

BRITISH TELECOMMUNICATIONS ENGINEERING



Included in this Issue

*Telecommunications in the
21st Century*

Service Management Systems

*Telecommunications Quality of
Service*



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 13 ■ PART 1 ■ APRIL 1994

Guest Editorial: The ICE Age Cometh Peter Cochrane and Bill Whyte	1
---	---

Telecommunications in the 21st Century

Customers Must be the Industry's Driving Force Sir Iain Vallance	3
The Bit Police—A View from the USA Nicholas P. Negroponte	8
Copper Mind-Sets Peter Cochrane	10

Service Management Systems: Supporting the Customer Interface Stephen Wilson	20
COSMOSS—An Innovative Approach to System Development Henri Van der Stighelen	25
Event and Test Management System Tim Young and Bob Holland	36
Electronic Trading: BT's Commitment to Changing the Way We All Do Business David J. Brunnen and Eric Wyllie	47
Telecommunications Quality of Service: Principles and Management Tony Mullee and Antony Oodan	53
A Whole-Life Approach to Management Decision Making Ian Napier and John Davies	62
Strategic Provision of Radio Nodes in the BT Access Network Jim Steggles	67
Internodal Signalling Peter G. Clarke	71

Regular Features

Field Focus	
Directly Buried Pre-Cabled Sub-Duct	86
Mobile Repeater Station for Restoring the Network	86
Power Drawings Database Within Southern Home Counties	87
BT News	88
Industry News	92
Book Reviews	94
Notes and Comments	94
IBTE Contacts	96

Board of Editors

Chairman John Prior, M.Sc.

Managing Editor Paul Nichols, B.Sc. **Secretary** Stephen Wilson, B.Sc. **Treasurer** Chris Fennell, B.A., F.C.M.A.

Alan Bealby, M.Sc., F.I.E.E.; Richard Beet, M.A.I.E.(DIP.), F.R.A.S., F.B.I.S.; Meredith Eddy; Martin Evans, B.Sc., Ph.D.;

Warwick Flury; Bernard Mallinder; Peter Martin; Ian Morgan, B.Sc., C.Eng., M.I.E.E.; Rodger Stevenson, I.Eng., M.I.E.E.I.E.;

Murray Trantor, B.Sc., C.Eng., F.I.E.E., D.M.S., M.B.I.M.; Keith Ward, C.Eng., F.I.E.E., M.B.I.M.; Syd Webster, I.Eng., M.I.E.E.I.E., M.B.I.M.;


Brian Wright, B.Eng.

Peter Cochrane and Bill Whyte

The ICE Age Cometh



Peter Cochrane



Bill Whyte

Information, Communication and Entertainment (ICE) stand poised to become the primary focus of society before the new millennium. This new ICE Age will not be cold and static, but hot and dynamic! Our ability to generate information and transport it around the planet on superhighways of optical fibre, to fixed terminals and mobile individuals through cellular and satellite radio systems, is about to change the way in which we communicate, work and live. There is not a single aspect of our lives that will not be touched by the communication, computing, and information revolution that is already upon us. The changes before us are comparable to the impact of the Industrial Revolution, physical transport and even the invention of the wheel. Unlike all other human activity, our ICE capabilities are doubling every year, but with a reducing use of energy and raw materials. Similarly, computer technology is providing a processing, storage and display ability for the office and home surpassing the mainframe computers of only 15 years ago. The new ICE Age will see a vast range of new facilities in the home, office, place of work, and in the pocket (or on the wrist) that will include access to vast data banks of information worldwide. The actual cost of moving information will be insignificant (and in many respects, it already is) as will be storage and presentation. We will be able to call for the latest publications and information on any topic. We will be able to drill down to a level of detail appropriate for our needs with interfaces that have been designed for human rather than technological convenience.

Conceptually, telecommunications is concerned with shrinking distance and compensating for the fact that 'remoteness' (from head office, shops, cinemas, family and friends) is a

barrier to decision-making or emotional contact. The fixed and mobile telephone network helps in part, but voice, text and data only offer a restricted range and facilities. New technologies offer the possibility of being virtually present at meetings while actually being at the other side of the globe. We could even wander through a shopping mall containing all the shops in the world without leaving home, and while in our favourite armchair, being given a health check by the networked autodoc. The hassle of travel will be minimised because we will concentrate on travel for pleasure or, where our journeys are absolutely necessary, the intelligent roads of the future will guide our vehicles around hold-ups to a pre-booked car park to which we are automatically navigated. Nor will we pay more for the privilege. Basically the same technologies that will enable these services will also help to reduce the costs of the suppliers. Vehicle navigation and control will allow 'just-in-time' delivery, as will home shopping. Routine customer queries will be handled by voice and image recognition computers while intelligent networks will provide easy hand over to live help desk operators for critical problems needing human attention. Communications-supported cooperative working will allow globally dispersed business teams to work as effectively as if they were in the same room.

A key difficulty with information technology today is that the interfaces effectively freeze out over 80% of the population by virtue of their unfriendliness. The humanisation of information, giving it characteristics that we can recognise, presenting it in a form that anyone from 3 to 90 years old can easily relate to is thus a major objective. The forms of these interfaces are already with us and based upon pen, hand, voice and

these developments are not mere incrementalism; they will require major changes to the BT business

screen. In their most advanced state they involve virtual reality and telepresence, where our senses are extended over distance to be embedded in data and information worlds, so we can view and interact with information in a new way. The use of multimedia and virtual reality in games is now established and becoming available at a rapidly reducing price. The combination of real and virtual worlds perhaps presents the most tantalising prospect.

Ultimately, information itself presents us with a major challenge. In Europe alone there are 6 million photographs of church windows, the Library of Congress in Washington requires 35 kilometres of book shelf each year to accommodate new volumes. The amount of information generated by mankind is doubling every 3 or 4 years. How are we going to find those articles, papers and snippets of information that are of interest to us? We already have difficulty in selecting a suitable TV programme when confronted with 100 cable channels, but with video-on-demand technology, the choice is likely to exceed any one of 10 000 videos at the push of a button! For us to cope with all of this will require some augmentation of the human brain—this has to be based on artificial intelligence. Software agents that roam the information fields and bring to the fore the most likely things of interest to us will be a vital ingredient in this technology revolution. However, they will need the ability to learn of our interests to tempt us into new fields as they develop.

The challenge we face is to accept that these developments are not mere incrementalism; they will require major changes to the BT business. Are we a bit carrier or an information processor? Are we a wholesale company and a separate service provider? How should we conduct and develop our business? Even the detailed structure of the network will have to be transformed; for example, one

cautious prediction suggests that ICE will result in network traffic growing by over 100 times that of today.

The major challenge with the future is the present: how easy it would be to get there if we didn't have to start from here. Even if we invent the cost-effective battery for the car of the future, how do we convert our petrol filling stations to battery charging and our mechanics to electricians and electronic technicians? How do we convince the doubters, the too-busies, the scared, that, this time, the futurologists have got it right? We must be able to show them that between the adequate Now and rosy Future is a reasonably trouble-free path and a route map that will take them on a cost-effective journey from here to there with familiar and friendly objects on the way. In truth, there will be some disagreement about the best route and there will even be some dead ends, but predict it and follow it we must.

Forecasting the future is extremely difficult and risky. So far mankind's efforts at technology forecasting have generally been pessimistic. This was particularly true of the telephone that was never conceived to be anything other than an instrument for entertainment! Trying to predict the development of the ICE Age and future information world is likely to be even more radically wrong—but predict we must!

Starting with this issue of the *Journal* we will take you on this journey into the future thanks to the contributed articles of numerous authors who have been invited to volunteer their predictions. These will cover the innovations that are on the horizon and also lead you well into the next century, covering the technologies, networks and applications that will determine how the inhabitants of the ICE Age will work, learn, keep healthy and enjoy themselves attended by robots, information engines and travel substitutes.

In collating articles for this, and forthcoming issues, we have enjoyed the full support and enthusiasm of the individual authors, and to them we extend our thanks for participating. We have also found the experience both enlightening and fun—we hope you, the readers, also enjoy a similar experience.

Peter Cochrane

**Head of Core Technologies Research,
BT Development and Procurement**

Bill Whyte

**Technical Division Manager Customer
Applications,
BT Development and Procurement**

Sir Iain Vallance

Customers Must be the Industry's Driving Force

Regulators and legislators indeed have tough issues to wrestle with in this transition to a highly competitive marketplace.

Introduction

Perhaps the most profound emblem the Modified Final Judgement (MFJ) has stamped upon industry is one simple word: competition. Before the MFJ changed our world, the preferred term in telecommunications was 'user' rather than 'customer', revealing overtones of commodity and a general disregard for market forces. That was understandable, given the industry's long heritage of regulated monopoly.

But ten years ago, the MFJ fired a shot, heard around the world, to change all that. Today, BT and other transnational carriers still feel its considerable effects in the markets as distant as New Zealand and as close to home as the UK. Indeed, the MFJ was far more than a watershed US domestic policy shift; it heralded a worldwide change. Moving the customer to the forefront, the MFJ set into motion trends that are gaining momentum even now.

Clear Trends from the MFJ

Most of those trends are good; the fruits of competition come to mind: for example, better quality, lower costs, and the astonishing explosion in customer services. In fact, I should confess, up front, my belief that the liberalisation of regulations surrounding telecommunications is clearly the best thing to happen for our customers for a long time.

This article was first published by the *IEEE Communications Magazine*, Vol. 31/12 in Dec. 1993. It is reproduced here with their kind permission and cooperation

At least one trend, however, is alarming. That is the trend toward investing regulators with greater power with no correspondent accountability. To address these two powerful and somewhat contradictory trends, and because I believe they both owe their current trajectories, at least in part, to the MFJ, perhaps it's fitting to begin with a look at how our world has changed since the early days of deregulation.

It is easy to forget just how different the telecommunications scene was ten years ago, just before BT was privatised. In 1984, my company had a few thousand kilometres of optical fibre in the ground; today, we have far more than a million kilometres, and by the close of the decade, we should have more than two million kilometres in the ground.

Back in 1984, no BT customer had the benefit of any digital technology. Today, 38 000 BT customers a week transfer to digital local lines, and the whole BT trunk network is digital—the first in the world, incidentally. At least as profound as the advent of digital and optical-fibre technology, however, is how our industry's deepest nature has fundamentally changed.

The structure of the industry in the early-1980s was virtually the same throughout the developed world. Each country had its national monopoly, usually state-owned, based on the notion that telecommunications formed a natural monopoly, and that the customer was better served by economies of scale than by competition. Those national operators often had a cosy relationship with national suppliers of equipment, a distorting

Barriers to competition are relentlessly breaking down everywhere in the developed world, and in much of the less-developed world as well

arrangement compounded by the system of national standards, which restricted the use of equipment largely to the country where it was developed.

To complete the picture, industry economics were grossly contorted by the social engineering of telephone pricing, with local service subsidised by a level of international and long-distance call charges bordering on the extortionate.

Seeds of Change

The seeds of change were, of course, sown by the introduction of competition, first in the USA in the late-1970s, leading to the MFJ, and then in Japan.

Competition came gingerly at first, because governments were inclined towards the natural-monopoly theory. In the USA, for example, local competition was long thought to be unachievable. However, after some initial confusion, customers quickly recognised the significant benefits (price cuts, improved service, better quality, and a wider range of products), so the overall benefits of competition were soon no longer in much serious doubt.

Today, customers are inclined to relinquish their rich new cornucopia of telecommunications benefits. They are, in fact, pressing for more.

Barriers to competition are relentlessly breaking down everywhere in the developed world, and in much of the less-developed world as well. The rate of change varies from market to market, depending on how well local regulators have learned (or not learned) from the MFJ's successes and shortcomings. To assess the progress of deregulation around the globe, let's take a quick world tour of the four major telecommunications markets, starting with the USA.

World of Change

United States of America

It's not surprising that in the first market to liberalise, competition in

the long-distance telecommunications market has proceeded furthest. This is obvious from the existence of three major aggressive inter-exchange and international service providers, offering some of the lowest prices and highest-quality services in the world.

Unfortunately, liberalisation in the USA has effectively stopped short of the local loop, and the Regional Bell Operating Companies (RBOCs) still enjoy a virtual monopoly over vast territories, some as large as the major European countries. Also, and this is particularly worrying for BT, liberalisation in the USA has not broken down the wall of protectionism. The xenophobic Section 310 of the Federal Communications Act still applies. With war brewing and Nazi propaganda in mind, it was passed in 1934 as a national security measure to protect the strategic value of the USA's radio and public broadcasting networks from control by foreigners (or aliens, as the Feds still like to think of us). It ought to have been consigned to the history books long ago.

But for reasons of political expedience, Section 310 is still with us, and raises a high if not insurmountable barrier to foreign companies owning more than 25 percent of any USA network in the States, and is a barrier to entry for companies such as BT, as witnessed in our brief relationship with McCaw Cellular. Aliens we remain.

Despite these significant shortcomings, the USA market is potentially one of the most attractive and open in the world, and it certainly presents a strong contrast with most of Europe.

Continental Europe

Here, where European Commission discussions have dominated the public debate for years, issues of competition and regulation are, of course, quite familiar. Apart from the UK, the European market for switched voice and telecommunications infrastructure is still largely closed, and controlled by state-owned Post, Telegraph and Telephone (PTT)

monopolies. The nonsense is that it is currently easier for a European company like BT to provide services across the USA, than across Europe.

Nevertheless, I think there are some good signs, signs that the state-owned PTTs understand that complete liberalisation of switched voice and of infrastructure is now inevitable; it is only a question of time. But the position so far is somewhat like St. Augustin's view of chastity, with the PTT's view of liberalisation an exact parallel: 'Oh Lord, let me be chaste, but not yet.' So we see a fair amount of prevarication. We see prevarication that is pushing the Commission into a position where it has to defend the basic tenets of the Treaty of Rome, namely, that there should be open and free competition across Europe, rather than the other way round, with the PTTs having to defend why it is that they should have a continuing special position of monopoly in the areas of infrastructure and switched voice.

Asia-Pacific

This third major world market is a high-growth area, attractive in terms of blue-sky possibilities, but a mixed bag in terms of market realities.

Australia and New Zealand are vigorously introducing competition and reducing protectionism; emblematic is BT's recent success in landing a \$300 million contract to build and run the government network in New South Wales. Japan, however, remains something of an enigma. To all appearances, legislation has opened the market, but those trying to do business there quickly realise that appearances can be deceptive. Foreign control of basic carriers is not permitted, and the regulation is an often impenetrable combination of ministerial guideline and carrier influence.

United Kingdom

My final stop on this world tour of liberalisation is at home. It is not generally recognised just how far the UK has come in the past ten years, but

I think it's fair to say we now have the most open market in the world.

Customer premises equipment has long been a free-for-all. Mercury and now Energis and other new entrants offer competing trunk and international networks. Increasingly, cable TV companies offer alternative local networks, while Ionica is poised to offer a wireless option to the home. Cellular networks—Cellnet and Vodafone—cover most of the country, as will the new personal communications network (PCN) services.

The MFJ contains specific restrictions on equipment manufacture. This is interesting from my perspective because the telecommunications supply industry in the UK was probably the first to witness competition when, in the early-1980s, telecommunications authorities around the world began to multi-source their buying. In response, national manufacturers have coalesced into multinational groupings: Northern Telecom with STC, Siemens with GPT, and so on.

It remains to be seen if network operators, like ourselves, will follow. There is a widely-held belief that telecommunications services provision will come to be dominated by four or five big players. There are embryonic signs this may be coming true, with BT's recent deal with MCI, and the emergence of groupings such as Unisource and EUNETCOM bringing together various European and perhaps American players to serve the multinational customers.

Such is the global regulatory landscape, as I currently see it, at various stages of evolution in various markets: clearly in a state of flux everywhere, but in general, relentlessly pushing toward the liberalisation first set into motion by the MFJ's influence in the world's largest telecommunications market.

The Role of Technology

It would be naïve to say all these changes are due to deregulation (or, more accurately, re-regulation) in the

wake of the MFJ. Much of the credit can go to the emergence of digitalisation and optical-fibre technology, which converged with the MFJ to change the fabric of our product and our market. Incidentally, the next step is imminent, as radio technology blossoms into the realm of widespread personal communications. Then the cord will finally, and literally, be cut—the umbilical cord tethering the customer to a proprietary network supplier.

But regulatory change has clearly hastened the development of these technologies and, in general, has pushed the global telecommunications industry toward a new reality of free-market competition. As we approach that ideal, customers, and transnational carriers in search of these customers, will continue to benefit.

Unfortunately, that mechanism puts the regulator in the driver's seat, precisely where the customer should be, and could be.

Unfortunately, at the core of it stands an inescapable fact that regulation and technology are, by nature, almost polar opposites. Technology is a dynamo that drives explosive change, and regulation—even bold policy like the MFJ—at its heart, attempts to control change to some extent. In effect, it usually inhibits its natural trajectory.

The Regulator's Troubling Rise to Power

Here lies the rub. The unavoidable appendage to regulation is the regulator. Of course, he has been with us since the beginning, but his role is changing and, to my mind, not for the better. He exists because of an irony embedded deep in our industry, a fundamental paradox of *laissez-faire*

capitalism: competitive success leads to dominance, itself the very enemy of competitiveness.

Our industry struggles to overcome that contradiction as we move from the realm of presumed natural monopolies into new relationships among competitive marketeers. And the preferred clutch mechanism to overcome that contradiction has always been regulation, engaged and disengaged judiciously, until the desired effect is achieved. Unfortunately, that clutch mechanism puts the regulator in the driver's seat, precisely where the customer should be, and could be.

The fact is, the best that could happen now, in face of the liberating technologies unfolding in our industry, would be for regulations, and the grip of the regulator, to relax even further. But the final word on that,

and indeed on this review of the past ten years, must inevitably rest with the single most important influence currently affecting our industry: not regulation *per se*, but the rise to power of the regulator.

In the UK, the model of regulation set up in the telecommunications industry—independent highly-discretionary, and without effective accountability—was generally copied in other privatised utilities. Today, looking back on our collective experience, I think it fair to say the resulting necessity of dealing with the regulation—at the national, European, and USA federal levels—has come to dominate the agenda and consume an undue amount of senior management time.

That original UK model was not a bad one, but after nearly ten years

As necessary as regulation is, it requires clearly-defined limits, or the result can be dangerous market distortions

facing up to the pressures of politics, media and indeed personality, it is beginning to creak. Like the MFJ, it could use further modification. But standing in the way is the nature of regulation itself, stemming the tide of change.

Three Forms of Regulation

What exactly are we dealing with? Stepping back a moment, I would distinguish three kinds of regulation: standardisation, antitrust and economic. Each of the three is proper and necessary, but the latter two, especially, can be carried too far. Why? Because once a regulator is appointed, the temptation and, indeed, the public and private pressure to over-regulate, becomes progressively irresistible.

And we must be careful about the dangers of over-regulation. As necessary as regulation is, it requires clearly-defined limits, especially in the realm of antitrust and economic regulation, or the result can be dangerous market distortion.

First, standardisation. The adoption of common standards is clearly vital to the industry. Technical developments, diversity of service and supply, and more sophisticated customer demand all make standardisation an increasing necessity.

To take a simple example, a traveller now needs several versions of a handset (not to mention a secretary to manage the bizarre regional variations in billing) to keep in touch over cellular network facilities while travelling across Europe. If we're serious about facilitating free movement in Europe and elsewhere, and if we're to really aspire to truly mobile technology, then all service and equipment suppliers must agree to common and binding standards to create that kind of roaming service.

And in fact, European operators and manufacturers have joined forces, under the aegis of the European standards body and the Commission, to draw up a digital cellular standard

for the whole of Europe, called *Global System for Mobile* (GSM). Eighteen countries have signed the memorandum of understanding on GSM, although there are some who believe it to be meddling over-regulation. But the essential point, and it is a major achievement, is that later on in the 1990s, this particular problem will have been cracked.

The second form of regulation, antitrust, is also basically concerned with giving customers the best service on the best terms. The trouble is, it is concerned with correcting failures of the market, outlawing agreements to distort competition, and abuses of market power.

All commercial activity in a free economy requires antitrust regulation; it will always be with us. But the historic monopolies in telecommunications have required it in a particularly intensive form—to protect customers from the high prices and poor services that a dominant player can inflict on them with impunity. Those are the legitimate specialised tasks of the FCC in the USA, or OFTEL in the UK.

The third form of regulation, economic regulation, is related but different: it is about allocation markets. And in my mind, overzealous regulation in this realm is clearly the most dangerous.

In the early days of telecommunications liberalisation in the UK, it was not enough to remove BT's monopoly, apply a rigorous antitrust regime, and then just stand back to let competition take off. Competitors needed a boost from Government. And that is precisely what Mercury, the American-owned cable companies, and the new mobile operators in the UK got. Mercury got favourable financial and other terms, and cable and mobile operators got exclusive access to certain technologies and services.

I recognised the need for those market-entry devices, though they didn't come free of cost, in many cases, to BT's own customers and shareholders. Creating competition

was in our interest; after all, BT's future lies with customers who choose to come to us. There is no future in customers who have no choice. But market-entry devices should be that and no more: a means to get competitors going, not to protect them from market forces once they're reasonably established.

In the UK, our current mix of antitrust and economic regulation is dangerously volatile. Perhaps readers will recognise a parallel in the USA. There is a danger that regulators who lack insight, or are impatient, or lack self confidence, or are swayed by short-term political considerations or personal ambition, may be tempted to confuse antitrust with economic regulation. They may begin to lose sight of the critical distinction between the two, and press on with promoting competitors rather than protecting competition. That would be the re-imposition of national state control, and it would be a disaster for any diverse and developing market.

I have a grave concern that we may be moving in this direction in the UK. With every change in the rules, we seem to see more regulatory intervention and more devices favouring certain competitors at the expense of others, and I confess I have BT's well-being in mind. I have no argument with the regulators that the rules need to change; we cannot manage a dynamic industry with regulation set in stone. However, we must be very clear about the criteria for change and its direction. Right now, we have no clarity of vision on either.

Five Regulatory Dangers

There are at least five causes of alarm. First, the human factor; regulators like to regulate. This is how they make their mark, and everyone wants to make a mark. This point is particularly troubling, given that the regulators' lack of accountability lets them intervene without giving a full account or explanation for their decisions.

Which brings us to a second concern: regulation in the UK and elsewhere has historically been a covert process, with no proper debate.

Third, I believe that the regulator has singularly failed to offer a clear vision of the marketplace. Often in the UK, regulators have shown little grasp of the telecommunications environment they believe will come about because of their regulations. Somewhere we have lost sight of the fact that the regulator is a surrogate for competition, not a manager of the industry.

Fourth, with this lack of regulatory vision and consistency goes a lack of clear direction for the industry.

And finally, regulation has not recognised its own limitations. It is one thing to impose a cap on prices, quite another to specify the number of fibre-kilometres to lay or the structure of a particular tariff. Regulation tends to involve itself in the running of the business, an area in which it has no expertise or focus.

However, I am encouraged by the current debate in the UK about regulatory governance, about who should regulate the regulators, about limits on their powers, and about the principles that should apply. In short, there's a growing feeling that regulators' powers require some safeguards, and I agree.

Early in 1993, the European Policy Forum published a report, 'The Future of Industry Regulations in the UK', which argued that privatised utilities are suffering from rampant over-regulation. Cento Veljanovski, the author of the report, argues that there is '... an urgent need for reform of the regulatory system to create a stable environment within which the management of the utilities can manage, the investors can be protected from high prices and poor service.'

I believe that an increasingly interventionist regulatory approach is doomed to failure, if only because the complexity and sophistication of the market will be such as to put it beyond the wit of the regulators to regulate.

This increasing complexity—thanks to the inevitable convergence through technology of various industries (telecommunications, computing, entertainment, education, etc.), some are regulated and some not—will, I believe, present the regulator with an insoluble problem. Modelling that complex interplay of industry will no longer be a practical possibility. And if you cannot model, you cannot regulate. You can intervene only in the knowledge that you have no means of knowing the consequences. Rational, responsible intervention in such a scenario is, in fact, impossible.

Everything points to the eventual, and some might say imminent, breakdown of regulation, even at the national level. Add to this the trend towards globalisation, and you compound the difficulty many times over. And finally, perhaps this is the time to remind ourselves that this is not an academic debate. In an age where we are essentially hardwiring the planet for on-line real-time thought exchange, we are in effect dealing with society's nervous system. Such neurosurgery is delicate.

The MFJ was the first attempt to manipulate the development of that nervous system, and like all groundbreaking surgery, it was in some ways crude. Nonetheless, the MFJ was revolutionary and necessary, and in the end it unleashed global market forces that show no sign of abating. It is our duty to channel those forces, minimally but properly, so that the rightful driver of the market—our customer, not the regulator—remains firmly in the driver's seat.

Biography



Sir Iain Vallance
Chairman, BT

Iain Vallance was educated at the Edinburgh Academy, Dulwich College, Glasgow Academy and Brasenose College, Oxford, graduating in 1965 with a BA in English. In 1972, he took an MSc at the London Business School, and was elected a Fellow of the School in 1989. His career started in the British Post Office in 1966, followed by a variety of appointments, initially in finance and procurement. He joined the Board of British Telecom at its inception in 1981. In October 1985, he was appointed Chief of Operations, becoming Chief Executive in October 1986, before taking up his present position as Chairman in October 1987. He is a member of the Confederation of British Industry President's Committee, the President's Committee and Advisory Council of Business in the Community, and the Governing body of the London Business School; a trustee of the Police Foundation; and chairman of The Princess Royal Trust for Carers. He is also on the international Advisory Board of the British-American Chamber of Commerce. He is a non-executive director of The Royal Bank of Scotland Group plc. He received honorary degrees of Doctor of Science from the University of Ulster, and Doctor of Technology from Loughborough University of Technology, in July 1992. His interests include hill walking, playing the piano, and listening to music.

Nicholas P. Negroponte

The Bit Police—A View from the USA

The problem with digital transmission is that it is just bits, and one bit is very much like another. How then do the politicians and the regulators know that their vision of the public good is not being subverted by the ubiquitous bit. In this article, Nicholas Negroponte examines the changing nature of the communications media and argues that, in the end, it will be the customer who decides which bit he wants and which bit he does not want.

A Licence to Radiate Bits

Imagine that you own a television station and the Federal Communication Commission (FCC) has just given you a licence to broadcast 20 million bits per second. Namely, you have just been given permission to become a local epicentre in the bit-radiation business. What would you do with your licence?

Let's face it, the very last thing you would do is broadcast high-definition TV (HDTV), if only because the programmes are scarce and the receivers few. So, this is what you would do. Firstly, with a little cunning, you would realise that you could broadcast four channels of digital broadcast-quality standard NTSC television, thereby increasing your audience share and advertising revenue. Upon further reflection, you might decide to transmit three TV channels, two digital radio signals, a news data channel, and a paging service. It continues. At night, when few people are watching TV, you might use most of your licence to spew bits into the ether for the delivery of personalised newspapers to be printed in people's homes. Or, on Saturday, you might decide that resolution counts (say, for a local football game) and devote 15 million of your 20 million bits to high-definition transmission. Specifically, you will be your own FCC for those 20 million bits, allocating and

multiplexing them as you see fit. That is, if the Bit Police don't stop you.

To be perfectly clear, this is not what the FCC had in mind when it allocated the HDTV spectrum among existing broadcasters, and the body politic will scream blue murder when it realises that TV stations have just increased their current television capacity by 400 percent for free!

The Message is the Medium

What will happen over the next five years is so phenomenal that it is difficult to comprehend. On the one hand, it is easy to state: 'We are in the process of leaving an analogue world and entering a digital one'. While this is simple to say, the consequences are not easily apprehended. For example, we once thought that audio, video, and data were different, but now we are told they are the same: they are all bits.

Sometimes, bits will be assigned to a medium at the transmitter (step one): this is the usual idea. But at other times, bits won't be in any medium, as such, but will be a model and string of bits which are transcoded at the receiving end into either audio or video or print (step two).

In the world before step one, we allocated specific spectrum to TV, radio, and various other applications, in part because (we thought) the highways of the sky required well marked and impenetrable median strips. One could easily determine, in advance, what would be found and where in the spectrum: voice, data,

The first version of this theme paper was published by *Wired Magazine* (San Francisco Vol. 2) in March 1993, and is reproduced here with their kind permission and cooperation

While there will be a fuss and, ultimately, regulatory legislation, in the end, all bits will be deregulated

video, and so forth. But now that ability has disappeared—consider the telephone company which has no idea if you are passing voice, electronic mail or fax over its wires. Will the FCC put a wireless tap on some part of the spectrum and say: 'Halt, this bit it is audio and should not be in this lane in the sky!'

Step two gets even more complicated. Imagine a relatively-simple example, such as weather information. Rather than transmit text, graphics and voice, consider the possibility of a digital model of the weather. Your receiver, no less a computer than a TV receiver, will digest and process that model (purists are welcome to call the model the medium). Your receiver will convert the model to sound or image, hard copy or soft copy, in greater or less detail at your discretion (or the receiver's inference of it). This is to say that the output medium is determined, after the fact, by the user!

What You Want: Bits of it Are and Bits of it Are Not

Probably most readers will assume that the existence of a Bit Police will result in the censoring of content. Not so! The consumer will do that by telling the receiver what bits to select. The Bit Police force will want to know and to control the medium: this, to my mind, makes no sense whatsoever. The problem, which is strictly political, is that the FCC has told existing TV broadcasters that they can have an additional 6 MHz for HDTV broadcasts, in parallel with and beyond their current allocation of spectrum (for the next 15 years, at which time they must give back their current spectrum). When the FCC made this allocation, it had no intention of creating a windfall. However, minority and special-interest groups will now complain because the bandwidth rich will get richer.

While there will be a fuss and, ultimately, regulatory legislation, in the end, all bits will be deregulated. There is simply no way to limit the freedom of bit radiation any more

than the Romans could stop the spread of Christianity, even though a few brave and early data broadcasters will be eaten by the Washington lions in the process.

Wired in a Wireless and Multimedia World

Take another example: this Journal. Like all journals, it is in a purely digital form during its creation. The text, obviously, is in a computer-readable form, the images are scanned and the layout produced on a computer-aided design (CAD) system. The style of creation is the epitome of both a digital process and a digital lifestyle (my contribution was written from the seat of an aeroplane and sent to the Editor by electronic mail). Only when the final pages are output to film for printing does the digital representation vanish.

Let's pretend that, instead of providing this Journal in hard copy, we could transmit it in bits. The subscriber could transcode them into this form or into another. We would then write the Journal very differently, of course. Among other things, we would provide varying levels of detail, our cutting room floor would be empty, and the Journal (if we still use the word) would be conversational. In a few years, we won't be pretending, but, for the time being, the reader must be content to hold physical paper shipped through time and space.


The message is that all information providers will be in a common business—the bit radiation business—not radio, TV, journals, magazines or newspapers.

There are five paths into the home: telephony, terrestrial broadcast, cable, satellite, and package media (compact discs, cassettes, flash cards, etc.). Note that the Bit Police have a descending mandate of control in the order listed—packaged media has none.

Each channel into the home has its own characteristics, timing and bandwidth, but each will be subject to

exactly the same convergence of media. Believe it or not, sophisticated consumer electronics companies don't understand this and are planning digital video cassette recorders (VCRs) for the home, to be released shortly, which accept only analogue input! In fact, why are they called VCRs anyway? They are bit collectors which vary only in bandwidth.

Biography



Nicholas P. Negroponte
MIT Media Laboratory

Nicholas P. Negroponte is the founder and director of the Media Laboratory of the Massachusetts Institute of Technology, where he is also Professor of Media Technology. He has written and lectured extensively about the promise of computer technology, human-machine interfaces and the changing nature of communication. Nicholas studied architecture at MIT, receiving his master's degree in 1966. Within a year he had founded the Architecture Machine Group at the same Institution. He has been a visiting professor at the University of Michigan, the University of California at Berkeley and Yale University.

Peter Cochrane

Copper Mind-Sets

Every few decades technological progress across a number of apparently disconnected areas presents an opportunity for radical change. The great benefit of these epochs is that they change the nature of problems and allow new solutions. Such an opportunity is now on the horizon with the dual realisation of very-low-cost telecommunications and computing.

Introduction

Much of our technological thinking and expectation is conditioned by a history of established practices and trends. In telecommunications, it is radio and copper cables with constrained bandwidths and high distance-related costs; for software, the acceptance of massive complexity to realise relatively simple tasks; for the interface between man and machine, it is technology convenient to the machine. In each case, the solutions become steeped in political and corporate correctness through massive investment programmes that perpetuate established trajectories. However, recent advances in technology are creating new opportunities for radical change. The prospects go beyond our established view of computing and communication, but to grasp them requires new, and not old, mind-sets.

Many systems and networks are still being designed as if the constraints of the copper past were still with us: compress the signal; save bandwidth; minimise hold times; continue with the old protocols although they are now the new bottle-neck; use the established design rules despite the disappearance of the original technology! The apportionment of operational and economic risks also seems to be vested in the past and requires realignment. The wholesale reductions in the hardware content of networks and machines relative to the massive expansion (explosion!) of software is a prime example, as is the poor human interface suffered on many teleporting services when the need to reduce physical travel is so apparent and bandwidth is so abundant.

Today we sit at an interesting crossroads where the advances in fibre-optic technology has (or will have soon) removed the bandwidth

bottle-neck and distance-related costs in telecommunications. Integrated electronics has effectively made information processing and storage limitless and effectively free, while software has introduced the potential for unexplained system and network failures on a grand scale. What should we do, or be doing? There would appear to be only a small number of key actions required to take the next major step. First, we need to start thinking and engineering **systems** and not hardware, software, networks and applications as if they were in some way disconnected. Second, we need to embrace new tools and techniques that might help us understand the non-linear nature of the world we have/are constructing. Third, we have to humanise machine interfaces and information presentation if we are to stand any chance of understanding the nature of our systems. In short, we need to be able to visualise before we can hope to conceptualise and understand. Finally, we should not shy away from starting again with a clean sheet of paper when incrementalism has driven us into a 'technological box canyon'.

Perspective

The growth and well being of civilisation is critically dependent upon rapid and efficient telecommunication for its organisation, control and stability. From the runners and beacons used in ancient times, to the Chappfe telegraph, electric telegraph, radio, and today's telephone and data networks, the need has always been for faster and more effective telecommunications. This is probably more evident today than at any other time with the emerging need to combine communication, computing and control. Tomorrow the list is likely to include robotics, telepresence and

variants of virtual reality. It is interesting to reflect that none of this, not even a ubiquitous telephone network, would be possible with wholly copper transmission technology. It is doubtful if there is enough copper and associated technology, or money available to pay for everyone to have just a telephone. Certainly the bandwidth limitations of copper and radio rule out any significant excursion into computing, control, telepresence and other future services. Only the timely arrival of optical-fibre technology during the last 15 years gives us the means to attain the dream of telecommunications for all at a price we can all afford, with the added bonus of a near infinite bandwidth and service expansion in the future.

It is perhaps understandable that the development of copper transmission and switching technology over the past 150 years has conditioned our industry to the network topologies, configurations, services and performance of the kind we currently enjoy. The reality is that the status quo will not do! Optical fibre is the only means of realising the next paradigm change, which may well be greater than the leap from the beacon to the electric telegraph. All we have achieved so far with optics is more of the same with better performance at a lower price. What we have to do is rethink the problem of systems, networks, operations and services to encompass the advances made possible by this young technology.

The exponential growth of telecommunications continues even during the present global recession and there is no sign of any downturn. However, it is sad to reflect that 85% of all the telephones are with the 15% of the population who have the money! Just 15% of us on this planet (about 825M telephones total and over 700M of these in the developed countries) have good communications—the rest have not. It would be nice to think that in the new millennium we could start to correct this imbalance. Indeed, long-term global stability is likely to

become almost wholly dependant on us achieving such an apparent ideal. Even so, in terms of growth potential a modest increase in industrialisation in the second and third worlds would see a massive growth in global communication.

Ecologically, IT has a significant part to play in the reduction of waste and raw materials. Unfortunately, much of this technology has seen an extension and expansion of earlier paper-based practices. In reality, all the base technologies to reduce drastically the consumption of paper (rain forest) and physical travel (barrels of oil) are available.

The cost of a telephone call, fax, computer communication session/file transfer and videoconferencing call in terms of raw material usage is insignificant compared to international travel costs by at least six orders of magnitude. Even in domestic terms, there are considerable opportunities for significant national savings. For example, the UK has an approximate population of 55M, a work force of 25M, with 23M dwellings and 25M telephones. The total transport cost, in energy terms, for the population to travel to work has been estimated at 1.8×10^9 GJ/year, while the entire telecommunication infrastructure consumes an estimated 5×10^6 GJ/year. If only 1% of the working population changed to full-time working from home, then an estimated 10^7 GJ/year could be saved. Globally this figure is probably of the order 10^9 GJ/year and could probably be increased significantly with enhanced telecommunications services.

A Selection of Limitations and Irritants

Much of our IT has changed little since it was originally introduced and leaves a lot to be desired from a user point of view. Some notable examples are as follows:

- a telephone that still has more or less the same bandwidth as it did in 1876;

- not being able to recognise people on the telephone, and not being able to see them;
- not having global telecommunication and computing mobility;
- brittle software in robust hardware, systems, networks and machines;
- coding systems that distort the voice and grossly distort moving pictures;
- trying to communicate with visual images of people who are the wrong size and colour and are distorted when they move;
- sterile teleconferencing environments that lack basic meeting-room facilities;
- having to travel so much when it could be circumvented by IT;
- having to send floppy discs and tapes via surface mail (*Frisbee Net!*) because it takes hours to send the data over the restricted-bandwidth telephone line;
- time delays in communication, control, computing and associated interfaces;
- a PC with a 25 MHz clock (and even that is getting old!) and a PSTN/ISDN only able to deliver 64 kbit/s (or low multiples thereof) as a switched service;
- PC screen that often only displays 75% of one page in an 'N' page document;
- inappropriate technologies surviving, and/or having their lives extended for the wrong reasons;
- e-mail addresses with enough characters to define the position of every atom in the known universe;

optical fibre has hidden properties that we have not discovered yet, and therefore affords a potential for exploitation well into the 21st century

- fighting through cascaded windows to achieve elementary objectives;
- having to communicate using an increasing amount of paper;
- voice command on car phones, but not on office machinery;
- a growing paper and electronic filing system that still has to be categorised, sorted and culled on a regular basis;
- living in a world that is both speeding up and removing delays in the decision-making process, while not providing new/adequate tools to cope; and
- radical technological opportunities being lost (and/or suppressed) through political/regulatory pressures and shortcomings in industry.

etc, etc, etc

New Networks

The longer transmission distances afforded by optical-fibre systems over their copper predecessors is predicated a drastic reduction in the number of switching nodes and repeater stations. This will soon be augmented by the optical amplifier and network transparency leading to numerous improvements across a broad range of parameters including reduced component count, improved reliability, reduced power and raw material usage, increased capacity and utility. A further (and revolutionary) development will see the concentration of more traffic onto single rather than several fibres in parallel. The advent of wavelength-division multiplex (WDM) makes such a proposal the more reliable option as the time to repair one fibre is far shorter than for multiple fibres!

Today we only access less than 0.005% of the fibre bandwidth available, and we might be able to approach 1% with currently available

technologies. But further, everything we have seen optical fibre do so far may be totally eclipsed by moving from the linear to the non-linear regime. It has hidden properties that we have not yet discovered, and therefore affords a potential for exploitation well into the 21st century. We already see experiments with bit rates of 10 Gbit/s over 10⁶ km of amplifying fibre.

As the cost of transmission continues to reduce we have to question the balance between transmission, switching and software. Radical reductions in the number of network nodes and repeater spans, consolidated switches, reductions in network management and constraints imposed by software look to be good targets. We may also have to move away from thinking that the local call does not just span the London area but the whole of the UK, expanding in time to encompass Europe and gradually the whole planet. E-mail is the only practical realisation of this concept that is presently available—and it is a subscription service!

Switching and transmission

Today we only use time and space for switching in telecommunications. But at the present rate of expansion it is becoming clear that we will not be able to realise the density of electronics to meet future growth by using these two parameters alone. It will become imperative to move into a third parameter to get the equipment density and energy concentration into existing buildings. For example, 10 years ago a typical UK repeater station floor had to accommodate some 2.4 Gbit of speech circuit capacity. That same floor today has 6.8 Gbit and projections indicate that 40 Gbit will be required by the end of the millennium. This packing density cannot be achieved with conventional electronics alone. Another degree of freedom is required—wavelength is the obvious choice. Recent developments have seen the demonstration of suitable technology such as contactless 'D type' (leaky feeder) fibres embedded in printed circuit

back planes. When coupled to an erbium-doped fibre which amplifies the light signal(s) when pumped optically, a loss-less multi-tap facility is realised which can distribute 10 Gbit/s and higher rates almost endlessly.

An interesting concept now arises—the notion of the infinite back plane. It could be used to link, for example, Birmingham, Sheffield and Leeds through the use of optically amplifying fibre that offers total transparency. Such concepts naturally lead to the idea of replacing switches by an optical ether operating in much the same way as radio and satellite systems today. The difference is the near infinite bandwidth of the optical ether. Demonstrators have already shown that a central office with up to 2M lines could be replaced by an ether system, but suitable optical technology is probably still some 15 or so years away. Systems of this kind would see all of the software, control and functionality located at the periphery of networks.

Signal format

Customer/service demand and technology are each providing the motivation to change the nature of telecommunications. In the midst of this revolution it is interesting to contemplate a further feature of communication history—the cyclic alternation between digital and analogue formats. Before the development of natural language, the 'grunt or gesticulation' was no doubt the simplest indicator of man's agreement or displeasure. We have since progressed to use the tally stick, smoke signals, drums, music, the written word, Morse code, telephony, and pulse-code modulation (PCM) (Figure 1). In parallel, computation has also followed an alternating analogue-digital history. We might thus presuppose that this alternation will continue into the future, and may even be encouraged by the inherent qualities of optical-fibre communications and photonic computing. There is certainly evidence that increasing

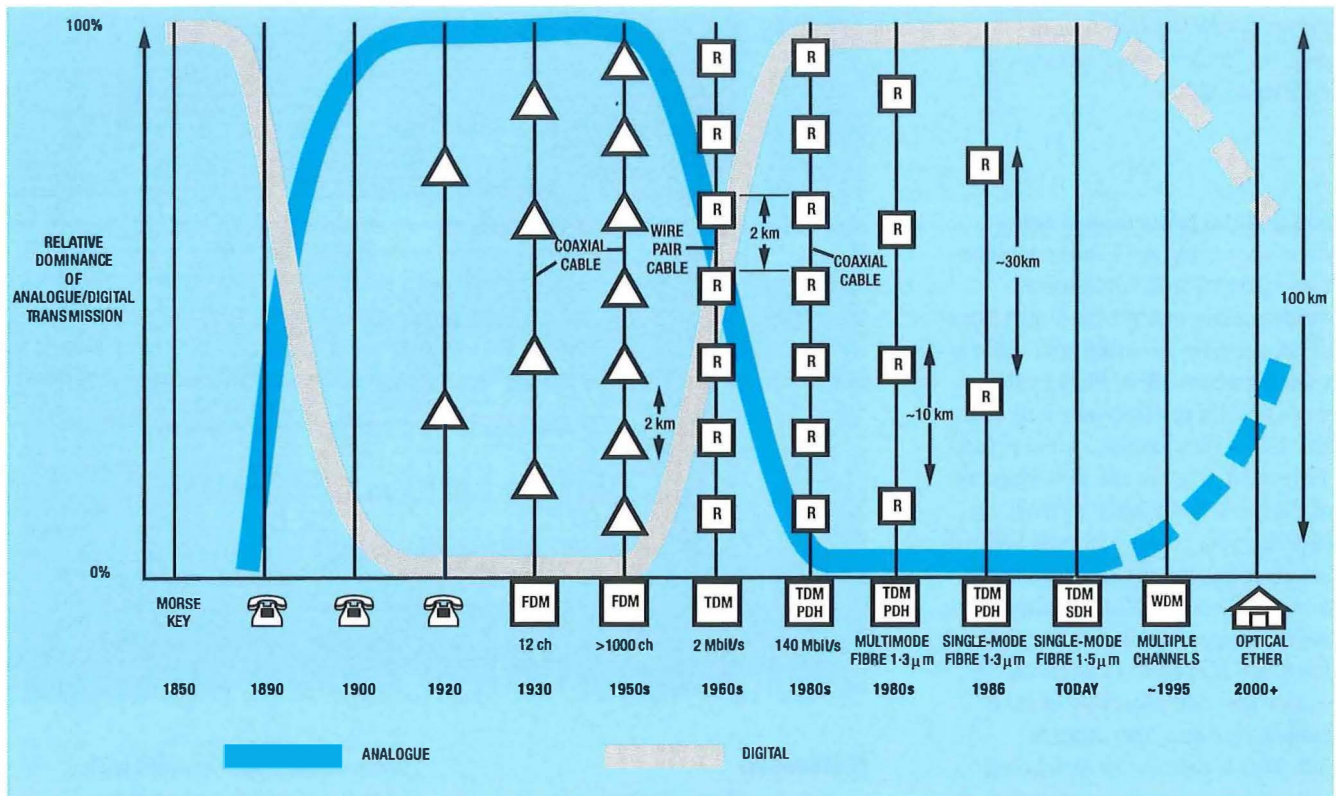


Figure 1—Cycling trend in analogue versus digital transmission

complexity and future requirements may necessitate the use of analogue techniques to solve the problems of achieving a useful level of artificial intelligence, associative computing and control systems. Optical amplifier technology looks ideally placed to meet this challenge and may even promote the move back to analogue signalling!

Mobility

While cellular radio technology is now well developed, and satellite mobile will no doubt follow, we suffer an apparent lack of spectrum, although we have yet to exploit microwave frequencies up at 180–300 Gbit/s and higher. But we also have the ability to take radio off air, up-convert it to the optical regime, feed it down fibre that is transparent through optical amplification, and get that signal to emerge in a distant cell undistorted. So, for multi-site working we could have access to the same facilities in all the buildings throughout an organisation, regardless of its geography/demography. We would effectively be sharing the same signal/work space, thereby giving all participants the illusion of being at the same location.

Another exciting prospect is optical wireless. The performance of

free-space optics in the home and office is very similar to microwaves, but with the advantage of a vastly greater bandwidth. Systems in operation in research laboratories are already providing pica-cellular illumination of the desk, personal computer, staff member, and offer the potential of active badges/communicators and inter-desk/computer links. Applications in the local loop might also be anticipated as an alternative to fibre, copper and radio.

For global mobility, the major challenges are likely to remain the organisational and control software necessary to track and bill customers—this represents the major system design challenge! Other issues include the ability to deflect/hold/store calls that are traversing time zones at unsocial hours.

Satellite and radio

Point-to-point satellite technology has now been eclipsed by optical fibre, carrying over 65% of all international communication. However, satellite now has a new role in broadcast, getting into difficult locations, mobility and service restoration. The path delay for geostationary satellites is problematic for real-time communication, but there are alternatives. There are proposals for low Earth

orbit (~1000 km) satellites, about the size of a football and launched by super guns, to form a cellular system in the sky using 70 or so low-cost units. Although they will periodically fall out of the sky, their replacement would be at low cost.

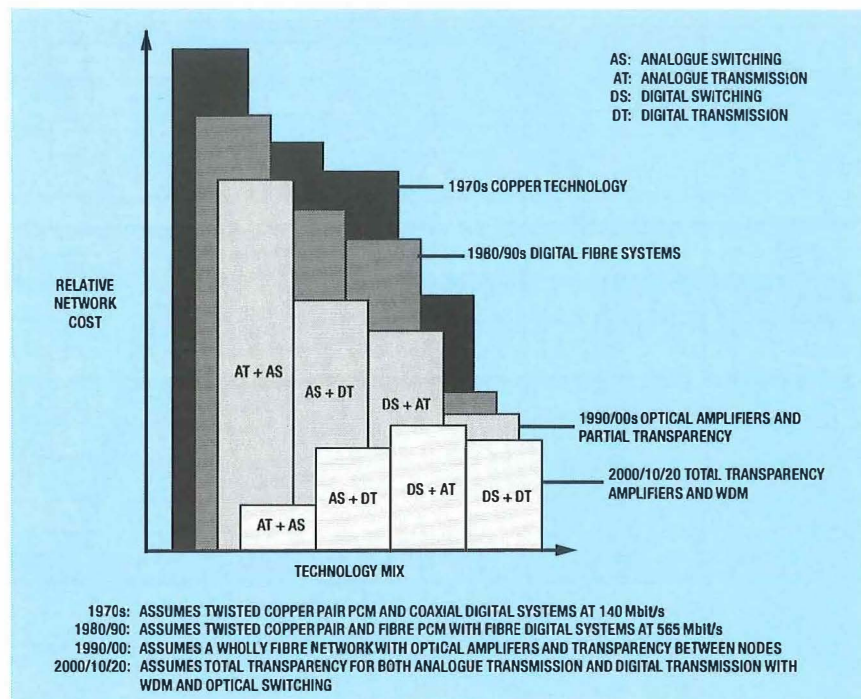
Other exciting developments include direct satellite-to-satellite links using free-space line-of-sight optics, on-board signal processing and switching, more compact and efficient coding schemes, individually positioned multiple microbeams, the use of much higher microwave frequencies than normal ground-to-satellite links, and even optical links from satellite to ground. All of this can be expected to extend the use of satellite technology by an order of magnitude (or so) beyond that of today. Fundamentally, satellite systems look set to migrate to mobile, and difficult/rapid-access applications. The technology is however a refinement of what we already have, and, other than using free-space optical links, it is hard to see any further major innovations.

The economics of analogue and digital

The development of PCM and transistor-based electronics jointly

Figure 2—Network costs with projected technology progression relative to 1970

enabled the realisation of time-division multiplex (TDM) transmission systems that could most economically combat the limitations of copper-pair crosstalk and attenuation. By about 1960, PCM systems were becoming established, and in the 1970s they became widespread. This in turn led to the development of digital transmission systems at higher rates which had well defined and internationally agreed standards by the mid-1970s. At this time, repeater spacings had fallen to 2 km for 2.048 Mbit/s on twisted-pair cables and 140 Mbit/s on coaxial cables. Systems operating at 565 Mbit/s rates were also being investigated but required repeaters at 1 km intervals. During this same period a number of administrations had completed economic studies that established the combination of digital transmission and switching as the most economic. This is reflected in Figure 2 which shows the options considered and their relative costs. Programme decisions made at that time are only now seeing their full realisation with the near full digitalisation of national and international networks in the developed world. A point to be recognised here is that historically the fixed assets of a telephone company (telco) are generally measured in £Bn, and any major network changes take 5–10 years to complete. Before fibre technology, the ratio for long lines would be of the order of 50% transmission and 50% switching. With the widespread introduction of optical fibre into modern networks, the transmission asset base may now be as little as 10% with some 70% of all resources expected to reside in the access network. This has significantly changed the economic balance away from the original figures. When repeaters are spaced at 50–100 km intervals and direct photonic amplification is introduced, an all-analogue system for national and international wideband services then looks far more attractive.



Software

The biggest single problem with software science and engineering is the almost total lack of any science and engineering. While we might use the excuse that it is a very young field, this does not give the full story. Indeed, should we even expect to achieve as mature an understanding as that enjoyed in mathematics and the physical sciences? In truth, software is a further level of abstraction beyond mathematics and may well defy real understanding for some time. A suitable analogy might be the mechanical engineer trying to design

some considerable way off, with structured programming, formal methods, requirements capture analysis (RCA) and object orientation being very unlikely to save the day. (Anyone who disagrees should contemplate the production of a bicycle using RCA—what chance is there that it will be right first time? Very little! In contrast, a series of rapid prototypes will certainly succeed.) All that can be expected is a tidying up of the software production processes we already have! Very soon we will be back to where we started, with an expansion in the size of software we are able to produce, and even a modest improvement in our

software may well defy real understanding for some time

a bridge on the basis of molecular descriptors, devoid of useful concepts such as 'Young's Modulus', or the electronics engineer designing circuits on the basis of quantum mechanics rather than 's' parameters. We are currently standing too close to the software problem—and we need to see through several more levels of abstraction before we are likely to form a full understanding and appreciation (Figure 3).

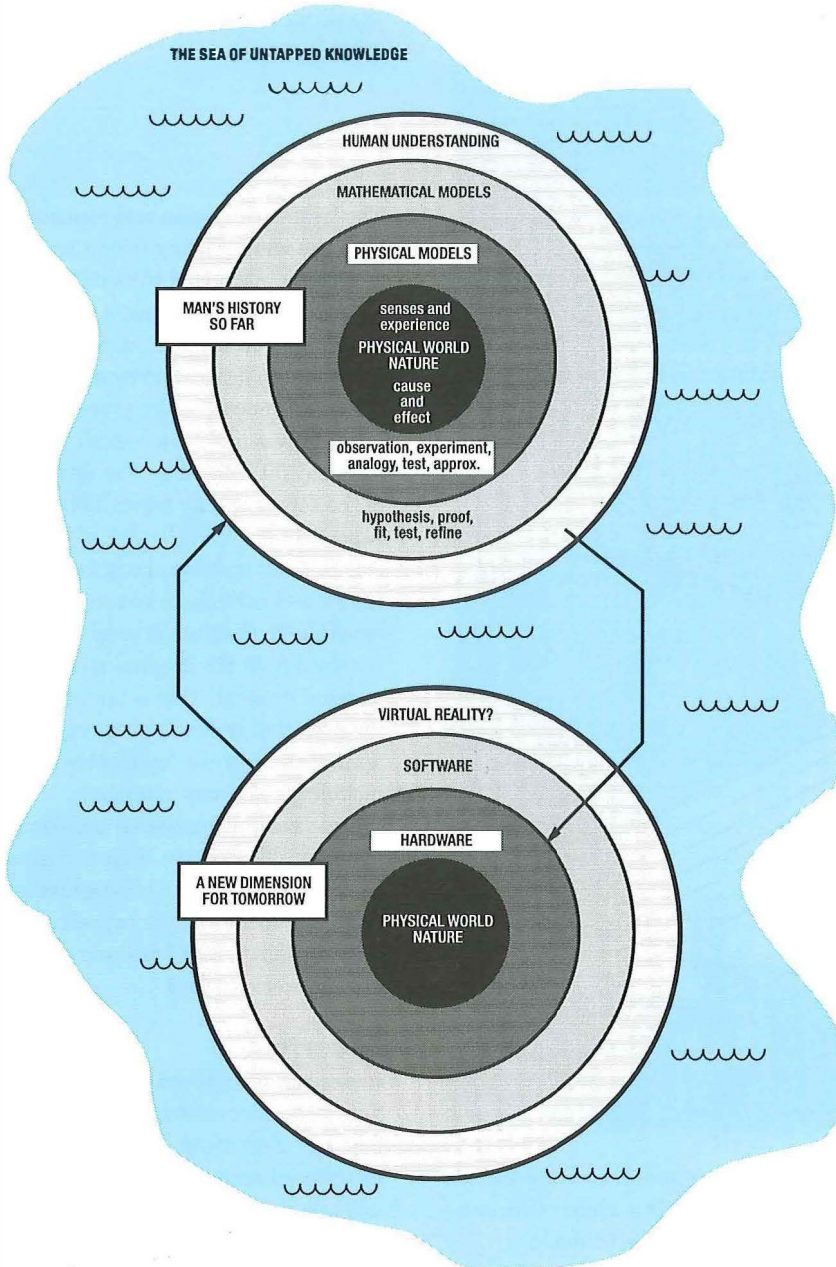
Unfortunately the above state of play in the development of software is little comfort to the designers, coders and users who are facing very real problems today. The solution looks

efficiency, but the level of confusion, uncertainty and engineering and operational risk will be at least the same, and may even be worse. The software factory approach might pay dividends, with tried and tested modules available in a reusable form. However, the chances are that history will repeat itself and the demand will be for custom software, for the same reasons we have had to produce custom and semi-custom integrated circuits.

Size and complexity

By any comparison of man's efforts, software is becoming a cause for

Figure 3—Software as an abstraction



things pose a considerable risk, which appears to be growing exponentially as we look to the future. We have to find new ways of solving this increasing burden of risk as the present trajectory looks unsustainable in the long term. Quite perversely the unreliability of hardware is coming down rapidly. It is becoming more and more reliable while software is becoming more and more unreliable, so much so that we are realising sub-optimal system and network solutions (Figure 4). From any engineering perspective, this growing imbalance needs to be addressed. If it is not, we can expect to suffer an increasing number of ever more dramatic (even global) failures.

Ants and ANTS—a viable alternative?

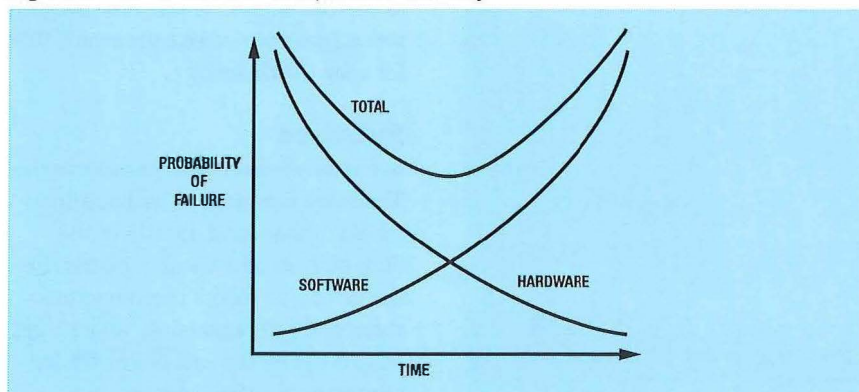
It seems remarkable that we should continue along a trajectory of developing ever more complex software to do increasingly simple things. This is especially so when we are surrounded by simple organisms (moulds and insects!) that have the ability to perform complex cooperative tasks on the basis of very little (or no!) software. An ant colony is one example where very simple rule-sets and a neural computer of only 200 (English garden ant) to 2000 (Patagonian ant) neurons is capable of incredibly complex social behaviour.

Numerous studies of ant and insect life forms have been reported with varying degrees of success. However, we have migrated to the ANT (autonomous network telepher) as a likely contender for the future control of telecommunication networks. Initial results from simulation studies have shown considerable advantages over the conventional software currently used. In one recent study of network restoration, only 400 lines of ANT code replaced the 1 000 000+ lines presently used in the real network. Software on this scale is within the grasp of the designer's full understanding, and takes only a few days to write and test by a one-man team! Some initial results are given in Figure 5.

significant concern in terms of its sheer scale. For example; the complete works of William Shakespeare take up about 450 m of paper; the line code for a small telephone switch is about 1 km and a central office is in

the 4-6 km range; network control centres are in the 6-10 km region; the Encyclopaedia Britannica is about 4.3 km. A full stop in the wrong place and the spacecraft misses the planet! In the software domain, very minor

Figure 4—Hardware versus software reliability






Figure 5—Network restoration with an ANT algorithm


The trial is run again and repeated with the winner being frozen and the rest being instructed to mutate. If after a few 100–1000 trials a competitor remains at the bottom of the league table it is culled by a reaper function. Repeating the process eventually realises an overall winner, after say 50 000 trials. The final form of the software may typically be expected to be half the size of the original, but operate about four times faster, and is likely to be incomprehensible (to the human) and far removed from the original seed from which it evolved. This is an interesting area that is developing rapidly for a number of diverse application areas including telecommunications, control, computing, signal processing and adaptive systems. Figure 6 gives a forecast for the rate of progress in this field, including the capacity development rates of desktop computers and workstations.

Evolution

Another radical approach that mimics nature is that of competitive evolution. In this case, a number of seed programs (say 10) are written to

address a given problem. They are let loose and timed. The winner (the one that takes the shortest time to complete the function) is then frozen. The rest are instructed to mutate!

Figure 6—Evolutionary hardware and software



Security

Security is an ever-increasing problem, and if we are going to realise distributed computing, telepresence and tele-action systems via telecommunication networks, certification at a distance will then become vitally important. To achieve this we can combine the scanning of the human face, eyes, lips, hands, thumbs, to provide a high degree of security. This can be augmented with voice recognition, handwriting, and ultimately, if we are still not convinced, we could ask for some obscure piece of information known only to the user. Using a small combination of these techniques, we can get error rates better than 10^{-10} for a few 1000\$ today.

Robotics

There are now around half a million robots in the world, mostly in the form of arms on assembly plants. In the future, as robots become autonomous and self organising, we may get caught out by the rate of growth for telecommunication capacity in a

one of our key problems is the distillation of core information from vast amounts of computer, network and system-generated data

similar manner to the recent growth in cellular radio. That is, we may not have radio stations in the correct locations ready for robots to do the things they have to do. They will need to know where they are, we will need

to know where they are, and when they get in trouble because of their limited ability we will need a capability for human interdiction at a distance. The forecasts for autonomous robot production from Japan

suggest the real possibility of machine-to-machine telecommunication/traffic densities rivalling those of the human race by 2020!


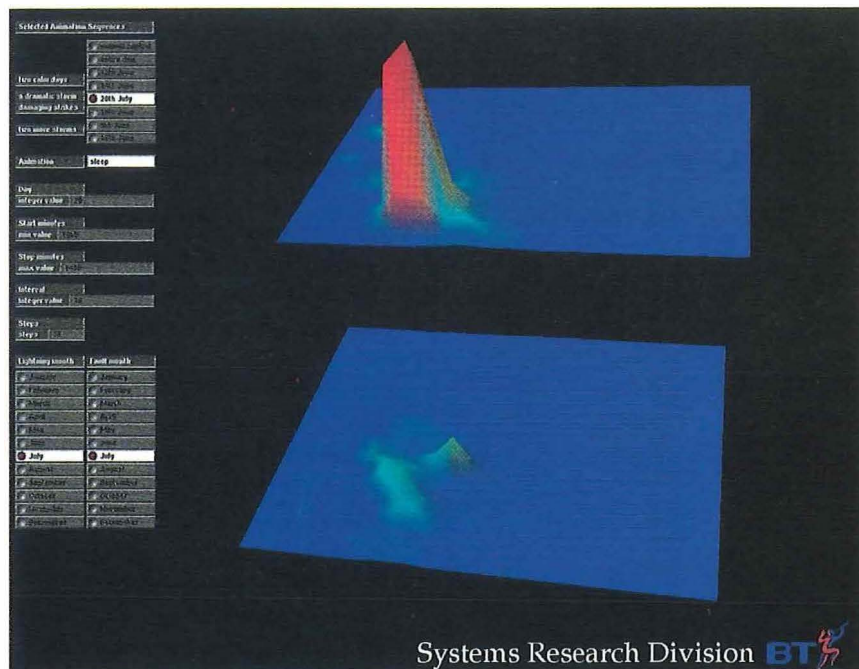
Visualisation

Today we are living in an increasingly complex, and non-linear, world. Unless we can develop new techniques to help us visualise future systems, so that we can conceptualise and understand/formulate solutions, our technological development may become stunted. At present, one of our key problems is the distillation of core information from vast amounts of computer, network and system-generated data. Virtual reality (VR) technology affords a new and novel opportunity to create new perspectives. Two examples of this approach are shown in Figure 7. The migration of a lightning storm across East Anglia and the resulting network failure reports are presented on two animated three-dimensional plates. In Figure 8 a portion of the BT network is shown in three-dimensional form. Here the user can 'fly the network' and gain new perspectives from inside and outside switching and transmission links.

Future Goodies—A Few Samples

- In 1993 a wrist-watch pager was launched with a price tag of several £1000. At the present rate of technological progress, a price below £100 should be realised by 1997.
- Around the year 2000 a wrist-watch communicator (with video?) will be available—Dick Tracey style—providing the challenge of new battery technology can be overcome. Prototypes are already under test.
- Using a very closed subset of voice communication with a machine, we can get satisfactory input/output today. But natural human voice

Figure 7—Lightning strikes in East Anglia



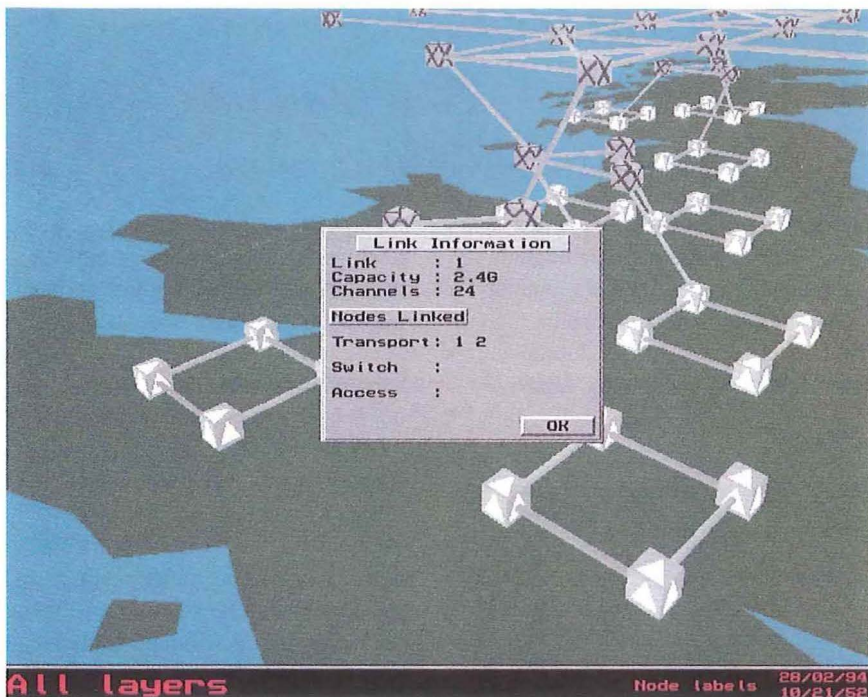



Figure 8—Flying the network—virtual reality as an engineering tool

on the office and IT than previously experienced through the increased expectations of the younger audience who are the workers of tomorrow. A few likely examples are:

- Virtual reality is poised to reach the home in 1994/5 in the form of games machines. Already you can link two Game Boy™ terminals via a cable and play each other across the room. A £100 modem allows players to communicate via the telephone network.
- Suppose it is Cup Final Day. The scene is Wembley Stadium; 100 000 people are in the stadium—minus 1—you! In your seat is a virtual presence terminal with a fish-eye lens giving full 360° vision back into your home VR terminal via a satellite broadcast link. Swipe the terminal with your Amex card and you can sit there in comfort and watch the game as if you were really there. You look to the left and right, front and back, and up and down: you have all-round vision giving you the perception of being at Wembley Stadium among the people around you.
- Already we see miniature cameras being strapped onto cars, cricket-stumps, quarter-back helmets. It is particularly interesting the different view that you get watching sporting events: it completely changes your perception.

Final Thoughts

In some respects we are currently suffering from technological indigestion, and it is likely to worsen as we approach the new millennium. We now have so much 'ripe technology' (and even more is coming) that we are spoilt for choice. What is required is a new approach to telecommunications, computing and IT that breaks the bonds with the 'copper' past and reaches out for new and novel solutions/applications.



communication (and full cognition) with a machine is probably another 15 years away (and perhaps 20), but by constraining the set that is being discussed, a workable level is then possible today!

- If the trajectory of electronics development continues, we can expect to see computers that will be 1000 times more capable than present-day machines within 10 years. In the new millennium, we can thus expect computing power 1 000 000 times greater and storage capacities that are far higher than we can even contemplate today.
- A simple question: when will a super computer equal a human from the standpoint of storage and sheer computational capability? A projection along our established trajectory suggests the answer to be the year 2010! At that time a super computer equals the human. By the year 2030 it will arrive as your friendly personal computer, requiring very much more than a 64 kbit/s ISDN port, yet no human intervention to establish communication with others of its kind!

Consumer/domestic technology can be expected to have a greater impact

Biography



Peter Cochrane
BT Development and
Procurement

Peter Cochrane joined the British Post Office in 1962 and was employed as a linesman and maintenance technician until 1966. In 1973 he graduated from Trent Polytechnic with a B.Sc. Honours degree in Electrical Engineering and gained an M.Sc., Ph.D. and D.Sc. from the University of Essex in 1976, 1979 and 1993 respectively. He is a fellow of both the IEE and IEEE, a visiting professor to Essex and Southampton Universities, an honorary professor at the University of Kent, and a visiting fellow to University College, North Wales at Bangor. He joined BT Laboratories in 1973 and has worked on a variety of analogue and digital switching and transmission studies. He has been a consultant to numerous international companies on projects concerned with systems, networks and test equipment development. In 1978 he became manager of the Long Lines Division where he was involved with the development of intensity-modulated and coherent optical systems, photonic amplifiers and wavelength-routed networks. He received the Queen's Award for Technology in 1990 for the production of optical receivers for TAT-8 and the PTAT-1 undersea systems. In 1991 he was appointed to head the Systems Research Division which is concerned with future computing and communications developments. During 1993 he became head of Research at BT Laboratories with 620 staff dedicated to the study of future technologies, systems, networks and services. He is also the BT Development and Procurement board member for technology.

Service Management Systems: Supporting the Customer Interface

This overview of service management systems provides an introduction to the articles on the COSMOSS and ETMS systems elsewhere in this issue of the Journal^{1,2}. It highlights the importance of information systems in supporting customer-facing operations and discusses trends in their development.

Introduction

Service management is a key differentiator in an increasingly competitive telecommunications market-place. Essentially it embraces everything concerned with customer-facing operations, ranging from procedural activities to more subjective matters, such as the customer's overall perception of how well a supplier and its employees respond to their requirements. For a telecommunications company, establishing a world-class network and product portfolio is not sufficient in itself; customers must also be left with a warm feeling, as if the supplier were in business purely for their benefit.

Service management operations are generally concerned with the marketing, provision, monitoring, maintenance and billing of products and services. Such operations are regarded as business processes, because they are initiated by a customer requirement and are only completed when the customer is fully satisfied that their requirement has been met. Until relatively recently, this involved paper-based, and hence labour-intensive, activities that took days or even months to complete.

This situation has been transformed within the last ten years or so by a considerable growth in the number of computer systems supporting service management operations, with some 200 now performing a wide variety of functions. However, many of these systems have been developed to automate individual components of an overall process or support the rapid introduction of new products and

services, without much thought being given to the end-to-end process or overall business objectives. As a result, systems can become inflexible and difficult to enhance when the need arises for further customer service improvements.

One example that has brought major benefits to both customers and service receptionists is the customer service system (CSS), described in an earlier issue of the *Journal*³ and shown being used in a typical customer service centre (Figure 1). This system has been very successful in integrating service and network management operations for telephony services and customer premises equipment but it cannot easily be developed to handle more complex products and services. Therefore, other systems have been developed to handle private circuits and advanced network services such as Linkline™ and Featurenet™. However, this makes it more difficult to integrate several diverse operations on a per-customer basis.

BT Applications Architecture

In order to address this problem and ensure that information systems are evolved in an integrated manner, BT has developed a single logical applications architecture which splits business operations into four components (see Figure 2). In addition to service management, there are three other components:

- *network management*, which includes applications and data required to control and manage the network, either on an individual element or end-to-end basis;



Figure 1—Customer service system

- *resource management*, which includes applications and data required to manage people and tasks, as well as assets such as vehicles and buildings; and
- *business management*, which includes supporting business-wide activities such as strategic planning, financial control and product management.

The architecture for network management applications and data is already well advanced along the lines described in earlier issues of this *Journal*^{4,5}, and subsequent articles in this *Journal* will review progress in this area. Although the architecture for the other applications areas is less mature, this is rapidly catching up as the experience gained from the development of

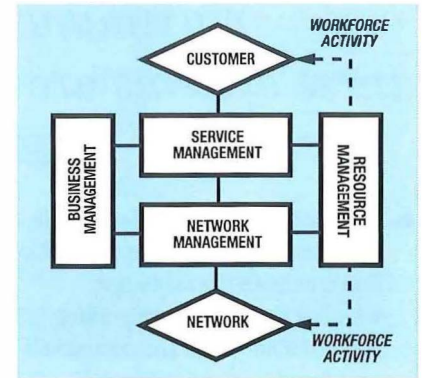


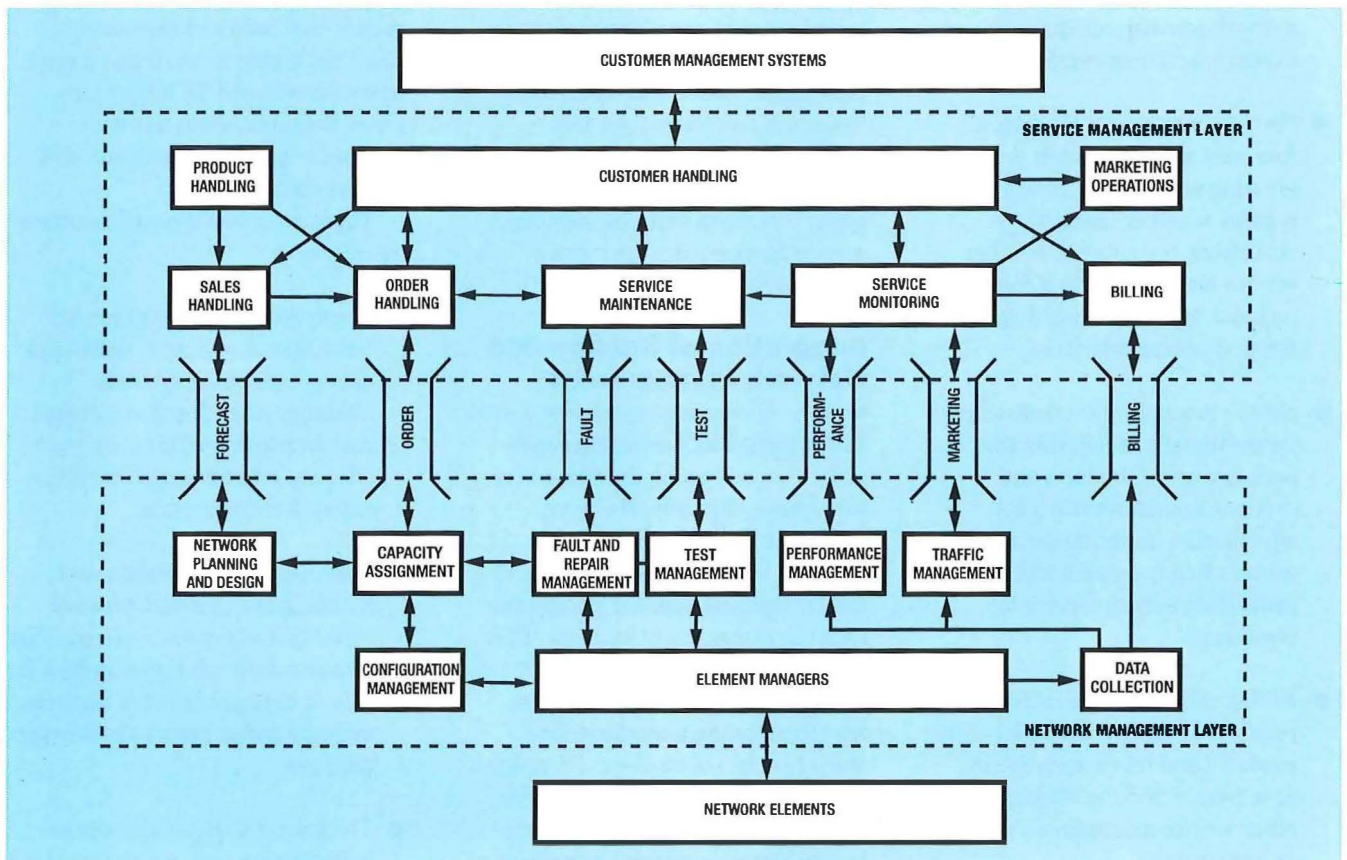
Figure 2—BT applications architecture

network management systems is carried through into other developments. The long-term objective is to integrate all applications and data management on a common distributed systems platform.

Service Management Architecture

Within the service management area, or layer, a logical architecture has been developed which splits all service management operations into eight discrete functional areas, as shown in Figure 3:

Figure 3—Logical architecture for service and network management



maximum benefit will only be achieved when these systems are also fully integrated with the appropriate network management systems

- *Customer handling* deals with all functions necessary to support the dialogue with the customer, whether this is through service receptionists or via the customer's own terminals and systems.
- *Products and services handling* maintains details of the products and services that are available, including functionality and specifications, for both enquiry purposes and selection at the order handling stage.
- *Marketing operations* includes the management of marketing campaigns and the accumulation of market intelligence.
- *Sales handling* includes pre-order activity such as contract and bid management.
- *Order handling* not only processes orders for new or changed products and services, but also manages requests for changes in the use of existing products; for example, network reconfigurations within an existing service agreement.
- *Service maintenance* includes all functions associated with progressing and clearing problem reports, whether initiated by customers, the network or other service management systems. This includes any supporting diagnostics and testing activities.
- *Service monitoring* is responsible for continually monitoring the performance of products and services against service level agreements, initiating corrective action when necessary and generating service reports to customers.
- *Billing* collates all the information required to bill customers for their products and services, gathering data from either the network or other service management systems as appropriate.

Grouping systems together within these functional areas enables future development requirements to be rationalised by identifying generic functions that can be applied across all customers and/or products. This will simplify the task of introducing new products and services or enable service management to be tailored to individual customer requirements by simply adding new applications to the appropriate area.

Each functional area represents both the applications and data appropriate to the processes involved. However, data may need to be accessed by more than one functional area, in which case either a master source can be designated or it can be stored in a separate global data area.

To give an example of how existing systems fit within this architecture, the CSS system discussed earlier carries out all functions except marketing operations and sales handling, although only for a limited set of products and services. By contrast, the two systems described in this issue of the *Journal* are specific to only one functional area but handle a wider variety of products and services. COSMOSS (customer oriented system for the management of special services) and ETMS (event and test management system) perform order handling and service maintenance functions respectively.

Integration of Service and Network Management

Integration of all service management systems would alone be a very worthwhile objective. However, maximum benefit will only be achieved when these systems are also fully integrated with the appropriate network management systems. This will enable seamless, 'hands-off' operations to be performed all the way from the customer interface down to network elements. Initially, such operations would generally be controlled by service receptionists, but an increasing degree of control

will be devolved to customers themselves to allow them to configure and manage their networks and applications as they require.

Figure 3 also shows how the service management layer will connect to the network management layer. Between the two layers are several logical interfaces, each of which can be regarded as a 'pipe' through which data can be exchanged. These provide a focus for the development of interface specifications between the diverse types of information systems contained in each layer.

Data flow between service and network management systems will generally be concerned with 'network features'. These in effect represent the functional capabilities available in the network and would be provided and managed by the network management layer according to a 'feature agreement'. In addition to details of the functionality of the network feature, the feature agreement will contain profiles, which specify key parameters for the management of the feature, such as provision and restoration times, performance levels and measurements required for billing purposes. Feature agreements will eventually form a key component of product specifications.

There are seven types of interface as follows:

- The *forecast interface* primarily links sales handling systems and network planning systems, although other functional areas may be involved. This contains advance information about future network requirements.
- The *order interface* links order handling systems and network capacity assignment systems. This contains data relating to orders for new or changed network features, including information about order progress.
- The *fault interface* links service maintenance systems and network

fault and repair management systems. This contains data concerning the management of faults, wherever they originate.

- The *test interface* links service maintenance systems and network test management systems. This is used for initiating test requests and communicating test results.
- The *performance interface* generally links service monitoring systems and network performance management systems. This conveys all data concerning the performance of network features against specified performance levels.
- The *marketing interface* primarily links service monitoring systems and traffic management systems, although other functional areas will be involved, especially network planning systems and marketing operations systems. This covers any data related to the usage of the network that may be required for marketing campaigns or customer service reports.
- The *billing interface* primarily links billing systems and network data collection systems. This contains all data that needs to be collected by the network and collated for billing purposes.

Taking this logical approach to architecture development does not necessarily mean that each function should be realised within discrete computer systems. For example, a function such as order handling for a particular product or service may happen to be performed in real time within network systems under direct customer control. However, it is still regarded as a service management function and requires a corresponding process to be invoked in response to the customer's requirement, even though this may be executed within millisecond time-scales.

In order to ensure that business processes can be managed in this way,

it is therefore essential that the functions and applications that support them are developed in a flexible manner. They will need to be capable of being run on a variety of computing and network platforms, but the key to all of this is an integrated information systems architecture.

Future Developments

In order to meet increasingly sophisticated customer requirements, the functionality of service management systems is expected to evolve rapidly. Developments over the next few years are likely to centre on the structure of the systems themselves and the customer interface.

As far as the systems are concerned, computing hardware will increasingly be divorced from the applications supported, enabling a single, distributed platform to be created. This will be capable of running a variety of applications, rather than just those specific to a particular customer, product or process. A flexible applications library would be created, with each application being selected as required during a business process according to a workstring template. When additional functionality is required, for example, when a new product is launched, new applications can then be created and added to the library, and templates revised, within much shorter time-scales. Similarly, databases will begin to evolve separately from processors to enable single, master data sources to be created and allow processors and databases to evolve separately.

As for the customer interface, service and network management will increasingly become a real-time operation for a wider range of products and services, controlled by service receptionists or customers themselves using terminals capable of managing any function at any location. Such terminals will have increasing access to artificial intelligence techniques to assist with the

customer dialogue or guide the customer through service control operations. There are also likely to be direct, high-level, interfaces between customer handling systems and the customer's own information systems, for automatic service control, reporting and billing purposes.

Summary

This article has given a brief overview of service management systems in order to provide some background to the articles on COSMOSS and ETMS later in this issue of the *Journal*. It has not been the author's intention to describe particular systems in any depth, but rather to encourage the reader to undertake further study of this very important subject.

Acknowledgements

The author would like to acknowledge that the information summarised in this article is based on the results of a wide range of studies and would like to thank the many people throughout BT who have contributed to this work.

References

- 1 VAN DER STIGHELEN, HENRI. COSMOSS—An Innovative Approach to System Development. *Br. Telecommun. Eng.*, April. 1994, **13** (this issue).
- 2 YOUNG, TIM, and HOLLAND, BOB. Event and Test Management System. *Br. Telecommun. Eng.*, April 1994, **13** (this issue).
- 3 GORTON, R. K. Customer Service Systems. *Br. Telecommun. Eng.*, April 1986, **5**, p. 18.
- 4 Network Administration for the 1990s. Special issue of *Br. Telecommun. Eng.*, Oct. 1990, **9**.
- 5 Network Management Systems. Special issue of *Br. Telecommun. Eng.*, Oct. 1991, **10**.

Biography



Stephen Wilson
BT National
Business
Communications

Stephen Wilson is responsible for service management strategy in the Customer Network Operations unit of BT National Business Communications. He is also chairman of a working group defining the service to network management interface and secretary to the Board of Editors of the *Journal*. He graduated from the University of Dundee in 1973 with an honours degree in Electronics and joined the British Post Office in the London South West Area, where he worked as a Training Officer and Exchange Maintenance Manager. He then moved to the Network Strategy Department in central London, where he worked on a variety of projects ranging from network design and cost modelling to long term network architecture studies. In 1991, he joined Business Communications Division where he had responsibility for product launch coordination and system evolution for service management centres before being assigned to his present post.

COSMOSS—An Innovative Approach to System Development

COSMOSS is more than just another new computer system, it is an innovative approach to system development. COSMOSS has been developed using a multi-division liaison between the business analysts, the system developers, their agents and the system users. This article gives a brief glimpse of the scale and complexity of a major system development and an appreciation of the objective of delivering a system that the customer wanted and the business needed.

Introduction

There had been attempts in the past to define and develop a system to manage BT special services (basically private circuits and telex services). The latest requirement was for a computer system for the management of special services for the Business Communications Division of BT having specific emphasis on the interface with the customer. The *customer-oriented system for the management of special services* (COSMOSS) was the final choice, and was originally targeted as a replacement for the computer automated private service system (CAPSS), that provided order-handling capability for special services.

CAPSS was being replaced because of difficulty in maintaining and developing the system. The problems stemmed from the complexity of the application language used and the non-BT-standard host machine.

The problems made it difficult to enhance the system to meet certain corporate objectives. These objectives

included the provision of a national system that was customer-centred and supported a single point of contact.

COSMOSS is a sophisticated order-handling system which references a unique customer management database and a products and services database. COSMOSS, which began a phased go-live in March 1994, has a modular architecture and is now one of the largest MVS/DB2 installations in Europe, if not the world.

Business need

The background to COSMOSS has its origins in BT's position as the country's most experienced and longest established provider of private telephony services. While this has obvious advantages, in the changing competitive environment, it also has its down side.

BT's processes have evolved over many years. They may not always match those of new companies in the market, who can capitalise on historic information and produce modern and efficient systems from scratch.

'Service management is a key differentiator in an increasingly competitive telecommunications marketplace...customers must also be left with a warm feeling, as if the supplier were in the business purely for their benefit...a considerable growth in the number of computer systems supporting service management operations...Therefore other systems have been developed to handle private circuits and advanced network services...'¹

'COSMOSS is a major system that pulls together key business areas to deliver this initiative'.

Figure 1—COSMOSS architecture

To cater for the existing customer base in private circuits during the 1980s, the company developed CAPSS. CAPSS was hosted on 27 regional databases that became tenuously linked from 1988 onwards.

Valuable as this system proved to be, it was order-orientated and not designed to be a customer-facing system. Thus, to identify the progress of a customer's order, each separate circuit would have to be accessed in turn to provide a view of the customer's end-to-end solution. This led to inefficient working practices and much duplication of effort.

In many cases, more than one BT person had to be involved to identify all the necessary customer information, often from more than one regional database and sometimes from paper records.

In addition, users of the system and BT's customers could sometimes be confused by variations in information derived from different parts of the country.

After a programme of market research among CAPSS users and BT customers, three major requirements were identified:

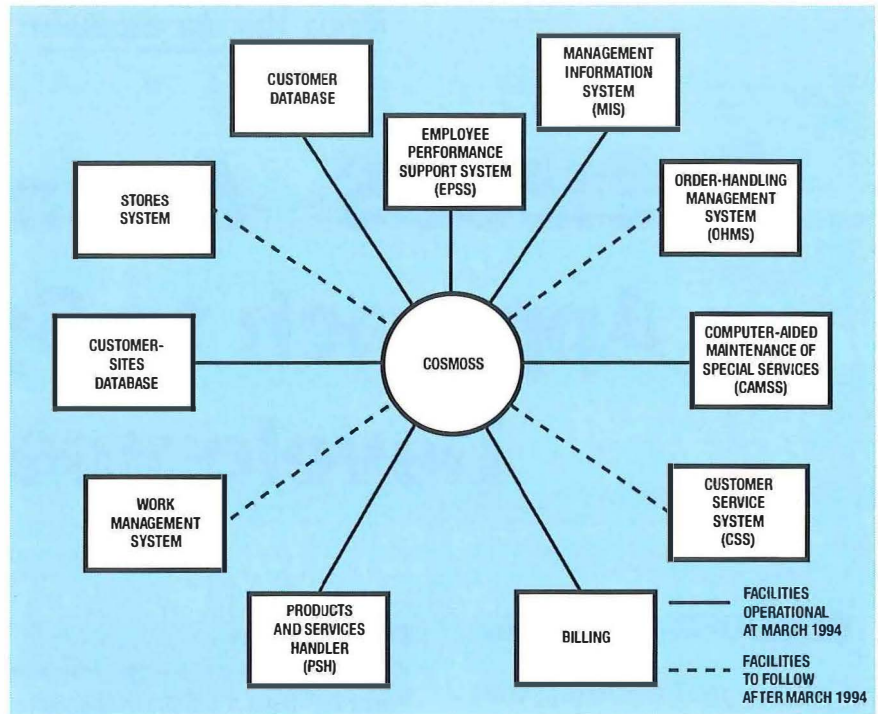
- standardised working procedures;
- a seamless national order handling database; and
- support for customer one-stop shopping.

From the start, COSMOSS, operating as a single national database, replacing the multiple CAPSS databases, to give a more flexible response to customer needs.

A brief overview

Users are allocated a single logon code, enabling access to all necessary parts of the database. This, in turn, allows the identification of the end-to-end progress of one, or all, of a customer's orders.

The new system initially has six interfaces (see Figure 1) linking to:



- a management information system (MIS), that governed the quality of service and performance information;
- the computer-aided maintenance of special services (CAMSS) system, used to store and progress repair information.
- a billing database;
- the products and services handler (PSH) products and services database;
- a customer database; and
- a customer-sites database.

COSMOSS was launched, in March 1994, with a national generic order-handling capacity using information drawn from these databases.

Later, in 1994/5, the link to the order-handling management system (OHMS) will be established and, in 1995, the system will be expanded to include links to the customer service system (CSS) for public switched telephone network and customer equipment orders.

System feasibility

In early-1990, BT's Business Communications Division commissioned a feasibility study for the

replacement of the CAPSS system, corporate strategy being a key driver for this initiative. A special business team was established to consider the options and produce reports and costings against business objectives for Business Communications' senior managers. The study was completed in October 1990. A consequence of this report was the COSMOSS feasibility study. By January 1991, the business survey and benefits study were produced.

Authority was given to proceed to the business case, which was signed off in May 1991. The COSMOSS programme was planned for three-stages with delivery beginning in October 1993. The second release would be March 1994, and the third and final release would be September 1994. The team produced the first scope document in July 1991. The current scope document was produced in September 1991. The scope was widened in response to Business Communications initiatives and direction from the integrated service management (ISM) initiative. The re-scope was to include single-point-of-order entry and a generic approach to products and services management.

When the business case was signed off, the special business team had authority to proceed to deliver the COSMOSS programme.

Project Implementation

People

Recruiting

The initial studies established the project constraints. The projected size of the team and the skill-set requirements gave the project manager and his team leaders a formidable recruitment task. The main resource area was established: the project manager and his team targeted commercial officers in Business Communications as the main source for recruits. Over a period of months, recruits were brought into the team. It became clear that the skill set would demand business analysis, management and technical depth. Consequently, the scope of resource was broadened to include the Business Communications engineering divisions.

Recruits to the project joined on a short-term basis, although the project manager envisaged that there may well be future opportunities for this highly-skilled team on a more permanent basis.

Establishing the roles

Early studies identified the key functional areas of COSMOSS. At the highest level, the COSMOSS team had established a front office, a back office, a project office and cross-functional teams.

The *front office* incorporated customer and site management, products and services, customer order handling and project management.

The *back office* incorporated job and activity management, resource management, engineering data and management information system (MIS).

The *project office* handled the project support, planning, documentation and reports. Additional responsibilities included budget control and personnel issues.

Cross-functional responsibilities were identified for data conversion, implementation, validation verifica-

tion and testing, and the user interface.

Team building

A programme of working events was rolled out to establish and develop the team skills of the COSMOSS people. The scale and importance of the project demanded a very high level of team values.

Skilling

Concurrent with the development of COSMOSS, each team member had to undergo a levelling process. This process ensured that everyone had a working knowledge of all aspects of the terms of reference. Furthermore, team members were expected to achieve a high level of expertise in those areas that affected them directly. The culmination of this lengthy skilling operation highlighted the foresight of the project manager and his team—the legacy of the COSMOSS programme was a pool of highly-skilled business analysts as a valuable asset to the company.

Planning

Overall plans for the programme were drawn up which identified the milestones and the associated deliverables. The scale of the operation dictated a phased delivery. Based on the scope, benefit study and business case, a release programme was drawn up for COSMOSS.

Terms of reference

The development of COSMOSS as a joint venture between Business Communications's business analysts and Development and Procurement's (D&P) application system developers required the establishment of inter-division terms of reference. From the outset, the project adopted a quality management system. In keeping with corporate policy, total quality management tools and techniques would be used extensively. The project adopted the D&P Information Systems Development (ISD) standard TELSTAR method for the development of the COSMOSS system. When

the project was under way, the decisions were taken to use computer-aided software engineering (CASE) tool technology, and SSADM/Excelerator was the choice. The volume and shareability of the documentation being produced demanded a standard. The decision was taken to migrate to the complete Microsoft Office system.

Managing the Project

COSMOSS, like all projects, has a definite life span. Essentially, good project management demands efficient juggling of limited resources over a set period to achieve the desired results. The application of project management was required to ensure goals were achieved on time, within budget and of best quality.

A project is unlikely to succeed if there is an inadequate strategy for tackling the basic problems. The BT Project Management Handbook provides key principles, techniques and procedures that can be applied to all types of BT programmes and projects. It has been written at a high level and provides a methodology for developing further procedures. TELSTAR is the BT project management tool for computer system development.

Project Management Handbook

The Project Management Handbook describes the principles, techniques and procedures applied to the management of projects, and programmes of projects, within BT(UK). It is a key document in BT's drive to achieve a consistent approach to programme and project management.

Good project management depends on:

- understanding and defining the organisation of the project and making it work effectively;
- ensuring that people understand their roles and those of others, and

Figure 2—Software development life cycle

that everybody concerned understands what is required of them; and

- ensuring that systems are clearly defined and are understood by all.

The procedures in the handbook apply to all projects, unless the client requests, or agrees, modified or additional procedures for specific projects. Thus, the COSMOSS programme adopted the TELSTAR approach to systems development.

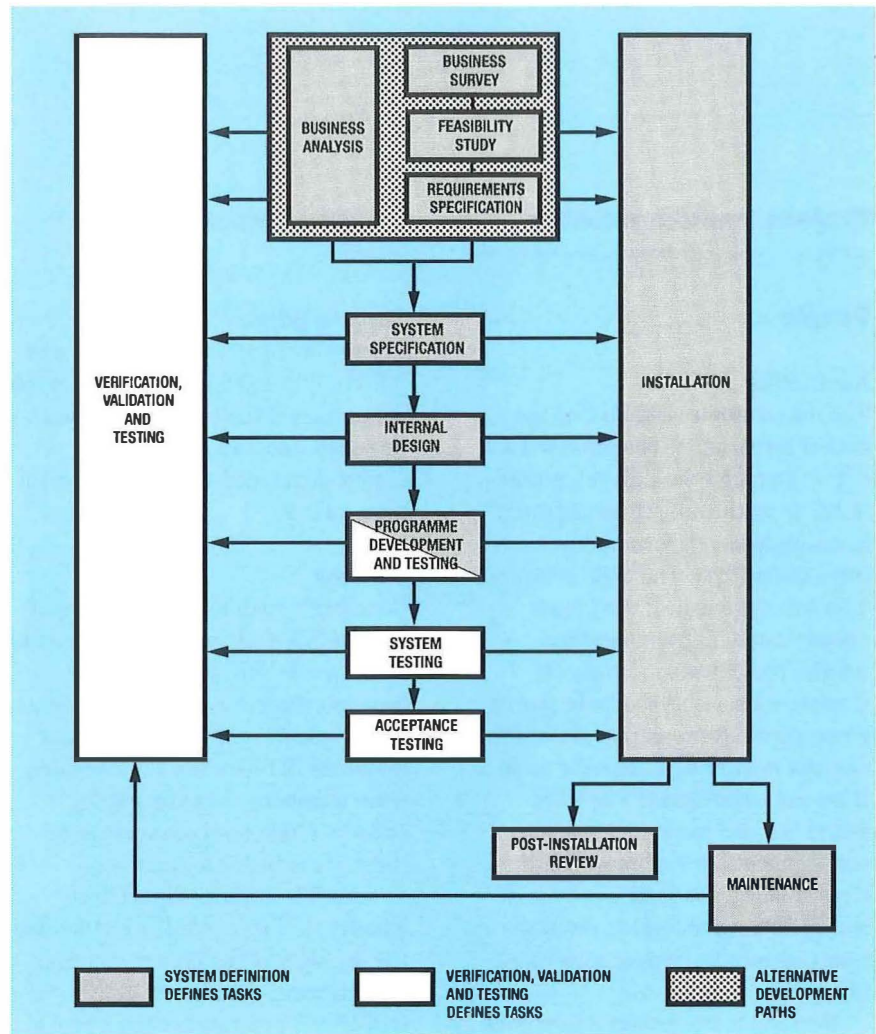
TELSTAR approach

TELSTAR is the recommended method for the majority of development projects within ISD. It was originally based on the structured development method, but has considerable additional features, such as project management, and verification validation and testing. It also considers areas such as non-functional requirements, performance engineering, and design reuse. Future releases will include configuration management, knowledge-based system development and business analysis.

It is generally a top-down development structure, in which models defining the whole scope of the system are developed in the early stages and refined in detail throughout subsequent stages. Some aspects of the method use a bottom-up approach, for example, current systems analysis, which leads to production of a description of an existing system. This approach is particularly useful where enhancements or closely-related systems are being developed.

Start-up costs for TELSTAR are comparatively low, but some investment had to be made to familiarise the team and the client with the procedures, techniques and documentation that are being used.

TELSTAR is a comprehensive methodology. The COSMOSS programme tailored its requirements, from the wide range of procedures, to fit the project (see Figure 2).



General objectives

The formal objectives of using system development procedures are to:

- prioritise and manage business requirements throughout the systems development life cycle;
- build developments in accordance with requirement priorities and within the scope of the business system models;
- use architecture and data management models to steer development;
- identify, agree, plan and manage the quality criteria for a project;
- manage, as a project, the progress of a discrete set of activities with defined end-products;
- use business requirements to drive the physical design, which drives the construction, and confirms compliance at each point of progression; and

- use engineering disciplines throughout the life cycle.

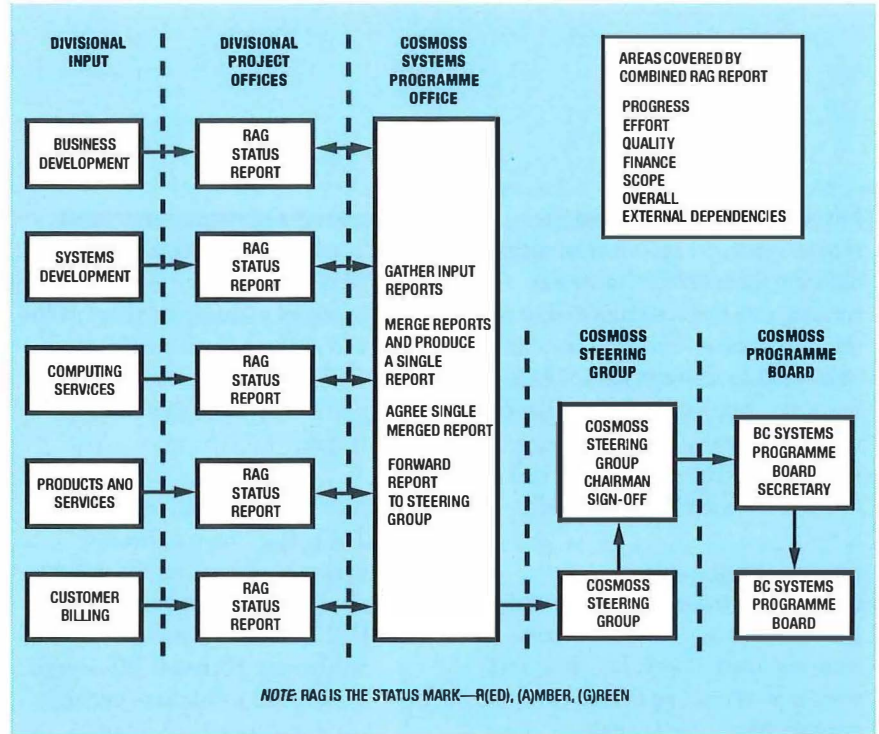
Quality management system

During the development life cycle, the COSMOSS programme was audited for British Standards Institution (BSI) BS5750 (ISO 9001) quality systems registration. This was an unforeseen mandatory requirement that placed contention on the project's resources. However, the project benefited overall from the tighter controls of the quality management system. The auditors' report included a comment that COSMOSS was one of the most impressively managed projects within BT.

Reporting on the programme

The COSMOSS programme adopted a reporting structure to senior management that is unparalleled within the company. Figure 3 explains the monthly flow of information that enables the Programme Board to maintain a very tight control of progress against plans.

Figure 3—COSMOSS high-level monthly project reporting procedure



The Development

The teams

Owing, in the main, to time constraints and the need to reduce any learning curves, it was agreed that the business analysts and specifiers (the customer) and ISD (the supplier) would work together in the capturing and documenting of the requirements (see Figure 2).

Requirements specification stage

The major deliverables from the TELSTAR requirements specification stage are shown in Table 1.

Table 1 Requirements Specification Stage

Deliverable	Owner
Problem and Requirement List	Business analysts
Requirement Definition	Business analysts
Data Model	ISD

While completing this stage, the team were being taught business-analysis skills and CASE tool management. The suppliers enhanced business analysis skills and understanding of the business environment. Interworking enabled the stage to be completed on schedule and progress to the next stage: system specification.

System specification stage

Moving into the system specification stage, the customer-supplier interworking was maintained. The major deliverables from the TELSTAR system specification stage are given in Table 2.

During this stage the need for strict change-management controls over the project was recognised and implemented. This ensured better control of the project's requirements and scope, both functionally and financially. This control enabled better decisions on prioritisation and budgets.

Table 2 System Specification Stage

Deliverable	Owner
System Specification	ISD (business analysts owned screen design, messages and processing and the user interface. ISD owned data design, management and performance)
System Security	Business analysts

Unfortunately, the strict control made every change, from major to minor, visible. The number of changes appeared alarmingly high. The Project Board felt that the scale and level of change within the project was making it unstable. Ultimately, this led to poor perception of the project as a whole, a review of its viability and consequently to the introduction of very close monitoring.

Internal design and code construction stages

Moving into the internal design and code construction stages, the customer and supplier teams split. The supplier concentrated on the areas with minimum customer input, other than quality-control review and verification validation and testing

work. Meanwhile, the customer team focused on combined system and testing. The test points, test designs and test scripts were produced in readiness for the application software. The major deliverables from this stage are shown in Table 3.

Table 3 Internal Design Stage

Deliverable	Owner
Internal Design Specification	ISD
Software Programs/Code	ISD

The major deliverables from the pre-system test or preparation stage are shown in Table 4

Table 4 Pre-Test Stage

Deliverable	Owner
Test Point Matrix	Business analysts
Test Design	Business analysts
Test Scripts	Business analysts

The major benefits of this splitting of the customer and supplier teams was the rapid move into both areas by experienced teams. ISD brought

forward their expertise from the requirement and specification stages, allowing them to move into code construction from internal design with minimal resource from the customer. The business analysts were able to move very quickly into the design of testing because of their detailed knowledge of the system, both in terms of functionality and design.

System testing stage

Moving into the system testing stage, the customer and supplier teams were reunited. Specialists from ISD were moved into the testing area to support the business analysts' core testing team.

The remainder of the ISD team were deployed on fault fixing. The development benefited from code modularisation and the object-oriented approach. Faults were identified, and fixes re-applied to the system-test environment without the worry or pain of major retesting.

A proprietary IBM product (TREC) was used for logging and progressing faults. Information was downloaded into an EXCEL spreadsheet to provide detailed statistics for better planning into the future and for continuous improvement.

Progress reviews

The continuous progress reviews enabled key decisions to be taken for the future of the programme.

- Optimise the CASE tool used. SSADM/EXCELERATOR was chosen. However, better results could have been achieved from a word processor and graphics package as the full benefits of the CASE tool were not fully realised until technical skills increased. The tool could have provided information that had been difficult to identify. This was due to insufficient attention, from the outset, to the internal model.
- Ensure that change management can differentiate between

'changes of requirement' and 'clarification of requirements'. This was not fully established and incurred a knock-on affect at the end of the system specification stage. Procedures were put in place to manage this in the future.

- Automate as much as possible. The project has a myriad of documentation, and changes are costly in time and money. The COSMOSS team decided to implement Microsoft Office with E-mail and a database management tool for documentation.

The close partnership between the customer and supplier produced a cross-fertilisation of ideas and ownership of responsibility that should become the accepted norm.

Hardware and the Network

Hardware strategy for COSMOSS

Scope of the hardware

The COSMOSS programme, in liaison with BT's Group Computing Services (GCS), looked at the needs of providing a national customer-focused database, offering up to 9000 user accesses and operating in the region of 800 complex transactions. A very large mainframe was needed. This mainframe was also required to execute the complex transactions within sub-second response times.

During the early stages of the project life cycle, only one machine offered the primary requirements: IBM's new ES9000/DB2 mainframe computer.

Network strategy for COSMOSS

Scope of the network

The network requirements for COSMOSS had four primary considerations.

- 1 Where appropriate, to utilise the existing network and terminals to provide COSMOSS users with access.
- 2 To support the changeover from the previous system to the COSMOSS system by permitting access to both for a three month period. This would include offering access to the other BT systems that the users required to execute their job.
- 3 To identify high-capability users of the COSMOSS system and to reduce their network access times.
- 4 To offer suitable network-management tools to ensure that network downtime does not interfere with access requirements for normal business operation.

Options available

BT has several computer networking strategies in operation throughout the company. These included:

- point-to-point networking, using Datel mux 5500 switches, which supported 30% of the CAPSS network strategy, supplied over Hewlett Packard (HP) terminals;
- BT's own open systems architecture (T-Net), which offers wide area networking in a region, and is linked together via network gateways using packet switching techniques (the limitation to data transmission across the network is the terminals that interface with network providers at 9600 bit/s);
- IBM's system network architecture (SNA), providing direct terminal access to the nearest IBM mainframe and passing information over the IBM network backbone, interfacing to the mainframe at 48 kbit/s; and
- personal computers (PCs) connected to a network provider via a local token-ring access method.

The high data transfer rates (10 Mbit/s) with the network provider ensures faster network response times.

Constraints

The network strategy for COSMOSS took into account the business needs and the needs of users to access other systems. However, budget constraints impacted on any move away from the existing networks.

In the case of the point-to-point network, users were restricted to HP terminals, with the resultant limitation to the HP mainframe. Alternative networking was required for those terminal users that could not interface with the IBM mainframe over this network.

The best network architecture for the users of the previous system was BT's T-Net, that offered access to both the HP mainframe and the IBM mainframe.

Users of existing BT systems that had access to other IBM mainframe-based systems could be offered network access via that system.

As part of the business need to record new business quickly, efficiently and on-line with customer calls, it was decided to connect special sales offices via PCs connected to a network provider.

Decisions

Owing to the wide range of networking options available to BT, and the varied needs of the users, it was impracticable to limit the COSMOSS network strategy to a single architecture. Multi-architecture ensured that BT was not restricted to a single supplier and could negotiate the best terms for network hardware. Additionally, T-Net offered its users generic access into systems that were otherwise incompatible in hardware or software configurations.

New Initiatives

The COSMOSS development programme has spanned a number of years. Therefore, it should not come

as a surprise that new initiatives came to light along the way. Several of these initiatives have been incorporated into the programme to ensure a state-of-the-art delivered system. In particular, COSMOSS has benefitted from object orientation, generic system-management concepts and an employee performance support system (EPSS).

Object orientation and COSMOSS

Followers of buzzwords in the software community will be familiar with the terms *object oriented* or *object orientation* or just *OO*. COSMOSS has become the first major user system in BT, and outside BT, to adopt and effectively use object-orientation principles. Object orientation is about a new approach to building systems. Object orientation promises a way forward that helps to deal with some of the perennial problems that have dogged the development of systems in the past. These problems range from not getting the requirements right in the first place through to expensive and heavy maintenance after implementation of the systems.

Object-orientation systems

Object-orientation systems are composed of *objects* which encapsulate both data and functionality that communicate with each other in strictly defined ways. Objects are real-world items that always exist in that area of the business; for example, objects are things like customer, order, product/service type and so on. These objects present a more stable view of the business world than programs and screens. So if, as is inevitable, there are changes in the type and manner of the business then these changes can be accommodated more quickly and with less overall impact than in the past. The tendency previously was for systems to be designed around their processing. This frequently resulted in transactions and programs accessing and updating the same piece of informa-

tion in different ways and in different places, resulting in complex and difficult-to-manage systems and an inability to implement even small business changes in reasonable time.

Although the object-orientation approach sounds common sense, it is a relatively new area, and its previous use tended to be in the more esoteric and small-user systems domain. COSMOSS has now brought the advantages of object orientation to mainstream large-volume transaction processing systems in BT, and has broken new ground overcoming technical and other constraints, and making object orientation work practically. The handling and updating of information through COSMOSS is done by the main transactions calling up *capsules*. These capsules contain all the code for performing a particular action on, or retrieving information from, customer or order objects. In object-orientation terms, the information is encapsulated with the processes that relate to it. The systems information in the objects can only be accessed by using these capsules, and so the information and the way it is handled is consistent across the system. It also means that there is only one area of the system to be modified if a change is required to a particular object. This improves the speed of changes and eases test and control. Furthermore, the way that the capsules are based around real-world objects make the impact of changes easier to understand and implement.

The object-orientation approach and the associated technical infrastructure adopted by COSMOSS also facilitates the distribution of systems and systems information. This is a crucial factor in servicing BT's customers not only in the UK, but also in the rapidly expanding global marketplace.

There is a large and growing body of opinion that object orientation is not just flavour of the month, but the way that all systems will need to be built in the future if they are to avoid

becoming dinosaurs soon after they are delivered. The adoption of object orientation should help COSMOSS remain a flexible and adaptable animal in a rapidly changing business climate.

For more about object orientation, see References 2 and 3.

Generic capability

There has been a tendency for system development solutions to be aligned with the needs of particular products or closely-knit product families; for example, where an order processing, a billing and a network system have been put together to service a particular product need. This vertical set of systems to address the requirements of a particular product set is sometimes known as *stove piping*. This can result in a great deal of similar functionality being re-invented each time for each new product set. The rationale for stove piping is often that there are commercial and time pressures that prevent the integration of these business-management systems in an effective way. Primarily, it is too lengthy a process to amend the existing systems to incorporate the needs of the new product.

Because of the wasted effort of stove piping, and the lengthy timescales for changes to existing systems, a key business driver in the development of COSMOSS was the need to respond quickly to the order-handling requirements of new and changed products. The requirement was for a *generic* capability that would embrace the needs of all types of products.

The COSMOSS approach was to bring into being a PSH that had the structure and capability to encompass change. The challenge was to hold all this information in an easily usable form to facilitate the order-taking function. This was achieved by a table-driven approach that consolidated information about products. The following information supported the order-taking process:

- specification,
- features and feature values,
- relationships, and
- prices.

The PSH brought together information that previously had been held, off-line and in disparate places.

The area of products and services is just one example of the COSMOSS development programmes' philosophy to go for flexibility and capability. The table-driven approach was used wherever possible to ensure that the data was independent of the programs that manipulated it. This ensured a faster reaction to change, and a safer, surer implementation of change for a system that must demonstrate the high reliability and effectiveness demanded by a customer-facing environment.

Employee performance support system (EPSS)

System performance is like good health: you always want more and better. Traditionally, performance has been a measure of how long it takes to achieve a response from the computer. However, improved system performance does not guarantee a comparable improvement in the effectiveness of the computer system. An important factor in any computer system is the user. When the performance boundary is widened to include the user, the result becomes more unpredictable. The EPSS targets this area to enhance the effectiveness of the computer system.

Traditionally, computer systems are developed and tested and a user guide is produced. A training-needs-analysis establishes the training requirement that is implemented on a finished-product computer system.

However, for a new single national system on which the client requires immediate service, there is contention between training and delivery. If the system has a large population of users that are geographically wide-

spread, there will be implementation-logistics problems. Short timescales, tight budgets and large numbers are poor ingredients for a quality training programme.

The solution was to commence training before system delivery and complete the programme just-in-time for going live. The EPSS provided simulation tools to deliver the look and feel of the real system. Computer-based training (CBT) could be delivered on a PC or on a mainframe terminal, and this provided an effective resolution to the implementation-logistics problems. Benefits were achieved in terms of time, cost and quality through consistency.

The scope of EPSS is not just the CBT. The EPSS introduces computer-based referencing (CBR), which is an electronic documentation system. CBR enables the upload, storage and retrieval of vast quantities of documentation. Documentation can be written by using conventional word-processing packages and automatically transferred to CBR using simple macros†. Once loaded, documentation can be managed cost-effectively and made instantly available to thousands of users. Maintaining the quality, integrity and availability of information held in one place is a simple procedure.

The strength of the EPSS is derived from an architecture which includes CBT and CBR. The EPSS can be invoked from the users' application program with the press of a HELP key. The EPSS provides sophisticated help by assessing* all the information on the user's screen and offering a range of options. These options can range from a simple explanation of a term to a browse through relevant documentation or a

† macro: a series of keystrokes that are captured to automate a process.

* assessing: the process where a user's screen is compared with the original blank screen to determine what new information is now present and may be causing a problem.

Figure 4—EPSS architecture

brief on-line tutorial. The choice is designed to suit the skill-level requirement of the user. An enhancement to the COSMOSS system added the external data access (EDA) facility. EDA enables the help system to access the host system for information such as valid value tables.

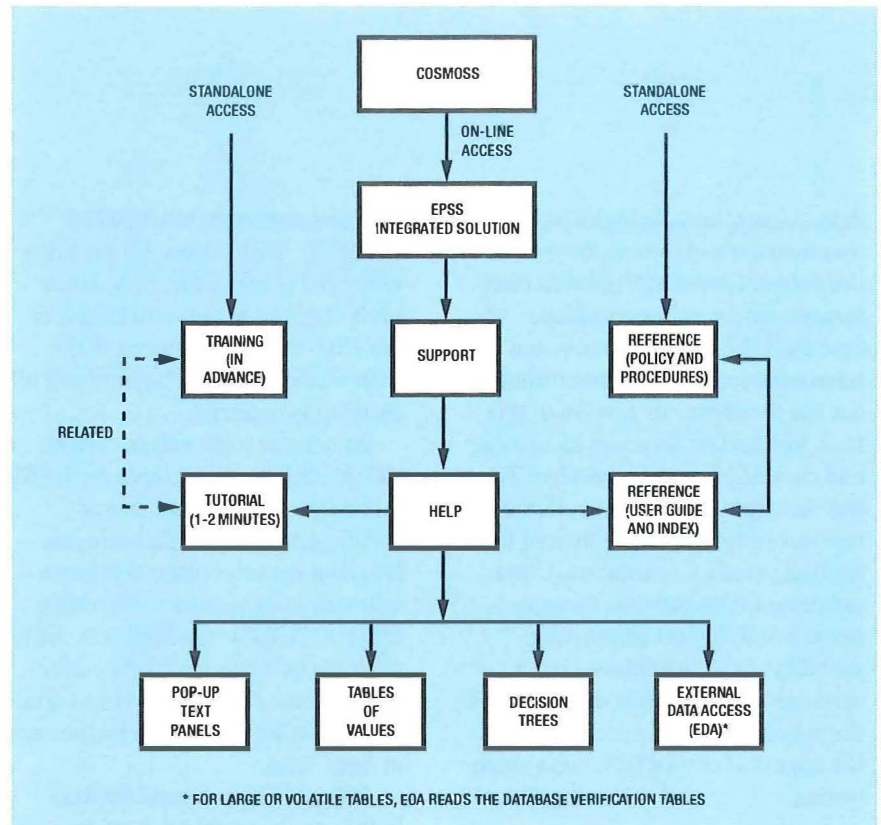
Figure 4, shows a typical EPSS architecture that includes all the above features.

The architecture of COSMOSS is non-rigid. Users have the flexibility to tailor system use to suit their requirements. The EPSS provides an on-line user guide to assist navigation within the system. The guide contains a range of expert scripts; that is, a mirror image of an expert user executing an operation. The script includes embedded dialogue panels that explain what is happening. Scripts can be readily enhanced and the range increased as required. Users with multi-sessions can run the scripts concurrently with the operation they are attempting. Additionally, the EPSS provides a user directory built on a hierarchical structure, which enables intelligent searching.

During the development life cycle of a computer system, changes will occur. Changes that affect the look and feel of the system will impact on the consistency of user education. The EPSS responds to this problem in two ways. Firstly, changes to EPSS productions are straightforward; they are made once and can be invoked immediately, for all users. Secondly, ongoing training and support can be made available, directly to the user. Users can access the EPSS directly from their terminals, for more lengthy training and document browsing, or via the HELP key for a just-in-time snapshot of support.

COSMOSS Application Testing

COSMOSS was developed and tested within the framework of TELSTAR. TELSTAR documents the life cycle of a system's development from the



initial conception through to delivery. System testing comes under the heading of verification, validation and testing.

Verification and validation

Verification and validation was undertaken at the end of the requirements-definition and system-specification stages of the project. Workshops were held at the end of each stage to test the product against business work strings and processes, and issues arising from these events were captured and progressed by using either the change-control system or the fault-management system. User group meetings were also held by the various functional specification and requirement teams. Selected audiences were invited to review and comment on the product.

Testing

A decision was agreed between the customer and ISD to have a joint system test for COSMOSS instead of the normal separate customer acceptance test. This was born out of the demanding timescales and the decision also to have an extensive user pilot trial prior to the go-live date. Testing within the development environment followed the normal path of:

- unit testing, against structured test plans, within the development environment;
- integration testing, within a separate environment in which structured link tests could be performed; and
- system testing, within an environment in which structured technical and business tests could be conducted.

Tools

An IBM proprietary product called *TREC* was used to track all system faults and anomalies during testing, and *XMODS* was used to control the migration of software between each test environment and onwards to the production machine located in Glasgow. A proprietary tool from Infoware, *CICS Playback*, was also used to capture test-script execution and to provide evidence of faults to the developers.

System test approach

For each system transaction, a number of business test points were defined; that is, specific business events that the transaction had to meet during testing, whether it be the capture of data or the creation of

data. A list of specific high-level tests was then defined to form the core of the system tests. Each test was then turned into a test design which described the business events and referenced to the test points throughout the document. In parallel to this task, system test data was identified and created for each test design. The test designs were then converted into run test scripts which referenced the test data ready for execution. Cross-reference tables between the test designs and the test points gave visibility of the percentage system coverage. Test execution maps showed the dependencies between tests and the impact of critical faults on system testing.

COSMOSS Security

COSMOSS, because of its usage and information content, has, like other computer systems within BT, been developed with security as a prime consideration. It seeks to meet its security goals by providing mechanisms that will safeguard the confidentiality, integrity and availability of the system and its data.

The requirements for security are defined in the COSMOSS Security Policy Document. This document was produced by a cross-divisional team, including representatives from the system users, D&P ISD, GCS and SecID. A member of the project team provided a full assessment of the perceived threats to the system. Key security policies and countermeasures are detailed. These must be developed as part of the system software or operating procedures required to counter any perceived threats. The policies included in this document provide a system-specific realisation of the corporate guidelines outlined in the Computer Security Manual and the Information Security Code.

Security is provided by a mixture of security features, included as standard within the products and software, that form the key building blocks of a special security subsys-

tem that has been developed by D&P/ISD. This subsystem provides enhanced security for those areas under highest threat or where it is felt that, due to the nature of the information stored, a higher level of security is required.

As a major high-impact system, COSMOSS is passing through the BT system security evaluation and certification scheme. This scheme, based on the international information technology security evaluation criteria (ITSEC) standard, sets out to evaluate how effective the security mechanisms are on the systems. Full evaluation was completed by the end of April 1994.

Security is a key consideration within any computer system. It underpins the way the system is used, operated and managed. As such, it is now becoming a key service/purchasing differentiator between BT and its competitors. In COSMOSS, BT has a security-conscious system that will help win this battle.

Training

COSMOSS is a single national system. It was developed in an environment of flux within BT; brought about by re-patterning of staff and the personnel release schemes to address rebalancing within company guidelines. COSMOSS introduced a user-interface team to tackle the problem.

Policy

The first task was to establish a training policy document. This document considered the problem, the environment, the options to proceed and a recommended solution.

Logistics

Early studies had identified a target population of about 6000 users. The users were geographically dispersed and covered a wide range of skills. The timing of delivery of training was crucial and the time off the job potentially impacted on service to BT's customers.

Implementation

The decision was taken for the Business Division to manage the implementation of training, and BT's own in-house training division designed, built and delivered the initial training courses.

EPSS

The COSMOSS system was designed to be flexible to meet the demands of its users and their customers. Unfortunately, the flexibility of use was a problem for training. However, the EPSS was able to provide a range of solutions to support the main thrust of training. Solutions ranged from on-line help to refresher training and expert scripts (see Innovations).

Support

When training was nearing completion and the COSMOSS system was undergoing the final customer testing, the support organisation was introduced. The role of support was to maintain the availability of the system and to respond to feedback from the users. The development team continued the task of delivering the future releases of COSMOSS.

System Implementation

The COSMOSS system was delivered to the customer in December 1993, three months before the scheduled go-live date. In the interim, the customer carried out customer acceptance testing (CAT). CAT involved full functional testing of the system with realistic data. Additionally, the customer was able to consolidate system hardware, the network and the user profiles.

The system went live in March 1994 as phase 1 of the phased release programme. Future planned releases are scheduled as follows.

- Phase 2 (September 1994):
Generic job management, enhancements to the system functionality and enhancement of

the link into the customer billing system.

- Phase 3 (March 1995): Enhancement of customer account management and the job management facilities. An initial link into the order-handling management system (OHMS) and a common interface to the customer maintenance system.
- Phase 4 (September 1995): Full electronic interface to CSS and the billing system.

Way Forward

Going global

Global customer services is part of the BT corporate vision to enhance customer service for the 1990s and into the next century. Global system strategy will provide an architecture for the future which includes single customer and site databases, a database for the management of products and services and a single customer billing initiative. The COSMOSS system supports the global initiative by providing an important and strategically-positioned back-end system.

Decision support

The EPSS has introduced a new era in user support. This new generation of user help is decision support, which recognises the system user as a performer. The EPSS for COSMOSS provides analysis of a user help request and offers a range of solution options to suit the needs of the requester. Into the future, the EPSS can support client/server/mainframe architecture. This support takes the form of information retrieval from a wide range of platforms over generic interfaces. The user can summon decision support from virtually any source, coupled with document retrieval, training and a wide range of multi-media products such as digital audio visual, spreadsheet or word processing.

Conclusion

The COSMOSS system was developed as a key enabler to meet the business needs of the company. The programme was, and continues to be, a partnership between the customers, their agent and suppliers, towards a common set of goals. The project adopted an end-to-end approach that was resilient to change and adaptable to the opportunities that emerged. The COSMOSS system has achieved its short-term aims and will evolve towards its full scope. Furthermore, COSMOSS provides an important component of the BT global system architecture.

Acknowledgements

The author would like to thank the COSMOSS development team and their associates for help and support in producing this article. Special thanks are extended to those colleagues who contributed directly: Alan Barnard, Seamus Reilly, Nigel Sykes, Iain MacMillan, Richard Brechin and John Greenstreet.

References

- 1 WILSON, STEPHEN. Service Management Systems: Supporting the Customer Interface. *Br. Telecommun. Eng.*, 13, April 1994 (this issue).
- 2 RUMBAUGH Object-oriented Modelling and Design. Prentice Hall, 1991.
- 3 COAD and YOURDAN Object-Oriented Analysis. Yourdan Press, 1990.

Glossary

CAMSS Computer-aided maintenance of special services
CAPSS Computer-automated private service system
CASE Computer-aided software engineering
CAT Customer acceptance testing
CBR Computer-based referencing

CBT Computer-based training
CICS Customer information control system
COSMOSS Customer-orientated system for the management of special services
CSS Customer service system
D&P BT Development and Procurement
EDA External data access
EPSS Employee performance support system
GCS Group Computing Service
ISD Information Systems Development
MIS Management information system
OHMS Order handling management system
PC Personal computer
PSH Product and services handler
SSADM Structured system analysis and design method
XMODS Transfer modifications

Biography



Henri Van der Stighelen
BT Business Communications

Henri Van der Stighelen was a member of the COSMOSS development team and has now moved into the future systems development team. He joined the COSMOSS team, in 1991, from the Quality Systems and Support Division of BT. Henri joined BT as a Youth in the Engineering Division and has experienced most aspects of telecommunications engineering. In 1984, he returned to studying and, focusing on mathematics and technology, he found an interest in computing. This culminated in a study of relational-database theory at university. Henri is an ideas person and brought to the team an enquiring mind that led to the introduction of the first employee performance support system for BT.

Event and Test Management System

BT's event and test management system (ETMS) is a major enabler in the quality of service and productivity of private services repair. This article outlines the background to the development of the system, and the practical implementation of cooperative network management architecture.

Background

In the drive to improve quality of service (in terms of time to repair) and productivity, the benefits of remote test management of telecommunications services have long been recognised. Remote testing allows a service to be tested and the fault localised, without incurring delay and cost of engineering resource having to visit remote exchanges to carry out testing. For technical reasons, private service testing in BT has evolved differently from the public switched telephone network (PSTN) testing, and is currently realised using the remote access and test equipment system (RATES) for analogue services and the digital access and remote test system (DARTS) for KiloStream services.

In addition to remote testing, digital services can detect problems through alarms generated by the managed private circuit digital network and MegaStream alarm schemes. This enables the maintenance engineer to manage problems as soon as they occur, rather than, in the case of testing, waiting for a customer to report a fault. KiloStream fault diagnosis is achieved by using a combination of alarm reports and remote testing via DARTS, while MegaStream currently relies solely on alarm reports.

History

From early experience of managing analogue and digital services over the London fibre access system (FAS), it became evident that the testing of digital and analogue services delivered over a switched fibre network technology using existing test systems (RATES, DARTS and alarm

analysers) created the following problems:

- up to five systems required for each diagnostic maintenance operation,
- complex operational procedures,
- slow repair times,
- high levels of skill to operate, and
- high cost.

Consequently, the cost and service delivery benefits of modern technology were somewhat compromised by the complexity of service management. With the development of new network technologies such as synchronous digital hierarchy (SDH) and fibre access network (FAN), coupled with new more complex private service products, a new approach to private circuit service management was required.

As part of the RATES programme, in 1991, BT carried out a study to propose new processes, procedures and systems to support analogue, digital and hybrid private services into the late 1990s.

The study carried out by BT Development and Procurement Human Factors, Network Management Department (NMD) and Management Science Consultancy Unit (MSCU) concluded the following:

- *Private service repair should be customer focused.* This would require processes and systems to support a customer-focused service management role, working into geographically focused network support organisations.

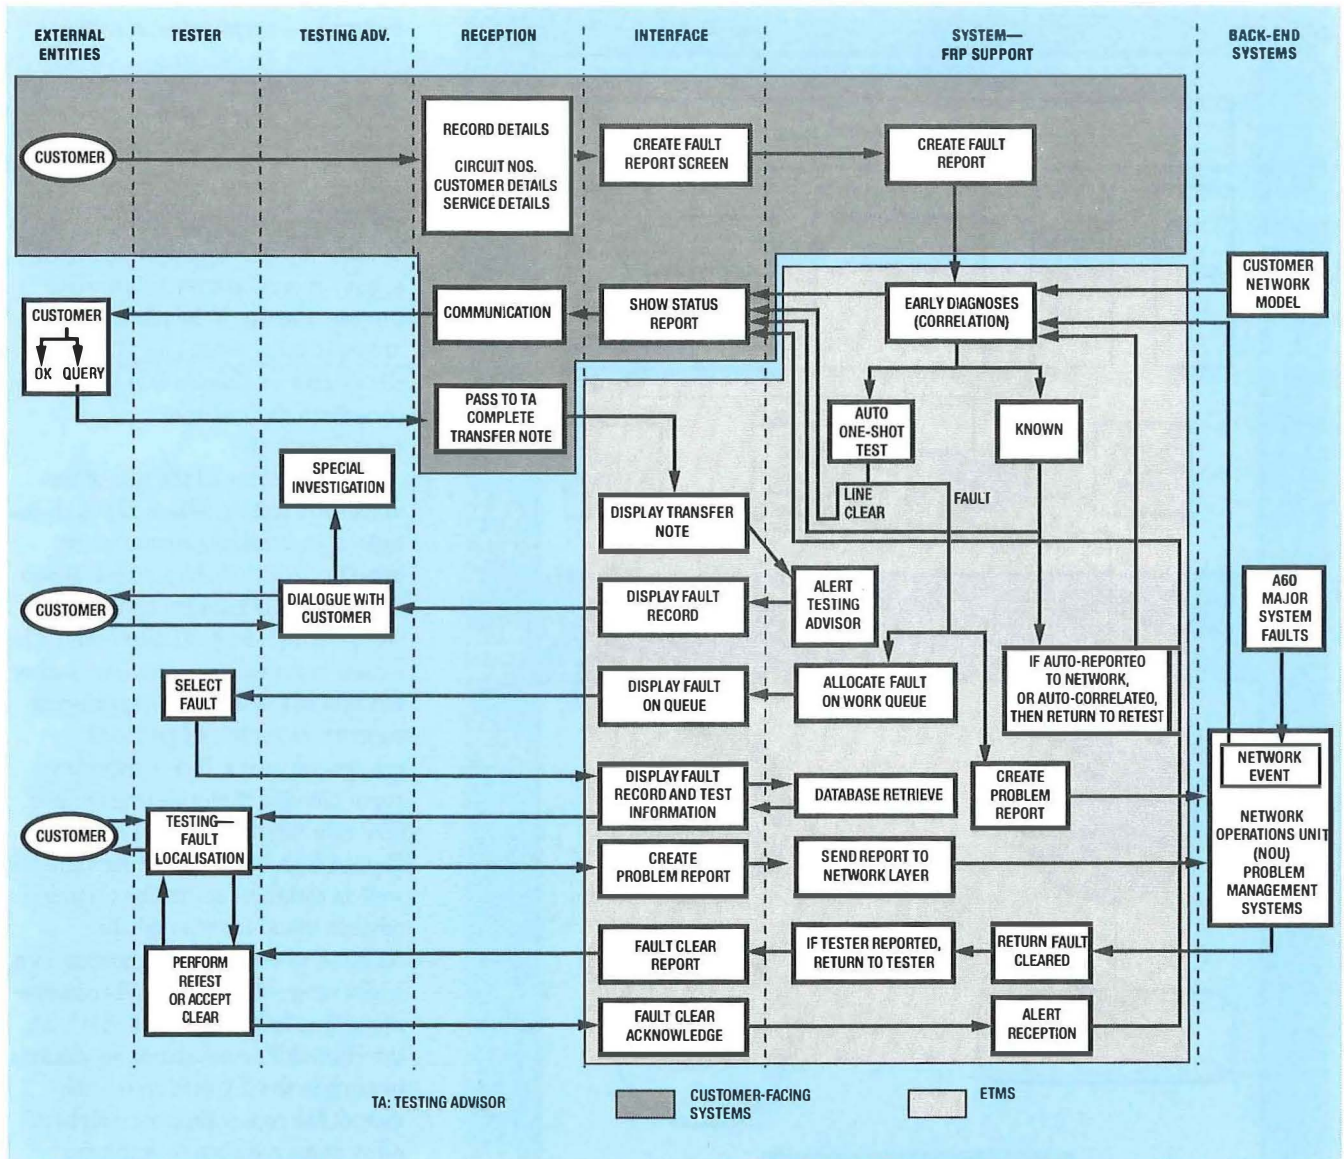


Figure 1—Task allocation chart

- *Integration of systems conforming to the cooperative network architecture (CNA).* Both the network support systems architecture (NSSA) and integrated service management (ISM) architecture conformed to the CNA management (CNA-M) architecture. This was therefore a fundamental requirement
- *Automation of repetitive processes.* During the study period, most manual processes (that is, data analysis, test, diagnostics and work distribution) were found to be repetitive and simple in nature; automation of these processes would improve process times and allow the human operator to focus on more complex diagnostic operations and customer interactive tasks. The MSCU identified

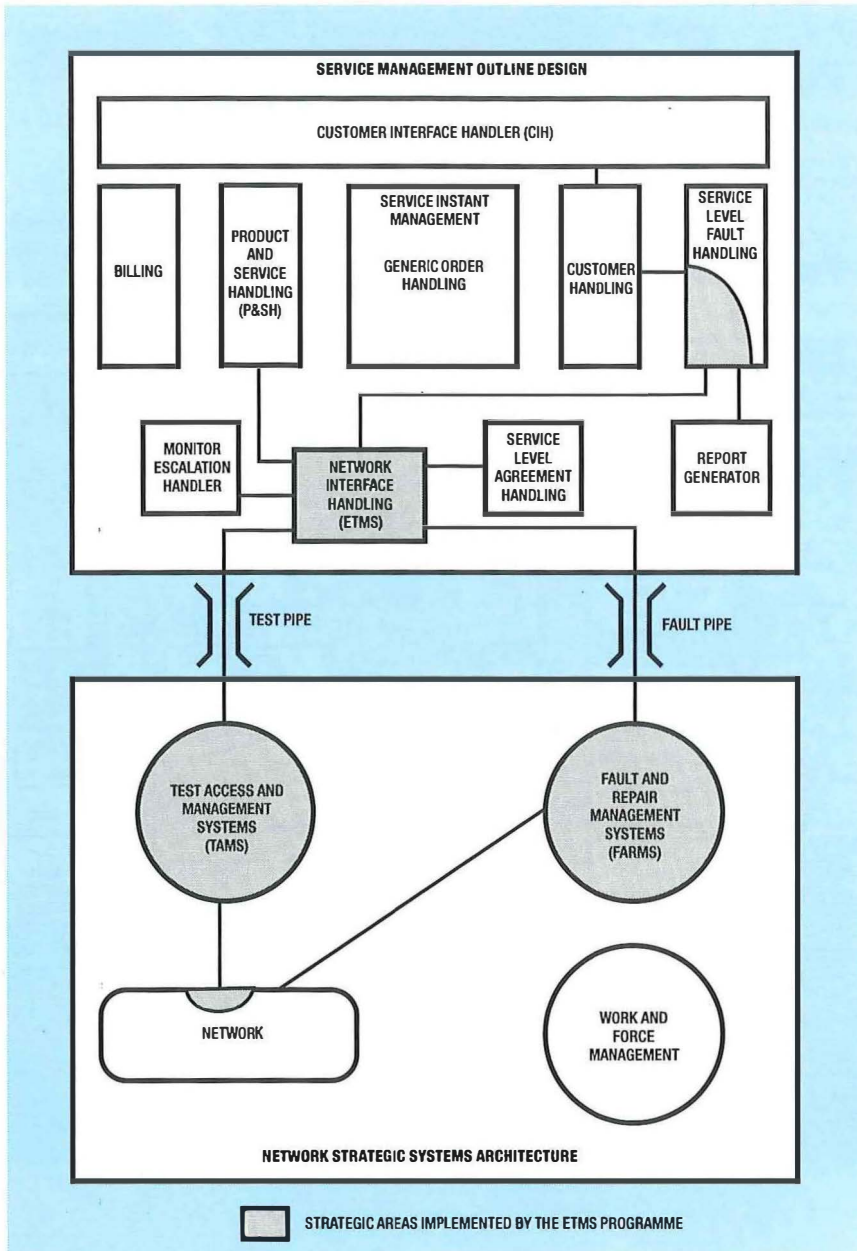
that significant tangible benefits would be achievable due to this function alone. They also forecast total fault volumes and where manual operation of the processes would be required. This identified the need to provide 230 work positions to perform technical service management functions; that is, to provide a technical interface to customers when required.

- *Provide the human operator with an interesting and rewarding role.* Study of existing processes showed that owing to the volume of reports and the skills required, the work distribution (controlled by the computer assisted maintenance of specialised services (CAMSS) system) had isolated testing officers to carrying out only a small part of the overall repair

process. This had taken away the more rewarding elements of this work—seeing the end result. It was evident that, in areas where fault volumes were lower and individuals had more responsibility for the end-to-end repair process, productivity and quality of repair were much higher.

- *Maximum reuse of existing systems.* Existing systems used for private circuit repair (CAMSS, RATES, DARTS, RENACE, GENETIC, CSS and INS) represented a considerable existing capital investment. Furthermore, during development of any new systems, customer services managed within the existing systems must be maintained and improved in line with quality-of-service targets.

Figure 2—Functional architecture



identify if a fault existed and, if so, allocate it to the appropriate repair agency or network reconfiguration process. Finally, if the fault had been automatically tested and allocated, the system would also carry out a test to confirm the repair or reconfiguration.

Having reviewed the role of the diagnostic testing officer (DTO) in the light of increasing automation, it became clear that the number of non-complex faults handled by the DTO would reduce. Instead, there would be a need to provide the customer-facing receptionist with direct engineering support for complex faults and customer queries. It was therefore recommended that a testing advisor (TA) role should replace the DTO. Having high customer-facing skills as well as technical skills, the testing advisor would increasingly be required to carry out diagnostics on a wider range of network and customer apparatus technologies. In line with the Human Factors recommendations to provide the TA with an overall view of the repair process, a sight of all systems required to manage a fault from end to end would be needed. Human Factors recommended a single screen with various terminal emulations be considered as a short-term solution only, the long-term solution being to provide a single logical system. This would provide the TA with a single set of screens designed to support the business process. The information (data) to build the screens would be selected from various systems identified as a particular 'data owner', such as CAMSS, CSS, RATES, GENETIC, INS, DARTS, RENACE and the alarm analyser, allowing the provision of consistent information.

Logical Architecture

The vision process was mapped against the existing ISM and NSSA functional architectures as shown in Figure 2 and then a target logical systems architecture was developed.

Process Development

Taking the key recommendations into consideration, a 'vision' process model was developed in the form of a task allocation chart shown in Figure 1. The major change to the existing fault management procedures was the use of event management systems to correlate automatically the customer fault report with known network 'events', prior to implementing a test. This course was recommended as private services fault diagnosis can be highly complex and requires testing of services with typically two or more circuit ends. Owing to intersystems networking, older technology and the need to monitor line protocols for specific periods of time to establish performance, for example, G.821, a single

service test could take up to 11 minutes to complete, whereas a query of an event database takes only seconds.

This would be an ideal situation if event management could support all services; however, a large element of BT's private services are routed over non-managed analogue networks (80% in 1993 reducing to 40% in 1998), which provide no event management facilities. Therefore, in the case of analogue technology based services, a 'one-shot' test was required to identify the status of the customer service. Where event and test diagnostics either failed to confirm the service status, or the customer did not agree with a right-when-tested (RWT) response, the problem had to be managed manually. Automatic testing would attempt to

Figure 3—ETMS logical architecture

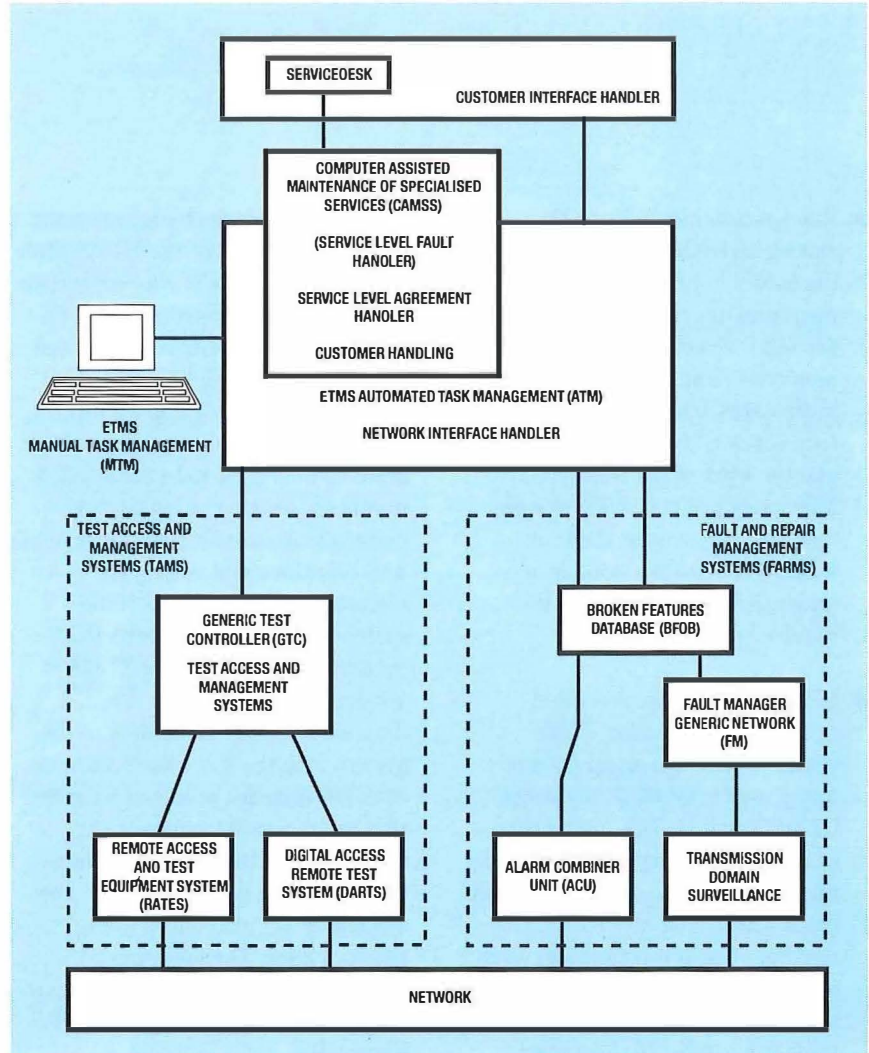
To meet the requirements of the 'vision process' within ISM and NSSA, the areas affected by the processes were studied. Considerable work within NSSA had defined logical test and event management systems, these being the generic test controller (GTC) as defined by the Test Access and Management Systems Evolution Study¹ and the broken features database (BFDB) defined by the Fault and Repair Management Evolution Study. Within the service management outline design, no logical system had yet been defined capable of supporting the business process. This was identified as a major area of development for the ETMS project. An outline design and a data framework conforming to CNA-M were therefore required.

The target logical architecture as shown in Figure 3 was developed by using existing systems and targeted developments in ISM and NSSA. However, a key functional area, automated management, did not exist within CAMSS as it did within CSS. To provide this, ETMS would provide CAMSS with a 'surround' system, providing the necessary network interface handling through the generic test and fault pipes identified at the network and service management layer interfaces.

Two major logical functions of the ETMS were identified to support this:

- automatic task management (ATM), automating the business process; and
- manual task management (MTM), providing the TA user interface to a single logical set of applications.

From the 'vision' processes and studies by the MSCU, a customer requirements definition (CRD) was produced by BT Business Communications Customer Network Operations (BC CNO) defining the business requirements to implement the vision process and the products supported. The key focus of the resulting specification of requirements was:



- an implementation plan to provide early business benefits, supporting existing products and services, with an enhancement programme to support the vision process by 1995;
- integration of support systems functions used for the end-to-end repair process as if a single virtual system but support the CNA-M architecture wherever possible; and
- pilot trials to ensure functionality would deliver quality-of-service and cost benefits forecast by the MSCU.

System Development

The requirements-capture study enabled the ETMS development team to scope in sufficient detail the definition of the system boundaries and the definition of requirements on external systems. The CRD dictated the viability of implementation

options and placed considerable constraints on the system design.

Before proceeding into detailed system design the following options were reviewed:

- *Could existing systems be enhanced?* CAMSS had been identified from the study as a key enabler and already the majority of the process was managed by CAMSS. However, CAMSS had already reached its maximum capacity and would require a major rework and larger main-frame capacity, which at the time was not available. CAMSS was also a single system and as such integrating access to the test systems would incur high communications costs and high complexity within the application. The functionality of CAMSS at the time was also under review as it provided service and network layer applications along with a service database. This option was rejected.

- *Were systems available on the market?* BT(NMD) was accepted as the leader in network management and had been developing products which were driving architecture and standards work in this area. This was particularly true of the CONCERT initiative and the work of the Network Management Forum. There were no known systems at the time which were either available or under development to offer the required functionality.
- *Single large system.* As stated above, the solution for ETMS would require the integration of many remote systems to a single CAMSS system. This option was considered to be over complex and impractical due to communications costs. This would also reduce the number of platform suppliers as it would need to be a mainframe installation, which would be a large capital cost to the project.
- *Distributed system.* The number of test systems that would require interfacing to was likely to be in excess of 20; this was further complicated by the number and diversity of physical locations. It was therefore considered practical to opt for an intermediate stage between CAMSS and the test systems. This intermediate stage would also deliver the need for local applications and facilitate the integration of terminal access to other support systems. The decision was made that the total requirement for ETMS would be distributed across a number of systems. This did not explicitly imply the need for distributed technology, but the flexibility offered by such technology was seen as a major benefit to the development and deployment of the project.

System context

From the task allocation charts and the logical architectures, a Yourdon

essential model was built to support the vision process for the ETMS. Full delivery of this model was impractical as the detailed requirements of the system were not sufficiently defined and the dependencies on other key initiatives would not allow a full definition and implementation of the systems interfaces to be achieved. A number of pragmatic builds were defined to ensure that the programme could continue and deliver the understood requirements of the system early. Figure 4 shows the context diagram from the ETMS Version 2.1 logical model. The diagram shows the boundary of the system with the key interfaces to its external systems, people and devices that are required to support the ETMS. It also highlights the transition of some of the interfaces to new systems being introduced to the network layer. The following describes the external systems that are required to support ETMS (refer to Figure 4).

- CAMSS is the system for receiving customer faults, retaining circuit data, fault tracking, fault progressing and fault despatching. It provides the necessary initial circuit data to allow the ETMS to decide which task and test scripts to operate.
- The user is shown to describe the manual task management facilities, which include the viewing and monitoring of alarm information and performing diagnostic test functions. External printer functions are also shown.
- The key network fault pipes are shown as the alarm combiner unit (ACU) and the broken features database (BFDB). The strategy for improving network fault reporting includes increased correlation and filtering of network events and the mapping of these events to network features which support customer services. The interface to the correlated information will be

via a single broken features management system; this will remove the need to support alarm reporting from the ACU. Ultimately, the ACU interface will be removed; however, for ETMS 2.1 the BFDB will only support Megastream services. The test pipe exists within ETMS as the GTC will not be deployed until a later build; however, the GTC functionality exists and is within ETMS.

- Fibre in the access network is now being deployed and private services must be supported. The telecommunications passive optical network element manager (TEM) is shown and provides the ability to loop back and reconfigure for digital private circuit testing.
- RATES is the analogue private circuit testing system and already provides end-to-end testing across a number of interconnected Hewlett Packard computers. The proposal to use cross-network technologies to support new services implies that, despite the off-load of analogue services to digital, RATES equipment will still be required to support local end testing of copper tails.
- RENACE manages the switching of digital private circuits and is used by ETMS to switch a faulty private circuit (PC) to DARTS test equipment. RENACE is also interrogated for link status data, which is used in the diagnostic algorithms.
- Once access is given to suitable test equipment via RENACE, DARTS is used to control loop-back and error monitoring to detect and localise faults on the circuit.
- Privilege user or NetView/6000 is shown as the system management functions required to operationally support ETMS itself.

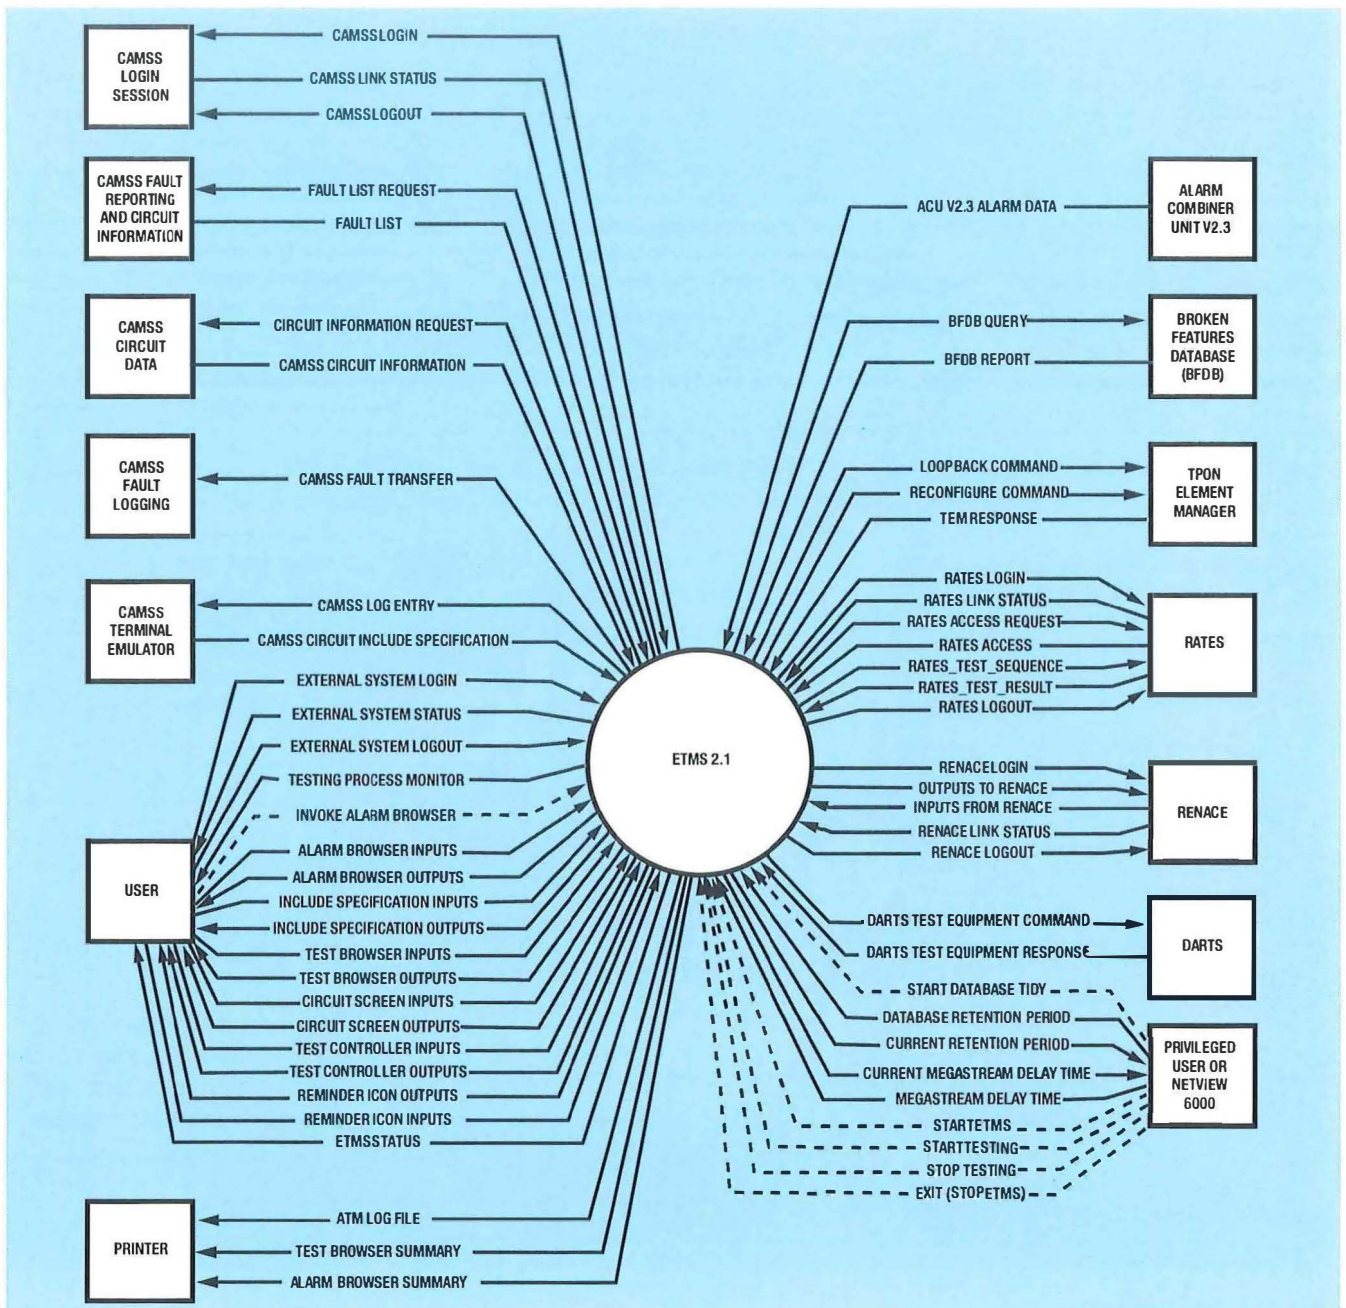


Figure 4—ETMS context diagram

Figure 5 shows a further decomposition of ETMS. Five functional groups are shown which summarise the main processing performed by ETMS.

Manage ETMS

This process has two main functions:

- system start-up and shutdown and the user interface to start and stop automatic testing by ETMS; and
- the facility to login to the external systems CAMSS, RENACE and RATES to allow automatic testing to begin. These logins are maintained for the duration of automatic testing.

Manage events

This process takes alarm events from the ACU and populates the local alarm database after validating and formatting the arriving data. This local database allows a fast look-up of more recent network alarms. Live alarms may be viewed by an end user.

Manage automatic tasks

This process takes faults off the CAMSS gateway queue when activated by the manage ETMS process. Each fault entry is analysed to determine if the faulty service is supported for automatic testing by ETMS. If the service is supported, then requests are sent to RATES and RENACE to obtain circuit routing,

access technology and test access information. This information enables the service test class to be identified and a test request is sent to the manage testing process.

In due course, the *manage testing* process responds with a test result which determines how the fault should be processed and onto which CAMSS queue the fault should be placed.

Manage manual tasks

This process, essentially, performs the same task as the *manage automatic tasks* process, but in an interactive manner, with the tester dictating which tests are run and in which order. The tester achieves this by

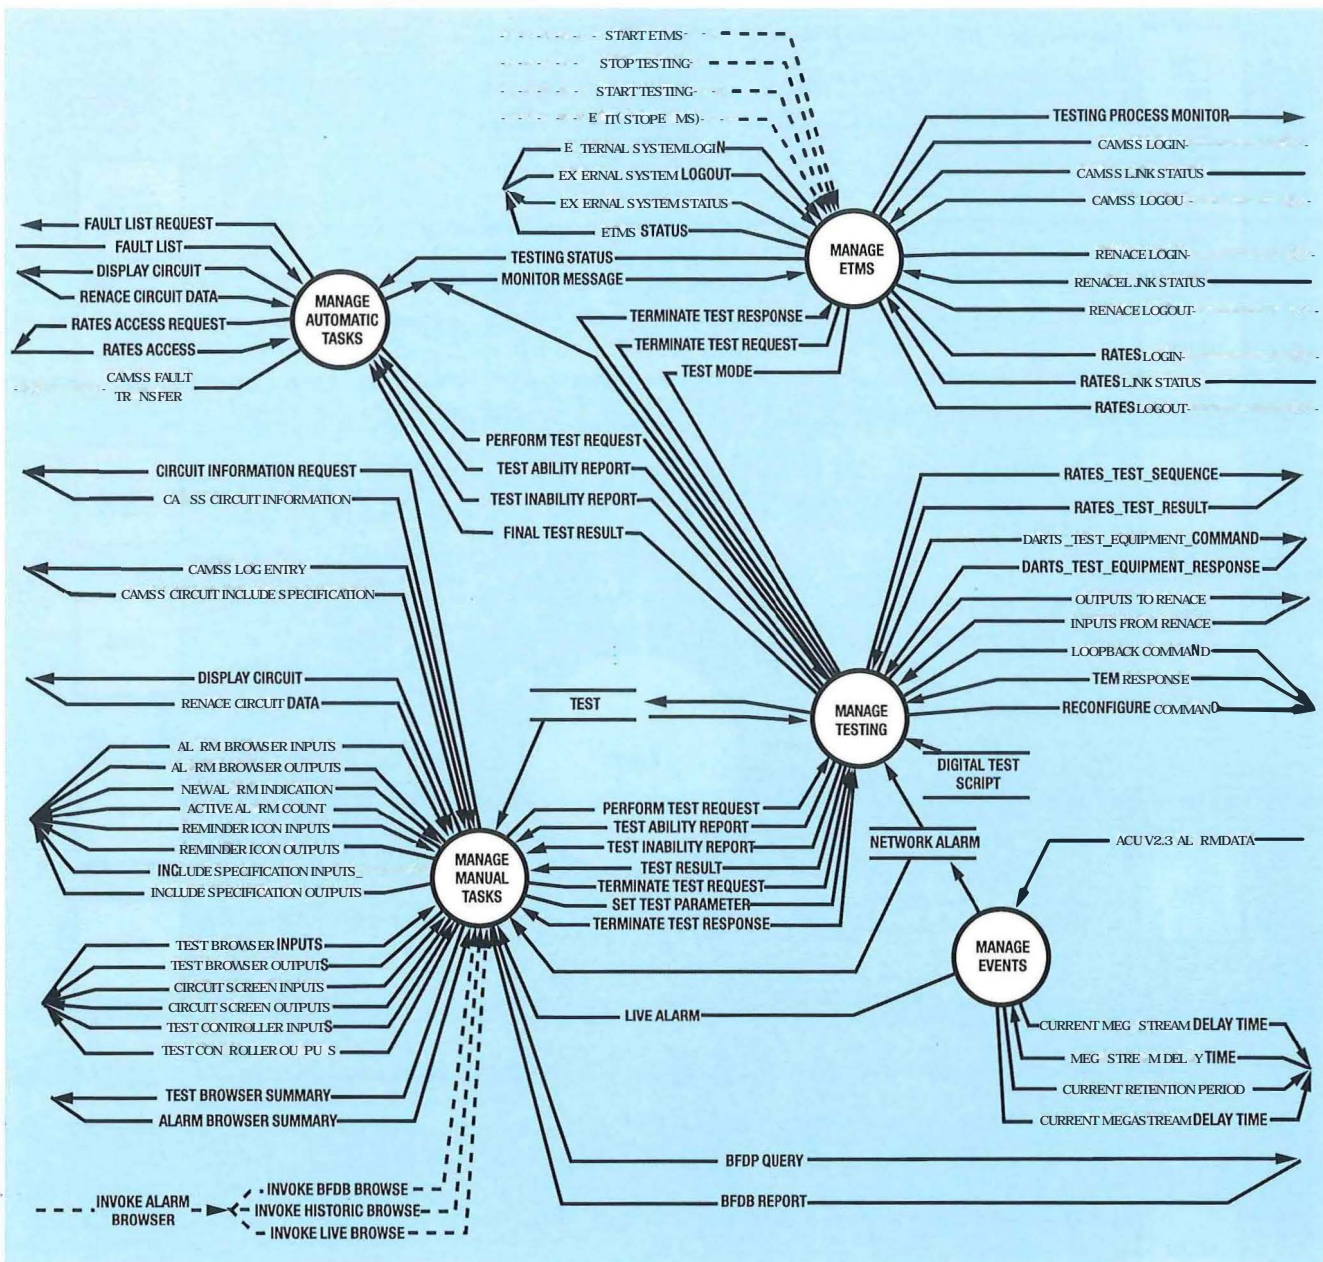


Figure 5—Main ETMS processing

using windows-based graphical user interface (GUI) tools to examine alarms, initiate tests and track test progress. This process also causes tests to be executed by sending test requests to the *manage testing* process.

Manage testing

This process is the 'testing engine' of the ETMS; it is activated by a *test request* message. The test class determined by *manage automatic tasks* is contained in the *test request* message and enables *manage testing* to determine which type of test sequence (test script) to run. Various test sequences are defined by test scripts, which are instructions in a tailor-made language interpreted by the generic test controller (GTC) which is the core

of *manage testing*. The GTC controls the sequence of alarm checking, circuit splitting and end-to-end testing by firing off 'child' processes according to the test script instructions. The GTC functionality was defined by the NCAB Test and Access Management System Evolution Study¹ and was incorporated into the ETMS until, at a later time, this functionality could be removed from the ETMS and a true test pipe implemented at the network layer. This will be implemented under the UK access plan (UKAP) over the next few years.

Development Strategy

The development strategy consisted of two major strands of work: the

integration of existing functions at the user terminal and the automation of existing manual tasks. The system was designed to support both; however, as workstation technology could already deliver a considerable level of integration through terminal emulations and local applications, these were considered as deployment-specific rather than mainstream application development. Ultimately, a single logical application would be developed to replace the terminal integration.

The development of the automatic function required that the user was the fall-back when automatic tasks had reached their capability. It was well understood from the beginning that ETMS could not fully automate

Figure 6—Typical trial configuration work position

the process, but considerable parts of the process could be automated, and the potential existed to extend this capability in the future once more experience had been gained in using the ETMS in the live environment.

This strategy offered several benefits to the development of the ETMS in that sufficient computing power would exist for both manual and automatic functions and hence a transfer from manual to automatic tasks would not require increased computing power. The manual functions were always available as fall-back during system failures, and local requirements could be supported to ensure that local variations in operation and organisation could be accommodated.

The development was to be UNIX-based and software would be written in C++ to facilitate the reuse of existing code developed in other programmes. This was a requirement placed on the project by the architecture boards and was seen as a key enabler for competitive tender and the beginning of a distributed platform for ease of extendibility and integration of future applications.

Implementation Plan

It was agreed that initial developments would be focused on trials to prove the business benefits and confirm development proposals. The trials would also be used to further capture detailed requirements for the automation process and the user interface.

The ETMS implementation plan was sub-divided into four major builds:

ETMS 1

A trial system to be used in several operational scenarios to prove that the business benefits were achievable and to finalise the system design. This was also the initial hardware deployment for desktop integration.

ETMS 2

National deployment of the ETMS platform and the introduction of automatic task management and



additional manual task management (MTM) functions. The Business Communications Technical Centre programme would be supported by this build and the introduction of the BFDB.

ETMS 3

Major upgrade to support new services and network technologies. New interfaces to emerging network and service layer systems to improve the quality of alarm information and the introduction of a network GTC.

System-to-system interfaces will be introduced to improve speed and to reduce the impact of individual system changes.

ETMS 4

Final enhancements to ensure benefits are achieved and to complete remaining features that were not implemented due to dependencies on other systems.

Trial System Architecture

The trial system was developed on a SUN platform. The configuration consisted of a single SPARC2 server and four IPX workstations. The desktop integration was achieved by the use of the network management workstation (NMW 2)² functionality with terminal emulators for CAMSS, RATES and RENACE access. Software was developed to perform similar functions to the existing personal computer based alarm analyser functionality to view live alarm information.

The most important aspect of the trial was the prototyping of the automatic task management. The focus for this was analogue PC testing as this only required interfaces to CAMSS and RATES. Early on in the ETMS programme, work was commissioned to provide application interfaces to these systems. The implementation of these interfaces was achieved by simulating user interaction to terminal emulators.

Trials of the automatic task management functions took place in Glasgow and London, and manual task management functions were trialled at Milton Keynes, Birmingham and Edinburgh. Figure 6 shows a typical trial configuration work position.

Trial results

The trials produced varied reports; the desktop integration was well received, but the alarm viewing facilities were not sufficiently well defined to support immediate user need. The automatic task management aspects of the trials proved that the technology could support the required functions; however, the interface to CAMSS and RATES was sensitive to minor changes of the remote systems and their configuration. The platform and system configuration proved to be unreliable and difficult to maintain and consequently many system problems reduced the effectiveness of the trials. However, from the results obtained, approval was gained to continue with the programme. The trial feedback

Figure 7—An ATM site configuration

and findings were used to improve the design and resilience of the deployed system and, in particular, suppliers were asked to provide a total systems management solution for the complete national ETMS.

Deployment System

Competitive tender

A competitive tender was issued for the procurement and deployment of hardware to support the ETMS requirement. The benefits of choosing an open platform were very apparent at this stage; the competition was high and IBM was awarded the contract. The hardware would be the IBM RS/6000-series of machines on the AIX operating system. The impact of changing platform at this stage of the programme was significant but the cost advantage compensated for any delays in the programme. The benefit of choosing UNIX was now a real business saving and justified adopting standards and open systems.

Deployment design

It was agreed that ETMS would be deployed in two forms as manual task management (MTM) and automatic task management (ATM). Both would be capable of supporting automatic and manual functions but availability and resilience would differ substantially.

ATM would be located at two sites each of which would contain three servers capable of being run independently; one of the servers would function as a stand-by system. Figure 7 shows the proposed configuration.

The MTM would comprise two servers, each of which could support all ETMS functions and act as the file server for the workstations. In the case of failure, any one of the servers could support the total workstation requirement per site and, if the ATMs were unavailable, MTM servers could be used as partial ATM (see Figure 8).

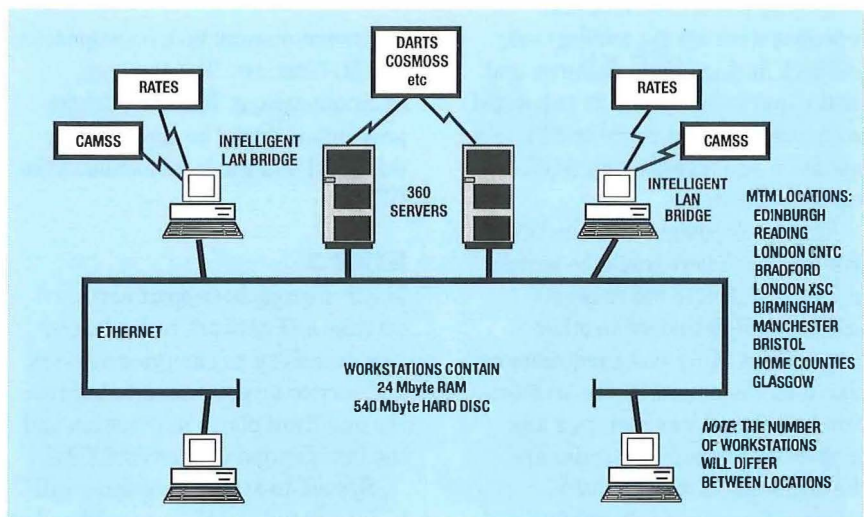
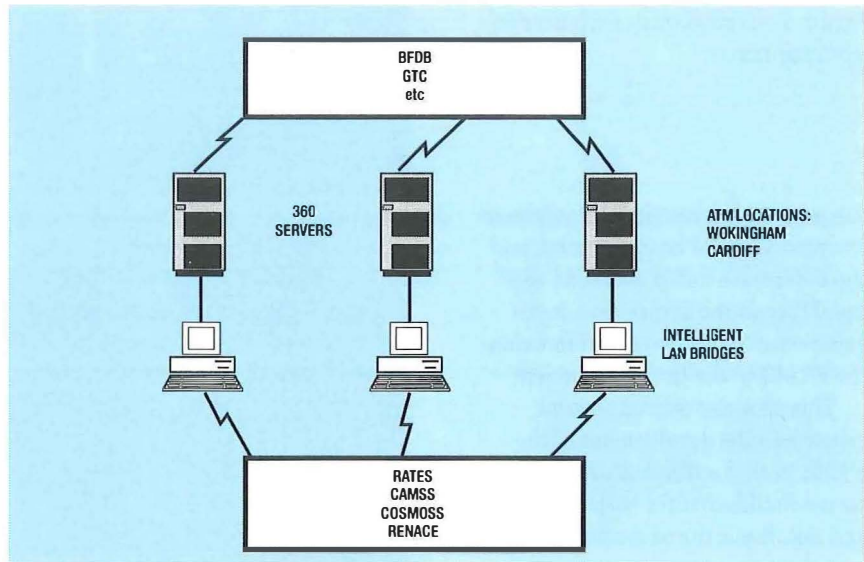
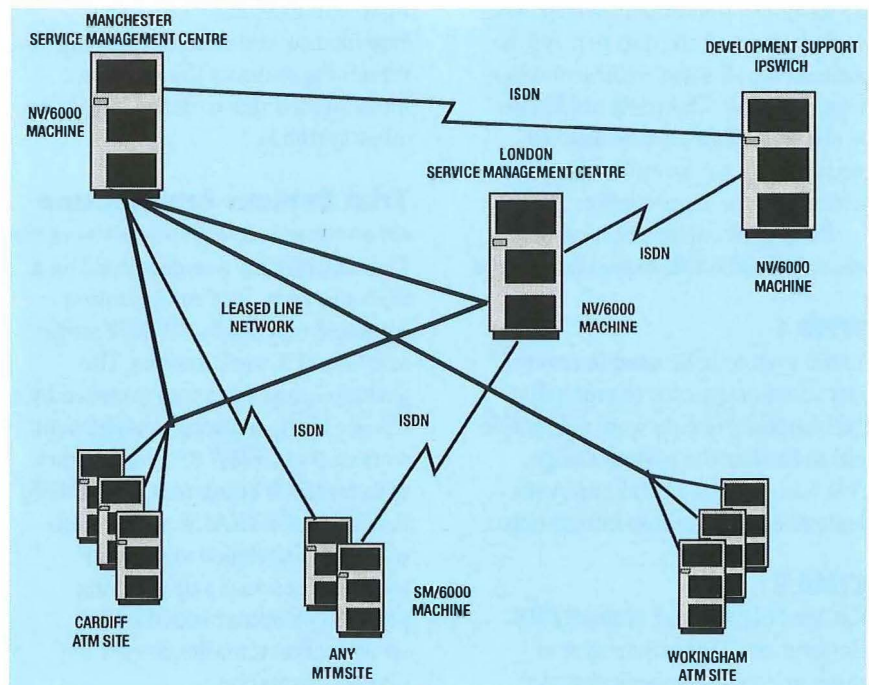


Figure 8—An MTM site configuration

Figure 9—ETMS managed network



The system management of all ETMS machines and applications would be managed remotely from a central management site using NetView/6000. A system model would be provided at BT Laboratories for support of applications and as a final disaster recovery point. Figure 9 shows the likely configuration and network of the system management solution.

Key deliverables

As already stated, ETMS will be delivered in four major builds, however each build has specific versions which are targeted at deployment and interface milestones. Table 1 summarises the key versions and the year in which they were achieved or have been proposed.

Table 1 Key ETMS Versions

ETMS 1.0	NMW 2 implementation	2Q/92
ETMS 1.1G	NMW 2 with test manager manual test functions and alarm browser facilities	3Q/93
ETMS 1.2	Automatic task functions. Extensive trials of analogue testing	4Q/93
ETMS 1.3	First deployment of MTM hardware (IBM) including new communications software (CITYNET)	1Q/94
ETMS 1.4	MTM site supporting ISDN services requiring local application support	2Q/94
ETMS 2.0 Phase I	Automatic testing of analogue private services with automatic alarm checking for digital services deployed on existing SUN hardware	1Q/94
ETMS 2.0 Phase II	Automatic testing of analogue and digital services and upgrade to coincide with major CAMSS upgrade. Still on SUN platform	2Q/94
ETMS 2.1 Phase I	First deployment of ATM on IBM hardware	2Q/94
ETMS 2.1 Phase II	MTM and ATM deployed with system management functions additional services types supported, TPN trial supported and MTM user tools providing the first release of the single logical vertical application	3Q/94
ETMS 3.0	Extension of GTC functions to support new network technologies and new service types defined by the private network services plan	2Q/95
ETMS 4	Programme complete	4Q/95

Future

Confidence is high that ETMS will deliver the forecast business benefits and that the system will provide opportunities to refine the repair process further by continuous improvements in the network alarm quality through programmes such as transmission domain surveillance, and by the ability to introduce new diagnostic capability through the use of the GTC and the GTI.

Already, many other potential business savings have been identified. The ETMS programme will provide BC with an early basic integrated platform that will bridge the gap from existing technology to the vision of an integrated platform with the ability to extend and

integrate new applications and facilitate the rapid deployment of new products and services.

ETMS is also seen as a key enabler to new programmes such as Advanced Fault Management and Business Intelligent Networks.

Project Management

The project has required that many parallel activities be progressed to ensure that the programme could deliver over a relatively short period, which over three years included capturing requirements, analysis and further refinement of the requirements, trial of suitable options and the development and deployment of the total requirement.

The project management team encouraged product prototyping and supported trials even when subject to a high risk of failure. It was apparent from the start of the project that full system test facilities would be impossible to create because of the unknown nature of the faults and the diversity of circuit types and configuration. Therefore, full system testing could only be achieved by controlled tests with live circuit problems. The project team has been multidisciplined and resourced from across BT and external companies. CNO, NMD GSEC, Human Factors, SUN, IBM and City Computing have all contributed significantly to the programme.

The complexity of ETMS and the number of external factors that could impact change to the programme has required that a high level of flexibility in scope and delivery be assumed. Parallel development techniques have to some extent achieved this, but ultimately the impact of changing an external system interface has been, and will continue to be a significant challenge to the project team. New processes are now being developed to manage and control deployment of internal and external system changes; that is, CAMSS, RATES, etc. The future introduction of standard interfaces will reduce the impact of

change and eventually lead to support cost savings.

The size of the deployment has also required a change to existing deployment practices, with IBM making a considerable contribution to this process. Good teamwork between BT and IBM has ensured minimum disruption to operational units and the successful introduction of the ETMS.

References


- 1 REED, C. F. Network Test Management Systems. *Br. Telecommun. Eng.*, Br. Telecommun. Eng., Oct. 1991, 10, p. 239.
- 2 VICKERS, A. D., and TRAYNOR, V. Network Management Workstation. *Br. Telecommun. Eng.*, Oct. 1991, 10, p. 254.

Biographies



Tim Young
BT Development and
Procurement

Tim Young joined the British Post Office in 1973. He worked in the Canterbury Telephone Area on exchange maintenance duties until transferring to BT Development and Procurement (D&P) in 1983 after receiving a B.Sc. Honours degree in Electronic Engineering from the University of Kent. He is currently technical group leader of the Test Management Applications development group in the Network Management Department of D&P and has been the system technical consultant and product manager of the ETMS and GTC products for the last 3 years.



Bob Holland
BT Business
Communications

Bob Holland joined the British Post Office in 1971 as a Trainee Technician Apprentice. He worked in London North Area on transmission and specialised services customer support duties. In 1985 after promotion to Assistant Executive Engineer, he took over responsibility for the day-to-day operation of CAMSS, and in 1988 moved to the RATES project team where he was responsible for the development of RATES. He is presently the diagnostic systems development manager for BC Customer Network Operations and is responsible for the development of the ETMS programme.

Glossary

ACU Alarm combiner unit
ATM Automated task management
BC Business Communications
BFDB Broken features database
CAMSS Computer assisted maintenance of specialised services
CNA-M Cooperative network architecture (management)
CNO Customer Network Operations
COSMOSS Customer orientated system for the management of special services
CRD Customer requirements definition
DARTS Digital access remote test system
DTO Diagnostic testing officer
ETMS Event and test management system
FAN Fibre access network
FAS Flexible access system
FM Fault manager generic network
GENETIC Generic network implementation and control
GSEC Glasgow Software Engineering Centre
GTC Generic test controller
GTI Generic test interface
INS Integrated network system
ISM Integrated service management
NCAB Network Control Architecture Board
NMD Network Management Department
MSCU Management Sciences Consultancy Unit
MTM Manual task management
NSSA Network strategic systems architecture
PSTN Public switched telephone network
RATES Remote access and test equipment system
RENACE Remote equipment network for ACE (automatic control equipment)
SDH Synchronous digital hierarchy
TA Testing advisor
TDS Transmission domain surveillance

Electronic Trading: BT's Commitment to Changing the Way We All Do Business

*When BT acquired Tymnet Inc. in 1989 it was able to enter the highly competitive UK electronic data interchange (EDI) market with the EDI*Net clearing house service. In an article in the Journal in January 1991, the basics of EDI and EDI*Net were reviewed. This latest review explores in more detail the essential elements that contribute to successful electronic trading schemes.*

Introduction

This article was originally conceived as an update to an article written in 1990 and published in the January 1991 issue of *British Telecommunications Engineering*¹.

Written shortly after BT's 1989 Tymnet acquisition, at a time of an intense internal debate to define the company's mission and essential objectives, the paper alluded to some of the strategic and marketing issues arising from new telecommunications capabilities. Electronic Trading—specifically the message switching service EDI*Net and the historical development of message standardisation—was described in some detail. But the article also described aspects of 'service delivery'—how it was being brought to market—and the wider implications for BT. This new review, if it is to update adequately the earlier, must therefore address the wider context and relate the service offerings to today's customer needs and BT's current strategic focus.

Definition

More than three years on, the first point to make is that we must start by restating what we mean by electronic trading and electronic data interchange (EDI); these terms are still not fully understood, or in everyday use, even among a technologically literate audience.

Electronic trading (or the more-recent North American term, *electronic commerce*) embraces a very wide range of networked computing and

messaging applications that are applied to business practice. The term can include facsimile, computer file transfer, electronic mail and EDI. This broad definition of electronic trading can be applied to any form of message or information flow and is even used by some commentators to include voice messages and tone-operated response systems such as those provided by banks for customer account enquiries.

A more-useful, narrow, definition limits the term to 'structured' messaging, thus excluding all manner of unstructured text such as electronic mail person-to-person memos and most fax messages. Structured messages—for example, an invoice, a purchase order, a bank payment instruction, the result of a blood test, or even an invitation to tender—can be designed to conform to an agreed message standard and, in that form, can then be exchanged directly between independent computer systems without manual intervention. Hence the formal definition of EDI as '*the computer-to-computer exchange of structured business data in agreed ('public') standard format.*'

The definition will, of course, change over time—the next obvious step being the removal of the 'business' qualifier as the applications widen to include personal and private use of these techniques for all sorts of transactions such as ordering and renewing books from the library. 'EDI purists'—champions of the widest possible delivery of EDI benefits across all trading communities—may insist on inclusion of the word 'public' in order to exclude proprietary standards agreed within

business designers are able to call on various service and design elements that together can provide a compound solution for the wider enterprise

specific user domains. This position is akin to a commitment to 'open standards', which are frequently in need of practical compromise to meet the realities of uneven development but still worthy as a long-term aim. Wherever the source and sink of the information is, or can be, computerised, there is scope for EDI, even to the extent of including the 'unstructured messages' within a standardised EDI message format or 'envelope'.

Purpose and Application

EDI is used:

- to overcome many time, cost and quality difficulties that are inherent in paper-based transactions;
- to address the issues of communications between applications systems that are independently designed and managed, operating on different types of computer, across a variety of communications networks using different communications protocols;
- to facilitate global trading across different time zones and operational schedules, particularly where the transaction-timing is critical to the enterprise; and
- to encourage the identification and development of *trading communities* where the benefits of mutual collaboration have a beneficial impact on economic growth and competitive advantages for the whole community or economy. The leading exponents of this are found in quick-response schemes in retailing and manufacturing.

To achieve all this, business designers are able to call on various service and design elements that together can provide a compound solution for the wider enterprise. The five most-significant elements are explained below.

Message standards

The invention and agreement of EDI message standards allow trading

partners to integrate the messaging process into their own computing environment without being forced to compromise design independence. There is therefore a growing level of participation in, and recognition of, the work of standards authorities and article numbering. The earlier article noted the increasing acceptance of UN/EDIFACT standards and the potential demise of industry-specific or national standards. There is now widespread acceptance that, although EDIFACT will, in the very long term, become dominant, the continued use of proven and reliable (but less-generic) standards, where large communities have made early and significant investments, is not in any sense life-threatening to the global trader and that any difficulties that may arise from this pluralist environment can be overcome by software that can cope with alternative designs.

Major players—the hubs at the centre of large trading communities—are more than ever aware that the collaboration of many small enterprises is not encouraged if the imposed investment costs of trading electronically cannot be shared across several communities. For the pioneers—for example, some major players in the UK's retail industry—the achievement of electronic trading has resulted from their mutual investment in time and effort over a decade. The commitment to making it happen is burned deep into the corporate soul and it is unreasonable (or even arrogant) to expect an overnight conversion to some alternative international superior wisdom.

Thus, major players, such as BT, may define entirely new messages—such as the 'One-Bill' message for telephone services billing of large multi-site customers—by using the UN/EDIFACT standards, and, at the same time, accept that the electronic bills it may receive for electricity, gas and water from the UK utility companies may conform to TRADACOMS while those in North America might be ANSI/X12 messages. Most countries (even within

Western Europe) are now more-prepared to accept that independent national or industry-specific message standards run counter to the ideals of free trade. UN/EDIFACT will benefit considerably from President Clinton's recent determination to switch all US federal government procurement of goods and services (valued at \$200 billion per annum) to electronic trading within 3 years².

The starting point for most organisations considering EDI for the first time has traditionally been a concern about the costs and quality problems associated with pieces of paper. Since it is estimated that at least 80% of paper-based 'business to business' transactions are now generated by computers, and a similar level of these invoke further input to computers, the delays and errors caused by human intervention represent a considerable cost burden and service inefficiency. Thus EDI pioneers sought to find a direct replacement for paper transactions within the context of existing business practice.

More recently, however, organisations have come to realise that it is better to first question the purpose, value and contribution of these transactions for the entire enterprise. If, for example, one has a wholly reliable trading relationship with a supplier, including visibility of his stock levels and delivery performance, what purpose is served by the interchange of delivery notes, acknowledgements and invoices? If one can safely assume that because the computer-generated order has been automatically dispatched the goods are received, then logic dictates that payment will inevitably be required. Why then impose upon the supplier the extra cost burden of creating and pursuing a demand for the money? Why impose upon your own organisation the cost of receiving and inputting this invoice when you knew all along that you were going to have to pay it, and had probably already agreed the time-scales in which it would be paid?

Why, an organisation might ask, do we have to devote so much space to

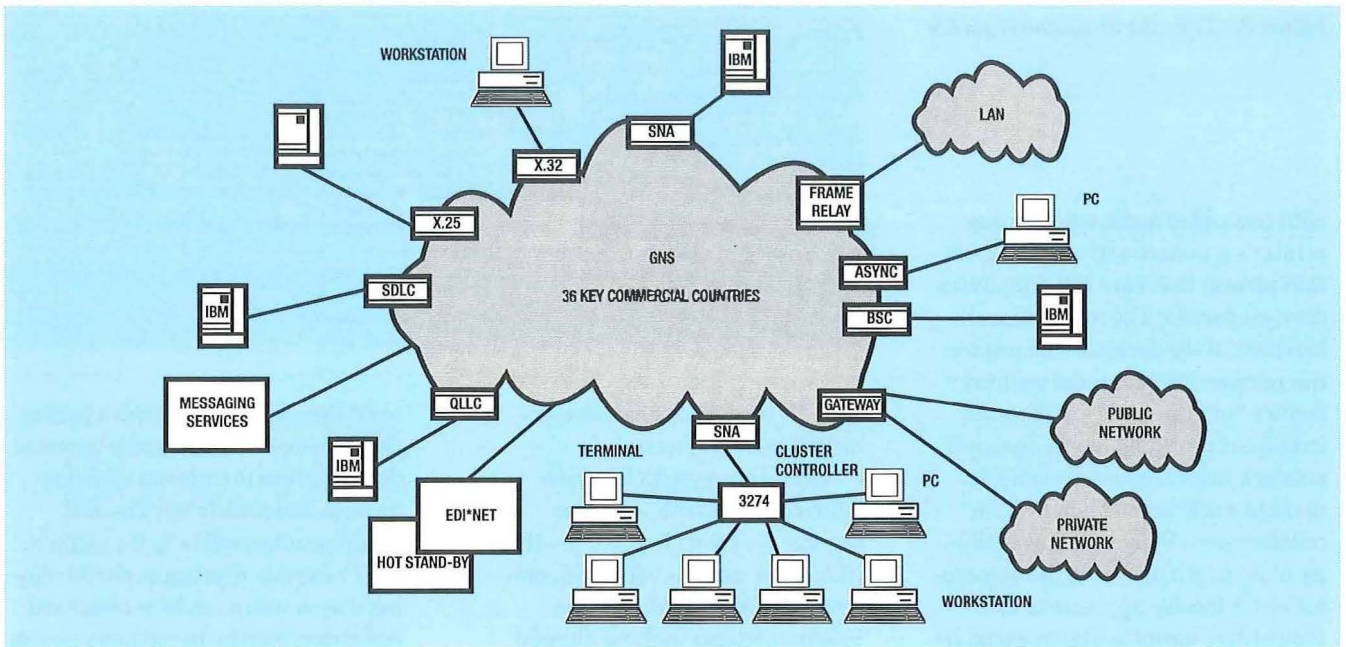


Figure 1—GNS capabilities

warehousing? Why do we have so much money tied up in goods we are waiting to sell? Why do we not know what to say to the customer who is waiting for his order? Why do we run out of the popular sizes and colours? Why are our goods more expensive than our competitors? 'Why do we do business this way?' is now the starting point that might lead to an electronic trading solution and, to the extent that this is difficult and is still more often invoked as a response to a crisis rather than an output of original business creativity, can partly explain why the conversion does not happen overnight. It also explains why the take-up of electronic trading has been boosted in the aftermath of a global economic recession.

Networks

The second major element in the realisation of compound electronic trading solutions has been the widespread availability of *managed network services* such as BT's Global Network Service, GNS. These networks, usually at their core based on X.25 packet switching techniques, help to solve the challenges of interconnecting different computer environments handling vastly different levels of traffic at globally dispersed locations (Figures 1 and 2).

Customers may connect to these networks either directly or via their local telephony service—a decision based almost entirely on traffic levels but occasionally dependent on the degree of real-time systems integra-

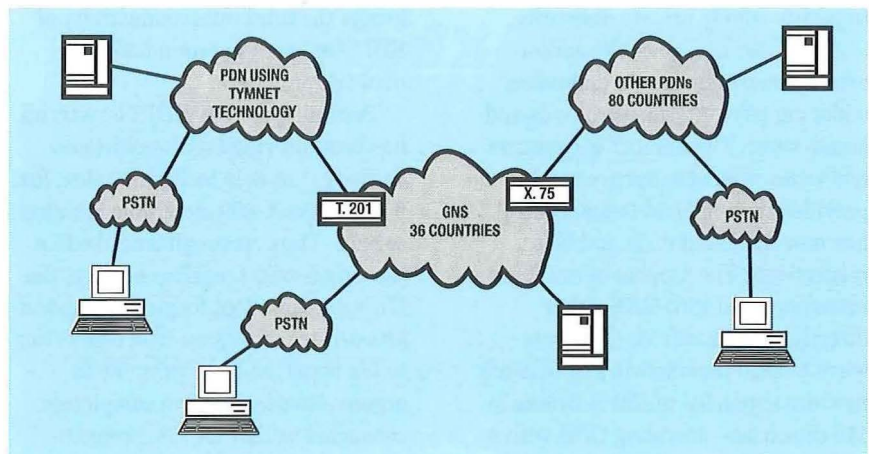


Figure 2—Global interconnection

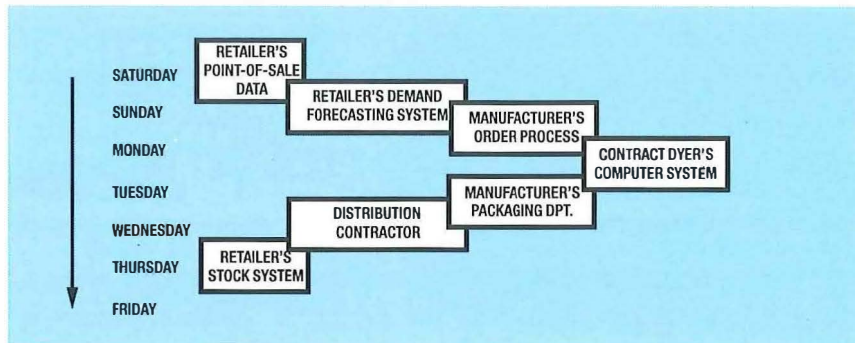
tion that is required. Some, but certainly not all, managed networks provide a facility for the completion of an X.25 transmission to the ultimate recipient via the local telephone network if that user is not directly connected to the network by a permanent circuit. The coverage of GNS within the UK is now so complete that the costs of short-distance circuits for packet-switched traffic are more easily justified than in some other countries and there is a growing awareness of 'virtual network' concepts.

An increasingly critical design feature is the extent of interconnection of ISDN lines to managed networks. As ISDN penetration accelerates and customers switch from inflexible and relatively expensive point-to-point circuits towards switched services whose cost is more closely aligned to traffic levels, the

needs for interconnection will increase, partly to handle bulk traffic for ISDN and non-ISDN locations and partly for the traditional managed network benefits of interface protocol conversion. Managed networks have also provided a convenient infrastructure for commercial information services companies who are now beginning to identify ISDN as a delivery mechanism for more-sophisticated types of information.

We can therefore anticipate a fusion of electronic trading techniques and messages that contain quite elaborate images, including encoded audio and video, that are generated by, and destined for, computer systems. An example might be the transmission from a designer to a printer of colour print specifications for production of a retailer's 'own brand' yoghurt pot complete

Figure 3—Typical transaction sequence



with embedded audio notes for the printer's guidance and an authorisation process to ensure full compliance from all parties. The international locations of the designer, the printer, the retailer, the shops, the yoghurt factory (and the cows) are of course irrelevant providing that a managed network infrastructure can enable them to work as one team in close collaboration. Without the availability of managed networks the opportunities for flexible approaches to competitive business design would be confined to the dwindling numbers of very large organisations who are able to justify wholly private networks.

BT's GNS has, since the earlier article, been enhanced to embrace a wider range of interface protocols and access rates. The number of countries and cities where the service can be provided with full end-to-end control has now increased to 36 and 900 respectively. The number of countries interconnected with GNS either directly or indirectly via X.75 gateways to local packet-switched network has also expanded to 169 networks in 116 countries—providing GNS with a leading position in terms of connectivity and global coverage. This degree of global coverage is an important factor in the design of electronic trading schemes where major organisations and their dependent trading partners—customers and suppliers as well as government tax collectors—need to be confident of the network's reliability and integrity. For this reason GNS and other networks now offer service level agreements that define the expected standards of performance and the processes for resolving any operational difficulties.

Clearing house facilities

The third element of compound electronic trading solutions involves the use of message switching services, such as EDI*Net, to overcome the logistical problems of time zones, schedules and a multiplicity of computer systems and alternative delivery media. The technical description of EDI*Net given in the earlier article is still valid but

the UK-based facility has been enhanced in several ways.

Firstly, the system has been interconnected with the other services operators in the UK—IBM, AT&T and most recently INS, now owned by GEIS. Adding these interconnections to those already available on EDI*Net systems in North America, Korea and Taiwan brings the total interconnectivity of EDI*Net users to around 25 EDI network providers.

Secondly, the UK EDI*Net service has been provided with additional delivery channels including telex, fax, post and an X.400 electronic message service. Thus, as an alternative to a *store-and-collect* mailbox service, the UK user can adopt forms of *store-and-forward* transmission. The connection to fax is particularly valuable to organisations who have completely converted to EDI for their conventional trading community but may need to deal with less competent organisations for short-term or infrequent or seasonal business. The addition of a postal capability enables large organisations to abandon their facilities for printing and posting where these may only be required for a small minority of the trading community—handing off the remaining work to an outside agency and realising operational cost savings.

Thirdly, EDI*Net in the UK has been enhanced by a 24-hour 'hot stand-by' capability that gives additional assurance to customers who are critically dependent on transaction times. Some organisations may have systems picking up instructions at 15 minute intervals in order to maintain work-flow through a production or distribution unit with severe consequences caused by communications disruption.

Figure 3 illustrates a typical transaction sequence. On a Saturday

night the sales statistics from a leading clothing store are reviewed by buyers to detect changes in customer spending patterns and stock levels. The EDI purchase orders will be in the manufacturer's computer system on the Sunday, but they in turn must issue colour and size instructions to the computer system at their contract dyers. The dyer's system will, in turn, send an EDI message back to the manufacturer's with progress information so that packing lines can be ready as the finished goods arrive. The delivery to the retailer's distribution depots (possibly involving yet another enterprise) is similarly time-critical and the end result is that, hopefully, the right colour and size woolly jumpers are on the shelves within five days of a purchase decision.

Finally the facilities available within EDI*Net for message status reporting have been made available to customers to enable them to verify that messages have been posted to, and picked up from, the recipient's mailbox. This important facility was included in the EDI*Net design at the time of Tymnet acquisition and is now proving invaluable in providing reassurance to major hubs who have previously not considered using more than one switching service but now find that substantial numbers of their trading community members prefer to have a choice of messaging service provider.

EDI user software and systems integration

The issues of end-user software were not covered in detail in the earlier article, but, from inception, a software certification process has been made freely available to encourage systems suppliers and present a wide range of software choice to the EDI*Net customer. Potential users of EDI can be roughly divided into two camps—hubs and spokes—or those that take and lead trading initiatives and those who find

the requirements thrust upon them. Early schemes left users with very little choice of EDI software and solutions were often imposed by large organisations whose buyer power was used to demand conformance to the hub's procurement practices. This environment is now in decline; partly because spokes are gaining experience in the context of several trading communities; partly because of a recognition of the need for thorough systems integration, and partly because spokes are themselves becoming hubs for their own trading community. Much of the value of EDI is lost if the recipient spoke is only equipped to convert the messages back into paper form. The influence of major hubs now only extends to recommendation of appropriate software packages and the definition of specific messages.

This position greatly benefits BT who, unlike other service providers, are not in the software business. The UK has a thriving EDI software industry and 13 of the leading software houses have joined with BT in a 'virtual enterprise'—EDI Business Team. Membership of EDI Business Team enables them to work closely with BT in developing joint proposals for customers. Recent significant customers for this joint approach include Next, House of Fraser, British Airways and Marathon Oil.

The creative and commercial input of the Business Team members complements the technical credibility of BT and provides solutions that are comprehensive and informed by experience. EDI Business Team was launched in October 1993 and the concept is now under consideration for extension into continental Europe.

The 1990 article anticipated the need for 'BT service intermediaries' operating on a wholesale basis. This is about to be realised by the appointment in 1994 of some EDI Business Team members to be full resellers of the EDI*Net service, relieving the network operator of much of the administrative and sales effort of dealing with many small and medium-sized businesses (and their computer system environments) and gearing up

to cope with a much faster rate of EDI penetration as large-scale conversion programmes get underway.

Electronic trading community management

The management of implementation programmes—the process of identifying and building collaborative trading communities—has been a major feature of the past three years' work by BT. Nobody should underestimate the effort required, even with support and encouragement by a major hub, to persuade trading partners to adopt a new way of doing business. BT's traditional competency in engineering project management had to be allied to a new commercial focus and position as a 'business facilitator'.

For many large enterprises, the notion that their suppliers (or customers) are *trading partners* belonging to a *trading community* has only recently arrived. The change in attitude that it represents stems from new management design practices that recognise the 'wider enterprise' and seek to gain competitive advantage by influencing the performance of factors beyond the conventional jurisdiction of the enterprise. This consideration often arises as part of a business re-engineering exercise where the interdependence of trading partners provides sufficient motivation to 'change the way we do business'.

The recent recession highlighted the extent of interdependence and may have focused greater management attention on these concepts. Large enterprises found it necessary to reduce their own operations to an essential core and, in doing so, have formed new relationships with agencies, outsourced service providers, or new suppliers to replace the activities that were once undertaken in-house. The old business model of *vertical integration*—seeking to gain control by trying to do almost everything—has been replaced by attempts to create and manage teamwork across a large number of notionally independent and smaller organisations.

This trend has been apparent for many years; one indicator being the positive job fertility rates of organisations with less than 50 employees in contrast to a decade of declining job fertility in larger organisations. One result of the new approach to business practice has been an increase in the number and significance of trade associations, federations and business clubs or groups who seek to address the issues of collaboration while still maintaining healthy competition. For example, the Retail Motor Industry Federation (RMIF) has taken an initiative on behalf of its 13 000 members to introduce an electronic trading scheme that is not subservient to any single large organisation. This is similar to the work of 'buying groups' such as Associated Independent Stores.

An important part of any electronic scheme is, therefore, an ability to manage the needs of the trading community. In this respect, BT is well served by its relationships with the Business Team members. The changes in business practice may mean that the primary influence in the market may not now be a major customer but a previously unrecognised trade association or agency who would not conventionally qualify for attention by account managers.

The Way Forward

The combination of all these elements provides the flexibility to overcome the difficulties that have inhibited the growth of electronic trading. Even after more than a decade of commercial electronic mail services the idea of a real alternative to postal services for many business transactions is only just becoming apparent to the world at large. The conceptual breakthrough and acceptance of what may seem to the information technologist to be relatively mundane has had to wait for experience and understanding of personal computing and local area networks to grow towards endemic proportions. Microsoft's Windows operating software has been called (by

Those that invest in IT to give them an advantage are 'telecompetitive' whereas those who just manage to get by are, at best, merely 'telecompetent'.

a UNIX purist) a computer virus, such has been its rate of growth. Now, with experience of using locally networked applications software, non-technical managers are beginning to realise how easily the concepts can be applied to more-distant transactions. A new class of executive or company is now emerging. Those that invest in IT to give them an advantage are 'telecompetitive' whereas those who just manage to get by are, at best, merely 'telecompetent'. If effective application of IT requires some experience then for many there is little chance of catching the leaders as they move on towards electronic trading systems that are not just confined to textual data but incorporate graphics and video and audio information. BT's position, particularly in the way it develops products and services and perceives market needs, will continue to be informed by its own experience as a user determined to be in all senses 'telecompetitive'.

The application of EDI is now widespread but far from universal. The finalists in the 1993 BT Award for Innovation in Electronic Trading included Texas Instruments, the NHS, HMSO, Rodenstock opticians and LIMNET—the London Insurance Market's community managers. When presenting the Awards, Peter Morgan of the Institute of Directors commented that 'the conversion of the UK economy to electronic trading is now well underway'. In some countries, notably in the Far East, electronic trading has been adopted as a cornerstone of national IT policy — policies which are concerned more about *use* than *production*. In the UK, this 'conversion' is now being described as a major infrastructural programme with huge cross-sector benefits. It is much less visible (and far less expensive) than investment in new motorways, the channel tunnel or North Sea oil, but the results are expected to be no less significant.

The relationship between economic growth and IT investment, between competitiveness and IT expertise has been the subject of many research projects³. On one hand

is evidence of failed IT projects and wasted investment, mostly, it would seem, where the technology was used merely to replicate the existing business practices. On the other hand, we have examples of imaginative initiatives that simply would not have been possible without the advent of electronic trading and EDI techniques. Successful electronic trading schemes look very much like a jigsaw puzzle that has no outside pieces. By providing some of the essential elements, supporting and encouraging the provision of others, and by using its imagination to create new solutions for its major customers, BT is increasingly being recognised as a force for economic growth with a significant role as a 'business facilitator'.

References

- 1 BRUNNEN D. J. Electronic Trading: The Development of Electronic Data Interchange Services for British Telecom. *Br. Telecommun. Eng.*, Jan. 1991, **9**, p. 235.
- 2 CLINTON, W. J. Streamlining US federal procurement by electronic commerce. Presidential Memorandum signed at The White House, 26 October 1993.
- 3 KRAEMER, PROF. K. L. University of California, Irvine. Lecture to PICT December 1993.

Biographies



David Brunnen
Advanced Business
Facilities Limited

David Brunnen is Managing Director of Advanced Business Facilities Limited (ABFL)—a marketing management agency specialising in the introduction of new information and communications products and services. His previous

experience in the telecommunications industry included the development of new ventures for BT. Before leaving BT in 1992 he became an Honorary Fellow of the M.B.A. Faculty at the company's staff training college. He is a Council Member of the Parliamentary Information Technology Committee (PITCOM), and a member of the Society of Competitive Information Professionals (SCIP). His new venture ABFL has been responsible for the creation and management of EDI Business Team—one of a series of new 'virtual enterprises' being developed to market novel information and communications services.



Eric Wyllie
BT Managed
Network Services

Eric Wyllie is the manager responsible for customer service delivery of BT's EDI and messaging services. He joined the then British Post Office in 1961 as a Youth-in-Training, moving to the Engineer-in-Chief's Office in 1965. While in telecommunications development he was responsible for the successful Herald/Pentara range of call-connect systems. In 1982, he moved to the BT Merlin large-PABX team where he coordinated the introduction of the Digital Private Network Signalling System No. 1 at Bank of America and the Palace of Westminster. Before moving to his present role, Eric was Business Manager of Control Systems, a BT venture into the intelligent buildings market. He joined Managed Network Services in 1990, as Head of the Active Support group for the introduction of EDI in BT and the commercial launch of the BT EDI*Net service in the UK and mainland Europe. He has also been personally responsible for several EDI retail sector wins at Debenhams, Iceland Frozen Foods, Associated Independent Stores, Argos, Sears, Asda and House of Fraser.

Telecommunications Quality of Service: Principles and Management

Quality of service is now a fundamentally important aspect in the retention and growth of telecommunications revenue. At present, the service provider's approach to quality has many shortcomings. This article identifies these shortcomings, provides definitions of quality and proposes a methodology for its management, thus benefiting both customers and service providers. The methodology could also be used as the basis for an architectural framework for quality-of-service studies.

Background

Recently the importance of quality in telecommunications (services) has increased. The following factors have contributed to this:

- The growing sophistication of customer premises equipment (CPE) is placing an increasing demand on the performance of the network.
- Customers' requirements are subject to rapid changes which can result in correspondingly rapid changes in quality.
- There is a general increase in the expectation among customers of the level of quality in all walks of life. With the rapid advance of technological applications in computing and the lowering of prices, customers are coming to expect similar benefits to be reflected in their telecommunications services.
- Different segments of the market have varying service-quality levels and price requirements. At one end of the scale there are many customers who do not require a high level of quality but who expect a lower price, whereas at the other end there are customers who require a high level of quality but who are also prepared to pay more for it.
- In countries with liberalisation and deregulation, telecommunica-

tion user groups are exerting more pressure for increased levels of quality.

- In competitive environments the scope for differentiation by price is diminishing rapidly. More and more frequently it is the service provider who can differentiate his services on quality who wins.
- There is a growing requirement from regulators for telcos to provide and publish achieved levels of service, and the regulator will be in a position to set the standards which must be met.

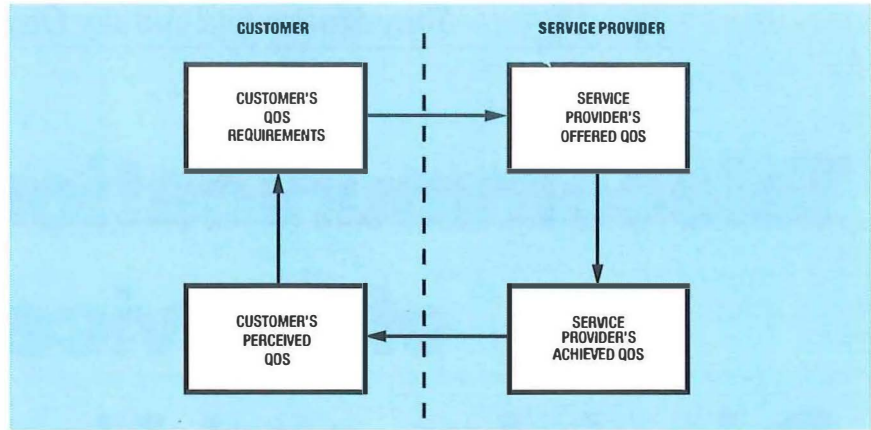
These factors sharpen the need to examine the effectiveness of existing approaches to quality of service (QOS) and to decide what changes are necessary to meet the growing demand by the customer for higher quality at an acceptable cost to the service provider. However, service providers will still need to learn how to achieve a viable balance between cost and quality.

Quality of Service

The principal characteristics of the existing approach to QOS may be summarised in the following statements:

- Most of the published work is the result of studies done on a cooperative basis with other service providers in the two principal international fora: CCITT and ETSI.

Figure 1—The four QOS viewpoints



- There is no commonly accepted QOS logical framework. This has resulted in an ad hoc, disparate and incoherent approach to the aforementioned studies.
- Customer requirements have not been given the consideration they demanded. The network was designed by technocrats who decided what was good for customers.
- Quality-of-service criteria are derived to align with the achievable performance of the network and not customer requirements.
- There is often no means of assessing whether a network is over- or under-provided in terms of quality.
- There is also a gap between the network performance statistics published by telcos and the QOS experienced by customers.
- Quality considerations are not effectively separated into customers' and service providers' viewpoints.

Despite these shortcomings, it is worth noting that QOS studies throughout the telecommunications world have accumulated a wealth of useful knowledge. However, it is necessary to address the above shortcomings in the light of the present telecommunications climate. This makes it necessary to understand the linkages between customers' requirements and service providers' offerings and the subsequent management of both.

Quality of Service Definitions

Before QOS can be managed, it is necessary to understand the differing perspectives and their definitions. The existing definitions of QOS lack the clarity required to express, separately, the service providers' and customers' viewpoints. The following

definitions provide the basis for identification and clarification of which aspect of QOS is being considered at any given time. These viewpoints are illustrated in Figure 1.

Quality of service required (by the customer)

Quality of service required by the customer is a statement of the level of quality of a particular service required or preferred by the customer. The level of quality may be expressed by the customer in technical or non-technical language.

A typical customer is not concerned with how a particular service is provided or with any of the aspects of the network's internal design, but only with the resulting end-to-end service quality. From the customer's perspective, quality of service is expressed by parameters which:

- focus on customer-perceivable effects, rather than on their causes within the network;
- do not depend in their definition on assumptions about the internal design of the network;
- take into account all aspects of the service including the network, network support and customer premises equipment;
- may be assured to a customer by the service provider(s); and
- are described in network independent terms and create a common language understandable by both the customer and the service provider.

It must be recognised that customer's QOS requirements can

sometimes be subjective. These requirements although subjective are extremely useful, and it is up to the service provider to translate them into something of objective use.

Quality of service offered (by the service provider)

Quality of service offered by the service provider is a statement of the level of quality which will be offered to the customer. This is the level of service which the service provider can achieve with the design of the network. The level of quality will be expressed by values assigned to network performance parameters, which not only cover the network but also network support. These parameters must be designed so that they can be understood by the customer. Each service will have its own set of performance parameters which cover the totality of the service, including all the service packaging.

Quality of service achieved (by the service provider)

Quality of service achieved by the service provider is a statement of the level of quality achieved by the service provider. It is an historical record of the levels of quality that have actually been achieved. These are expressed by values assigned to the parameters specified for the offered QOS. These performance values are summarised for specified periods of time; for example, for the previous three months and/or on an annual basis.

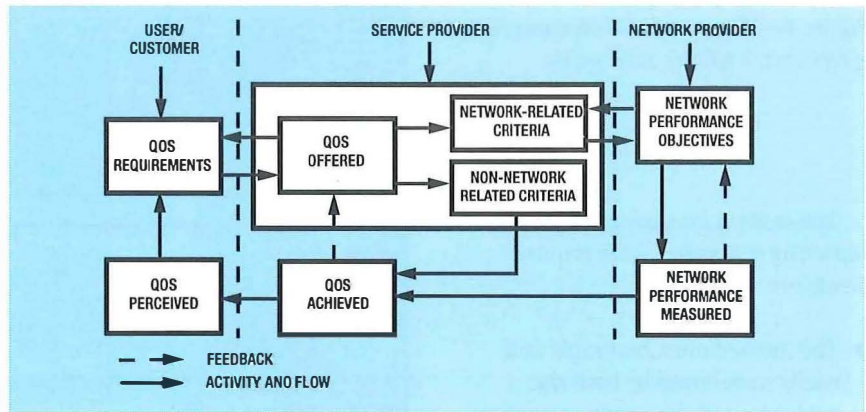
Quality of service perceived (by the customer)

Quality of service perceived by the customer is a statement expressing the level of quality experienced by the

Figure 2—Interrelationships between QOS viewpoints

customer. The perceived QOS is expressed usually in terms of degrees of satisfaction and not in technical terms. Perceived QOS is assessed by various methods including customer surveys, customer's comments and customer complaints. Figure 2 shows how the various QOS viewpoints interrelate with each other. The service provider and the network provider have been separated in the diagram to illustrate the fact that the service provider need not always be the network provider. However, the service provider must always take full responsibility for the QOS offered to the customer.

From the intra-relationships between the QOS viewpoints in Figure 2, it can be concluded that both the customer's and service provider's quality interest must be in a state of equilibrium in order for successful business relationships to be achieved. Therefore, it is necessary to manage the activities and relationships associated with the QOS viewpoints to obtain the optimum



quality levels commensurate with the price the customer is willing to pay. The methodology described in the subsequent section was developed to achieve this goal.

The Management of Quality of Service

A methodology for the management of QOS activities, termed the *quality cycle*, is shown in Figure 3. The principal stages in the quality cycle are:

- capture of customers' QOS requirements;
- establishment of the offered QOS by the service provider;

- formulation of performance specifications for the network and guidelines to carry out the associated tasks;
- measurement and computation of the achieved QOS by the service provider; and
- assessment of the customers' perceived QOS.

The management of each principal stage is now examined.

Customers' quality-of-service requirements

The starting point of the quality cycle is the assessment of customers' QOS requirements.

Figure 3—The quality cycle

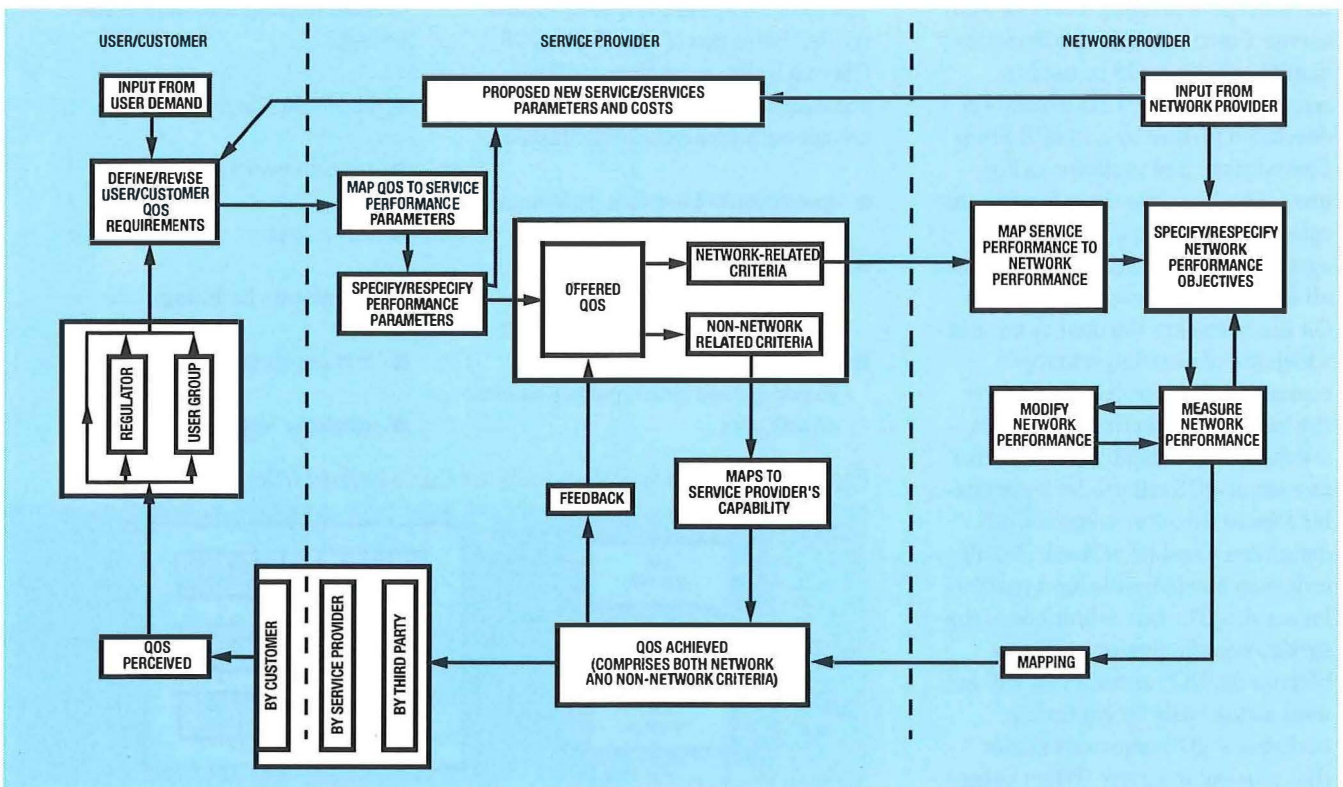


Figure 4—Matrix to facilitate capture of customer's QOS requirements

The criteria for successfully capturing customers' QOS requirements are:

- The method must be simple and easily understood by both the customer and the service provider.
- The method should cover most, if not all, QOS requirements (that is, network and non-network related criteria) of any telecommunication service.

- The method should be reasonably easy to administer and not make undue demands on resources.

Various frameworks for requirement capture have evolved over the history of telecommunications. Many of these have been service specific; that is, suitable only for a particular service. Others have concentrated only on the technical quality of the service offering. What is needed is an approach which combines all the requirements above and can also be applied to any service. A framework was described by Richters and Dvorak in 1988¹ which showed how the concept of dividing a service into service functions and identifying the quality criteria could be used to construct a matrix. This matrix was developed further by a FITCE Study Commission² and is shown in Figure 4. On the y-axis the telecommunication service is divided into 11 specific service functions, which cover all functional aspects of any service. On the x-axis are the quality criteria which are of most importance to customers. The resulting cells form the basis for QOS criteria. All cells are defined generically, but to arrive at a set of QOS criteria for a particular service the service-specific cell definitions must be derived. Not all cells may be applicable for a particular service. The cell definitions of the service-specific version will then become the QOS criteria and will be used as the basis for capturing customer's QOS requirements for that particular service. When using

SERVICE QUALITY CRITERIA		SERVICE FUNCTION						
		SPEED 1	ACCURACY 2	AVAILABILITY 3	RELIABILITY 4	SECURITY 5	SIMPLICITY 6	FLEXIBILITY 7
SERVICE MANAGEMENT	SALES 1							
	PROVISION 2							
	ALTERATION 3							
	SERVICE SUPPORT 4							
	REPAIR 5							
	CESSATION 6							
CALL TECHNICAL QUALITY	CONNECTION ESTABLISHMENT 7							
	INFORMATION TRANSFER 8							
	CONNECTION RELEASE 9							
	BILLING 10							
	NETWORK/SERVICE MANAGEMENT BY CUSTOMER 11							

the matrix the following factors must be given consideration:

- The sample size and customer segment must be chosen to represent the customer base.
- The frequency of administration of the framework should reflect changing requirements.
- Customers' responses may need to be translated into terms more meaningful to service providers.

The Offered Quality of Service

The second stage of the quality cycle is the determination of the offered QOS followed by the management of the associated tasks. The following tasks are normally associated with this stage:

- specification of network performance;
- specification of non-network related performance;
- prioritisation to deliver QOS requirements which cannot be met on day one;

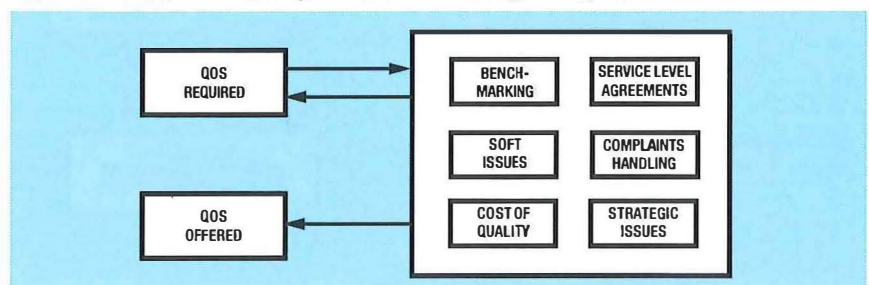
- specification of the monitoring systems;
- specification of the logistics to deliver the offered QOS; and
- publication of the QOS targets.

Determination of the offered quality of service

The service provider decides on the level of QOS to be offered to the customers after assessing customers' QOS requirements and mapping this to business capability. The process of arriving at the offered QOS is shown schematically in Figure 5. The following issues must also be addressed during this stage of the process:

- benchmarking,
- service-level agreements,
- soft issues,
- complaints handling,
- cost associated with QOS, and
- strategic issues.

Figure 5—Illustration of the process to arrive at offered QOS



the cost of providing the level of service required by the customer will be a vital issue in deciding what the offered QOS to the customer should be

Benchmarking

Benchmarking is the process of comparing the performance of one company with that of others³. Benchmarking may be carried out on a service-by-service or product-by-product basis. Such comparisons are made when a service provider wishes to compare performance data with that of a competitor, on the basis of a mutual exchange of information.

The obvious criterion for meaningful comparison is the availability of performance information on the same set of parameters. In cases where this is not available some form of 'tailoring' and interpretation may be necessary for meaningful comparisons. The most effective use of benchmarking information is to enable the service provider to decide where performance should be pitched in the quality league.

Service level agreements

A service provider could enter into an agreement with customer(s) or other service provider(s) to provide an agreed level of quality for a particular service. This is known as a *service level agreement* (SLA). Examples of SLAs are:

- Service provider agrees to provide a minimum level of service for a particular service or group of services to any customer, with a penalty clause for rebate if the quality falls below the published target.
- Bilateral agreement between the service providers of two countries to provide a particular level of service.

The principal advantage of such an agreement is the opportunity for the two parties concerned to focus on a particular quality need.

Soft issues

There are a number of soft issues related to quality which are not obvious in the QOS requirements captured using the matrix in Figure 4. Examples of soft issues are:

- politeness and helpfulness of operators,
- customer goodwill,
- service and repair to suit customer's convenience,
- ethical, societal and aesthetic values of the service provider, and
- public image.

The service provider should be aware of the current soft issues, and consider these when arriving at the offered QOS.

Complaints handling

The following extracts from Telecom's Handling of Customer Complaints published by the Australian Government Publishing Service, Canberra⁴, illustrate the importance to be given by service providers to complaints.

'para 6.8: Complaints should not be seen primarily as a by-product of quality of service improvements covering installation, repair, billing, proportion of payphones in operation and so forth. A complaint covers more than the subject matter. It also includes customer perceptions as to how they are being treated. This emphasises the need to treat complaints as a separate part of quality of service.

para 6.9: ..poor treatment of complaints by a monopoly is inexcusable; in a competitive situation it can be downright foolish.

para 6.12: .. basic requirements for .. (complaints handling) .. process to operate:

- improved quality-of-service measures that attack the causes of complaints,
- a system that deals effectively with complaints at first point of contact,
- close monitoring of why complaints escalate.'


The profile of complaints should be analysed on a service-by-service basis to help identify the relevant inputs. However, problem localisation and resolution may require taking a cross-service view.

Cost associated with quality of service

Inevitably the cost of providing the level of service required by the customer will be a vital issue in deciding the offered QOS. As most service providers have been operating under monopolistic conditions there has been no strong commercial reason for efficient costing systems to have been developed. This has resulted in a lack of understanding of the relationship between cost and quality.

Figure 6 shows traditional curves associated with the cost of preven-

Figure 6—The cost of quality



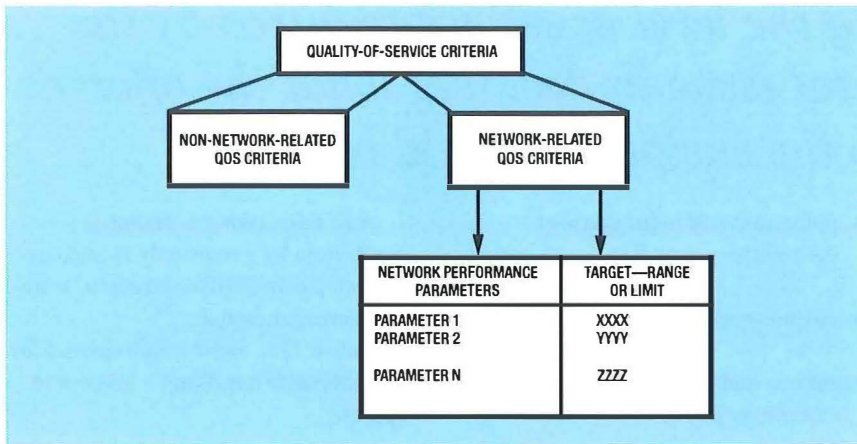


Figure 7—Relationship between quality of service and network performance

tion and the cost of failure with total quality cost being the sum of the two. The total quality cost curve is U-shaped with respect to quality levels, and the level that gives minimum cost may at first sight be the level of service to be aimed for by the service provider. However, the relationship between cost and quality should also bring in the additional relationship of revenue. When this relationship is established, the elasticity of quality versus revenue when competition is present may be attempted. The customer will be looking for value for money. Customer interest in this area may be expressed by the equation:

$$VFM = B - P,$$

where *VFM* is value for money, *B* is benefit and *P* is the price for the service. It will be necessary for a service provider to have a grasp of the cost-versus-quality relationship in order to control and manage quality. Precise quality-versus-cost relationships are likely to be a long way off, but every effort to expedite the arrival of this state of affairs would be beneficial to the service provider.

Strategic issues

Examples of strategic issues are:

- What is the product life cycle of this service and the resulting implications on the quality?
- What is the opportunity cost of providing a level of service better than asked for?
- What factors are likely to influence QOS?

The relevant issues should be identified and considered in the process of arriving at the offered QOS.

Publication of quality-of-service targets

The offered QOS should be published within the service-provider's organisation for the information of all concerned personnel, and action planned to remedy any subsequent shortfall in performance should also be detailed.

The service provider may publish this information for the benefit of customers. Lack of confidence on the part of the service provider to carry out this exercise may be perceived as a weakness or as an act of concealment.

Specification of network performance

The offered QOS must be translated into performance statements in order that both the network- and non-network-related QOS performance targets may be set. Figure 7 illustrates schematically the separation of the network- and non-network-related QOS criteria. Network performance parameters are derived and design targets set to meet the offered level of quality.

To illustrate, let us assume that the offered QOS for the accuracy of information transfer for basic telephony is that not more than 1 call in 100 should experience moderate difficulty in call clarity. When mapped into network performance parameters, this would result in network performance measures including transmission loss (on analogue circuits), noise of all types, echo, delay, sidetone, clipping etc. To fix the

target for end-to-end performance it is necessary to know the combined effect of these characteristics and/or components on call clarity. Since such a relationship is not linear in practice it is necessary to arrive at these relationships empirically.

Having specified an end-to-end performance measure, this must then be broken down into elemental performance levels and targets which will need to be monitored in order that their combined effect will meet the end-to-end performance required by the customer. As far as possible CCITT Recommendations should be used for specifying network performance measures. The above process should then be repeated for all the principal services supported by the network. It is quite possible that differing performance levels may be required of the same measure for different services. Then it is the responsibility of the service provider to decide upon the performance levels of the network. The service provider will need to consider the following points:

- meet the most stringent requirement for all services,
- provide an overlay network for the most stringent requirement, and
- separate the high performance parts of the network for the higher quality users.

This is an area where the ingenuity of the service provider will be challenged to the utmost.

Specification of non-network-related performance

The specification of non-network-related measures could be carried out by standard operational analysis for the quality requirements.

Specification of monitoring systems

Monitoring systems will be needed to track the achieved performance against the offered performance. For

Customer opinion is valuable because it gives the service provider a finger-on-the-pulse knowledge of what the customer thinks of the quality of the services.

technical performance this may involve measurements on individual elements of the connection. Monitoring systems for non-network-related performance would normally involve recording information from which achieved performance can be computed.

There are two objectives in the selection of monitors: firstly, to ascertain how well the achieved QOS compares with that offered, and secondly, how the network performs. The choice of measuring systems will reflect these objectives. The following guidelines indicate the principal points to be observed in the selection and specification of monitoring systems:

- Wherever possible the monitoring systems should reflect directly the network performance specified.
- Where it is necessary to break down network performance to elemental performance the achieved end-to-end performance should be estimated from the measurements taken from the monitoring systems.
- Careful analysis will identify the necessary non-network activities to be recorded.
- Measurements may be taken on a sample basis. The sampling ought to be decided by the service provider based on their own needs and those of customers and the regulator.
- The advantages and disadvantages of monitoring live traffic as opposed to artificial traffic should be carefully evaluated before a monitoring system is specified.

Specification of the logistics to deliver the offered quality of service

An inventory for the implementation of the offered QOS and the associated monitoring systems ought to be made. Types of resources are:

- Physical resources such as information brochures, customer premises equipment, associated exchange equipment, test equipment, etc.
- Resources to record customers' fault reports and systems to process these within the service provider's organisation.
- Resources to produce achieved QOS levels from observed performance data.

The Achieved Quality of Service

The achieved QOS is the statement of performance actually delivered by the service provider over a specified period of time. The performance must be expressed as the end-to-end performance as experienced by the customer.

These are based on measurements taken and information recorded during the time frame in question from the monitoring systems. Both network- and non-network-related performance ought to be considered. The achieved performance may be published in two principal categories: one for the customer and the regulator (that is, to go outside the business) and the other for use within the business. However, it will be far more beneficial in the long run to produce one set of published performance data for use both internally and externally. The achieved QOS will be compared against the offered and the perceived QOS. Any discrepancies or changes must be the subject of investigation.

The Perceived Quality of Service

The perceived QOS is the level of satisfaction expressed by customers and users. The satisfaction level is expressed by an opinion score and is ascertained usually through customer surveys.

Customer opinion is valuable because it gives the service provider a finger-on-the-pulse knowledge of what the customer thinks of the quality of the services. The perceived QOS has on frequent occasions been found not to correlate with the achieved QOS. Some of the reasons for this may be:

- The customer's perception, being subjective, is liable to vary with time and the recorded value is thus accurate only for the instant when it was recorded.
- The customer may be influenced by recent good or bad experiences with the service provider.
- The media can influence the customer's perception.
- The customer may be ill-informed, (for example, a recent survey showed that the customer thought they were paying much more for a fixed duration local call than the actual tariff).

While comparing the perceived QOS with that of the achieved QOS all possible reasons for the mismatch should be considered. Corrective action should then be taken, such as:

- improvement in the achieved QOS,
- improvement in the offered QOS,
- education of the customers,
- revision of the correlation techniques (between the achieved QOS and the perceived QOS), and
- review of customer surveys.

The opportunity cost of not following up the feedback between the perceived and achieved QOS may be high because in a competitive environment the customer will have the opportunity to buy from an alternative service provider. Figure 8


Figure 8—Two levels of service provision

illustrates two cases. In case 1, a higher percentage of customers are supplied with a number of services in the first five days, but at the tail end there is a longer wait. In case 2, the tail is shorter, but in the first five days a lower percentage of customers than in case 1 are supplied with the services. Which is perceived by the customers as providing the better level of service? The answer lies with the customer. Service providers could be mistaken if they arrive at their own interpretation without considering their customers' perception. The above curves also illustrate the need for tail management. The average performance may appear respectable but the customers at one end would experience unacceptable levels of quality.

Role of the Regulator

A regulatory body is normally appointed in countries where competition in telecommunication services exists. In the UK this body is called OFTEL. Its responsibilities are described in the Telecommunications Act 1984.

The principal role of OFTEL, as far as quality is concerned, is to ensure that realistic interpretations of achieved quality of service are published and that where improvements are necessary these are carried out. The performance statistics should be published in terms meaningful and understandable to the customers. To date (Spring 1994) the performance measures on which achievements are to be published have not been finalised by OFTEL. Yet another regulatory body is the European Commission. For the principal telecommunication services they produce directives. Embodied in the directive is a section on quality of service. The service provider will be required to publish achieved results on a regular basis. As an example, ten measures have been proposed for basic telephony in the Open Network Provision (ONP) directive. The precise relationship between OFTEL's list of performance measures and that



of the EC has not yet been publicly stated.

International Comparisons of Quality of Service

Comparisons of quality of service achieved by different service providers in other parts of the world would enable service providers to assess their own standing in the performance league.

International bodies such as the OECD (Organisation for Economic Cooperation and Development) and ETNO (European Telecommunications Network Operators) have attempted to persuade the service providers from their respective member countries to agree on a common set of definitions to

The service provider who offers credible quality at the lowest price wins.

enable true comparisons of performance to be made. Published performance on a common set of measures is still awaited.

Performance Engineering

Performance engineering embraces all engineering activities associated with the management of the quality of service. These activities include traffic dimensioning to give the planned availability, transmission performance to give the required call clarity etc. Also included are the monitoring systems to measure the achieved level of quality. When the engineering activities of the network reflect the philosophy described in the quality cycle, the network may be considered to be designed for optimum quality commensurate with the price a customer is willing to pay.

Conclusion

In the present business climate it will be disadvantageous for a service provider to ignore the demands on quality made by customers. This article proposes a methodology, which could simplify the identification and management of tasks associated with quality of service. There are no short cuts, and rigorous adherence to the methodology will bring rich returns for service providers. The service provider who offers credible quality at the lowest price wins.

The methodology described in this article could be considered to form the basis for an architectural framework for the study of QOS in the interna-

tional fora. The nature of some of the issues in QOS is so fundamental that no one service provider can work in isolation; these issues have to be addressed in the international fora and mutually acceptable solutions found. The overriding advantage of an internationally-agreed framework is that the cumulative effect of QOS studies will complement each other to the benefit of all parties. The time now is ripe for a significant shift in the approach to QOS and its management.

Acknowledgements

The authors wish to thank both Keith Ward and Robert Faulkner (BT Products and Services Management) for their contribution to this article. Thanks are also expressed to the Federation of Telecommunications Engineers of the European Community (FITCE) for Figures 2, 3, 4 and 7.

References

- 1 DVORAK, and RICHTERS. *Communications Magazine*, IEEE, Oct. 1988.
- 2 Final Report of the FITCE Study Commission on Study of Network Performance Considering Customer Requirements, 1991–1993.
- 3 CHERRETT, PAUL. Practical Benchmarking. *Br. Telecommun. Eng.*, Jan. 1994, **12**, p. 290.
- 4 Telecom's Handling of Customer Complaints. Report from the House of Representatives Standing Committee on Transport, Communications and Infrastructure, Australian Government Publishing Service, Canberra, May 1991.

Glossary

CCITT International Telephone and Telegraph Consultative Committee

ETSI European Telecommunications Standards Institute

FITCE Federation of Telecommunications Engineers of the European Community

ONP Open Network Provision

Biographies



Tony Mullee
BT Business
Communications
Global Marketing

Since joining BT in 1972, Tony Mullee has held a number of management posts within both the technical and the commercial sides of the business. He was also one of the triad responsible for the development of BT's Master's Programme, designed to develop BT's future senior managers. Tony chaired the two-year FITCE European Study Commission 'Study of Network Performance Considering Customer Requirements'. Currently, Tony is employed as BT's Strategic Global Marketing Manager, with responsibility for BT's Advanced Managed Services Portfolio including Concert technology. Tony has a B.Sc. in Applied Physics, an M.Sc. in Telecommunications Management and an M.B.A. from Essex University.



Antony Oodan
Consultant

After graduating in physics, Antony Oodan worked for over a decade on the design and manufacture of thermionic tubes. His contribution included improvements in the manufacturing techniques of submarine cable repeater tubes and in the production of long-life tubes for Government use. He moved to telecommunications, joining BT about two and a half decades ago, and has had extensive experience in transmission, switching and service provision. His life-long penchant has been towards standards and quality. A recent spell as Secretary of a Study Commission under FITCE enabled him to contribute towards his favourite subject. Antony now heads his own consultancy on quality of telecommunication services. He is a Member of the Institute of Management.

A Whole-Life Approach to Management Decision Making

The need to reduce costs in the face of increasing market pressures is widely recognised by managers throughout industry. The whole-life approach described here can be used to improve the information available to decision makers, using a simple process and friendly computer tools for a range of business applications, including supplier and technology selection.

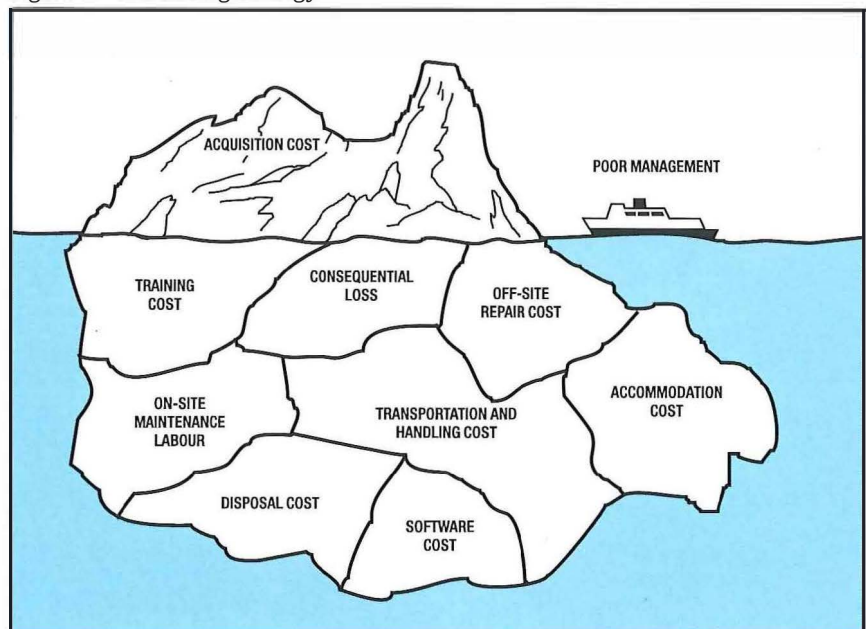
Introduction

When taking decisions which affect the specification, design and selection of hardware or software, it is important to take a long-term view of costs. The danger of concentrating on up-front costs only is often illustrated by the iceberg analogy (Figure 1). The technique of whole-life costing was initially applied to the procurement of military equipment by the US Department of Defence in the late-1960s¹. This followed the realisation that the costs of ownership were far greater than the initial purchase price. It was estimated that only 26% of the \$126 billion US defence budget was spent on the procurement of new products, compared with 61% on operation and support². Care must be

taken, however, not to assume that all these costs could be scaled down if equipment were more reliable or maintainable; many of these operating costs in the military arena arise from the need to plan for active battle service, not just routine peace-time maintenance³. This should also be considered before applying military analogies to commercial situations.

Attempts were made during the 1970s to apply whole-life costing more widely. In the UK, the Department of Trade established a National Terotechnology Centre in 1975 with the objective of promoting the adoption of whole-life costing to asset management in British industry. However, even as late as the 1980s, there was little evidence that UK managers had modified their view that initial cost

Figure 1—The iceberg analogy



WINA allows users to perform complex whole-life analyses in a windows environment without specific programming or spreadsheet knowledge.

was the most important consideration when acquiring capital assets⁴.

In the late-1980s, a whole-life costing group was established in BT Group Procurement Services. Procurement policy was changed to include mandatory formal consideration of whole-life costing when selecting suppliers for certain product types, and several commercially available models were evaluated. The US Department of Defence model, CASA (cost analysis and strategy assessment), was selected as the most appropriate, and was used extensively until the development of an in-house model, WINA (see below), in 1993. From 1990 to 1993, around 60 studies were completed in support of procurement decisions. Products assessed include core and access transmission systems, computers and customer apparatus.

Whole-Life Costing and Financial Analysis

Giving consideration to downstream costs in decision making is obviously nothing new. When analysing a proposed project or purchase, the traditional approach would normally have two stages, technical cost estimation and financial analysis:

1. Estimates are produced for the annual costs of items such as purchase, maintenance, repair, power consumption etc. These could be arrived at using 'hard' historical data, established 'rules of thumb', expert engineering judgement or sophisticated causal models. Normally someone with detailed knowledge of the system under consideration would make these estimates, using a combina-

tion of these methods. Estimates of revenue must also be made for investment appraisals.


2. The estimates are then processed by a financial analyst, who ensures that such things as inflation, the time value of money, and tax implications are considered. The output would normally be a net present value (NPV) for the project. An excellent description of these financial issues is given by Lewis Pilgrim in *Telecommunications: A Structured Information Programme*⁵. A PC-based model, RAPID (rigorous appraisal of investment decisions) is widely used in BT to produce NPV estimates.

Possibly several iterations will be required, due to changes in the basic estimates or the predicted inflation rates. It is important to remember that the aim is to predict accurately the future with uncertain data.

The whole-life approach is intended to improve this process:

- by providing a standard methodology for technical cost estimation by modelling the link between basic cost drivers (such as system reliability) and the resulting cash flows;
- by combining the two stages in a single computer model, with in-built financial logic, thereby bridging the technical and financial issues; and
- by providing extra facilities in a user-friendly model such as graphical output, fast 'what-if' analysis, a system hierarchy builder, population profiling and life-cycle templates for a range of products. Figure 2 shows the flow of data through WINA.

Figure 2—Simplified block diagram of WINA



The BT WINA (whole-life investment appraisal) model, developed by BT Development and Procurement, facilitates the above.

The Whole-Life Approach and WINA

The BT WINA model was developed by the Intelligent Systems Unit at BT Laboratories and Group Procurement Services (GPS) Finance to satisfy a set of requirements agreed by representatives of all the major divisions in BT. WINA allows users to perform complex whole-life analyses in a windows environment without specific programming or spreadsheet knowledge. The programming language used was Smalltalk, an object-oriented language⁶. WINA runs under the Windows™ operating system on a PC with a minimum specification of a 386 processor and 4 Mbyte RAM.

A simple, seven-stage process for whole-life investment appraisal is shown in Figure 3. The WINA model

Figure 3—The seven-stage process

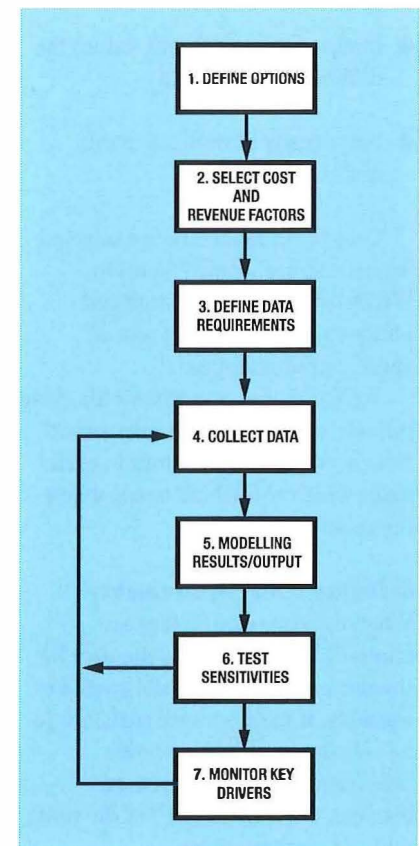


Figure 4—WINA component hierarchy example

can help with most stages by providing a modelling structure and templates where appropriate.

1. Define options

The first stage is to define the options under consideration. This involves setting the system boundaries, which can range from selecting a supplier for a simple product to comparing two fundamental network options. In addition, quantities and purchase profiles need to be defined (for example, 20 000 fax machines per month for two years, five transmission systems per year, etc). WINA provides facilities for population profile entry and a powerful system hierarchy builder (Figure 4). Any number of options can be compared, but it is important to consider the 'do nothing' possibility when supporting business cases.

2. Select cost and revenue factors

For decision support, factors should be selected which are:

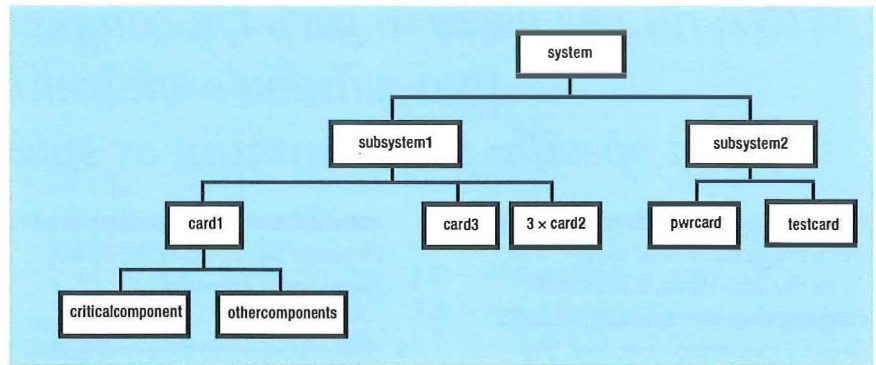
- likely to be of significant financial value,
- likely to have different values for different options, and
- not already committed (sunk costs).

Cost/revenue factors are selected from a comprehensive menu in WINA, which can also store templates to model the life cycle of specific product types.

WINA divides a product's life cycle into six stages; the user may select from a range of cost factors in each stage which contribute to the whole-life cost/revenue.

3. Define data requirements

When cost/revenue factors are entered, WINA prompts the user for the data required to model each. For example, if maintenance costs are to be calculated, WINA requires information on the mean time between failures (MTBF) of the units making up the system.



4. Collect data

Input data should be gathered from equipment manufacturers, existing field databases and expert opinion as appropriate.

5. Modelling results/output

WINA produces overall results in tabular or graphical form (Figure 5), after processing the input data.

6. Test sensitivities

WINA allows 'what if' tests to be performed easily. Inputs can be varied and new graphical output produced on screen within seconds. In this way, the 'key drivers' can be identified. When key drivers are identified, effort can be concentrated on sourcing accurate data for the items which really matter.

7. Monitor key drivers

Key drivers identified by WINA should be monitored during the lifetime of the equipment if possible. Real results can then be produced to


check against the original prediction. This is a vital stage which should improve accuracy of future studies and enable supplier performance to be evaluated.

WINA allows individual buyers and engineers to perform their own cost analyses in the first instance without assistance from a specialist analyst.

In Defence of Time Preference for Money and Discounted Cash Flow

Discounted cash flow techniques are used as a formal method for expressing the preference for spending money in the future as opposed to spending now. Future expenditure is preferred because money in the bank has interest-earning potential; for example, if £100 is required next year, only £93 would need to be invested this year (assumes 8% interest). Discounting works on a

Figure 5—Typical WINA results



It has been estimated that up to 80% of the whole-life cost of a typical product is committed in the design stage of the life cycle

compound basis, resulting in expenditure more than five years out being drastically reduced when considered in today's terms.

Engineers sometimes view the application of discounted cash flow (DCF) as financial short termism, standing in the way of 'spend to save' measures and overemphasising the initial cost of a project. This view is strongly refuted by Philip Pugh⁷, who defends time preference for money and DCF for two reasons:

- The first duty of financial management is to pay the bills as they arrive on a day-to-day basis. If that is not accomplished, there will be no future to be long sighted about.
- Expenditure delayed is money saved, because this will either reduce borrowing or liberate funds to be gainfully employed elsewhere. The application of DCF is valid if and only if the test discount rate used represents the real cost of capital.

What Factors Should Be Considered and When?

Whole-life investment appraisal can be used to support a range of business decisions. These range from the relatively simple selection of suppliers to meet a well-defined requirement, through consideration of different technical solutions, to selecting the actual functionality required. A range of different cost/revenue factors can be considered, depending on the possible options. The WINA model provides sample templates relating product types to likely cost and revenue factors.

Design to Lowest Whole-Life Cost: The Ultimate Goal?

It has been estimated that up to 80% of the whole-life cost of a typical product is committed in the design stage of the life cycle⁸. Expenditure on things such as repair and maintenance may well be required during

the in-service period, but is an inevitable consequence of decisions made much earlier on by the designer. This is demonstrated in Figure 6. The obvious conclusion is that system end users such as BT have a vested interest in the design work which will normally be carried out by the equipment manufacturer. Two possible strategies that a purchasing organisation can adopt to take account of this are:

1. Leave the design work to the suppliers, and then select the lowest whole-life costing option on a competitive tendering basis. This method is currently in use by many large organisations, and involves a formal transfer of technical information in the tender documents, to allow the procurement analysts to assess the whole-life costing implications.
2. Work more closely with suppliers so that they have a fuller understanding of their customers' cost drivers, and can influence their design work accordingly to reduce them. This would involve an iterative, two-way transfer of information.

Option 2 initially seems preferable, although it requires a high degree of trust between customer and supplier. The supplier needs to be

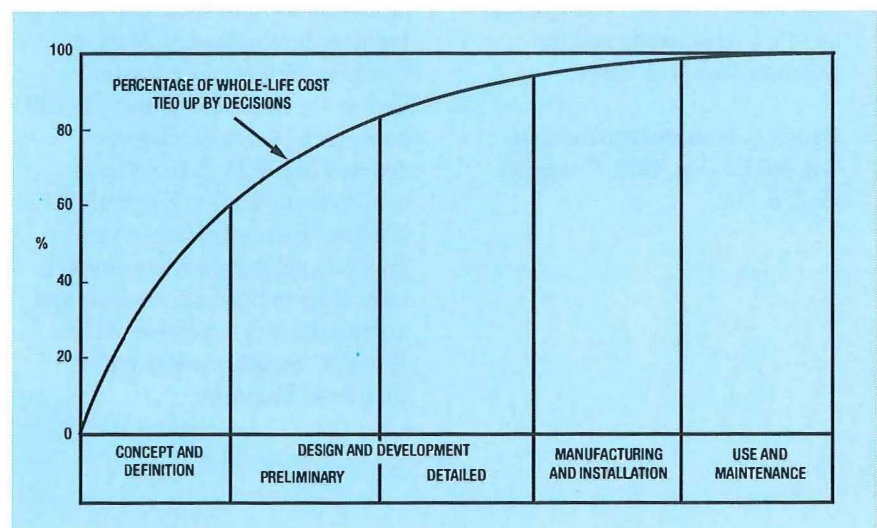
certain that the customer will consistently use whole-life costing as the financial criterion for selection so that his efforts to reduce costs of ownership will not lose him business because the customer reverts to a purchase price based decision. Conversely, the customer needs to be assured that the supplier cannot exploit the relationship to increase purchase prices without producing the desired reductions in ownership costs. In addition, internal overhead and labour rates could be sensitive, particularly when a supplier may also be a competitor.

Conclusions

A whole-life approach to management decision making has been described, based on a simple seven-stage process and the use of the BT WINA (whole-life investment appraisal) model. The importance of applying the correct financial logic to cash flows has been emphasised, as has the need to take a consistent approach to technical cost estimation.

The dependence of equipment costs of ownership on the initial design stage of the life cycle has been discussed. It is vital that purchasing organisations make designers aware of the long-term financial consequences of their work if whole-life costs are to be minimised.

Figure 6—Commitment of cost by life-cycle stage



Acknowledgements

The authors would like to thank Simon Paul of GPS Finance for his significant input on the history of whole-life costing.

References


- 1 PAUL, S. C. N. Whole Life Costing. BT Internal report, 1993.
- 2 GIBBS, J. Basic Principles of Whole Life Costing. Paper from Life Cycle Costing seminar, Jan. 1988.
- 3 PROFFITT. Life Cycle Costs, Paper from IEE Colloquium on Life Cycle Costing and the Business Plan, Feb. 1994.
- 4 WOODWARD, D. Lifecycle Costing for Optimum Asset Management. *Accounting World*, Feb. 1990.
- 5 PILGRIM, L. Investment Appraisal. *Telecommunications Engineering: A Structured Information Programme* (supplement to *Br. Telecommun. Eng.*), July 1993, Issue 9.
- 6 DAVIES, J. and NAPIER, I. New Directions in Investment Appraisal. Scheduled for publication in *BT Technol. J.*, Oct. 1994.
- 7 PUGH, P. G. and V. Spend to save measures: Their financial justification. Paper from IEE Colloquium on Life Cycle Costing and the Business Plan, Feb. 1994.
- 8 DRURY, C. *Management Accounting*, 3rd Edition. 1992, Chapman Hall, p. 809.

Biographies



Ian Napier
BT Development and
Procurement

Ian Napier joined the Post Office in 1977 as a telecommunications apprentice. He was sponsored to study for a B.Sc. in Engineering Electronics at Warwick University. On graduation in 1985 he joined BT Materials and Components Centre to work on component quality and reliability. He obtained an M.Sc. in Microelectronic Technology from Middlesex University in 1992 and is presently studying for an M.B.A. He led the development of whole-life costing in BT procurement from 1989 to 1993, and is a member of the Institution of Electrical Engineers.



John Davies
BT Development and
Procurement

John Davies graduated in Computer Science and Physics at the University of London in 1981. After two years in industry, he obtained an M.Sc. in Computer Science from Essex University. After four years with GEC he returned to Essex where he received his Ph.D. in the area of artificial intelligence. He joined BT in 1990 and currently leads several projects applying decision- support technology to financial analysis and appraisal. He is a member of the British Computer Society and a Chartered Engineer.

Strategic Provision of Radio Nodes in the BT Access Network

The strategic provision of radio nodes has become a major method of distributing digital network services for many UK operators, while BT has continued to rely heavily on its duct and cable network. As service delivery time-scales continue to reduce, there is a need to exploit every method of customer service provision, and digital microwave radio in the BT access network must be used to complement its optical fibre partner, rather than competing with it.

Introduction

Provision of customer digital wideband services by radio is not new. Neither is the concept of customer radio nodes—you only have to look across the skyline of any major city to see an array of lattice steel and dish aerials adorning its most predominant buildings (see Figure 1).

You will be looking at what is probably the most common method of providing network services for many UK operators.

BT has, for many years, used radio as well as optical-fibre cable in its core network, but has continued to exploit its existing and extensive access duct and cable network as the most economic solution for service delivery, thus keeping customer costs to a minimum.

Digital circuits have even been sold on the basis that 'BT will provide optical-fibre cable, rather than that unreliable radio!', a fact that worked wonders to enhance the profile of radio when it was actually needed for engineering reasons or speed of provision in the early days.

The fact is, modern digital radio has a reliability and quality of service at least as good as its fibre partner, as well as often being quicker to install, and sometimes more appropriate. BT account managers have also become much more aware of the advantages of using radio to meet customer delivery requirements, and the number of single point-to-point radio installations in BT's access network is now growing rapidly. With a few notable exceptions in major city centres, however, BT radio nodes are still relatively rare.

Figure 1—Mercury Communications Limited radio node structure on the Leeds skyline



In today's marketplace, prospective customers are expecting faster and faster provision of service. Their own business needs are changing quickly, and they quite rightly expect network service providers to keep pace. They are able to shop around, and those lattice steel structures on the skyline suddenly become a dark formidable challenge to the ability of duct and cable, or even to reactive radio provision, to compete in the demanded time-scales.

Until recently, BT's MegaStream provision followed a 60 working day process for new circuits (20 days for successional circuits), and a survey completed in mid-1993 showed that customers were, on average, requesting completion in 45 days. The probability of failure on new circuits was, therefore, already being built in, even before an order was taken. The average provision time in the United States is in the order of 15 days.

In December 1993, BT reduced its MegaStream process to 30 working days (10 days for successional circuits), and is still providing many in less than this to meet customer demand. A rigorous drive towards further reductions in standard process time-scales is currently under way.

North East Radio Node Project

In September 1993, BT Worldwide Networks North East Zone commissioned a study on the use of access digital microwave radio links, for the rapid provision of digital network services in an urban environment.

The study considered:

- recent changes in the General Development Order (GDO) affecting BT's legally permitted development rights for the provision of radio structures (planning permission requirements);
- the likely consequences of the new Road and Street Works Act;

- changes in time-scales being demanded by customers for digital network services;
- consideration of how other operators provide their services; and
- previous North East studies on radio deployment in the outer core network.

The study recommended the strategic pre-provision of radio nodes in selected cities throughout the North East, and following financial approval, the first tranche of these sites has now been constructed.

The product portfolio available from these node sites will be the complete raft of digital network services available using existing 18 GHz radio technology, including MegaStream and other similar digital services (for example, ISDN 30), with normal response times of between 6 and 15 days; that is, comparable with that now available in the United States.

Radio Node Sites—Definition

A radio node is, in essence, a concentration of microwave dishes, affording line of sight to a pre-determined area of potential customers up to 15 km from the node site. This can be achieved from either a structure on the roof of a building, or, as an alternative, a ground-based mast. Each node structure is capable of supporting between 15 and 30 microwave radio dish aerials.

Associated with these structures is suitable accommodation to house radio, multiplex, and transmission line terminal equipment, together with main cable feed(s) back to an exchange building, giving access to the BT network. Coaxial intermediate frequency (IF) cables to the roof structure are also pre-provided.

Dishes and associated radio units are not pre-provided. They are both expensive and readily available, and are, therefore, provided reactively on demand.

Provision of customer service is subsequently made quicker and easier by the removal of any major holding factors or labour intensive work at the node site. Effort can then be concentrated at the customer end, where an ever-increasing range of antenna mounts has been developed to simplify installation.

Circuits provided using radio node facilities may remain fed by radio, or later be transferred to cable when convenient; for example, when the customer site becomes part of a passive optical network catchment area.

Short-term temporary radio service may also be given for speed of provision, while normal duct and cable techniques are deployed.

Selection of Sites

The initial study recommended a first tranche of node sites in five key cities in the North East, with judgement being made on the basis of both commercial importance and geographical spread. The actual number chosen was based on central guidance on likely roll-out throughout the UK.

Selection of target areas within the selected cities was then made on the basis of:

- heavy concentrations of existing and potential digital network service customers;
- 'fibre lean' areas where penetration of existing optical-fibre cables was low or non-existent;
- prime sites (for example, business parks etc.);
- top ' $N \times 100$ ' customers; and
- any location where coverage would be advantageous; for example, to counter a direct competitive threat.

Local tactical planning guidance was obviously essential in these first two stages of node site location.

Figure 2—Radio node provision at a third-party site

A detailed site search was then carefully made by using a radio planning terrain database (which contains Ordnance Survey contour information for determining probable line-of-sight coverage), and local detailed planning knowledge of possible sites and buildings.

This was followed up by a physical survey at selected sites, to confirm suitability of the site and visible evidence of maximised customer coverage, at lowest possible cost, within the possible 15 km footprint.

It soon became apparent that although city centre sites were often best situated, they would, in many cases, be at greatest cost to rent accommodation and roof space. The best city centre sites were also, in many cases, already occupied by other operators' node structures. BT's own buildings and exchanges, although not always in city centres, often proved to be very favourable alternatives (that is, from the outside of cities looking in).

A final choice was made by using a selection of both BT and third party buildings.

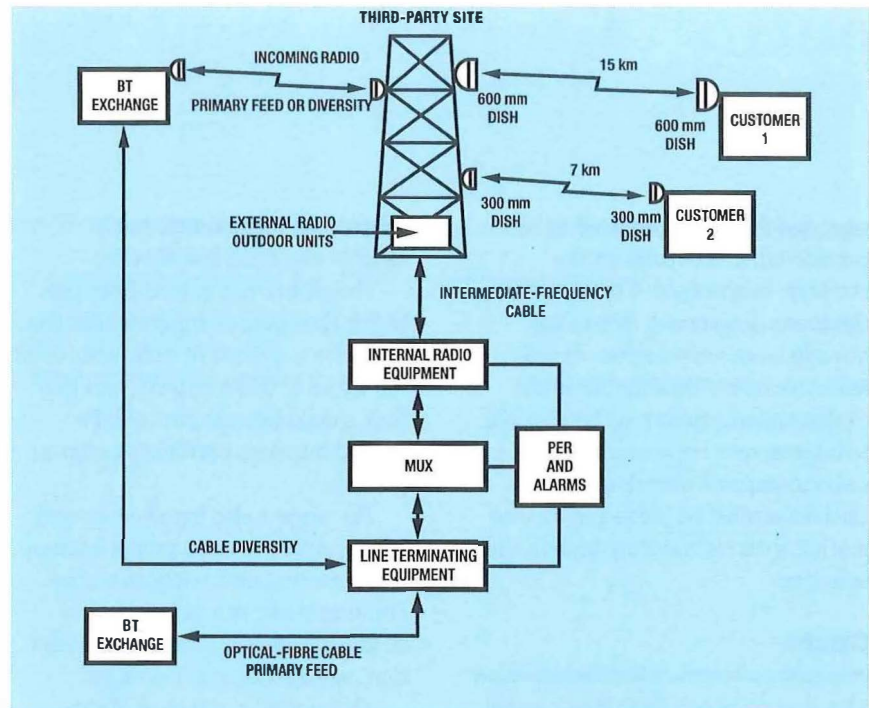
Technical Aspects— External Structure

The external structure must be capable of carrying an array of both 300 mm and 600 mm diameter dish aerials, for providing customer service at varying distances from the node.

A radio system using 300 mm diameter dishes at both ends has a range of up to 7 km, while 600 mm diameter dishes can reach customers up to 15 km away. Smaller dish aerials with a range of 2–3 km are currently on trial.

In simple terms, the strength requirement (and hence size and cost) of an external structure is proportional to the sum of the total surface area of the dishes that it is designed to carry, due to their weight and the total wind resistance.

Height, orientation, and geographical location are also taken into



account in the complex calculations involved.

As most nodes will be located in or near major cities, where most customers will be within a few kilometres of the node site, 300 mm dish aerials (or smaller) are expected to predominate. This maximises the number of dishes capable of being erected on a pre-designed structure, the systems that the node can then accommodate, and hence the overall cost per circuit of the installation.

The external structure may require structural alteration to the roof or building fabric to which it is attached, and expert structural design for the installation is essential. Planning permission may be needed, depending on the status of the structure in terms of the General Development Order, and third-party wayleaves may be required. These activities can be time-consuming, and may considerably increase the hidden costs involved in constructing the node site.

Technical Aspects— Internal and Network Requirements

If the site is an operational BT building, spare digital network capacity and power facilities are likely to be readily available.

At non-BT sites, separate accommodation is required to facilitate the provision of appropriate radio, multiplex, and line system terminal

equipment. Ancillary services such as PER power plant with stand-by batteries, alarm reporting, and MegaStream reporting, will also be required.

Although the equipment room can be up to 1 km from the external node structure, it is desirable for it to be located as closely as possible to minimise cable and installation costs.

The method of connecting each radio node back into the BT network needs to be carefully considered. As mentioned above, all existing BT exchange sites, except some in more rural locations, will already have a digital presence, and allocation of spare network capacity may be available immediately at minimum cost.

Non-BT sites may need the provision of optical cable and line systems, or alternatively an incoming high-capacity radio system may be used. Both of these options will be more costly than the BT exchange site, and may be a determining factor in site selection.

A decision needs to be made on the type of incoming line or radio system to be used. A 16 × 2 Mbit/s system may be used at small sites, with single feeds, to reduce costs, while a 34 Mbit/s system with 2/34 MUX may be used at larger sites, with diversity, to facilitate patch out (for fast change over in case of service failure).

The questions of preferred network configuration, service protec-

tion, and diversity also need to be considered, and depend on the strategic importance of the site and the customers served. Sites may already have some degree of resilience because of existing network configuration, and it may be possible to interconnect between nodes by radio to improve overall security. Each case must be judged on its own merits, with cost figuring high in the equation.

Costs

The original North East study costed several site options in detail, and concluded that the installation of radio nodes, and subsequent provision of customer service, compares favourably with the average costs of normal MegaStream provision on a circuit-by-circuit basis.

The ongoing reuse of radio equipment and dish space on the external structure reduces the whole-life cost of the installation, and the penalties of losing orders because of otherwise not meeting customer response times make the provision of radio nodes even more financially attractive.

Conclusion

In today's competitive market-place, the ability to respond quickly will, in many cases, be a deciding factor in the ability to compete at all for many major network services, and it is now more important than ever before to exploit every available technology in the quest to achieve shortest possible time-scales for service provision.

BT's traditional methods of duct and cable in its access network must be enhanced by the radio solution, which should be seen to complement its fibre partner, rather than competing with it.

Not just the right technology, but where possible its proactive and pre-tested provision, by the right people, with the right skills, are essential ingredients towards ensuring that the customer both perceives, and is

offered, not just a quick, but a reliable and fault free service.

The experience gained from the North East project suggests that the proactive provision of radio nodes will satisfy all of these criteria, and that they should become part of BT's normal business portfolio as soon as possible.

The project also involved several BT departments, and people in many of its internal processes, including building work, site negotiation, structural design, planning, installation, and commercial activities.

The mutual realisation of the importance of the project culminated in all of the above working closely together to achieve completion in extremely tight time-scales—an excellent team effort by everyone involved.

Acknowledgements

The author would like to acknowledge all Tactical and Detailed Access Planning teams in North East Zone for their contribution to this project; the Radio Planning team in North East for their contribution, assistance and technical advice. The Radio Engineering and Systems team in the Policy, Performance, and Planning directorate for their support and advice; the project team involved in the current project installation stage, for their determination to make it happen; and Norman Hails (Skillbase) who completed the original North East study.

Biography



Jim Steggles
BT Worldwide
Networks

Jim Steggles is Transmission Network Planning Manager, Worldwide Networks North East Zone. He joined BT as an apprentice in Sheffield Telephone Area in 1963, and worked on Switch and Transmission Planning duties before taking his first management role on External Planning in Middlesbrough, in 1972. He moved to Leeds, in 1979 as Executive Engineer, External Efficiency, before moving back to Transmission Planning and Works duties in the Regional and later Trunk Network offices. He became Trunk Network Regional Works Manager for Switch, Cable and Transmission, in 1985, and moved to his present duties on the company's reorganisation in 1991. His academic qualifications include a B.A., degree and membership of the British Institute of Management. He is also North East Zone Liaison Officer for the IBTE Associate Section.

Internodal Signalling

Common-channel message-based signalling is one of the key features of the modern telecommunications system. There is a need, however, for the signalling system used to evolve in a manner that allows new features to be added easily. Internodal signalling between exchanges uses CCITT Signalling System No. 7. As part of a series on the development of standards for ISDN, this article outlines the recent enhancements to this signalling system to accommodate new supplementary services and features such as intelligent network operation.

Introduction

Since the publication in *British Telecommunications Engineering*, in April 1988, of the series of articles on CCITT Signalling System No. 7 (SS No. 7), many enhancements have been made to the CCITT recommendations, and ETSI specifications have been produced to document the application of SS No. 7 in Europe. The enhancements consist mainly of the addition to the integrated services digital network (ISDN) user part (ISUP) of support for new supplementary services, the addition of message flow control and compatibility features to facilitate future enhancement, and the introduction of the new intelligent network application part (INAP) to support intelligent network (IN) operation.

Internodal signalling is the means by which telephone exchanges communicate with each other in order to set up calls between customers. Without it, even simple calls cannot be made. Internodal signalling may be considered to comprise two parts: the signalling network (which is, in effect, a data network dedicated to the control of exchanges), and the signalling messages that are conveyed across it. The ideal system is one in which the signalling network is so reliable and fast that messages always reach their destinations and with negligible delay, and in which the repertoire and content of mes-

sages is unlimited, so that new features can be easily added. SS No. 7 goes a long way towards meeting these ideals, and enhancements are being made to improve it even further.

In essence, SS No. 7 is a common-channel message-based signalling system, using 64 kbit/s channels in 2048 kbit/s transmission systems in which each channel used for SS No. 7 (that is, each signalling link) can carry the signalling for up to about 2000 speech circuits. Up to 16 signalling links, comprising a signalling route, can be provided between a pair of exchanges, giving an ultimate route traffic capacity of about 30 000 erlangs. Where this is insufficient, additional signalling routes can be provided.

SS No. 7 general architecture model

Signalling System No. 7 is specified in the form of a number of modules known as *parts*, most of which are optional. The approach is similar to the Open Systems Interconnection (OSI) Protocol Reference Model¹, in which the protocol is conceptually divided into seven *layers*. The object of this is purely to simplify the specification process, by allowing the protocol within each layer to be specified independently of that in the other layers. It also allows different upper layers to use the same type and the same instance(s) of underlying

General documentation

CCITT Recommendation Q.700, entitled *Introduction to Signalling System No. 7*, provides a general description of the principles of SS No. 7 and of SS No. 7 networks, and was revised during the 1988–92 study period.

CCITT Recommendation Q.1400, entitled *Architecture framework for the development of signalling and OA&M protocols using OSI concepts*, is a new recommendation. It provides a framework for the common development and evolution of access and internodal signalling protocol specifications using OSI concepts, and provides guidance on techniques that should be applied to the detailed process of specifying protocols.

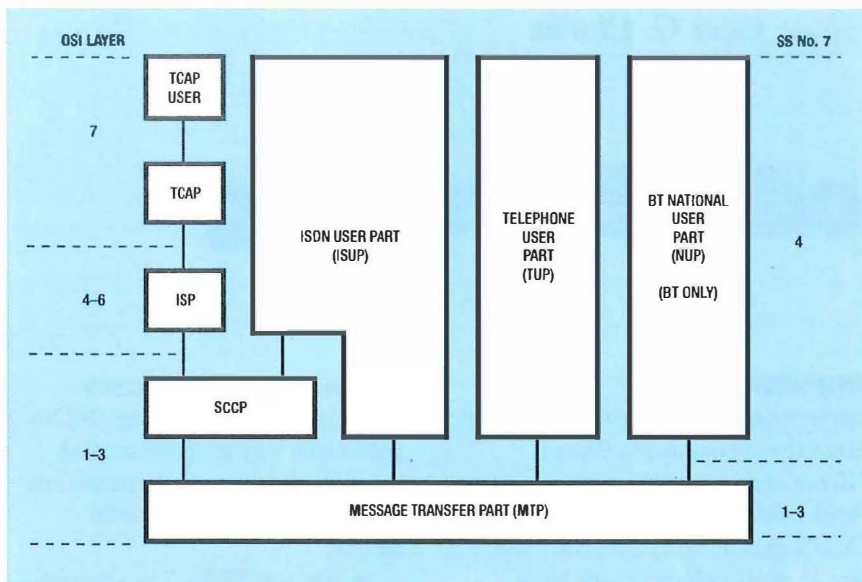


Figure 1—SS No.7 general architecture model

layer(s). Figure 1 shows how the main parts of SS No. 7 fit together to form a layered protocol, and how those parts and the SS No. 7 levels relate to the OSI Protocol Reference Model and its layers.

Technical overview

Message transfer part (MTP)² This provides a reliable message transfer between nodes, and is the only part that is mandatory.

Telephone user part (TUP) This is a CCITT standard protocol that uses the MTP to communicate with the TUP at other nodes for the purpose of establishing (and releasing) telephone connections.

National user part (NUP)³ This is BT's version of the TUP, and provides also the procedures to support a number of ISDN supplementary services, over the connections established.

ISDN user part (ISUP)⁴ This is a CCITT standard protocol for the purpose of establishing connections supporting ISDN services.

Signalling connection control part (SCCP)⁵ This provides enhanced addressing and translation features that allow the transfer of signalling messages between two indirectly connected nodes when no speech connection is required; that is, non-circuit-related signalling. In conjunction with the MTP, it provides a global message transfer service that approximates to the OSI Network layer service. It supports both

connection-oriented* and connectionless* protocol classes.

Transaction capabilities (TC)⁶ This is a simple protocol that provides features for the establishment of a signalling dialogue between nodes when no speech path is required between them. Strictly, it is considered to consist of two smaller parts: the transaction capabilities application part (TCAP), and the intermediate service part (ISP). However, in practice, the ISP has few functions and no protocol associated with it. TC uses the connectionless service of the SCCP.

Intelligent network application part (INAP) This is a TCAP user whose purpose is to provide non-circuit-related signalling capabilities between network nodes; for example, between an exchange (service switching point (SSP)) and a database (service control point (SCP)).

Operation, maintenance and administration part (OMAP) This is an example of a TCAP user which generates the non-circuit-related protocols for the management of signalling networks.

Summary of Implementation of SS No. 7 in the BT Network

In the BT national network, the existing inter-exchange signalling system uses only two of the parts of SS No. 7 mentioned above: the MTP, to provide a reliable means of conveying messages between call processors, and the NUP (BT's own version of the TUP), providing the specific call control protocols between

them. This is used between all of BT's digital exchanges (System X and AXE10) in the national network. TXE4 exchanges have also been modified to use SS No. 7.

BT has not yet implemented the SCCP or TC in the national network, but their provision will be necessary for the support of advanced IN operation.

The digital derived services network (DDSN) uses the NUP for both 5ESS-5ESS and 5ESS-System X DMSU interconnection. Between the 5ESS nodes and databases, an AT&T proprietary signalling system, comparable to early versions of SCCP and TC, is used.

The advanced services units (ASUs) employ, at Phase 1, the American National Standards Institute (ANSI) version of the ISUP for communication directly between the ASUs. At Phase 2, this will change to the ETSI ISUP. For interworking with System X DMSUs, the NUP is used.

Message Transfer Part (MTP)

History

The CCITT MTP recommendations first appeared in 1980 (Yellow Book). Many minor amendments have been made since then to enhance the operation of the system, and to clarify the recommendation text. A major enhancement to increase the standard message length capability from 62 octets (1 octet = 8 bits) to 272 octets was agreed in 1985 and published in 1988 (Blue Book).


During the 1988-92 study period, the only significant changes were the revision of the restart procedures, and refinement of the user part availability/unavailability procedures. No further changes of substance are expected.

BT implementation

The BT MTP is almost identical to the CCITT MTP and can carry user part messages up to 62 octets long

* See Glossary of Terms

Figure 3—General format of SIF portion of TUP message



In anticipation of the need to operate signalling over asynchronous transfer mode (ATM) networks, for the support of broadband services, work has already started to specify a signalling ATM adaptation layer (SAAL). Recommendations on this are likely to be produced in the near future.

User Parts

Telephone user part (TUP)

History

The CCITT TUP recommendations first appeared in 1980 (Yellow Book). Amendments have been made since then to add support for a limited set of supplementary services, and to clarify the recommendation text. The 1984 (Red Book) version supported the following supplementary services:

- closed user group,
- user access to calling line identity,
- user access to called line identity,
- call redirection,
- malicious call identification, and
- 64 kbit/s connectivity.

During the mid-1980s, the EEC analysis and forecasting group (GAP) produced a report on the coordinated introduction of ISDN in the community. The report recommended the phased introduction of ISDN service

in Europe from 1988 to 1993. However, the ISDN user part (ISUP) was not considered sufficiently stable for implementation at that time. Consequently, the decision was taken to specify an enhanced version of the TUP, referred to as the TUP+, for the support of limited ISDN service between European countries. In addition to ISDN services, this provided support for the following supplementary services:

- closed user group,
- calling line identification,
- sub-addressing, and
- user-to-user data transfer.

BT implementation

BT has not implemented the TUP *per se* in the national network, but it is employed on a large number of links to other countries worldwide. BT has not implemented the TUP+, and has no plans to do so.

Technical summary

In common with all user parts, the TUP operates from switching node to switching node, to set up connections on a link-by-link basis. As a Level 4 user of the MTP, the TUP is responsible for the contents of the signalling information field (SIF) in any messages it passes to the MTP. Figure 3

shows the general format of the SIF portion of a TUP message.

Although the destination point code (DPC) and originating point code (OPC) are used by the MTP, they must be determined by the TUP, and so are common to both Levels 3 and 4. The service information octet (SIO), value 00000100, identifies the TUP as the Level 4 user to which the MTP must deliver the message at its destination. The circuit identification code (CIC) identifies the internodal traffic circuit which is being used for that particular call.

TUP messages are identified by a pair of heading codes—H0 and H1—each of which consists of four bits, giving a potential maximum number of $16 \times 16 = 256$ message types. Table 1 shows the allocation of these codes to the messages specified for the TUP.

National user part (NUP)

History and BT implementation

The NUP⁴ was evolved from the early CCITT TUP, but modified to suit BT inland needs, in collaboration with GEC, Plessey and STC. Further divergence took place as the TUP developed somewhat slowly, and the NUP now differs significantly from the TUP. The NUP has been enhanced to support national ISDN services, in respect of which it currently supports the following:

BEARER SERVICES:

- circuit-mode 64 kbit/s unrestricted bearer service,
- 3.1 kHz audio bearer service,
- circuit-mode speech;

SUPPLEMENTARY SERVICES:

- multiple subscriber number (MSN),
- call forwarding (CF),
- calling line identity presentation (CLIP),
- calling line identity restriction (CLIR),
- connected line identity presentation (COLP),

TUP Documentation

The CCITT Recommendations concerned with the TUP are as follows.

- Q.721—Functional description of the SS No. 7 TUP; provides a general description of the TUP.
- Q.722—General functions of telephone messages and signals; lists the messages and parameters, and indicates their purposes.
- Q.723—Formats and codes; specifies the detailed format and coding of each TUP message type.
- Q.724—Signalling procedures; specifies the procedures for the establishment, supervision and release of normal telephone calls using the TUP.
- Q.725—Signalling performance in the telephone application; specifies the performance requirements of the TUP.

During the 1988–92 study period, minor changes were made to Recommendations Q.723, 724 and 725. Q.721 and 722 remain as published in the Blue Book.

The TUP+ was published in 1986, as a CEPT recommendation, T/SPS 43-02, in the form of an exceptions document to the CCITT TUP recommendations.

Table 1 TUP Message Types

HO H1	0	1 FAM	2 FSM	3 BSM	4 SBM	5 UBM	6 CSM	7 CCM	8 GRM	9	10 CNM	11	12	13	14	15
0	R						ANU									
1	R	IAM	GSM	GRQ	ACM	SEC	ANC	RLG	MGB		ACC					
2	R	IAI			CHG	CGC	ANN	BLO	MBA							
3	R	SAM	COT			NNC	CBK	BLA	MGU							
4	R	SAO	CCF			ADI	CLF	UBL	MUA	R	R	R	R	R	R	R
5	R					CFL	RAN	UBA	HGB	R	R	R	R	R	R	R
6	R					SSB	FOT	CCR	HBA	R	R	R	R	R	R	R
7	R					UNN	CCL	RSC	HGU	R	R	R	R	R	R	R
8	R					LOS			HUA	R	R	R	R	R	R	R
9	R					SST			GRS	R	R	R	R	R	R	R
10	R					ACB			GRA	R	R	R	R	R	R	R
11	R					DPN			SGB	R	R	R	R	R	R	R
12	R					MPR			SBA							
13	R								SGU							
14	R								SUA							
15	R					EUM										

ACB	Access barred	EUM	Extended unsuccessful backward set-up	NNC	National network congestion
ACC	Automatic congestion control	FAM	Forward address message group	RAN	Re-answer
ACM	Address complete	FOT	Forward transfer	RLG	Release guard
ADI	Address incomplete	FSM	Forward set-up message group	RSC	Reset circuit
ANC	Answer, charge	GRA	Circuit group reset acknowledgement	SAM	Subsequent address
ANN	Answer, no charge	GRM	Circuit group supervision message group	SAO	SAM with one signal
ANU	Answer, unqualified	GRQ	General request	SBA	Software generated group blocking acknowledgement
BLA	Blocking acknowledgement	GRS	Circuit group reset	SBM	Successful backward set-up information message group
BLO	Blocking	GSM	General forward set-up information	SEC	Switching equipment congestion
BSM	Backwards set-up message	HBA	Hardware failure oriented group blocking acknowledgement	SGB	Software generated group blocking
CBK	Clear-back	HGB	Hardware failure oriented group blocking	SGU	Software generated group unblocking
CCF	Continuity failure	HGU	Hardware failure oriented group unblocking	SSB	Subscriber busy
CCL	Calling party clear	HUA	Hardware failure oriented group unblocking acknowledgement	SST	Send special information tone
CCM	Circuit supervision message group	IAI	Initial address with additional information	SUA	Software generated group unblocking acknowledgement
CCR	Continuity check request	IAM	Initial address	UBA	Unblocking acknowledgement
CFL	Call failure	LOS	Line out of service	UBL	Unblocking
CGC	Circuit group congestion	MBA	Maintenance oriented group blocking acknowledgement	UBM	Unsuccessful backward set-up information message group
CHG	Charge	MGB	Maintenance oriented group blocking	UNN	Unallocated number
CLF	Clear forward	MGU	Maintenance oriented group unblocking	R	Reserved
CNM	Circuit network management message group	MPR	Misdialed trunk prefix		
COT	Continuity	MUA	Maintenance oriented group unblocking acknowledgement		
CSM	Call supervision message group				
DPN	Digital path not provided				

- connected line identity restriction (COLR),
- malicious call identification (MCID),
- subaddressing (Sub),
- terminal portability (TP),
- three-party service (3Pty),
- user-to-user signalling (UUS).

Currently, the NUP is used within the BT national network and for its interconnection with the networks of other licensed operators (for example, Mercury Communications Ltd. (MCL)), and other administrations, (for example, Hull).

BT, as a network operator, has signed and committed itself to the ETSI

NUP Documentation

The specification for the NUP is contained in BTNR 167.




Figure 4—General format of the SIF portion of NUP message

Memorandum of Understanding (MoU) on the Implementation of European ISDN Service by 1992. MCL has also signed the MoU. Consequently, the UK network, along with those of the other member states, is required to provide at least the services identified as having priority (*starred services*) in the MoU, by December 1993 at the latest, and preferably by December 1992. Publication of the ETSs for the support of these

services, and of the basic call, by ISUP is now complete.

Technical Summary

The general format of the SIF portion of a NUP message is shown in Figure 4.

The SIO, value 10000100, identifies the NUP as the Level 4 user to which the MTP must deliver the message at its destination. The use of

Table 2 NUP Message Types

H0 H1	0	1	2	3	4	5	6	7	8
0	IAM	*	*	Address complete	Answer	Circuit free	*	Confusion	
1	IFAM	*	*	*	Clear	Blocking	*	ISDN composite service information	Protocol negotiation
2	SAM	*	Send 'N' digits	Congestion	Re-answer	Unblocking	*	Send service	Enveloped ISUP
3	FAM	*	Send all digits	Terminal congestion	Release	Blocking acknowledgement	*	Service	Enveloped TCAP
4		Additional set-up information	Send additional set-up information	Connection not admitted	Coin and fee check	Unblocking acknowledgement	*	Additional call information	
5				Repeat attempt	Operator override	Overload	*	Operator condition	
6				Subscriber engaged	Howler	*	*	User-to-user data	
7				Subscriber out of order	Extend call	*	*	Swap	
8				#	*	*	#	#	
9				R	R	*	#	#	
10				Call drop back	R	*	#	#	
11					R	#		Nodal end-to-end data	
12					R	#		User-initiated suspend	
13					R	#		User-initiated resume	
14						#			
15						#			
16						#			
130									Enveloped ISUP segment

IAM Initial address message
 IFAM Initial and final address message
 SAM Subsequent address message
 FAM Forward address message group

* Allocated for intra-exchange use only
 R Reserved for UK use by suppliers
 # Previously allocated to BT use, now redundant

H0 codes 9-255 not allocated
 H1 codes 17-129 and 131-255 not allocated

all other parameters is as explained for the TUP.

Although BT plays a leading role in the development of the CCITT recommendations, there are some messages and procedures that are peculiar to the BT network (for example, for operator services), and to expect their explicit inclusion in an international standard would not be reasonable. NUP messages are identified by a pair of heading codes: H0 and H1, each of which consists of one octet, giving a potential maximum number of $256 \times 256 = 65\,536$ message types. Table 2 shows the allocation of

these codes to the messages currently specified for the NUP.

A detailed description of the operation of the NUP may be found in earlier issues of the *Journal*³.

ISDN user part (ISUP)

History

The CCITT evolved the ISUP⁴ to support ISDN services, and it first appeared in 1984 (Red Book).

The CCITT standard SS No. 7 user part for the support of ISDN services is the ISDN user part (ISUP).

The CCITT 1988 (Blue Book) ISUP contains some options not required in Europe, and supports only six supplementary services:

- direct dialling in (DDI),
- calling line identity presentation (CLIP),
- calling line identity restriction (CLIR),
- call forwarding (CF),
- closed user group (CUG), and
- user-to-user signalling (UUS).

It omits four supplementary services and some other features that are required in Europe to

support the requirements of the European ISDN MoU. ETSI has therefore produced an exceptions document to the CCITT 1988 ISUP to support further ISDN services, and to eliminate the options not required for international connection in Europe.

The results of the ETSI work were submitted to the CCITT, and although, in principle, CCITT recommendations are potentially applicable for both national and international implementation, this led to the production of CCITT Recommendation Q.767 explicitly for use on international

interconnections. This is a subset in that it eliminates some options and supplementary services not required, and a superset in respect of some additional supplementary services.

However, in respect of the basic call (that is, not supplementary services), the CCITT international

ISUP Documentation

The CCITT Recommendations concerned with the basic call as supported by the ISUP are Q.761–Q.764, and Q.766. The exceptions applicable to international interconnections using ISUP are specified in Recommendation Q.767. The supplementary service protocols are specified in Recommendations Q.730–Q.737. The content of each recommendation is as follows.

- Q.761—Functional description of the ISDN User Part of Signalling System No. 7; describes the applicability of the ISUP, specifies the primitives to and from the message transfer part (MTP), describes the version compatibility features and identifies the messages supported.
- Q.762—General function of messages and signals; describes the purposes of the messages used by the ISUP, and of the parameters contained within those messages.
- Q.763—Formats and codes; specifies the formats of ISUP messages including the binary coding of parameters.
- Q.764—Signalling System No. 7 ISDN User Part signalling procedures; describes the ISUP signalling procedures for the establishment and release of 64 kbit/s circuit-switched national and international connections.
- Q.766—Performance objectives in the ISDN application; specifies performance requirements for the ISUP.
- Q.767—Application of the ISDN User Part of CCITT Signalling System No. 7 for international ISDN interconnections; describes the services to be supported by international connections established using the ISUP, specifies exceptions to the 1988 ISUP recommendations when used for international interconnection, and contains additional information of relevance to this application of the ISUP.
- Q.730—ISDN supplementary services; provides general information relating to the support of supplementary services by the ISUP, and specifies the generic procedures, including error procedures, relating to supplementary services.
- Q.731—Stage 3 description for number identification supplementary services, using SS No. 7.
 - §1—Direct Dialling In (DDI)
 - §3—Calling Line Identification Presentation (CLIP).
 - §4—Calling Line Identification Restriction (CLIR).
 - §5—Connected Line Identification Presentation (COLP).
 - §6—Connected Line Identification Restriction (COLR).
 - §8—Sub-addressing (Sub).
- Q.733—Stage 3 description for call offering supplementary services, using SS No 7.
 - §2—Call Forwarding: Busy (CFB).
 - §3—Call Forwarding: No Reply (CFNR).
 - §4—Call Forwarding: Unconditional (CFU).
 - §5—Call Deflection (CD).
- Q.733—Stage 3 description for call completion supplementary services, using SS No 7.
 - §1—Call Waiting (CW)
 - §2—Call Hold (CH).
 - §4—Terminal Portability (TP).
- Q.734—Stage 3 description for multiparty supplementary services, using SS No 7.
 - §1—Conference calling (Conf).
 - §2—Three Party service (3Pty).
- Q.735—Stage 3 description for community of interest supplementary services, using SS No 7.
 - §1—Closed User Group (CUG)
 - §3—Multi-Level Precedence and Pre-emption (MLPP).
- Q.737—Stage 3 description for additional information transfer supplementary services, using SS No. 7.
 - §1—User-to-user signalling (UUS).

During the 1988–92 study period, Recommendations Q.730, Q.761, 762, 763, 764 and 766 were revised; and new Recommendations Q.731, 732, 733, 734, 735, 737, and 767 were produced.

Version 1 of the ETSI ISUP has been published as ETS 300 343. Reference should also be made to the relevant ETSI specifications for Version 2: ETS 300 360 (basic call) and 300 356 (supplementary services).




Figure 5—General format of the SIF portion of an ISUP message

ISUP subset specified in Recommendation Q.767, is virtually identical to Version 1 of the ETSI ISUP. ETSI ISUP Version 2 has now been drafted, and aligns closely with the CCITT 1992 ISUP.

BT was instrumental during the 1988–92 CCITT study period in enhancing the basic ISUP to include an essential but simple message segmentation procedure, and forward compatibility features to provide a solid but flexible platform for future enhancement.

As far as the protocol is concerned, support of any supplementary service is optional. Also the protocol is designed to operate properly even if not all nodes in a network support all of the supplementary services. During the 1989–1992 CCITT study period, the additional protocols to support some further supplementary services (including all except four of the MoU services), and enhancements to existing services, were added to the ISUP for publication in the 1992 recommendations.

The services supported by the 1992 ISUP are as follows:

- direct dialling in (DDI),
- multiple subscriber number (MSN),
- calling line identity presentation (CLIP),
- calling line identity restriction (CLIR),
- connected line identity presentation (COLP),
- connected line identity restriction (COLR),
- subaddressing (Sub),
- call forwarding: busy (CFB),
- call forwarding: no reply (CFNR),
- forwarding: unconditional (CFU),
- call deflection (CD),
- call waiting (CW),
- call hold (CH),

- terminal portability (TP),
- conference call (CC),
- three-party service (3Pty),
- closed user group (CUG),
- multi-level precedence and pre-emption (MLPP),
- user-to-user services 1–3 (UUS1–3).

Technical summary

With regard to the format of ISUP messages, if all of the parameters that might be needed for all supplementary services were included in the *initial address message* (IAM), the message would be very long. Furthermore, on any particular call, use would be made of only a few of these additional parameters. This would render the signalling system rather inefficient. For this reason the ISUP uses a message format consisting of three distinct parts:

- mandatory fixed part, containing a sequence of mandatory parameters of fixed length;
- mandatory variable part, containing a sequence of mandatory parameters of variable length; and
- optional part, containing a set of optional parameters of variable length.

Figure 5 shows the general format of an ISUP message.

The ISUP messages are specified in a single group of up to 256 message types. So far, only 56 messages have been specified, and they are listed in Table 3.

The format of the mandatory fixed part is conventional (that is, like the TUP) with coding fields of finite size, and in a predetermined order.

In the mandatory variable part, the start of each individual parameter is indicated by means of a pointer, which

indicates the number of octets between itself and the parameter concerned. Because these parameters are mandatory and the order is predetermined, the name of each parameter is implicit in the order.

In the optional part, each parameter is coded using the 'name, length, value' technique, in which each parameter to be sent in a message is assembled in three parts; the first is the parameter name, the second is the length in octets, and the third is the essential content of the parameter. These parts of the whole parameter are then transmitted in that order. Other parameters may then be similarly transmitted. The order in which the set of parameters is transmitted is immaterial.

A detailed description of the operation of the ISUP may be found in earlier issues of the *Journal*⁴.

In order to support broadband services, recommendations for a broadband version of the ISUP (B-ISUP) have been prepared in draft form by the CCITT. These are undergoing refinement, and final versions to support the initial service offering (Release 1) were completed at the end of 1993.

Separation of call control protocol from bearer connection control protocol

Some new supplementary services, for example, completion of call to busy subscriber (CCBS), could be supported more efficiently than at present, by simply sending a message to the remote node without concurrent establishment of a speech connection (particularly as it will not be used if the called party is busy). This requires 'separation' of the call

Message Type Code	Message Type Name
1	Initial address
2	Subsequent address
3	Information request
4	Information
5	Continuity
6	Address complete
7	Connect
8	Forward transfer
9	Answer
10	Reserved (used in 1984 version)
11	Reserved (used in 1984 version)
12	Release
13	Suspend
14	Resume
15	Reserved (used in 1984 version)
16	Release complete
17	Continuity check request
18	Reset circuit
19	Blocking
20	Unblocking
21	Blocking acknowledgement
22	Unblocking acknowledgement
23	Circuit group reset
24	Circuit group blocking
25	Circuit group unblocking
26	Circuit group blocking acknowledgement
27	Circuit group unblocking acknowledgement
28	Reserved (used in 1988 version)
29	Reserved (used in 1988 version)
30	Reserved (used in 1988 version)
31	Facility request
32	Facility accepted
33	Facility reject
34	Reserved (used in 1984 version)
35	Reserved (used in 1984 version)
36	Loop back acknowledgement
37	Reserved (used in 1984 version)
38	Reserved (used in 1984 version)
39	Reserved (used in 1988 version)
40	Pass along
41	Circuit group reset acknowledgement
42	Circuit group query
43	Circuit group query response
44	Call progress
45	User-to-user information
46	Unequipped CIC
47	Confusion
48	Overload
49	Charge information
50	Network resource management (new for 1992 version)
51	Facility (new for 1992 version)
52	User Part test (new for 1992 version)
53	User Part available (new for 1992 version)
54	Identification request (new for 1992 version)
55	Identification response (new for 1992 version)
56	Segmentation (new for 1992 version)
57-223	Spare
224-255	Reserved for national use

Table 3 ISDN User Part Message Types

control protocol from the bearer control protocol, specifying them in separate modules. This modularity principle could also be usefully extended to the specification of separate protocol modules for basic call control and supplementary services. Additional modules may also be specified for the control of additional bearers; for example, broadband, $n \times 64$ kbit/s.

The principle of protocol separation is being applied to the later releases of internodal signalling for broadband, although recommendations are currently only in draft form. Separation could be applied to the narrowband (64 kbit/s) network by enhancing the INAP to provide call control, in conjunction with the ISUP for bearer connection control. These aspects are under discussion in the ITU-TS (formerly CCITT) and ETSI.

Non-Circuit-Related Signalling (NCRS)

The signalling procedures of conventional user parts such as the ISUP, NUP and TUP are based on the assumption that a circuit-switched connection is being established concurrently with the signalling. For this reason it is referred to as *circuit-related signalling*, and carries the information needed to identify the transmission circuits used to make the connection. IN operation will require the transfer of information between the SSP and the SCP, where no speech path is provided or required. The IN therefore introduces the need for signalling between nodes in the absence of a speech connection. Also, because this request for information and the response will generally need to be completed in the interval between the customer's call request and establishment of the call, the time taken will be critical. The use of circuit-related signalling for this purpose would be very inefficient. For this reason, the SCCP has been specified by the CCITT to provide NCRS.

In contradistinction to the user part signalling, the nodes between

which NCRS is used may not be directly connected by a signalling link. The originating node may not know the physical destination address, and intermediate translation may therefore be required to route NCRS messages to the correct destination. Also, the SCP may be duplicated for security, and consequently the specific address of the active SCP may not be known to the originating SSP. Under both of these circumstances, the SCCP provides the correct routing translation facilities. In OSI terms, the SCCP uses the services of the MTP for the lower layer functions, and provides services to higher layers that approximate to a Network layer (3) service. This service is in turn used by transaction capabilities (TC), which is a simple protocol specified by CCITT to provide dialogue control and a means of combining requests for information between the same pair of nodes.

NCRS will be needed for efficient support of advanced IN operation, mobile telecommunication services, and some supplementary services. The specific application protocols to support IN-based services will be embodied in the INAP. This protocol standard is specified by the CCITT, and presently supports IN Capability Set 1 (CS1). However, several versions exist and the standard is not yet sufficiently refined for confident implementation. Further work is planned to refine the standard and to add further capability sets.

The adoption of international standards for NCRS poses no particular problems for BT as far as the protocols themselves are concerned, beyond the provision of the TC, the SCCP, and a 272-octet MTP at all exchanges using the feature. However, the behaviour of NCRS traffic is not limited by the dimensioning of the PSTN/ISDN in the way that circuit-related signalling traffic is. Appropriate signalling network structure, monitoring and control will be needed to ensure that NCRS traffic does not overload the signalling network. These aspects are under study.

Signalling connection control part (SCCP)

History

In order to support the remote interrogation of centralised databases, updating of visitor location registers in mobile telecommunication, and remote activation of supplementary services, the need was identified for a flexible data transfer mechanism for use within telecommunication networks. Such applications do not require the concurrent establishment of a speech connection. Therefore, the use of conventional user parts, such as the TUP, NUP or ISUP would be inefficient for these purposes.

The MTP can route messages only to the physical location indicated by the DPC of the message. Within that node, the MTP can distribute messages to any one of a maximum of 16 entities. The limit of 16 will be insufficient for future applications; for example, to IN databases. IN databases will need to be duplicated for security (for example, to allow updating of software), and duplicates may not be at the same physical location.

Generally, signalling point codes have significance only within their own signalling network; for example, the UK network. The Pan-European Cellular Radio (PECR) system will require means to send visitor location updating messages between location registers in different countries, and therefore crossing signalling network boundaries.

The SCCP recommendations were first published in the CCITT Red Book (1984). Further refinement followed, resulting in the Blue Book version (1988). During the 1988–92 study period, BT realised that use of the connectionless service of the SCCP to convey messages to support the mobile telephone service would require a message length capability exceeding that currently specified. The addition of new messages and procedures to allow message segmentation in support of this requirement

was eventually agreed. Procedures to specify SCCP restart, and other refinements were also introduced.

BT implementation

A subset of the SCCP functions (translation and routing) will be needed in the BT national network to support advanced IN operation and some new supplementary services (for example, call completion to busy subscriber (CCBS)), in which it is necessary to address messages directly between nodes which are not connected by direct signalling links. It will be introduced as and when required to fulfil these purposes.

Technical summary

The SCCP⁵ was conceived with the objectives of transferring data in a connectionless mode, establishing temporary connections, and/or using permanent signalling connections, and, incidentally, aiming to provide an OSI Network layer service. The services provided by the SCCP to its users are specified in the form of four protocol classes as follows:

- Class 0—connectionless, message sequence not guaranteed;
- Class 1—connectionless, message sequence guaranteed;
- Class 2—basic connection-oriented; and

- Class 3—connection-oriented with flow control, and detection of message loss and mis-sequencing.

In the two connectionless modes of operation, the SCCP transfers messages independently of each other, and consequently provides no means of correlating one with another, or of verifying the ability of the destination user to receive them. The inability of the SCCP to deliver a message to the destination user may be indicated to the originating user by means of a further connectionless message. Some applications, particularly those involving the transfer of large amounts of data, benefit from the establishment of virtual, or logical, connections. This mode of operation benefits both the network and the user because it verifies the ability of the destination user to receive the data at that time. The SCCP protocol class is determined dynamically by the SCCP user by the use of the appropriate request primitive.

The SCCP regards all of its local users as subsystems, and they are addressed using subsystem numbers (SSNs). Provision is made for translation from an appropriate global title (global address), to the signalling point code of the destination SCCP (DPC) plus an SSN to identify the destination user. The global title may

SCCP Documentation

The following CCITT recommendations are those concerned with the SCCP:

- Q.711—Signalling Connection Control Part; provides a general description of the SCCP.
- Q.712—Definition and function of SCCP messages; describes the functions of, and specifies the primitives associated with, the SCCP (including SCCP management) messages.
- Q.713—Formats and codes of SCCP messages; specifies the formats and codes of the SCCP (including SCCP management) messages.
- Q.714—SCCP procedures; specifies the SCCP procedures including those of SCCP management.
- Q.716—SCCP performance; specifies the performance requirements of the SCCP.

All of the recommendations listed above were revised during the 1988–92 study period.

The ETSI standard for the SCCP (connectionless protocol classes only) is contained in ETS 300 009. The BT requirements are contained in BTNR 145.

consist of a national or international telephone (or ISDN) number, a Telex number, a data network terminal number, or a number from any other specified scheme. Which numbering scheme(s) is/are supported in a particular SCCP instance is determined at the planning stage. The translation facility is of particular benefit in an IN, where databases are duplicated for security. Where these operate in a main/stand-by mode, the identity of the operational database is not known at the origin of the enquiry. In this case, the enquiry is routed by global title to an SCCP that is in close proximity to the database pair, and has knowledge of their status. This SCCP then may supplement or replace the global title with the DPC and SSN of the currently operational database.

The general format of an SCCP message is similar to that of an ISUP message as shown in Figure 5, each optional parameter being coded using the same 'name, length, value' technique.

The SCCP messages are specified in a single group of 16 message types, as shown in Table 4.

In addition to the messages listed in Table 4, five messages are specified for subsystem management.

A detailed description of the operation of the SCCP may be found in earlier issues of the *Journal*⁵.

Transaction capabilities application part (TCAP)

History

The transaction capabilities (TC) protocol was developed by the CCITT as a simple general-purpose application protocol for non-circuit-related signalling (NCRS). CCITT recommendations for TC first appeared in 1988 (Blue Book). The main change in the 1988–92 study period concerned the addition of dialogue handling procedures to allow the negotiation of application context. The addition of these procedures and the consequential enhancements to the TC messages are optional and will not cause

Table 4 SCCP Message Types, and their Applicability to Protocol Classes

Message Code Type	Message Name	Protocol Class			
		0	1	2	3
1	Connection request (CR)			*	*
2	Connection confirm (CC)			*	*
3	Connection refused (CREF)			*	*
4	Released (RLSD)			*	*
5	Release complete (RLC)			*	*
6	Data form 1 (DT1)			*	
7	Data form 2 (DT2)				*
8	Data acknowledgement (AK)				*
9	Unitdata (UDT)	*	*		
10	Unitdata service (UDTS)	*	*		
11	Expedited data (ED)				*
12	Expedited data acknowledgement (EA)				*
13	Reset request (RSR)				*
14	Reset confirm (RSC)				*
15	Protocol data unit error (ERR)			*	*
16	Inactivity test (IT)			*	*
17	Extended unitdata (XUDT)	*	*		
18	Extended unitdata service (XUDTS)	*	*		
19–255	Spare				

compatibility problems when interworking with Blue Book versions. The Blue Book recommendations specified only a connectionless mode of operation, leaving open the question of whether a connection-oriented version should subsequently be

specified. No need has been identified for connection-oriented operation, so all references to connection-oriented operation have now been deleted. Other, less significant, refinements were also made during the 1988–92 study period.

TC Documentation

The following CCITT recommendations are concerned with TC.

- Q.771—Functional description of transaction capabilities; provides a general description of TC.
- Q.772—TC information element definitions; describes the functions and specifies the primitives for TC users.
- Q.773—TC formats and encoding; specifies the formats and codes of TC messages.
- Q.774—TC procedures; specifies the procedures for TC.
- Q.775—Guidelines for using TC; provides guidelines to users of TC for specifying applications of TC.

The latest version of the ETSI standard is contained in ETS 300 287. The BT requirements for TC are contained in BTNR 140.

BT implementation

TC will be introduced into the BT national network, concurrently with the SCCP, in conjunction with applications of IN operation as and when required.

Technical summary

The TCAP protocol⁶ has been developed to provide end-to-end dialogue control, and is used in conjunction with the SCCP to satisfy the above needs. It is a simple, application-independent protocol ideally suited to the transfer of non-circuit-related information. It may be considered to be structured in two sub-layers, the lower sub-layer being the transaction sub-layer (TSL), and the upper being the component sub-layer (CSL). Five types of message are specified for the TSL:

- *Begin* This initiates a transaction with a remote node.
- *Continue* This is used to send more information relating to an established transaction.
- *End* This terminates an established transaction that has served its purpose.
- *Unidirectional* This is a 'one-shot' message that does not result in the establishment of a transaction, and requires no acknowledgment at the TC level.
- *Abort* This terminates an established transaction that did not reach a satisfactory conclusion.

Messages relating to the same transaction are correlated by the use of references similar in principle to those used in business correspondence.

Five types of information element (component) are specified for the CSL:

- *Invoke* This requests action by, or information from, a remote user.
- *Return Result—Last* This confirms completion of action requested, and/or contains all information, or the last segment of information, requested in a previous *invoke* component.

- *Return Result—Not Last* This contains an intermediate or partial response to an *invoke*.
- *Return Error* This indicates that the *invoke* was not successful, and the reason can be indicated in associated parameters.
- *Reject* This indicates rejection of the *invoke* request because of a protocol error, for example.

TSL messages are formatted using the 'name, length, value' technique. The CSL protocol data units are sent as parameters of the TSL.

Examples of the applications of TC will include the support of:

- IN-based services;
- local-exchange-based supplementary services;
- credit-card validation services;
- mobile telephony, for conveying location updating;
- universal personal telecommunications (UPT) and number portability;
- SS No. 7 operations and management functions (checking routing tables).

Intelligent network application protocol (INAP)

History

Rapid progress has been made in CCITT in the 1988–92 study period to specify a protocol to support initial implementations of INs. The initial version of the intelligent network application protocol (INAP) has been specified to support the set of services included in Capability Set 1 (CS1). However, it is recognised that the recommendation is rather immature, and that there are ambiguities that could lead to interworking difficulties

between implementations developed independently.

BT implementation

It is intended that BT will implement the INAP when the need arises to provide advanced IN capabilities in the BT national network.

Technical summary

It is intended that the operation of the IN, including a description of INAP will be the subject of a future article.

Internodal signalling for mobile communications

History

An attempt was made to establish a global standard for the protocols to support the communication needed between location registers and the mobile switching centres. This resulted in the development of the 1988 CCITT Blue Book Recommendation Q.1051, Mobile Application Part (MAP) of SS No. 7, that describes the application protocol for the transfer of mobile telephone location information between mobile switching centres (MSCs), home location registers and visitor location registers to support the mobile telephone service. However, this was a first attempt, and some inconsistencies have been identified. Also, the three regional implementations have diverged, and the need for a global standard is by no means overwhelming. Consequently, it was agreed in the 1988–92 study period that Recommendation Q.1051 should be deleted, and replaced by references to the three regional standards. Further work in the ITU on protocols for the support of mobile communications is likely to be based on the INAP.

IN Documentation

Q.1218—IN Interface Recommendation; describes the SS No.7 INAP. Reference should also be made to ETS 300 374-1, the Core INAP, for implementation throughout Europe to support CS1.

MAP Documentation

The current MAP signalling interface standard of relevance to the European Global System for Mobile (GSM) communication is the ETSI ETS 300 303.

BT implementation

The GSM service in the UK is not operated by BT, but by Telecom Securicor Cellular Radio (TSCR) and by Racal Vodafone. BT *per se* therefore has no implementation of the MAP protocol.

Technical summary

Description of the signalling to support the mobile telephone service would not be appropriate without an associated description of the operation of the mobile telephone service as a whole. This may be the subject of a future article.

Signalling System No. 7 Management and Testing

History

Significant progress has been made in CCITT Study Group XI in the 1988–92 study period in developing the SS No. 7 management philosophy and in its integration with the telecommuni-

cation management network concept of Study Group IV. The Blue Book recommendations concerned with monitoring and measurements (Q.791) and the OMAP (Q.795) have been deleted, and replaced by a new set of recommendations providing a more comprehensive view of SS No. 7 management as a whole.

BT implementation

BT has adopted a rigorous approach to the testing of implementations of SS No. 7 that was demonstrated in the bilateral MTP testing that took place between the UK and Belgium in 1984. This approach has been continued with new implementations and new interconnections, and as a consequence, serious escalating failures have been avoided in the BT national network. The CCITT signalling network management standards are relatively new and incomplete, and consequently none of the testing performed by BT to date has explicitly conformed to these

standards. However, it is intended that management of the BT signalling network and the CCITT signalling network management standards will converge.

Technical summary

The potential scope and complexity of signalling network management (including testing) is considerable, and the existing recommendations are relatively immature and incomplete. Therefore, no attempt is made in this article to provide a technical summary. The topic is sufficiently comprehensive to become the subject of a separate article when mature standards have been established.

Future

Arising from competition, BT, in common with many other network operators worldwide, is experiencing increasing pressure to offer new features and services in ever short-

SS No. 7 Management Documentation

The following recommendations concerned with the management of SS No. 7 networks were produced by CCITT during the 1988–92 study period.

- Q.750—Overview of Signalling System No. 7 Management; gives an overview of the functions, procedures and entities for managing SS No. 7, and shows the relationship between TMN and SS No. 7 management.
- Q.752—Monitoring and measurements for SS No. 7 networks; identifies the measurements to be made on the signalling network.
- Q.753—SS No. 7 management functions; contains the informal text descriptions for the MTP routing verification test, the SCCP routing-verification test and the circuit-validation test.
- Q.754—Signalling System No. 7 management ASE definitions; provides the formal definition of the MTP routing verification test and the SCCP routing verification test (SRVT) parts of OMAP.
- Q.755—SS No. 7 protocol tests; describes the operation of the MTP tester for validation of an implementation or compatibility testing of two implementations, and the simulation of a user part for the generation of MTP test traffic. A section covering the SCCP remains to be added.

The following recommendations are the test specifications for particular parts of SS No. 7.

- Q.780—SS No. 7 test specification general description.
- Q.781—MTP Level 2 test specification.
- Q.782—MTP Level 3 test specification.
- Q.784—ISUP basic call test specification. This specifies a detailed set of tests for validation and compatibility testing of the 1988 version of the ISUP (Q.761–764). Most of these tests are applicable also to the international interface (Q.767) version of the ISUP.
- Q.785—ISUP protocol test specification for supplementary services. This specifies a detailed set of tests for the supplementary services supported by the 1988 version of the ISUP (Q.730), and by the international interface (Q.767) version of the ISUP.
- Q.786—SCCP test specification.
- Q.787—TC test specification.

Recommendations Q.780, 781 and 782 were revised by CCITT during the 1988–92 study period; Recommendations Q.784, 785, 786 and 787 are new recommendations produced by CCITT during the 1988–92 study period.

ening timescales in order to meet market windows. 'Quick-fix' solutions have been successfully employed to meet the more significant of these demands. However, to sustain and enhance this market position, it will be necessary to exploit the advantage BT has of an existing network, by upgrading its infrastructure to support new features in the most efficient manner, and introducing flexibility mechanisms that will allow rapid enhancement to meet future market demands. A key element of network flexibility is the IN concept, and this is seen as a vital aspect of the evolution of telecommunications networks throughout the world. Concerning signalling, BT has played a leading role in the international standards fora in ensuring that flexibility mechanisms are incorporated into signalling protocols. This has resulted in signalling standards that allow earlier versions of protocols to interwork directly with new versions without loss of capability, and allow additional capabilities to be added to the new versions. These capabilities, together with the new INAP, and the emerging broadband protocols, should ensure that internodal signalling is able to provide all that is required of it well into the twenty-first century.

References

- 1 JENKINS, P. A., and KNIGHTSON, K. G. Open Systems Interconnection—An Introductory Guide. *Br. Telecommun. Eng.*, July 1984, **3**, p. 86.
- 2 LAW, B., and WADSWORTH, C. A. CCITT Signalling System No. 7: Message Transfer Part. *ibid.*, Apr. 1988, **7**, p. 7.
- 3 MITCHELL, D. C., and COLLAR, B. E. CCITT Signalling System No. 7: National User Part. *ibid.*, Apr. 1988, **7**, p. 19.
- 4 DAVIES, C. G. CCITT Signalling System No. 7: Integrated Services Digital Network User Part. *ibid.*, Apr. 1988, **7**, p. 46.
- 5 CLARKE, P. G., and WADSWORTH, C. A. CCITT Signalling System No. 7: Signalling Connection Control Part. *ibid.*, Apr. 1988, **7**, p. 32.
- 6 JOHNSON, T. W., LAW, B. and ANIUS, P. CCITT Signalling System No. 7: Transaction Capabilities. *ibid.*, Apr. 1988, **7**, p. 58.

Biography

Peter Clarke joined the Post Office Engineering Department in 1962. He graduated from the University College of North Wales in 1965 with a degree in Electronic Engineering, and is a member of the I.E.E. Following several years of involvement with the procedures and equipment for the maintenance of Strowger switching and signalling systems, he joined the BT-Industry Advisory Group on Systems Definition (AGSD), contributing to the establishment of the basic requirements for transmission, switching and signalling for BT's integrated digital network. When the AGSD was replaced by the Telecommunications Systems Strategy Department (TSSD), he became responsible for the production of Strategic Network Plans specifying the performance requirements for the integrated digital network. He then spent some time in Network Strategy, Network Planning and Works, and Network Systems Engineering on the specification of BT's requirements for internodal signalling. He has been directly involved in the negotiation of international standards for digital transmission, switching and signalling since 1977.

Glossary of Terms

The purpose of this section is to serve as a reminder of the meaning of some of the basic terms used in connection with message-based signalling (and in many cases with data networks, which are similar in principle of operation). A sound grasp of these is an important prerequisite to the understanding of operation of signalling systems.

(Telecommunications) signalling

Information transfer between telecommunications network nodes concerned (directly or indirectly) with the establishment, supervision or release of telecommunication calls or connections.

Circuit-related signalling

Signalling between switching nodes (exchanges) in a telecommunications network for the purpose of establishing, maintaining and releasing a circuit-switched connection comprising a concatenation of traffic circuits interconnecting those nodes.

Non-circuit-related signalling

Signalling between switching nodes and other types of node in a telecommunications network, that takes place independently of any traffic circuits.

Channel-associated signalling

This is circuit-related signalling in which the means of information transfer has a fixed physical or logical relationship with the circuit being used to carry the telecommunications traffic.

Common-channel signalling

Signalling in which the means of information transfer has no fixed physical or logical relationship with any individual traffic circuit, but is shared dynamically with other signalling information.

(Communication) protocol

A formal set of rules specifying the functions and procedures required to facilitate communication between two nodes in a network. Its specification consists of the definition of messages, primitives (see below) and procedures.

Layered protocols

Consider two nodes with processors between which communication is required. Protocols are specified for this purpose. A feature of the layered approach to protocols is that each layer is specified separately, and

considered to provide services to the layer above, which is referred to as a *user* of those services. An aspect of particular importance is the method of distinguishing internode messages (which have to be standardised both semantically and syntactically), from signals between protocol layers within each communicating node (which need only be semantically standardised). For the purpose of writing the specification, each of these signals is referred to as a *primitive*, and each piece of essential information conveyed by it is referred to as a *parameter*.

Messages

Messages are the quintessence of modern signalling systems, being the sole means by which information is transferred from one node to another. Because each node in a network or between networks may have been obtained from different suppliers, it is necessary to specify messages precisely in terms of both semantics and syntax. Consequently, protocol specifications include detailed message formats, and coding of the individual bits representing the values of all parameters included in each message. At the sending node, each protocol layer adds its own header or envelope as the information passes down the protocol stack, rather as a letter for secure delivery might pass through several stages of enveloping, addressing and registration.

Primitives

The nature of primitives is entirely the concern of the designer of the node, and should consequently not be constrained by the protocol specification. A primitive comprises information or instructions passed between protocol layers within a node. On the assumption that the entire node is designed and built by the same supplier, there is no need for the purchaser of the system to specify, or be aware of, the electrical nature of the primitives. Only the semantics of the primitives need to be specified, including the parameters which are

essential for conveying that meaning to the other node. Primitives are usually described by the name of the layer of the underlying functional block which supports them. Each generic primitive associated with a particular function may exist in up to four basic forms, depending on the nature of the function. These are: *Request*, *Indication*, *Response* and *Confirmation*. They are used in the following way:

- *Request* This type of primitive is generated by a user to request that a specified function be performed.
- *Indication* This is passed to the user to inform that user that a specified function is being performed. The same primitive is used whether execution of the function was initiated by the network or by another user invoking the corresponding Request primitive.
- *Response* This is passed from the user to the network to indicate that the user acknowledges reception of the corresponding Indication primitive.
- *Confirmation* This is passed from the network to the user to indicate to the user that the action required as a result of passing an earlier Request primitive has been carried out.

(OSI) Protocol Reference Model

The model¹ was devised primarily so that the process of specifying communication protocols could be modularised and simplified, and so that as much use as possible could be made of any potentially common data transfer principles. The model is only a framework of principles relating to the external behaviour of a node, and is not intended to place constraints on the internal design of any implementation. For example, it is not necessary for the interface between individual layers to be realised, either in hardware or software.

Connection-oriented protocols

The means of information (or data) transfer in which an association is

established between sender and receiver before any attempt is made to transfer the data. It avoids the risk of flooding the network with undeliverable messages if the intended recipient is unable to receive them at that time. The association could take the form of a physical connection, but this is not essential. The most usual form of association is that of a logical (or virtual) connection. In this, the originator sends a request to the intended recipient of the data to ask whether the data can be accepted. If the intended recipient responds with confirmation that it is able to receive the data, the logical connection is considered to be established, and data transfer may start. A complementary procedure is used to release the logical connection. The packet-switched service uses this mode of operation. Connection-oriented procedures are useful when a large amount of data needs to be transferred.

Connectionless protocols

The means of data transfer without prior confirmation that the intended recipient is able to receive it at that time. Where small amounts of data are to be transferred, this method can be efficient, because it avoids the overhead of the additional messages which would be required to establish and release a logical or physical connection. A disadvantage is that each message needs to contain full addressing information.

Directly-Buried Pre-Cabled Sub-Duct

The cost of providing resilient optical-fibre cable routes to small rural exchanges can often be very high because of the cost of providing new duct. One means of reducing the cost in these situations is the introduction of an aerial optical-fibre cable. While this can, especially where there is an existing pole route, dramatically reduce the cost of providing a new cable, it can have the disadvantage of poor visual impact especially in areas of outstanding natural beauty and national parks.

In some circumstances, to overcome this objection while still retaining the ability to achieve economic cable provision in rural areas, a method of directly burying a pre-cabled sub-duct using a modified moleplough with a vibrating plough head (Figures 1 and 2) has been introduced.

Where route conditions are suitable, this would typically be a grass or unmade verge of peat to mixed rock, it has been shown that the cost that would have been involved had conventional duct laying techniques been employed. The speed of provision can be up to four times faster than when conventional ducting methods are used.

A typical example of using this technique was a 9.3 km 8-fibre spur

from the existing Lochgilphead-Tarbet No. 3 cable in Argyll, West Scotland. No duct existed along the route, which follows the line of a 30 mile single-track road serving scattered communities with farming and forestry interests. Apart from the cost of providing 9.3 km of conventional duct, the envisaged frequent road closures expected as part of the operation were a cause of concern to both local residents and Strathclyde regional council.

In the event, the cable was installed by using the directly-buried pre-cabled sub-duct technique and ready for service in four weeks from commencement as against an estimated 15 weeks. A saving of 46% in cost was achieved and a letter from the regional council congratulated BT in achieving minimal disruption to the environment, daily lives of local residence and compliance with both the letter and spirit of the new street works act. No closures were necessary.

A further bonus was achieved when a second pre-roped sub-duct was installed over 2 km of the route at a minimal extra cost. This is to be used to install a copper cable for the local residence and allow the existing pole route, which is prone to frequent failure in the prevailing harsh winter weather conditions, to be recovered.

Maintenance of cables installed in this manner may be carried out using

the same technique as would be employed for cable in a conventional duct. Where the effort required to recover or install a new length becomes excessive, it is a simple procedure to dig down, expose the sub-duct and cut it to enable recovery and installation in two directions. The sub-duct can be rejoined by using a conventional connector.

After successful execution in Scotland, fibre planners in mid- and North Wales are currently investigating two routes which would lend themselves to this method of construction. The terrain has similar logistic problems to Scotland. Execution is expected in the earlier part of 1994.

The pre-cabled sub-duct is a conventional sub-duct mono bore 3A extruded over a standard COF 181 type cable. The cable can be of any size from 8-96 fibres. The sub-duct is buried at a depth of 600 mm and joints are every 1.8 km. Standard JRF11 joint boxes are used at jointing points. The moleplough comprises a CASE four-wheel drive and steer tractor unit mounted with a vibrating plough head capable of maintaining a fixed depth relative to the surface and a chain cutter for opening carriage-way surfaces.

*Richard Avery and Colin Barrett
on behalf of IBTE North Wales
Centre*

Figure 1—Vibrating head



Figure 2—Vibrating head working in rocky location



Mobile Repeater Station for Restoring the Network

The Scottish Restoration Planning Group in Edinburgh was given a project to design a mobile transmission repeater station (TRS) which would be used in the event of a catastrophic failure of a node site in the public switched telephone network (PSTN). The mobile TRS would quickly restore the PSTN transmission requirements of 80% of sites in the UK before augmentation was required. In addition, the remaining 20% of 'super-node' sites could use the facility as a base for restoration. The mobile TRS would fit into a container 2.4 m by 9.0 m approximately; that is, an ISO container size. A design was submitted to BT's Core Transmission

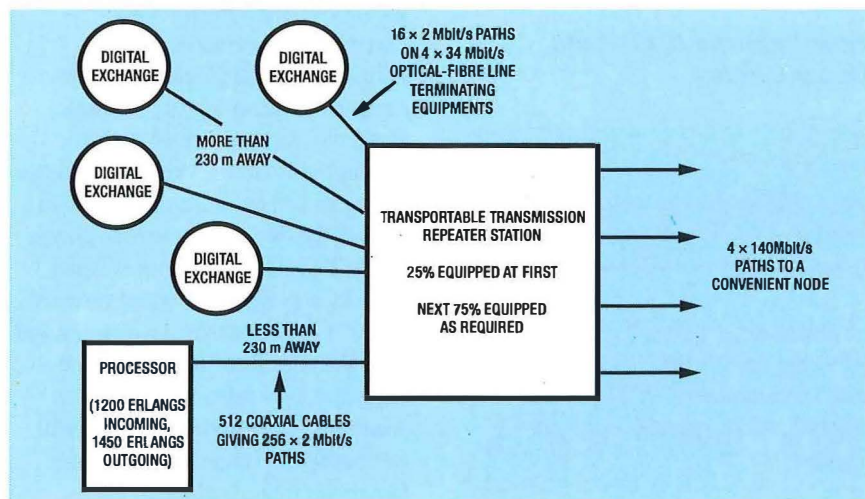
Plant Team and is now being procured by them and Group Logistics Services (GLS). It is anticipated that a number of these will be available in Autumn 1994.

It is impossible to replace a typical TRS. For example, an ordinary Scottish site serving a town's total transmission needs has an average of around 13×565 , 13×140 , 9×34 and 5×8 Mbit/s optical line systems alone, let alone the attendant muldexes. So the effort was put towards fulfilling the needs of a typical processor site, satisfying the 80% criteria mentioned above. The Exchange Dimensioning Group in Scotland suggested that a typical processor with 1200 erlangs of incoming traffic and 1450 erlangs of outgoing traffic would need 252×2 Mbit/s paths. From this information, the equipment required could be designed.

To meet the demand of 252 Mbit/s paths with the least equipment, the TRS is equipped to give 256×2 Mbit/s paths. Spare optical-fibre line terminating equipments (OFLTEs) are included to be taken away to the other end(s) to overcome the well-known manufacturer equipment incompatibility problem which has dogged plesiochronous transmission systems.

To connect attendant digital exchanges more than 230 m away (the distance at which the 2 Mbit/s signal begins to have errors), $4 \times 16/2$ OFLTE MDXs are included, again with mirror image systems 'parked' waiting to be taken to a distant end. See Figure 3 for a diagram of the connections.

Figure 3—Schematic of exchange connections



Enough rack space and cabling are provided to cater for four processors, but only 25% of the equipment is provided on cost and changing-technology grounds. An agreement has been reached with GLS to provide the 75% extra needed in an emergency. Other spare racks are provided which can accommodate MegaStream, KiloStream etc. for important customers, and a 62-type audio rack with a Krone-type frame to accommodate audio cables.

This mobile completes the 'family' of transportable equipment planned and gathered by the Restoration Planning Groups in the UK at present. They service the PSTN needs of sites facing disaster situations beyond 'business as usual' procedures. The Central Operations Unit at Oswestry controls a range of kit, from a mobile digital main switching unit (DMSU) down to a small 600-line local digital exchange, all of which can be interconnected thanks to this repeater station.

In summary, the transmission mobile:

- can carry traffic of four processors to either one site or any multiple of four;
- can act as a central point for fast onloading of four RCU/RSSs or outward to four sites, 16×2 Mbit/s per site;
- is 'manufacturer compatible' with regard to line systems;

- has facilities for fibre and audio terminations;
- provides a basic service to get sites up and running;
- answers a need highlighted in every exchange emergency exercise conducted to test procedures; and
- allows for further onloading as specific needs are identified.

Graeme Lyall for IBTE Scotland East Centre

Power Drawings Database Within Southern Home Counties

Records have always played a substantial part in any planning office, and with screen-based planning fast becoming a reality, drawing records are now an essential part of that process.

Drawing offices throughout the country have changed to computer-aided design (CAD) and many new drawings are now available on computer, although numerous older drawings still have to be converted. These are readily available to the drawing offices while planners still have paper copies.

In BT Worldwide Networks (WN) Southern Home Counties (SHC), the first steps have been taken to produce a system that is interactive between drawing office and planning office. Figure 4 shows three drawing offices based at Tunbridge Wells, Brighton and Bournemouth with the file server in the Bournemouth office. The offices are linked across BT's own multi-protocol router network (MPRN) using transport control protocol (TCP) and uses an Intergraph file management system. The MPRN works with both Unix- and DOS-based products, which is essential for this integrated system.

The six power planning offices within WN SHC have been provided with a PC and a plotter and are linked into the server at Bournemouth. The planners are able to access the drawings database, and view or print a power record for any operational building within the zone. Planners are also able to show any required amendments that they would like the

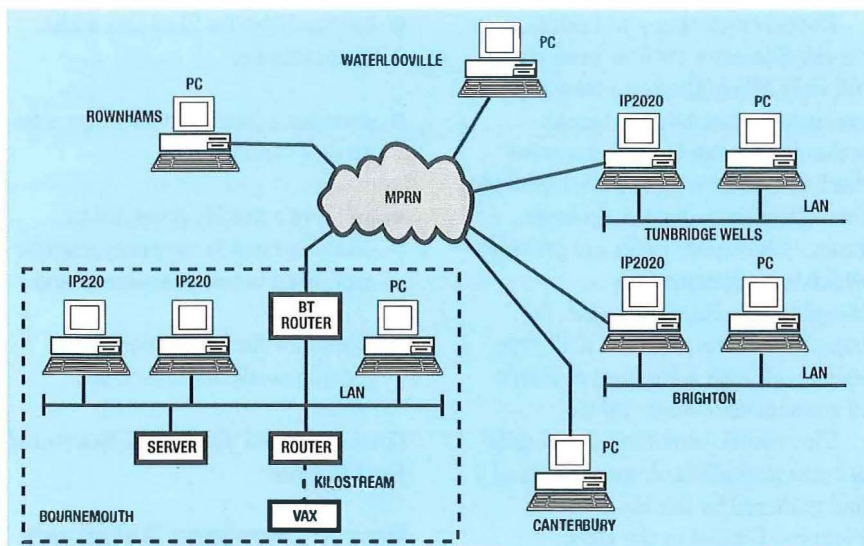


Figure 4—Drawing office and power planning network

drawing office to make without altering the existing drawing. This is achieved by drawing on a different layer to that occupied by the main drawing. This is called *red-lining*.

The advantages of the system are:

- instant access to the most up-to-date records and drawings,
- time savings in the amendment of drawing,
- balanced workload for planners and draughtsmen,
- no paper records, and
- a system that is easily monitored for quality and quantity.

At present concentration is being given to electrical distribution drawings, as there is a statutory obligation to do so under the 16th Edition of the IEE Regulations, but

the screen-based planner will require all the other types of drawing if they are to be effective in their role. One way to overcome the deficiency is to scan the paper records onto the system. These may not be to the required quality but would overcome the problem in the short term. This is being pursued.

The drawing offices controlled by Internal Planning at Brighton and Southampton will shortly be connected to the system to enable them to produce floor-plan drawings. Access to these drawings is required by power planners and the advantages of their connection to the system can be seen.

Any system such as this is only as good as the information that it contains; teamwork between the drawing and planning offices will ensure the accuracy of the information held.

Peter Baker for IBTE South Downs Centre

Field Focus Articles

Are you aware of a local BT project that would make the subject of an interesting article in the Journal's Field Focus? For example, novel solutions to field problems could make the bases of very interesting items.

If you feel that you would like to contribute an article, then please contact the Managing Editor, British Telecommunications Engineering, Post Point G012, 2-12 Gresham Street, London EC2V 7AG (Tel: 071-356 8022; Fax: 071-356 7942). Alternatively, please contact your local IBTE Centre (see p. 96 for details).

Virtuosi—Support for Virtual Organisations

A consortium of UK industrial and academic partners, led by BT, is collaborating on a project that aims to create a computer-generated environment where users can work with each other, regardless of location. Using virtual reality to support cooperative work across dispersed groups, Virtuosi will allow people to join discussions and problem-solving activities at home or in different offices or factories across the world.

Two pilot services will be developed for the project. One will be concerned with improving communication between a number of BICC cable factories throughout the world. The aim is to enable managers and workers to cooperate as if they worked in a single factory. For example, experts will be able to 'visit' shop-floor staff on another continent to resolve technical and production issues at the remote factory. A compact three-dimensional representation of the total organisation will be generated on the screen of a desktop computer. Users will be able to move around this visualisation to locate individuals and sources of expertise at remote sites. The system will then integrate video windows and computer-based data into this virtual world, to support conference calls between the participants.

The other pilot service will be based on the textile and clothing industry to allow cooperation between a number of small and medium-sized enterprises involved in the design and manufacture of garments. This application will be part of Nottinghamshire County Council's programme to support industry in the area and to provide competitive advantage to its users. Designers will use virtual-reality techniques to create a visualisation of a garment. The designer will be able to discuss the design with remote buyers who will also see the garment being worn by a model on a 'virtual catwalk'. Users will be able to make changes and view the effects during their discussions and, once agreement has been reached, the details of the design will be transferred immediately to the factory for manufacture.

Expertise from the UK company, Division Ltd, will contribute to the core framework of the project, providing powerful tools to allow users to interact within the virtual environment. BT Laboratories and GPT Ltd will be involved in developing the telecommunications aspects of the project. GPT's contribution will also be supported by GEC Marconi Hirst Research Laboratories which has considerable experience in distributed virtual-reality applications.

A strong theoretical basis will be provided for the project by Nottingham, Lancaster and Manchester Universities. These universities are all involved in UK and European research projects to establish the fundamental principles for interacting and cooperating in virtual environments that will be set up across future information superhighways. In addition to developing models and systems, they will also be involved in user-requirements capture and assessing the effectiveness of the two application pilots.

BT to Trial Interactive Services to the Home

BT has announced that it has started technical trials of a new service which promises to create a revolution in the home. The new service brings together the telephone and the television to enable customers to choose and order entertainment and information services from a menu on an ordinary television set. The material is then transmitted from a central database over the telephone network to the television set.

Potential applications are extensive and range from movies and television programmes to home shopping, educational and training material, electronic publishing, information services and video games.

Paul Reynolds, Director of BT's Information, Communication and Entertainment (ICE) Programme, explained that, because the system was interactive, customers would be able to call up a movie when they wanted it. Eventually they would be able to order products seen on the screen, or change a standing order with their bank, just by pressing a button. All this would be

possible while still being able to use the telephone.

Called *interactive multimedia services*, the new concept is undergoing technical trials at Kesgrave, near Ipswich, and will involve up to 70 BT employees. In the technical trial, video material in digital format will be sent along existing copper wire using a technology called *asymmetric digital subscriber loop* (ADSL). Fibre-optic links will also be tested as an alternative to copper. Following successful completion of the trial, BT will conduct consumer trials with 2500 households in the autumn and is currently evaluating potential sites.

Various major companies are working with BT to supply a range of programming for the technical trial. These include the BBC, Carlton Communications, Granada Television, London Weekend Television, Picture Music International and Thames TV. BT is also working with suppliers to obtain movies from the Hollywood studios.

BT is working closely with Northern Telecom, which has been awarded the contract for ADSL equipment; Alcatel Network Systems, which will supply much of the fibre technology; Oracle which is providing the multimedia server software; nCUBE which will provide the multimedia server hardware; and Apple which is supplying the set-top box.

New British Technology Brings Tomorrow's Information Superhighways Closer

BT has pioneered a world-beating device at its Martlesham Heath Laboratories, which will increase the capacity of current optical-fibre links from around 1900 simultaneous telephone conversations to a staggering 70 000. And the technology has the potential to handle 150 000 calls over 150 km of continuous fibre-optic cable.

The new product will play a key role in tomorrow's information superhighways in allowing existing 1.3 micron fibre optics—of which there are more than three million kilometres in BT's UK network—to achieve far greater data rates over much longer distances. These superhighways could carry a mixture of analogue and both

compressed and uncompressed digital signals to make it technically possible to deliver such services as three-dimensional television, interactive multimedia services (sometimes referred to as *video on demand*), high-quality audio and computer data to the office or home.

Previously, BT researchers played a key role in the development of 1.5 micron amplifiers, which led to increases in capacities in the longer wavelength 'window'. The new breakthrough increases the signal strength by 1000 times at 1.3 microns in the region of the spectrum where distortion of the pulses is about 10 times lower. The new optical amplifier is compatible with the huge installed optical-fibre infrastructure and will allow the full potential bandwidth of BT's existing fibre-optic investment to be unleashed in a crucial move towards ubiquitous broadband services.

The amplification process is based around transferring energy from a local energy source to the weak 1.3 micron signal entering the amplifier. This is achieved by doping an advanced type of optical fibre, known as *fluoride fibre*, with a chemical element called *praseodymium*. The energy transfer is mediated by the praseodymium ions which are incorporated into the core of the fluoride optical fibre. The energy source which provides the energy needed for amplification is a small solid-state laser which emits light at a wavelength of 1.047 microns, a wavelength which can excite the praseodymium ions in the fibre core.

Light at a wavelength of 1.047 microns is combined with the weak 1.3 micron signal at the amplifier input and both optical signals are coupled into the praseodymium-doped fluoride fibre. As these two optical signals travel through the doped fibre, the 1.047 micron light is absorbed by the praseodymium ions and a proportion of its energy transferred across to the 1.3 micron light. In this way the 1.3 micron signal grows as it travels through the amplifier while the 1.047 micron light decays. Typically, 20 m of doped fluoride fibre is sufficient to produce amplification factors of $\times 1000$.

Sixth Queen's Award for BT Laboratories

BT Laboratories (BTL) has again won the Queen's Award for Technological Achievement. This year's award goes to the Access Networks Division for blown-fibre technology in the telecommunications network.

This technology is now well-established but is set to play a major role in enabling BT to take optical fibre directly into people's homes and businesses to deliver a new and exciting range of multimedia services.

Blown-fibre technology allows lightweight optical fibres to be 'floated', using a small air compressor, into plastic tubing which has already been installed in underground ducts. The process eliminates the risk of fibre damage and allows it to be quickly, safely and economically installed in the most complicated routes.

Optical fibre, with virtually unlimited transmission capacity is well-established as the preferred medium for long-distance communications between major exchanges. Optical fibre connected straight into the home or business is the next logical step and will offer communications capabilities which will transform business practices and home entertainment—and may in fact change the way many people conduct their lives.

Chris Wheddon, Director of Network Services at BTL, said: 'Blown fibre is an excellent example of BT investing in the future, and it will allow BT to install in the UK the intelligent multimedia broadband network, popularly known as the *digital information superhighway*.'

Blown-fibre technology will also play a vital part in BT's passive optical network trials and video-on-demand trials—at the 'leading edge' of fibre entry into the access network.

BTL has led the world in optical-fibre technology and has received three previous Queen's Awards for Technological Achievement that recognised the work in the field of optical fibre and optoelectronics devices. Blown-fibre technology has been patented throughout the world by BT and licences have been granted in Europe, USA and the Far East.



BT technician at the kerbside uses blown-fibre equipment to install optical fibre into the telecommunications network

BT Global Network Services Helps to Contain Nuclear Disasters

BT's Global Network Services (GNS) will be providing key networking and electronic messaging elements of the Government's Radioactive Incident Monitoring Network (RIMNET), the second phase of which was launched in April.

RIMNET is a £12 million nuclear emergency response facility which will provide a basis for Government decision if an overseas nuclear accident affects the UK. RIMNET Phase 2, which replaces the interim Phase 1 system, will collect data about levels of radioactivity, display it to provide a clear picture of the developing situation and communicate Government advice rapidly to the public, media and those involved in the response to the accident.

In essence, RIMNET Phase 2 consists of a network of 92 gamma radiation dose monitors, sited at Met Offices throughout the UK. The monitors are linked to the main central database facility (MCDF) in London, and a back-up installation in Poole, Dorset, via BT's GNS. Every hour, the MCDF polls the monitors for their latest readings. The data is processed at the MCDF ready for access and assessment by a team of experts located at the DoE's Technical Coordination Centre (TCC). The TCC team is then responsible for making the analysed information available to

various government departments and authorities, and other approved users.

In addition to the monitor system, RIMNET Phase 2 also receives supplementary data on other forms of pollution (such as mercury or nitrate levels) from approved data suppliers via BT's Messaging Services. The approved data suppliers, which might include universities, or local government or environment offices, record their own local findings onto a standard questionnaire, which is then transmitted directly to the MCDF. In return, they are able to access and download information from the central database.

BT Global Network Services Launches ISDN Dial Back-Up Service

BT Global Network Services (GNS) has launched a back-up service, based on the integrated services digital network (ISDN), for digital network customers. The service is called ISDN Dial Back-Up (DBU).

The growing trend is for managed data network services to be used for business-critical applications where even the most minimal failure rate would be unacceptable. The introduction of ISDN DBU now provides the ultimate in network resilience.

The ISDN DBU services is currently available to UK customers of the BT GNS LAN Interconnect service at speeds of 64 kbit/s. Service is to be expanded to cover X.25 customers and be rolled-out to European LAN Interconnect customers later in 1994.

To take advantage of the service, customers need an ISDN2 connection at their premises. An ISDN terminal adaptor is provided and maintained by BT, and each customer connection has a dedicated channel at the installation. Should a network access failure be detected, the ISDN DBU service is automatically initiated and continues until the fault is rectified. As part of a managed service, the ISDN call charges are borne by BT, and the reliability of the service is guaranteed by regular polling of the ISDN DBU link. As an ISDN2 link provides two 64 kbit/s channels, and

only one is required for a single back-up link, an additional benefit is that the customer can use the spare circuit for another application such as file transfer.

Safeway Selects BT's EDI*Net Service to Improve Its Supply Chain

Safeway Stores Plc, one of the UK's leading supermarket chains, has introduced BT's electronic data interchange (EDI) service, EDI*Net, to enhance the relationship with its 2000 plus suppliers.

Safeway is committed to working more closely with its suppliers to streamline the supply chain; and sees EDI as a tool to help build such relationships through the provision of accurate data in the most effective way.

EDI has already been introduced for Safeway's ordering process with many of its suppliers. Starting in February 1994, a number of other applications were introduced, including order confirmation, advance delivery notifications, proof of delivery, invoices, product and pricing information and payments.

Neil Lawrence, BT Commercial manager of EDI*Net explained that the first phase of Safeway's roll-out programme during 1994 would concentrate on setting up EDI links with its remaining non-EDI suppliers. Once this had been achieved, the company would focus on providing those suppliers already using EDI with the added resilience of EDI*Net.

Safeway's choice of EDI services is designed to open up and encourage competition in the market. According to Nick Ford, Director of Supply at Safeway, this made BT the natural partner to provide such a service as it also advocated competition based on customer support and value for money, rather than locking customers into a single network or single software solutions. Safeway was also impressed with the high level of support which BT could provide to enable the faster and more effective roll out to its suppliers.

BT is working in partnership with EDI Business Team members Kewill-

Ketal and Meadowhouse Bar-Laser to provide flexible end-to-end solutions for organisations of many different sizes and capabilities.

Top Marks for BT's ISDN Quality-of-Service Results

A recent test programme, designed to monitor the reliability and quality of BT's ISDN service, showed a 100% success rate in the connection of calls.

The nationwide test, which was conducted over a 3 month period, evaluated both local and national calls in three key areas of performance: namely, call-establishment delay; data transfer; and loss of established calls.

- *Call-establishment delay:* 100% of calls were connected within 2 seconds, the average being 1.5 seconds, which is vital for time-critical applications such as LAN-to-LAN connection or accessing information services.
- *Data transfer:* 60: second tests were used to measure whether data transmission was error free once a call was connected. The results showed that more than 98% of transmissions were without any errors. The very small numbers of affected calls would not be noticeable to the majority of customers because the use of error-correction communications software makes minor errors virtually transparent.
- *Loss of established calls:* No calls were lost during data transmission which is an indication that re-dialling, which can prove costly and disruptive, is virtually eliminated for companies using BT's ISDN for file transfer or videoconferencing.

The network monitoring system has been developed at BT Laboratories based on the guidelines issued by the International Telecommunications Union—Telecommunications Standardisation Sector (ITU-T). BT uses numerous test sites throughout the UK to measure and collect

information on ISDN data calls placed both locally and nationally.

BT and Norwegian Telecom Form Alliance

A strategic alliance enabling Norwegian Telecom to market and distribute BT's data and voice communications solutions in Norway has been announced. Norwegian Telecom will have access to the full range of BT's end-to-end managed services for business customers.

Tormod Hermansen, chairman of Norwegian Telecom, said: 'The main beneficiaries of this alliance will be our customers, as access to BT's top-class international telecommunications services will improve their companies' performance.'

For BT, the alliance represents the next stage towards achieving its declared ambition of becoming the world's leading telecommunications group; the company recently announced alliances with MCI in the US and Grupo Santander in Spain.

Alfred Mockett, managing director of BT's Global Communications division, said: 'This strategic alliance reinforces BT's commitment to putting customers first.'

'Norwegian Telecom's in-depth knowledge of the Norwegian business community, in particular the oil, gas and shipping industries, coupled with BT's expertise in the field of design, installation, management and maintenance of international communications networks, will provide the Norwegian business community with first-class telecommunications solutions.'

The alliance paves the way for potential joint activity to extend geographical coverage and develop technical solutions, based on the combination of skills, knowledge and expertise of Norwegian Telecom and BT.

BT Secures ITV Sport Contract

ITV Sport has awarded a one-year outside-broadcast contract, valued at £100 000, to BT's Visual and Broadcast Services division. Using its

broadcasting expertise, BT will provide a tailored range of packages to meet ITV Sport's requirements at major UK sporting, ranging from the Coca-Cola Cup to athletics events and major boxing championships.

BT's flexibility was key to securing this contract. Whereas TV outside broadcasts have normally been quoted on an ad-hoc basis, BT is providing ITV with a made-to-measure, packaged service enabling individual events to be costed and planned in advance. Roger Philcox, technical producer at ITV Sport, explained that BT had tailored a specific broadcast service to meet their needs. The service offered was convenient and cost-effective, and allowed them to plan precisely their own budget and resource requirements.

Steve Maine, Director of BT's Visual and Broadcast Services division, commented: 'The ITV Sport contract illustrates BT's responsiveness to customer needs and our ability to meet these requirements cost-effectively. By installing a fibre network, which is already linked into over 100 of the UK's premier venues, BT is demonstrating its commitment to the future of British broadcasting through investment in new products and services.'

BT's expertise and extensive infrastructure will be used to provide ITV with a broadcast solution based on a combination of fibre, microwave and satellite news-gathering technology. BT is the only broadcast-services provider capable of offering such a solution.

BT's broadcast capability has recently been demonstrated at events such as Test cricket in the West Indies, the Whitbread Round-the-World Yacht Race, the French Open Tennis and European Championship football.

Tenth OFTEL Annual Report

In OFTEL's tenth annual report published at the end of March, Director General Don Cruickshank has highlighted key issues in 1993 and outlined his plans and ambitions for the future. He says: 'The focus must be on the customer ... customers should get the best possible deal in terms of quality, choice and value for money. The interests of customers as the first priority is closely followed by the national economic interest—although that is consumers again in the end. While an open competitive market attracts inward investment, UK providers geared to respond to customers' demands will compete more effectively in the global telecommunications market.'

The Director General emphasises throughout OFTEL's goal of getting a better deal for customers. He continues: 'What means does OFTEL have for achieving our goal? The most effective means I have is promoting competition. This will lead to real choice—by which I mean there should be three or more service providers knocking at the door.'

He discusses the terms of interconnection between BT's and other operators' networks. He reviews the key events in 1993 and outlines the next stages which were fully discussed in OFTEL's recent statement *Interconnection and Accounting Separation: The Next Steps*. (see separate item).

Numbering issues are another high priority in the minds of competing operators. The Director General comments: 'In a competitive market the allocation of numbers must be fair, and seen to be fair. It is for this reason that OFTEL is taking over the administration of UK numbering from BT and I have put plans into action to achieve this. We approach numbering issues with one guiding principle. Numbers are a national resource which should be used for the benefit of customers rather than operators.'

The Director General discusses the progress made towards OFTEL taking over numbering and towards number portability being introduced. He also highlights OFTEL's enquiry into the provision of directories and

directory information. He emphasises that consumers' concerns about the way their directory information might be manipulated will be fully taken into account.

Looking ahead to 1994 and beyond, the Director General emphasises his commitment to breaking down barriers to competition—including technical and regulatory barriers. He comments: 'The Network Interfaces Coordination Committee (NICC) ... has the vital task of reducing technical barriers. Many services such as number portability require cooperation between operators on technical specifications for interfaces. I have asked NICC to draw up a new, more focussed, programme for 1994.' He continues: 'I am very much alert to the danger that regulatory procedures themselves may form a barrier to development and innovation. For instance, an area we will be looking at closely in 1994 is the licensing process ... we will be exploring ways to streamline the procedures.'

He outlines the progress he expects to see with interconnection and numbering issues. He comments: 'we will also be more proactive in identifying and dealing with anti-competitive behaviour.'

Three longer-term issues the Director General addresses are

- the convergence of telecommunications, information and entertainment,
- the growing internationalisation of the telecommunications market, and
- the complex question of social provision.

The Director General concludes: 'The future prospects both for competition and for customers look encouraging. But there is still a long way to go and much work to be done. My team at OFTEL and I are determined that 1994 will be a watershed when increased competition will be much closer to a reality in the market, and will be delivering substantial benefits for many, many more customers.'

Statement on Interconnection and Accounting Separation

Don Cruickshank, Director General of Telecommunications, has announced a three-stage plan to bring in better interconnection arrangements. His Statement, *Interconnection and Accounting Separation: The Next Steps*, sets out the framework for progressive changes aimed at promoting fair competition and getting the best possible deal for customers.

The first stage, which is now in effect, builds on the Director General's determination of how much Mercury should pay BT for interconnection. It provides a common basis for calculating the charges paid by other operators to BT for connection to its network. It should give operators confidence that these charges will be both fair and available to all without undue discrimination.

The second stage demands the implementation, by January 1995, of a new list of some 74 standard interconnection services. Charges for these services will be determined by an open and thorough investigation of how BT allocates costs to them. It also allows for the separation, for accounting purposes, of BT-Network, which provides interconnection services to other operators, from BT-Retail, which sells services direct to customers. Plans for this stage are already well advanced.

Stage three will include further consideration of a range of complex, and often conflicting, issues such as alternative concepts of costs assessment, charging structures and the range of services which should be offered. When this analysis shows that changes would be beneficial to customers they will be introduced as soon as practical.

Satellite Small-Talk Success

INMARSAT has successfully made the world's first satellite voice connection to a cellular-sized telephone. Trials of a future global satellite handheld telephone system tested the power margins needed to provide the voice quality required for

the INMARSAT-P service. These trials were designed by INMARSAT and Australia's Telstra Research Laboratories.

The telephone was a combination of a Nokia GSM-standard cellular telephone with a special antenna designed for the trials by the Indian Space Research Organisation (ISRO) connected to an NEC INMARSAT-M briefcase telephone electronics. The INMARSAT-P telephone, which will incorporate all of the electronics within the handset, is expected to be the same size as the smaller cellular handhelds on the market today.

ISDN Makes an Impact

Reed Exhibitions' National Telecommunications Survey for 1993 indicated a positive response to ISDN and EDI, perceived high costs of mobile and the fact that most people expect teleworking to play a part in their business. The survey, based on replies from more than 600 telecommunications or IT managers who responded to a questionnaire sent to 4000 UK companies, was drawn from a National Computing Centre's register.

More than 87% of respondents expect ISDN to make a strategic contribution to their organisations over the next two years. Other highlights from the report showed:

- Only 28% of organisations with a turnover of less than £10M had an official telecommunications strategy as compared with almost 75% of those with a turnover exceeding £250M.
- Nearly 75% expect to be using ISDN voice technology in two years time and almost 50% expect their usage of it to be significant. Other technologies which are expected to grow rapidly include voice over data, ACD and voice messaging, all of which are expected to be used in around a third of the organisations within the next two years.
- Teleworking is expected to grow considerably, from 37% now to 62% in two years.

- Satellite data transmission was reported by over 15% of the utilities and communications organisations.
- Broadcast data technology was reported by over 20% of the finance organisations and 16% of the utilities.
- Internal and external e-mail, EDI and networked fax are expected to be in use in around 80% of the responding companies in two years time.
- Two thirds of the respondents use cellular telephones with 39% believing that they would be using cellular data technology in some form within the next two years. Forty per cent expect to be using PDAs or other portable communications/computing device.
- Most users believe the main benefits of mobile communications is improved efficiency, better service and higher productivity. Few identified cost savings as a potential benefit, rather, more than half of the respondents cited cost as a serious drawback to implementation.

Energis Catches the Tube

Energis has signed an agreement with London Underground to install fibre cable over 65 km of London Underground's rail network. Cable will be laid around the Circle Line and along connections from Kings Cross to Tottenham and from London Bridge to Wimbledon. Energis sees the move as complementary to its existing development and as a relatively cheap way of gaining access to central London. Energis has already completed its link between London and Leicester.

National Cable Network Proposed

Six regional networks to cover London, the Midlands, the north, East Anglia, the south coast and Scotland are being formed by cable

TV companies which plan to implement a national backbone to rival those of BT and Mercury by mid-1995. The move will bring together the combined cable networks that are being set up in 62 urban areas by 20 operators. Optical-fibre and microwave links will be used to bypass Mercury and BT's networks.

Licence for MFS Communications

MFS Communications has been granted a national PTO licence by the DTI to offer public and private organisations a range of switched and non-switched, voice and data, domestic and international services. MFS is to begin construction of a 16 km fibre network within the City of London.


European Workshop Discusses Environmental Performance of Telecommunications Products

In March, more than 70 senior representatives of telecommunications companies across Europe and

the USA joined a two-day workshop at The Royal Society for the encouragement of Arts, Manufactures and Commerce to share ideas for the improvement of the environmental performance of their products. The workshop—the first of its kind in the telecommunications sector—introduced the concept of *eco-design*, which is central to new business ideas of environmental sustainability as well as an issue for competition in the marketplace. Eco-design concepts make environmental factors the foundation of product design. Findings of the workshop will be published by the RSA later in the year.

SDH for London

Esprit Telecom, a provider of value-added international services, has signed up with City of London Telecommunications (COLT) to use its synchronous digital hierarchy (SDH) network for local-loop connections. Esprit provides services to companies in the European Union for worldwide connections, routing traffic through its own network onto international lines.



Broadband: Business Services, Technologies and Strategic Impact

by David Wright

The problem with any new subject like broadband is that there is a delay before any book on the subject is published. Some authors in their rush to be the first, write before the standards are finalised, often resulting in errors. Credit is due to David Wright who has produced a timely book which is accurate and is the most comprehensive book I have seen on the subject.

The book is aimed at managers and professionals with an interest in public broadband networks. It does not concentrate just on technologies but includes sections on services and strategies. These last two sections are important because without them there will be no users. The strategies basically involve finding businesses which require high bandwidth and are prepared to be early adopters of new technology, (such as the insurance and publishing industries) or can act as a show-case, like health care establishments. These strategies are well known and may only be appropriate to a particular country. For example, the health service in the United Kingdom is unlikely to be a market leader for broadband services, while this is a growing market in North America. The book makes it clear that customers will only buy the service if they are genuinely useful and help customers with their business, either in terms of

notes and comments

Journal Articles

Contributions of articles to the *Journal* are always welcome. Anyone who feels that he or she could contribute a telecommunications-related article (either short or long), which may embrace technological, commercial and managerial issues, is invited to contact the Managing Editor, BTE Journal, Post Point G012, 2-12 Gresham Street, London EC2V 7AG (Tel: 071-356 8022; Fax: 071-356 7942). Guidance notes for authors are available and these will be sent on request.

New Subscription Rates—External Customers

New subscription rates for the *Journal* for external customers (that is, not employees of BT) will take

effect from July 1994. The new rates are as follows:

Price per copy (companies, universities, libraries and other bodies): £9.25 (£10 including postage and packaging UK; £11.25 including postage and packaging overseas).

Price per copy for private individuals (to be paid for by personal cheque): £5.00 (£5.75 including postage and packaging UK; £7.00 including postage and packaging overseas).

One year's subscription (companies etc.): £40.00 UK; £45.00 overseas.

One year's subscription for private individuals (to be paid for by personal cheque): £23.00 UK £28.00 overseas.

Overseas customers can pay by sterling drafts drawn on London. Because of high bank charges, payments in foreign currency cannot be accepted.

cost reduction, revenue generation or managerial effectiveness. I particularly liked the case studies at the end of the chapters where companies had invested in telecommunication services to improve business and beat competition.

In contrast the services section appears to be a mixed bag. The bulk of the section concentrates on audio, data, fax, image and video applications viewed from both a service and technology perspective. The arguments given for choosing a technology were realistic, but inconclusive. SMDS and multimedia looked out of place. It was good to see SMDS defined as a service, but frame relay and B-ISDN are described quite rightly in the technology section. However as network operators will be offering services based on SMDS, Frame Relay and B-ISDN, it would have been good to compare them. Although expanding the remit of the book, sections on residential services and the increasingly popular TCP/IP would have been useful.


The first section of the book is on technologies and includes frame relay, SONET, DQDB, ATM and B-ISDN. By leaving out LAN protocols like Ethernet, token ring and FDDI the book is clearly aimed at wide area public networks. The inclusion of SONET belies the North American pedigree. I feel the inclusion of SDH would have given the book wider appeal within Europe and could easily be modified in a second edition since SDH and SONET are similar. The ATM architecture cube which shows ATM working over FDDI, DS3 and DS1 is also specific to North America. The description of the technologies is good. It answers the question 'Why' as well as 'How'. For example, why is the ATM cell 53 octets long and why is DQDB better than Ethernet for wide areas? The book is well structured giving an overview of the technology before going into more detail. However, some of the technical descriptions are incomplete, either because of too much detail or had not been defined by the ITU. An example of this is the generic flow control (GFC) parameter in the ATM UNI. More detailed self-explanatory diagrams would have

been helpful, although there are plenty of diagrams to complement the text.

The preface says the book is suited to MBA and MIS programmes and I would agree with this. I feel it is also suitable for consultants or anyone who wants a good grounding in broadband both in breadth and depth. For specialists there are better books on individual technologies, but none which cover such a broad range. People already working in the area may feel that it confirms their knowledge rather than substantially adds to it. As with any book there are weaknesses, but these are minor. It is fairly easy to read and gives a business perspective of broadband and I have no hesitation in recommending it.

*Published by Artech House.
£63.00. xvii + 476 pp.
ISBN 0-89006-589-6.*

Reviewed by Trevor Johnson



The Basics of Electronic Data Interchange

The Basics of Electronic Data Interchange in Manufacturing

The Basics of Electronic Data Interchange in Smaller Companies

by Keith Blacker

These three titles are the first in a series of remarkably well-written and concise guides for readers encountering electronic data interchange (EDI) for the first time.

Keith Blacker, a highly experienced 'electronic trading' expert, has resisted the temptation to write big books. By adopting an almost 'Ladybird' format the author has set himself a challenge—to distil the essence of a large subject accurately and clearly. For any business manager or student who imagines that EDI is difficult, dull or impossibly deep, these books will be the perfect antidote.

The author manages to cover more than the essential elements of EDI. By page 23 in the first title we are treated to the clearest-ever 90-word explanation of packet switching, and by page 31 the mysteries of, and need for, systems integration are perfectly revealed with the help of a simple diagram and 'When integrating a computer application in one company with another application in a different company, it is necessary to ensure that the information which is transferred means the same thing to both the sending and receiving organisations.'

The first title is clearly intended as a primer for the rest of the series where examples of EDI in various types of business can be explored in slightly more detail. Because it has taken more than 25 years for EDI and electronic trading to get off the ground; because there are many experts who have a vested interest in making it all complicated; because EDI is quite simply one of the great enablers for changing the way we do business; and because it is so central to BT's global mission and the aspirations of BT's business customers; these are books that are badly needed.

*Published by EDISTONE Books
(+44 (0) 564 823 707). £4.95.
ISBN 1-897815-04-2
ISBN 1-897815-01-8
ISBN 1-897815-02-6*

Reviewed by David Brunnen

**BRITISH
TELECOMMUNICATIONS
ENGINEERING** ISSN 0262-401X

Published in April, July, October and January by *British Telecommunications Engineering Journal*, 2-12 Gresham Street, London, EC2V 7AG. (Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in this *Journal*, unless any such statement is made specifically by the Board.

© 1994: The Institution of British Telecommunications Engineers.

Printed in Great Britain by Unwin Brothers Limited, The Gresham Press, Old Woking, Surrey, GU22 9LH.

Subscriptions and Back Numbers

Price: Single copy—companies, universities, libraries and other bodies £5.00 (£5.75 including postage UK, £7.00 including postage overseas); private individuals (paid by personal cheque) £3.50 (£4.25 including postage UK, £5.50 including postage overseas). Annual subscriptions (including postage)—companies etc. £23 UK, £28 overseas; private individuals (paid by personal cheque) £17 UK, £22 overseas.

Overseas customers can pay by sterling drafts drawn on London banks or by sterling Eurocheques.

Price to BT staff: £2.25 per copy.

Back numbers can be supplied if available at the single copy rates given above.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 2-12 Gresham Street, London EC2V 7AG.

Remittances for all items (except binding) should be made payable to 'BTE Journal' and should be crossed '& Co'.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: 071-356 8050. Fax: 071-356 7942.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: 071-356 8050. Fax: 071-356 7942.)

Copyright

The entire contents of this *Journal* and the *Information Programme* are covered by general copyright and special permission is necessary for reprinting long extracts, but editors are welcome to use not more than one-third of any article, provided that credit is given at the beginning or end, thus: 'From British Telecommunications Engineering'.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by the *British Telecommunications Engineering Journal* 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/94 \$5.00 + .00



**THE INSTITUTION OF
BRITISH TELECOMMUNICATIONS
ENGINEERS**

IBTE Local-Centre Contacts

The following is a list of Local-Centre contacts to whom enquiries about the Institution may be directed. Where no Membership Secretary details are shown, membership enquiries should be directed to the Local Secretary.

<i>Centre</i>	<i>Local Secretary</i>	<i>Membership Secretary</i>
Aberdeen	Jim Martin (0224 752688)	George Cameron (0224 752857)
East Anglia	Terry Birdseye (0702 373330)	Keith Phillips (0206 717244)
East Midlands	Dave Bostrom (0691 666501)	Ian Bethell (0602 478587)
Lancs & Cumbria	Trevor Hughes (0254 666237)	John Nicholson (0254 666237)
Liverpool	John Johnson (0925 415705)	Clive Maude (051-229 4074)
London	Mike Burgess (0908 693141)	Peter Martin (071-492 6138)
Manchester	John Halton (0663 763556)	
Martlesham Heath	Peter Shaw (0473 644378)	Ian Gilbert (0473 645500)
North Downs	Nick Smith (0227 474823)	
North East	Richard Jackson (091-261 3560)	
North Wales & Marches	Peter Clay (0743 274205)	
Northern Ireland	George Adams (0232 328293)	Brenda Sharvin (0232 320092)
Scotland East	Graham Neilson (031-668 5878)	Lewis Shand (031-668 5482)
Severnside	Ian Davis (0272 206701)	Dean Shephard (0272 206457)
South Downs	Graham Lindsay (0273 222234)	
South Midlands	Tony Atkinson (0234 277874)	Jim Coley (0234 274043)
South Wales	Peter Coleman (0222 514524)	Keith Voyce (0222 555075)
Staffordshire	Ray Piper (0785 813483)	Paul Field (0785 762419)
Thames Valley	Stephen Graham (0252 337104)	Dave Prout (0734 501244)
West Midlands	Richard Harwood (021-230 3954)	Paul Jervis (021-233 8162)
West Scotland	Graham Wellby (041-220 3997)	Jill Maxwell (041-220 5837)
Westward	Rob Rand (0392 212681)	Chris Gould (0392 212663)
Yorks & Lincs	Steve King (0532 378124)	Paul Horncastle (0532 760242)

Associate Section Zone Committee Contacts

The following is a list of Associate Section Zone contact points to whom enquiries about the Associate Section should be directed.

<i>Zone Committee</i>	<i>Contact</i>	<i>Telephone No.</i>
London	Terry McCullough	081-456 6798
Midland	John Sansom	0533 513155
North East	Geoff Jenkinson	0532 378163
Northern Home Counties	Andy Edmonds	0473 645171
Northern Ireland	Brian Walker	0232 232565
North West	Phil Holt	061-600 8111
Scotland	Bill Lindsay	0382 302629
Southern Home Counties	Martin Harvey	0424 754385
South West	John Horton	0272 329303
Wales	Ferg Brown	0952 249886

IBTE National HelpLine — 071-356 8008

NEW BOOKS FROM ARTECH HOUSE

Asynchronous Transfer Mode Networks: Performance Issues

Raif O. Onvural

This comprehensive new title takes you from the application integration problems of today's networks to the virtually limitless realm of tomorrow's ATM networks, citing practical methods and simulation models useful in understanding and coordinating ATM networks. It provides useful insights into protocols, reactive and preventive congestion control mechanisms, virtual paths and virtual circuits, traffic characteristics of new applications, bandwidth allocation and switching requirements.

Hardcover. 230 pages. 1994.
Order Book BT662 **£52.00**

Teletraffic Technologies in ATM Networks

Hiroshi Saito

An invaluable reference for engineers and researchers developing ATM nodes, ATM networks and ATM LANs, this new title outlines the latest advances in teletraffic technologies in ATM. Providing valuable information on internetworking between the telecommunications and teletraffic technologies, it surveys the various switching technologies, shows how to evaluate the performance of ATM networks and demonstrates implementation of teletraffic control via connection admission control (CAC).

Hardcover. 190 pages. 1994.
Order Book BT622 **£55.00**

Secure Data Networking

Michael Purser

Networked data is especially vulnerable to deliberate malicious threats in the form of hackers. What specifically can you do to ensure protection from threats against your network using logical rather than physical means? This practical guide is designed to protect the smallest to the largest of networks from unauthorised access, explaining how to counter security threats and make optimum use of encryption, integrity checks, uniqueness mechanisms and security management. An authoritative, comprehensive and practical treatment of a sensitive and mission-critical subject.

Hardcover. 241 pages. 1993.
Order Book BT692 **£53.00**

Broadband Telecommunications Technology

Byeong Gi Lee, Minho Kang and Jonghee Lee

Synchronous, integrated, broadband, high-speed, fast-packet, connection-oriented, video, optical, multimedia ... these increasingly common terms are characteristic of the various aspects of emerging telecommunications technologies. This essential reference includes a comprehensive technical overview of current trends in technology development, including synchronous digital transmission (SDH/SDNET), BISDN and the associated ATM, high-speed data communication networks and broadband video services.

Hardcover. 577 pages. 1993.
Order Book BT653 **£67.00**

Transmission Networking: SONET and the Synchronous Digital Hierarchy

Michael J. Sexton and Andrew B. D. Reid

This invaluable reference provides comprehensive coverage of the principles of transmission network architecture and traces the development of the supporting standards in the CCITT. In general terms, the text presents the deployment scenarios available to a network operator and their impact on equipment specifications; and covers interactions with other networking technologies including MANs, frame relay and BISDN.

Hardcover. 360 pages. 1992.
Order Book BT551 **£59.00**

HERE'S HOW TO ORDER:

By Fax:

Fax your purchase orders and credit card orders to
+44 (0)71 630 0166

By Mail:

Send your order with payment to:
Artech House Books, Portland House, Stag Place,
London SW1E 5XA.

By Phone:

Call +44 (0)71 973 8077

GUARANTEED RIGHT OF RETURN

If not completely satisfied, you may return any book in saleable condition with the invoice within 30 days for a full refund.

 Artech House Books BOSTON • LONDON

TELECOMMUNICATIONS

Contents

Guest Editorial: The ICE Age Cometh	1
Customers Must be the Industry's Driving Force	3
The Bit Police—A View from the USA	8
Copper Mind-Sets	10
Service Management Systems: Supporting the Customer Interface	20
COSMOSS—An Innovative Approach to System Development	25
Event and Test Management System	36
Electronic Trading: BT's Commitment to Changing the Way We All Do Business	47
Telecommunications Quality of Services: Principles and Management	53
A Whole-Life Approach to Management Decision Making	62
Strategic Provision of Radio Nodes in the BT Access Network	67
Internodal Signalling	71
Field Focus	86

