

VOLUME 14 PART 4 JANUARY 1996

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

Major Series:

*The Information Industry and its
New Technologies*

New Feature:

pcochrane@btlabs



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 14 ■ PART 4 ■ JANUARY 1996

Guest Editorial: Towards the Information Age Bob Foster	273
<hr/>	
<i>The Information Industry and its New Technologies</i>	
Convergence—Synthesising a New Industry Bonnie Ralph and Andy Reid	274
Electronic Commerce Kieran Levis	281
Service Providers for Electronic Commerce David Giddings	286
The Web of Circumstance Interview with Keith Teare	291
The Fourth Information Revolution and its Impact on Some Public Information Services Julian Stubbs	294
Digital Signatures: A Solution for Security in the Information Society Nigel Hickson	301
Community Networks David Wilcox	310
<hr/>	
<i>Telecommunications in the 21st Century</i>	
Image Processing for Telecommunications in the 21st Century Mike Whybray and Mark Shackleton	315
The Virtual Science Park Richard Drew, Peter Dew, Christine Leigh, David Morris and Jayne Curson	322
<hr/>	
<i>Architecture</i>	
Transport Network Architecture Alan Jackson	330
<hr/>	
<i>Human Factors</i>	
Usability Engineering of Speech-Based Services Charanjit K. Sidhu and Gerry Coyle	337
<hr/>	
Access Network Support Systems—An Overview Stephen Potts	341
Access Networks Evolution and Preparation for Implementation EURESCOM Project 306 Alan Quayle, Jerry Trigger, Simon Fisher and Dave Faulkner	345
A View of a Future Intelligent Network Parminder Mudhar, Peter Glen, Carla Capellmann, Ronald Janmaat and Jørgen Nørgaard	356
<hr/>	
<i>Regular Features</i>	
pcochrane@btlabs	362
Book Review	363
BT News	364
Industry News	366

Board of Editors

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

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

Assistant Editor Richard Kemp, B.SC., C.ENG., M.I.E.E.; Alan Bealby, M.SC., C.ENG., F.I.E.E.; Martin Evans, B.SC., PH.D.; Warwick Flury; David Greenop, B.SC., M.SC.; Fred Landery, B.A.; Sam Quinn, B.A., DIP.M.; Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.J.MGT.; Keith Ward, C.ENG., F.I.E.E., M.I.MGT.

Theme Editors

The Information Industry and its Key Technologies Bonnie Ralph **Architecture** Andy Valdar


Telecommunications in the 21st Century Peter Cochrane and Bill Whyte **Human Factors** Martin Cooper

Published by: The Institution of British Telecommunications Engineers, 2–12 Gresham Street, London EC2V 7AG

ISSN 0262-401X

Bob Foster

Towards the Information Age



This is the second of three issues of the *Journal* including articles dealing with the new technologies that make possible multimedia telecommunications and computing services and the effects these services will have on our lives as we enter the *information age*.

We live by communicating; the very function of society depends upon it; we communicate to educate, to trade, to set and apply laws, and for pleasure and entertainment. As such, any major change in communication technology always affects the way we live.

The current predominant form of telecommunications, the telephone, has made interactive communication possible over distances. Prior to that, people could only communicate over distance by writing things down and having them carried by courier.

However, the telephone is limited to person-to-person spoken communications that falls well short of the richness of communication possible when people meet in groups such as in a schoolroom, a market, a court or a business meeting. In such environments, people can see one another, hear one another and share documents.

Only now is low-cost technology becoming available that would allow the reproduction of rich broadband multimedia communications over a distance, and it will still be many years before this will be as ubiquitous as the telephone is today.

However, the rapid growth of narrowband multimedia services like

the Internet already gives an indication of the changes that broadband multimedia communications (*information superhighway*) will bring.

In conjunction with multimedia computing services, the 'superhighway' will enable a new 'information age' where information can be held in digital form, manipulated, analysed and transmitted to any part of the world.

In this issue we look at some social and commercial implications of this progress towards the information age.

Unusually for the *Journal*, many contributors are from outside BT. This is because we wanted to get the views of people who would use, rather than provide, the technologies of the information age.

All the contributors foresee radical and mostly beneficial changes arising from the widespread availability of broadband multimedia services, but they also point out the many barriers that still exist, such as cost, ease of use and interworking.

A summary of the contributions reads like a school report: 'shows promise but could do better'. This is a salutary point because it is the companies that 'do better' in offering 'better' services that will be the long-term winners in the information age. To do better, we really need to understand not only the technology that makes these services possible but how these services will be exploited and the commercial and social changes they will bring about.

In the first article, on p. 274, Bonnie Ralph and Andy Reid, both from BT Networks and Systems, examine the different roles that companies can play in the information age and how this not only offers better ways of doing things we do today but also makes entirely new services possible.

Then, Kieren Levis of Cortona Consulting describes the considerable opportunities for electronic commerce. He points out that this is still in its infancy and there are still problems of bandwidth and security to be overcome.

In his article, David Giddings of BT National Business Communications (NBC) points out a number of areas where we can expect to see technical and commercial developments in the next few years. He also identifies the need for companies to collaborate to make the information age a reality.

The article on p. 291, a record of an interview with Keith Teare of Easynet, Cyberia Cafe and Cyberia Records, shows how an entrepreneur can exploit multimedia technologies to develop new markets. It gives a taste of the sorts of businesses that will develop in the future as the technology becomes easier to use and more widespread.

Julian Stubbs, also of NBC, describes how the eventual widespread availability of multimedia services will impact on the role of libraries and museums. The change will be dramatic, though, as he points out, it will take a few years yet before the technology is good enough or ubiquitous enough to compete with existing library and museum services.

Nigel Hickson of the Department of Trade and Industry describes the work required to develop some of the tools needed to support electronic commerce, digital signatures and encryption. This is one of the few areas where government intervention may be required to ensure electronic commerce has the same legal status as conventional commerce.

David Wilcox, whose experience as a journalist includes some time on the *Evening Standard*, describes how new communication services can be used to provide community networks. Such networks already exist in the US based upon little more than e-mail and bulletin boards, and the growing use of the World Wide Web Internet protocol will make them more attractive and easier to use.

Bob Foster

Manager, Products and Services
BT Networks and Systems

Bonnie Ralph and Andy Reid

Convergence—Synthesising a New Industry

The convergence of the content creation, computing and telecommunications industries into a new communications industry is having a profound effect on our lives. This article looks at some of the implications of this new industry.

Introduction

What **is** convergence? Is it a functional synthesis of the telephone call and the TV broadcast? Is it a raft of technologies which alter, forever, the economic environment in which the ‘unconverged’ industries of content creation, computing and telecommunication used to operate? Is it the clashing of commercial tectonic plates, as major players in all three areas collide like so many icebergs in a polar sea? Or is it the infrastructural support of total human individualism, as the market strives to serve segments of one person, via mass customisation?

Has it happened? Is it happening? Or is it **going** to happen? Above all, what does it mean to us, to the cosy world of telephony-based products and concerns about efficient transmission and switching which we **used** to inhabit? The analogy of world maps at various historical periods can be used to explain what has happened. If you look at the map in Figure 1, you will see that some bits are missing: the Americas and the Antipodes! It’s not that America didn’t exist at this date, or that it wasn’t inhabited. But even if the map-maker knew it existed, it wasn’t something of which he had to take account. It was just too far away. Consider the map at Figure 2. Today, not only do we have to take account of the fact that one can **go** to America, quite easily, we also have to take account of it in our lives. We eat American food, use American computers, watch American TV and use American names for things.

A few years ago, when we came to work in the morning, we might be aware (dimly) that we inhabited the same world as TV companies, publishers, educationalists or museums. But we didn’t have to think about them at

work. At best, they were in a large, rarely-opened box marked ‘Major Customers’. But mostly, they had nothing to do with us at all.

Now they are part of the same industry as we—the new information and communications industry. Our world map has altered, and people whose business we only dimly understand are changing our world, our jobs, and our lives.

This article examines some of the implications of the existence of this new industry.

Industry Structure

First of all, what does the new industry do? How is it structured? In 1995, the European Telecommunications Standards Institute (ETSI) ran a Strategic Review Committee (SRC6) on the European Information Infrastructure. SRC6 developed a model for use in determining significant technical interfaces where standards were needed to ensure industry development.

Once upon a time, telephone companies built and operated a public switched telephone network and offered one main universal service—*telephony*. Customers subscribed to this service, which they would use sparingly, as it was not just expensive, but someone more important might be needing to use this national resource at the same time. Large computer companies built large, monolithic computers which could hold and process data, either to run the corporate accounts or carry out scientific calculations. Public service broadcasters, in the words of Lord Reith, produced the best programming material for everyone. Film makers produced feature films which people would sit and watch in the cinema and then might see several years later on television when the public




Figure 1—Ptolemy map (by permission of The British Library)

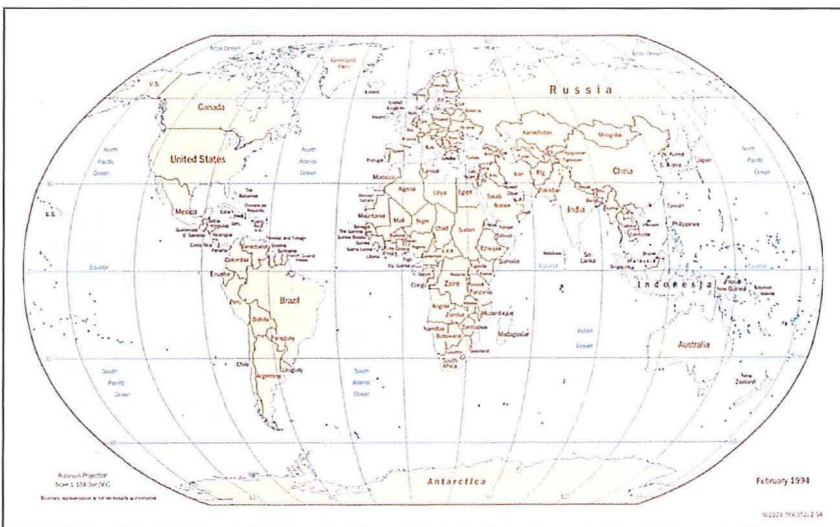


Figure 2—Modern world map (by permission of The British Library)

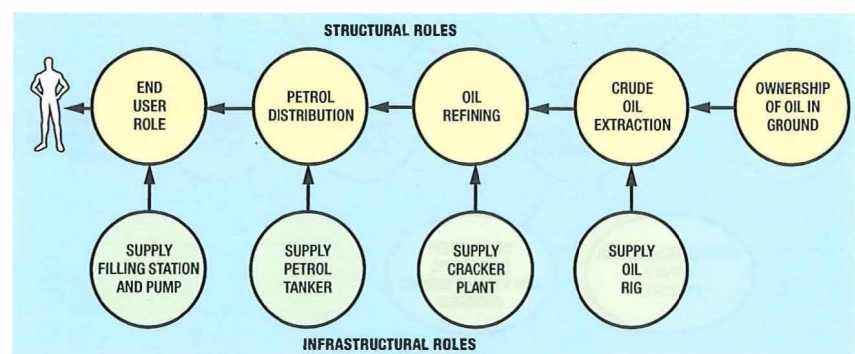
service broadcaster had bought broadcasting rights.

Each industry had a self-contained world and there was very little interaction between them. Each industry formed its own way of working. Both telecommunications and broadcasting were based on national public service organisations. The computing industry was based on a few very large companies, while the film industry, although dominated by the big Hollywood companies, did allow some room for smaller players. However, all these industries did have one important thing in common—they all knew what their customers wanted, indeed better than the customers did themselves!

As the major players in these industries now start to realise the importance of one another's domains,

it is vital that we form a new understanding of how the converged industry works. As players from these industries pass information between each other, we need to understand and define an entirely new range of interfaces.

Figure 3—A value chain



Roles and players

The first factor which ETSI/SRC6 recognised was the differentiation of roles in an industry from the players who undertake these roles.

Roles are well-defined processes within an industry. They should be reasonably long-lasting and their nature should be broadly agreed by everyone in the industry. A role normally forms part of a chain of roles, which, when combined, represents the industry. A role has its own economic value.

A player is a company or other organisation which can undertake one or more roles. The more roles a player undertakes, the more integrated the player; the fewer the roles, the more specialised he is. Within the information industry there is no strong indication, as yet, whether we will see more larger, integrated players or more smaller, specialised players. There is currently a great mixture and this seems likely to persist for some time.

Structural roles and infrastructural roles

Within most industries, and the new information industry is no exception, we can differentiate between roles which are associated with the main value chain of the industry. This value starts with the raw materials needed for the industry, works through roles associated with manufacturing, includes the retailing of the main products of the industry, and finishes with the end user. This is illustrated in Figure 3.

Both the telecommunications industry and the computing industry must leave behind much of their history and even some treasured principles in order to form this marriage.

In the information industry, the value chain starts with the *ownership of information sources* role. This role is associated with the ownership of intellectual property rights, statutory and civil rights, and even the ability to create useful information. Players include, for example, libraries, museums, art galleries, teachers, schools and universities, doctors and health organisations, and each of us as individual citizens are owners of large amounts of very useful information.

This information content is then packaged with other information content and made available to end users by information service providers. Next, the information from these sources needs to be assembled into basic information content; for example, a film, or a distance learning package. This is carried out by the *information provision* role in the ETSI/SRC6 enterprise model.

The relationship between the information provision role and the information service provision role needs to be brokered and managed, which is itself a role. This is also true of the relationship between the information service provider role and the end user role. In the ETSI/SRC6 enterprise model, these are called *information brokerage* and *information service brokerage*, respectively.

These roles are the *structural* roles of the information industry. However, in order to operate successfully, they require an infrastructure. In the information industry, this infrastructure includes the devices to store and process information, telecommunications networks to transport the information, terminal equipment through which the information can be requested and displayed, and systems which support the development of the information content and services. All these pieces of infrastructure are supplied by *infrastructural* roles. As we can see, the supply of telecommunications services is simply one infrastructural role in the information industry. This is shown in Figure 4.

Finally, we note that there is an infrastructural role which is the communication and networking of information. This role supplies the *information infrastructure* and is variously known as an *open systems environment*, *distributed processing* (or *computing*) *environment*, *information network infrastructure*, and even the *information superhighway*. This infrastructural role, drawing together both telecommunications and computing infrastructure, is central to the development of the new information industry.

The Information and Communications Infrastructure

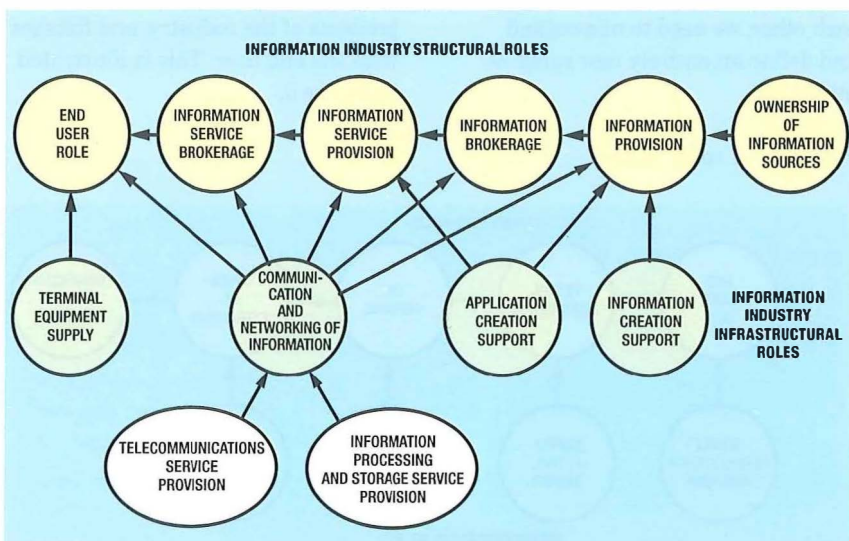
'For this reason, a man will leave his father and mother and be united to his wife, and they will become one flesh.' In the same way that a married couple are more than two people and are usually treated as a single item, so it is with information infrastructure. The marriage of telecommunications and computing results in a new being, information infrastructure, which draws heavily from each partner's expertise. However, it cannot be simply put together by installing telecommunications ports on computers. The 'leaving' part of marriage is as equally appropriate as the 'cleaving'. Both the telecommunications industry and the computing industry must leave behind much of their history, and even some treasured principles, in order to form this marriage.

The job of the information infrastructure is to form a set of interconnected computers, including terminal equipment, which can support information industry applications. An application is usually broken down into components, and the information infrastructure must support each component and all the messaging and other information transfer between the components of the application. This is illustrated in Figure 5.

The application needs a homogeneous infrastructure which can allow it to run in such a way that it does not need to understand the details of the way the information storage, processing, and messaging takes place. The application should be capable of design at reasonably high level after which it can be set to run on the information infrastructure.

The telecommunications industry has a history of trying to specify all possible applications within the telecommunications interfaces. This history must be left behind. The information infrastructure needs to be general and allow all sorts of applications to develop, most of which cannot

Figure 4—Structural and infrastructural roles in the information industry



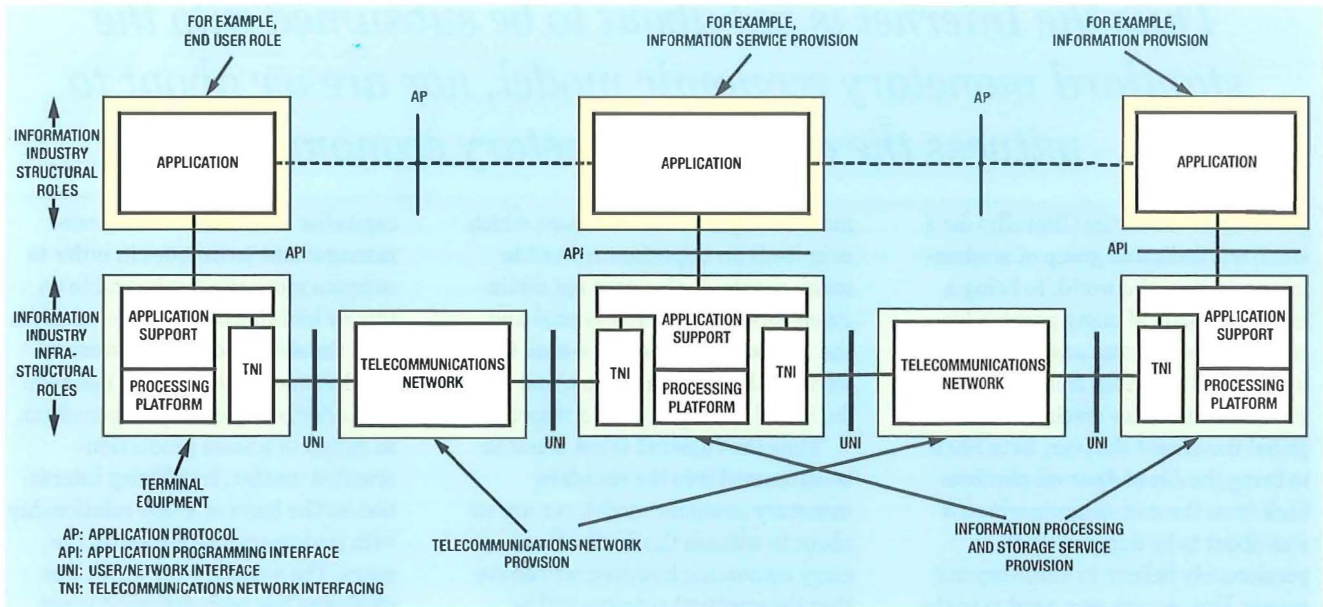


Figure 5—Information infrastructure supporting information industry applications

even be properly envisaged as yet, let alone specified. Understanding the limited scope of telecommunications within the information industry, and that there are other players, not traditionally in the telecommunications industry, is essential. Telecommunications players can no longer agree, and then dictate, standards.

The computing industry has a history of monolithic and proprietary solutions. This too must be left behind. The information infrastructure is diverse and includes many elements, all of which cannot be supplied by one player. Moreover, the applications are expensive to develop and the information industry is not going to support the development of the same application many times over, once for each player's infrastructure standards. Open standards are very important to the information infrastructure, and the computing industry has a very poor history in this area. The best example today of an information infrastructure is the Internet. It forms a single network for interconnected computers on which information industry applications like information browsing (World Wide Web) can be supported.

Most marriages also have their difficulties and rely on a level of commitment and hard work in order to survive. Both sides are often guilty of ignoring the importance of the other and taking them for granted. To the computing industry, telecommunications is often dismissed as 'just the

physical layer', while the telecommunications industry is often guilty of dismissing computing and data altogether as peripheral to the 'main' telephony service. Neither industry said 'obey' in their marriage vows, nor are they very likely to! However, there is now a growing commitment in Europe and throughout the world that 'for better, for worse, for richer, for poorer', the two industries must unite in order to produce the information infrastructure on top of which the new information industry can grow.

The Internet—Radicals Meet Conservatives

Economists start from the observation that the world's resources, both physical and human, are limited, but that inherent human demand for those resources is unlimited. Economics is the study of the mechanisms by which these unlimited demands are matched to the limited resources to satisfy those demands. Our Western economy is based on the idea that this dilemma is best managed by a market economy where things are all measured in a monetary value. Money allows us to compare the value of things, trade things, keep an amount of value until such times as we want to buy something (save), and even provide a means by which we can pay for things at a later date (have credit). Money and monetary value is the true currency of our Western economy. Governments, financial institutions,

corporate institutions, as well as individuals (or tax payers!) all form part of this great system where everything ultimately comes down to monetary value. Anyone within a large organisation who has tried to prepare a case for a new product or a new investment will know this to be so. Likewise, anyone who has been to a bank manager, seeking to set up a business, knows that 'brass tacks' are measured only in pounds and pence.

Many members of the Internet community, and especially those who have been responsible for its early success and growth, do not support this view of life. There are many things in life which are valuable but which cannot be measured in money. The Internet was formed and grew on the idea of free sharing and exchange of information so that groups of like-minded people could debate and discuss matters important to them independent of distance, nation, or even wealth. The Internet represents a new community, divorced from the monied economy. Even today, the great majority of information on the Internet is provided free in a global information 'swap shop', and there are no pounds or pence anywhere in sight. This ethos runs deep. Indeed, for many, the rejection of monetary economics marks the Internet as the first successful experiment in anarchy; no government, no police, no banks, just a vast global community.

It unlikely to be coincidental that at the same time as the Internet has

Thus the Internet is not about to be subsumed into the standard monetary economic model, nor are we about to witness the death of monetary economics

grown from the forum (literally) for a small but dedicated group of academics throughout the world, to being a significant part of many people's lives, environmental issues and the green movement have risen from a small group dismissed as 'crackpots' to a global movement that can force Shell to bring the *Brent Spar* oil platform back from the mid-Atlantic where it was about to be dumped. Both passionately believe in value beyond money. How do you give a value to the ability of a group of academic specialists in cancer research to swap freely ideas and data? How do you give a value to the annual migration of whales down the Pacific coast of America?

As the Internet grows into a more general information infrastructure, and more and more commercial organisations become involved, this clash of economic thinking will be seminal. It would be easy to say, as many do, that the Internet has simply ridden on the back of education and defence budgets of many governments and now it must join the real world of monetary economics. Several factors make things more complicated than that.

In general, monetary economics is best suited to 'widgets'. Each widget has a value which is its price, and this price, through the pricing mechanism, broadly reflects the costs of making the widget. Information is generally expensive to create. In the past, it has also been expensive to distribute, either in the form of books, films, records and tapes, pictures, or even word of mouth. Information infrastructure makes the distribution of information very cheap. This means that there is very little 'per-usage' cost and little justification for charging on a per-usage basis, which means it is hard to fit to the 'widget' model. This is not just true of information. There are many aspects of wealth, education, and even defence which do not fit the 'widget' model. We should expect to see many imaginative ways for rewarding the labour required to create information. The Internet

makes recognition much easier, which is in itself an important reward to many people. It also does not distinguish between the professional and the amateur. Anyone can e-mail a world authority on a subject, and they do, in full anticipation of a response.

Thus the Internet is not about to be subsumed into the standard monetary economic model, nor are we about to witness the death of monetary economics; however, we believe that the eventual outcome will be closer to the latter than the former!

The Internet—A Clash of Paradigms

Determinists may claim that the Internet was always going to emerge as the shaper of the information age. It clearly is not the case that interactive TV will be available before, or even at the same time as, Internet access. Whether or not the Internet is an inevitable development, it is certainly a genie which, now it has emerged from the bottle, will effect radical change in the industry, even if it were to fail tomorrow.

The Internet is used to describe three very different things. It is a set of robust data protocols which enables information to pass across any telecommunications transport more sophisticated than a telegraphy network. It is an **actual** 'network of networks' provided by **actual** providers, with **actual** services, price and facilities. (This is the Internet of the legendary growth rate and the corner in the obscenity market. It is also the Internet of newspapers, shopping malls and that key medium, the e-mail.) But the Internet is also an international community of inventors—of technologies, content, new ways of doing business, new ways of communicating and having fun.

It is also one of the most frustrating environments in the world!

Because the Internet is the way it is, the information age will develop in a way it might not have done otherwise. There was already a requirement for industry in general to

capitalise on computer-integrated management techniques in order to compete more effectively, and to do this by building on a mass-production base the ability to 'mass customise'. The Internet's contribution has been to challenge the broadcast paradigm, so suited to a mass production-oriented market, by offering interaction as the basis of a new relationship with customers, the audience, the users. The weapon with which this challenge has been delivered is not the technicolor compressed image, but the humble e-mail, that product of the convergence of telecommunications and computing, of the computer file and the telephone call. The e-mail is poised so indeterminately between speech and text that if you defame someone in an e-mail it is classified as slander rather than libel, just as though it were live speech.

So what does e-mail do for the information industry? Very simply, it gives the user the ability to 'bite back'. In fact, e-mail is so compelling; it is very hard to read one without clicking on the **REPLY** button. And when you visit your favourite Internet site, and the feature you specially like has been removed, you don't just sit there, you complain, vociferously, by e-mail! *The Daily Telegraph* gets about 100 letters every week, but that is nothing like as many as the *Electronic Telegraph*, which gets between 1500 and 2000 e-mails.

Practitioners of the interactive paradigm do not eat omelette on a Friday because the supermarket eggs were all broken and it's too much bother to take them back and complain. They do not buy clothes they don't like because there aren't any they do like. Practitioners of the interactive paradigm e-mail the shopping mall, asking why they can't design the clothes themselves! In a word, they are not passive.

Figure 6 gives a few (not terribly serious, but hopefully thought provoking) stigmata by which the interactive world and the broadcast world can be distinguished from one another. If you disagree with us about

BROADCAST	MEETS	INTERACTIVE
<p>Characteristics: 'quality', professional, large audience</p> <ul style="list-style-type: none"> ● TV ● 'Movies on demand' ● Mass production ● Julie's Pantry ● The Telegraph ● 'Encarta' ● 'Nicole' ● Party political broadcast 		<p>Characteristics: often uncertain quality, amateur, specialist audience, or audience of one</p> <ul style="list-style-type: none"> ● E-mail ● Yahoo (search engine on Internet) ● Mass customisation ● Mongolian Barbeque (a restaurant chain in South West London where you make up your own recipes) ● Hot Wired/Electronic Telegraph ● 'Engines for Education' (a 'learn by doing' method available on Internet) ● 'What's New' (on Netscape) ● White House server

our examples, you know what to do—e-mail us!

Where Professional Meets Amateur

The tendency of the mid-to-late twentieth century has been to depend on even narrower specialisation, allied with extreme professionalism. People do not get jobs as biologists, they are appointed to study the human genome. They aren't electronic engineers, they are experts in distributed computing environments. You don't go into advertising, you go into media-buying, or trade-mark searches. The products of the broadcast paradigm are the glossy monolithic results of tribes of specialists. The interactive paradigm, on the other hand, has tended, at least so far, to produce tightly-focused products in an environment where many more varied players can tout their wares. If you don't like the idea of searching for a publisher, you can advertise your **own** book on the Net. Why pay an expensive agency to manage your advertising when you can do it yourself on your own home pages, taking on the role (not always with total success, it must be admitted) of graphic artist, copywriter, media-buyer and so on, all by yourself?

Part of this explosion of 'non-professional' activity comes about because the roles of the older environment are either blurred, or migrating, in an information-networking environment. Publishers suddenly aren't publishers any more, they are ex-systems engineers who know how to manage sales over the Internet, and builders of Internet access and services. Their criteria for success

don't necessarily match those of Faber, or HMV.

In its February 1994 issue, *Wired* magazine featured an article on the 'Death of Advertising'. What **does** happen to advertising in a world of interaction and mass customisation? It is, after all, only an instrument of marketing, and is currently uniquely fitted to the world of mass production. What is presumably needed in the world of mass customisation is an equally professional means of getting each customer to tell you what he or she wants, rather than the existing mechanisms for making him or her want what you've got to sell.

The movement of players among the roles of the industry leads to experimentation on a grand scale. Keith Teare's new contract for recording artists (see 'The Web of Circumstance' in this issue of *British Telecommunications Engineering*) is revolutionary in the fullest sense. If it succeeds, some existing publishers may topple—the new relationship with creators will not sustain them at their existing level, using their existing methods.

In speaking about the professional meeting the amateur, what we are really talking about is radical experimentation—exciting, risky, enabled by technology, and driven by the same market forces that are pushing towards customisation.

The Need for Myths and Metaphors

It is always possible to tell when a really important technical change hits humanity. It elicits the response 'This is not allowed. It is dangerous (and, probably wicked).' The players in, and users of, information networking,

Figure 6—Distinguishing the interactive and broadcast worlds

should not be surprised that sober citizens are doing everything in their power to turn off the new technologies, by burying them under waves of social anxiety. The industry will need to respond with propitiatory myths. History and prehistory are full of these, each used to ease us into the use of some frightening new technology. Examples include 'Prometheus has taken the blame for fire, so it's okay to use it.' (Greek) 'It's fine cooking food provided the holy men don't touch it.' (Maori) 'You won't shake to death on trains provided someone walks in front with a flag.' (British) 'We already have one on privacy. (It's safe to answer these questions, provided you tick the little box that says the answers won't be passed on to other organisations.)'

Perhaps *British Telecommunications Engineering* should run a competition for the best information networking myths! It's a problem that only the industry itself is in a position to solve.

PC Meets TV, Music Meets Pictures, and Spielberg Meets Gates!

BT is experimenting with interfaces where images and music combine to manage human response to information input. It's a new role for the composer—helping business people to run their jobs. A recent *New Scientist* article on interactive films suggested that true progress in this field depends on object orientation, not something Disney has hitherto needed to think about. The PC with a TV card will soon be commonplace, and one day, perhaps, the TV set-top box will be as powerful as a personal computer.

Convergence means applying the techniques of one of the contributing industries to new and creative forms in the others.

At the World Conference on Computers in Education in Birmingham in Spring 1995, delegates saw children perform brilliantly, singing madrigals, and dancing, supported by computer software that works with

These developments are vital to the future supply of information content, without which the builders of information networks are soon going to find themselves in a seller's market.

midi systems. You can now compose on the PC, juggling any combination of voice and instruments as you write, and altering the result easily and quickly.

These developments are vital to the future supply of information content, without which the builders of information networks are soon going to find themselves in a seller's market. A wide variety of new artistic genres are emerging (some, it must be said, more artistic than others). First off the starting block have been interactive games, supported by both traditional and new players in this corner of the industry. Now, there are a few interactive films, in which the 'viewer' (nearly as horrible a word as 'Infotainment') gets to decide what happens next. Notice in passing that this is an interesting example of customisation.

The interactive book has been pioneered, but only in print so far. Consider, for a moment, the barriers facing a writer (content creator) who wishes to write an interactive book for the PC screen. Until the skills of the C++ writer are brought to bear on his problem, there is no software package that enables branching to take place on rules applied only at run-time. If he wishes to use still images, video or music he has to handle the copyright issues or produce his own. The question is, will it be the conventional publishers who take charge of this problem, or new players with the requisite software skills; for example, players to whom object orientation is second nature, and who have been using it for years to design other kinds of network?

Meanwhile, the common or garden 'story' is becoming more and more a commodity, and the recent secession of publishers from the Net Book Agreement suggests that the pricing of the conventional book may well be about to reflect this trend in no uncertain manner. There is clearly a wide value gap between Milton's vision of a book as 'the precious lifeblood of a master spirit' and the 6000 manuscripts sent to Mills and

Boon each year, which is not to say that Mills and Boon do not meet a genuine market need, and meet it excellently.

Of course the genres of the interactive age may well be quite different from those of the broadcast age—using terms like *game*, *film*, *book* may well not do them justice, in the event.

Summing Up

This article has attempted to consider information industry convergence from as many angles as we could identify—and we have undoubtedly merely scratched the surface. Our main message is that convergence leads to **synthesis**, to the 'truly new'. We hope we have demonstrated that, and something else that goes with the 'truly new', which is the opportunity to introduce innovation in all our activities, and to make the 'truly new' work for us. And of course, if you don't agree, you can always send us an e-mail!

Biographies



Bonnie Ralph
BT Networks and
Systems

Bonnie Ralph spent the first two years of her working life as an Inspector of Taxes (failed). She joined what was then the General Post Office in 1968. The contributions she remembers with the most satisfaction include the Post Office's first ever computer-assisted training course, the first cardphones and trainphones, an arduous struggle to build an acceptable case for a networked Centrex product for BT's large business customers, and the Information Product Lab, now subsumed into a more widely-scoped Customer Applications Bureau. Although most

of her career has been spent in product marketing, product management and product development, she currently can be found clinging to a niche in Networks and Systems Future Platform Department.

E-Mail:

ralphb@btwan04.oldpaul.bt.co.uk



Andy Reid
BT Networks and
Systems

Andy Reid went to Dundee University and did his honours thesis in software engineering. He joined BT in 1982, and after a diversion into international transmission equipment and then leased line marketing, he joined BT's team on SDH standards. This covered areas from the design of the world's best clocks to his original theme of software engineering. He made many important contributions to the network management standards for SDH. He now has two bosses, one for the strategy of transport networks and one for the strategy of information networks. In the past year, he has been an influential member of ETSI's activities on information infrastructure.

E-mail:

reida@btwan04.oldpaul.bt.co.uk

Kieran Levis

Electronic Commerce

Electronic commerce is probably a more important consequence of the convergence of telecommunications and computing than the much more speculative multimedia services, yet very few people understand what it is, never mind its implications. This article highlights some of the findings from a recent study by Cortona Consulting on the developments of new forms of communications between businesses.

What Is It?

The term *electronic commerce* (EC) is mainly associated with two particular applications of it: electronic data interchange (EDI) and, particularly recently, buying and selling over on-line services. This is misleading:

- EDI is a set of standards for exchanging structured messages between computers and is not in itself a form of EC. It has mainly been applied in streamlining the logistics of purchasing, invoicing and delivery, just one facet of EC.
- On-line buying and selling is, as yet, in its infancy. The volume of goods actually being sold over the Internet and services like CompuServe is still tiny.
- There are several other well-established examples of EC—electronic funds transfer, travel reservations systems and real-time foreign exchange dealing.
- Many new forms of EC, such as collaborative product development, are growing rapidly in importance.

EC embraces all forms of interactive business transaction which are facilitated by networks of computers. It is becoming particularly important now because of the much greater number of businesses and individuals who are able to use these networks, and the growing number of ways in which businesses can conduct transactions electronically, both with other organisations and directly with consumers. It is becoming as diffuse as commerce itself. However, it is the business applications which will probably take off faster than the consumer ones.

Business-to-business electronic commerce can be divided into seven main categories:

- **Automated purchasing and logistics** is carried out between trading partners with well-established relationships—this is what EDI is mainly used for. Intelligence and valuable information is normally concentrated on the computer systems of the participants rather than the EDI channel between them.
- **Exchange systems** enable large numbers of buyers and sellers who do not necessarily know each other to establish contact, exchange information and do business with each other. An exchange system is a market place in itself. Exchange systems for consumers will tend to be open to most buyers; those for vertical business markets will often be restricted to accredited participants, such as brokers or travel agents.
An electronic exchange system, typically includes a single computer system which all parties can access and which holds very large amounts of volatile information about products and processes. The intermediary who manages this process therefore plays a critical role.
- **Payments and banking** are closer to exchange systems than automated purchasing and logistics in that specialised intermediaries are indispensable, but the role is essentially one of transmitting information securely. Some facets of payments such as teller machines and credit card authorisation are already automated. New forms of payments and banking are an important enabler for EC in consumer markets.
- **Information delivery** describes transactions where networks are used to distribute commercial

Figure 1—Categories of electronic commerce

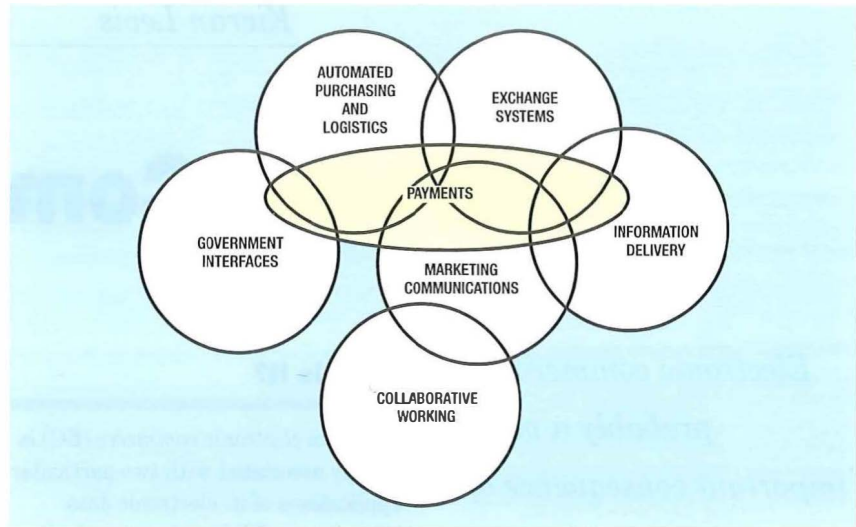
information products from seller to buyer. The longest established are on-line text services, but increasingly software, training and other multimedia material can be delivered over networks. In many cases, intermediaries are not involved.

- **Marketing communications** is mainly about direct communications between companies and their customers and will generally not involve intermediaries other than network providers—indeed it can be a means of bypassing them. Marketing communications embraces a wide range of business activities from customer support, electronic publishing and advertising, to gathering information about customers and prospects and selling to them directly.
- **Collaborative working** describes the sharing and exchanging of information between companies who are working closely together on developing new products or on other complex projects, using groupware platforms like Lotus Notes. Increasingly, the shared information needs to be held on network-based computers.
- **Government interface** describes transactions between businesses and government departments such as tax returns and customs forms.

There are several long-established examples of the first four. The last three are new and immature.

These are not hard-and fast divisions as the overlaps in Figure 1 make clear, but they provide a framework for analysing different kinds of market. For each of these categories, the market drivers, business models and the role of intermediaries will tend to be different. Understanding the difference between them is fundamental to an assessment of how they will develop.

A small number of companies are currently committed to EC as a



general way of doing business. Their most important applications fall in the marketing communications circle, but spread into several of the others. Most dabblers in electronic commerce, however, will tend to start in one of the seven areas for a specific reason; for example:

- because a major customer or government department requires them to (automated purchasing and logistics or government interface);
- to obtain access to specific information or trading opportunities (information delivery or exchange systems);
- to develop products with trading partners more efficiently (collaborative working); and
- to cut costs and restructure supply chains (automated purchasing and logistics or marketing communications).

The categories of EC should not be confused with the tools which may be used in different applications. As Figure 2 shows, some of these, such as e-mail and on-line information, can be relevant to most categories.

Why Is It Important?

Electronic commerce contributes to economic efficiency in five important ways:

- (a) by shrinking distances and timescales,
- (b) by lowering distribution and transaction costs,
- (c) by speeding product development,
- (d) by giving more information to buyers and sellers, and
- (e) by enlarging customer choice and supplier reach.

Figure 2—Key tools for electronic commerce

CATEGORIES OF BUSINESS-TO-BUSINESS ELECTRONIC COMMERCE	KEY TOOLS FOR EC						
	E-MAIL	EDI	INTERNET	VPNs	ELECTRONIC CATALOGUES	ON-LINE INFO	ENCRYPTION
AUTOMATIC PURCHASING AND LOGISTICS	●	●	●	●			
EXCHANGE SYSTEMS	●		●	●	●	●	●
PAYMENTS AND BANKING		●	●	●		●	●
INFORMATION DELIVERY	●		●	●		●	●
MARKETING COMMUNICATIONS	●	●	●	●	●	●	
COLLABORATIVE WORKING	●		●	●		●	●
GOVERNMENT INTERFACE	●	●	●	●	●		

● TOOL CAN PLAY A SIGNIFICANT ROLE

● TOOL IS GENERALLY IMPORTANT

● TOOL IS ESSENTIAL

Different combinations of these represent powerful market drivers for the adoption of electronic commerce in most markets and industries. They can also result in radical changes in the way markets work, the structure of industries and the shape of national and international economies.

(a) Shrinking distances and timescales

Faster and cheaper telecommunications are rapidly reducing the importance of location in the global economy. EC specifically enables many ways of compressing distance and timescales:

- One of the key factors in the globalisation of financial markets has been the ability to conduct transactions across the world, in real-time and 24 hours a day. Foreign exchange, the largest of these markets, is largely electronic—others are becoming so.
- In manufacturing industries like automotive, electronics and textiles, the trend towards subcontracting work to countries with lower labour and unit costs has been accelerated by high-speed data networks.
- Customer service and back-office work is following a similar trend.
- India's thriving software development industry is largely dependent on satellite links which not only carry orders and specifications from companies in Europe and North America but actually deliver the finished product.
- Small knowledge-intensive organisations (like Cortona) can operate globally, using networks for electronic dialogue and to exchange work-in-progress between units in several countries.
- Other small companies can use networks, notably the Internet, to find trading partners and customers in countries and regions where

they do not have, and could not afford, a physical presence.

(b) Lower distribution and transaction costs

EC can cut distribution and transaction costs:

- by eliminating paperwork and the associated labour costs;
- by streamlining delivery processes—supplies can be sent just-in-time, orders can be received from anywhere but shipped from the closest location;
- by reducing stock holdings—inventories account for over 20% of gross domestic product in countries like Britain and France;
- by reconfiguring supply chains to reduce and simplify the number of supplies and bypass middlemen who no longer add value; and
- reducing further the number of payments which are made by labour-intensive cheques and cash.

The cost benefits of EC are not equally distributed. They accrue particularly to banks and to major buyers at the end of long supply chains, such as motor manufacturers and food retailers. EC is a threat to many small suppliers and particularly to wholesalers.

In some industries there is scope for large cost reductions. It is here that EC is being adopted fastest and where there could be significant restructuring. Most forms of retailing, including banking, have dedicated a large proportion of physical distribution infrastructure to disseminating information from suppliers, capturing it from customers and holding large amounts of stock along the chain. All of this could in principle be short-circuited by electronic means.

(c) Faster product development

Shorter product life cycles mean that getting new products to market

quicker is now a competitive necessity in many industries. With the decline of vertical integration and the growing concentration on core competencies, the process is increasingly becoming a collaborative one between trading partners.

One of the most important ways of reducing product development time is by sharing and exchanging detailed information in near real-time over networks. Groupware tools and (mainly local area) networks have yielded big productivity gains in intra-company applications. They are now increasingly being adopted for intercompany working—in industries ranging from automotive, aerospace and electronics to fashion, advertising and publishing. In some cases, this can mean reducing time-to-market from several months to a few weeks.

(d) Better-informed buyers and sellers

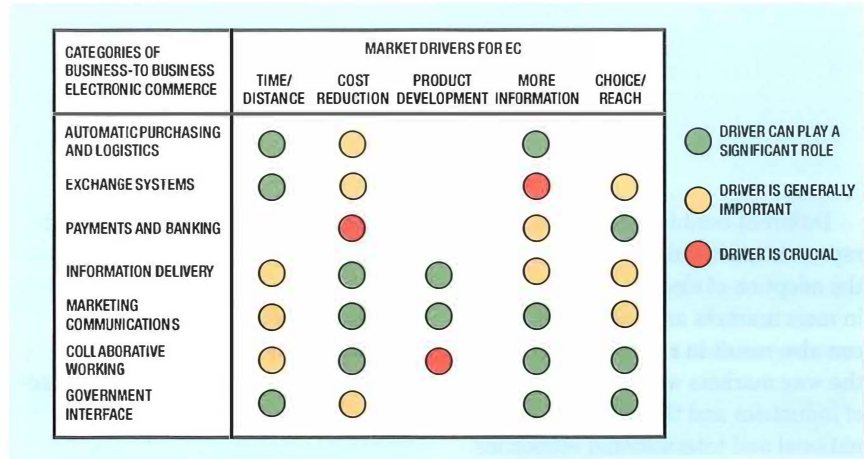
In most industries, more precise targeting of prospects and tailoring of product offerings are critical to competitive success. Computer networks enable suppliers to capture and analyse large amounts of information about customers and to develop integrated information systems. These can both automate purchasing and payments and, even more valuably, update customers' databases.

Different forms of EC also enable buyers of all kinds to obtain more information, more easily about a wider range of products and suppliers, and conduct transactions with a growing number of them directly. This has been most notable where specialised exchange systems such as computer reservations systems for travel agents have been established. New forms of network computing allow more types of buyer to do this, using, for example, the World Wide Web on the Internet or on-line services like CompuServe and Microsoft Network.

(e) Enlarging customer choice and supplier reach

Access to more information and more suppliers empowers some buyers—

Figure 3—Market drivers for EC



they can choose from a wider range and even compare vendors' responses to specific requests.

Suppliers can also establish contact and do business with a wider range of trading partners and customers. Many more organisations can do business internationally, as noted above under 'Shrinking distances and timescales'. Other examples of greater reach include:

- more direct contacts from prospective customers, via the World Wide Web for example, rather than through the medium of wholesalers and retailers; and
- the growth of new exchange systems.

Networks can also help suppliers to communicate more deeply with customers and prospects and build stronger relationships. More information can be exchanged and between parts of organisations which did not previously have contact with each other.

Figure 3 shows the relative importance of these five drivers for each of the seven categories. Three—time/distance, cost reduction and more information—apply in varying degree to virtually all applications. However, there is frequently one dominant driver in each category:

- cost reduction in payments and banking,
- more information in exchange systems,
- faster product development in collaborative working.

The impact of the drivers also varies according to industry sector. Few are entirely immune. Where there has been obvious scope to cut distribution costs, as in retailing supply chains, or transaction costs, as in the airline and banking industry, this will be the most

compelling driver. In both automotive manufacturing and software development, however, reducing product development times has been the key factor.

What's New About EC?

Several additional factors make the adoption of EC much more likely now:

Digital connectivity

Digital connectivity is growing exponentially and reaching critical mass in many communities of interest—both business and consumer. The key elements in this are:

- the presence on most knowledge workers' desks of communicating personal computers which are ever more powerful, affordable and easy to use;
- the widespread adoption of new tools for exchanging and sharing information, in particular local area networks, e-mail and groupware;
- the doubling of the size of the Internet every year since 1988—a faster rate of growth than that of any other communications medium, faster even than the PC.

Some see a combination of two new laws at work: Moore's (computing power doubles every 18 months) and Metcalfe's (the value of a network equals the square of the number of users). In some markets, digital connectivity is becoming comparable with that of voice telephony.

Accessibility

Most early examples of electronic commerce were based on mainframe computers, information was stored centrally, and both it and the technology were too difficult and expensive to be used by all but a small number of specialists in large companies. The new era is based on millions of individuals, including many consumers and small-to-medium-sized enterprises, using communicating PCs both to access and to share information. Much of that information is more widely distributed—for example on the World Wide Web—but the big breakthroughs have been graphical user interfaces like Windows, which make personal computers much easier to use, and browsers, which make information easier to find.

Unlike highly structured, old-style EDI computer-to-computer transactions, which often took months if not years to implement, trading partners can now proceed incrementally and experimentally, starting for example with simple e-mail and building up to more structured exchanges. New-wave digital communications are defined less in terms of technology, and more on business needs, and are frequently initiated by end users rather than central IT departments. They are intrinsically less hierarchical. They encourage dialogue, the sharing of information and even creativity. They can also be real-time, interactive and multimedia.

Organisational change

Changes in the structure of organisations and their relations with each other are also encouraging new forms of EC:

Electronic commerce is already changing fundamentally the way business is being done in some industries and will eventually affect nearly every business and every aspect of the economy

- Business process re-engineering and other fundamental reviews of corporate efficiency are not confined to internal operations. They need to embrace the gamut of a company's operations, in particular the automation of purchasing and streamlining of supply chains, where big efficiency gains can often be achieved.
- Most companies now seek to concentrate remorselessly on core activities and competencies. Vertical integration is being abandoned in most industries, and many are actively disaggregating key functions and operations.
- Monolithic organisational structures are giving way to more joint ventures, outsourcing and even virtual organisations. These and other forms of cross-company collaboration need better, more integrated information flows.
- Many more companies, both large and small, now have an international perspective and seek a global reach.
- More areas of the economy, such as the health service, utilities and government departments, are now subject to commercial pressures, and are particularly ripe for radical improvements to their external communications.

The most notable place where new applications have been developing has been the Internet and in particular the World Wide Web. Many thousands of companies have established a Web presence which they are using to communicate directly with customers and prospects to make available accurate up-to-date information, to provide customer service and technical support, and in a growing number of cases to make direct sales. The volume of sales, particularly to consumers, is likely to grow fairly modestly in the short term—very few yet have the neces-

sary bandwidth for Web applications and there are still important security problems on payments to be resolved. Eventually, however, on-line sales should take off in a big way.

The significance of the Internet is much broader. The fact that any business can communicate directly, easily and interactively with any customer, supplier or partner, anywhere in the world is a major advance in digital connectivity. It takes EC out of the back room and into the business mainstream.

Summary of Conclusions

Electronic commerce is already changing fundamentally the way business is being done in some industries and will eventually affect nearly every business and every aspect of the economy, from social security payments to workplace geography. It can shrink geographical distance, restructure supply chains, cut distribution and transaction costs and improve the efficiency of markets by giving more information and choice to both buyers and sellers.

Most applications of networked computing and data communications to date have been intra-company, confined to large organisations and the preserve of a few specialists. Most early forms of electronic commerce were complex, expensive and specialised. Now millions of knowledge workers, small businesses and consumers are achieving digital connectivity and many more organisations are seeking to communicate more effectively with their customers, suppliers and trading partners. Most businesses can find some facet of electronic commerce relevant to their needs. A few are committed to it as an overall way of doing business.

Of the seven categories of EC, the largest in the long-term, in terms of volume though not of value, will be payments and banking, automated purchasing and logistics and government interface, because they could be relevant in some form to so many businesses. However, only in pay-

ments and banking is there a significant role for an intermediary.

Nor will intermediaries play a particularly big role in most information markets, where new technology is enabling customers with specific needs to get the information directly from source and to manipulate it themselves. However, when intermediaries combine the role of information provider with that of organiser of an exchange system, they add considerable value. Demand for these, however, is generally specialised.

The areas of EC which will grow fastest in the short-term are those which require virtually no intermediary at all: collaborative working and marketing communications. Both can be tried out easily and affordably in simple forms like e-mail and Web pages. These can form the building blocks for more ambitious subsequent ventures. These will often require new kinds of networks which can support the sharing of large amounts of information.

Biography



Kieran Levis
Cortona Consulting

Kieran Levis is Principal of Cortona Consulting, which he founded in 1992.

Cortona specialises in evaluating markets for new information and communications services. Recent consultancy assignments he has led include large-scale reviews of evolving markets for networked multimedia, technology-based learning and electronic commerce. He also consulted in IT marketing and business strategy at the PA Consulting Group. As an executive at BT and at British Aerospace he started up and ran a number of new businesses—in on-line information, narrowcast television and audiotext. These included the first telephone information and private satellite services in Britain.

David Giddings

Service Providers for Electronic Commerce

Many types of service provider are emerging as the data networking industry evolves towards on-line commercial services – electronic commerce. This article reviews the hierarchy of service providers that provides the facilities required to bring a service, over a data network, to the end consumer.

Introduction

As the data networking industry evolves towards on-line commercial services, or *electronic commerce*, a number of different types of service provider are emerging. The most well known are the traditional telecommunications companies (telcos) and the Internet service providers, which have sprung up in their literal hundreds in the past two to three years. In this article, which approaches the subject of service provision from the point of view of the telecommunications and data processing components of the new industry, the concept of a service provider is taken further to include those companies that provide each facility

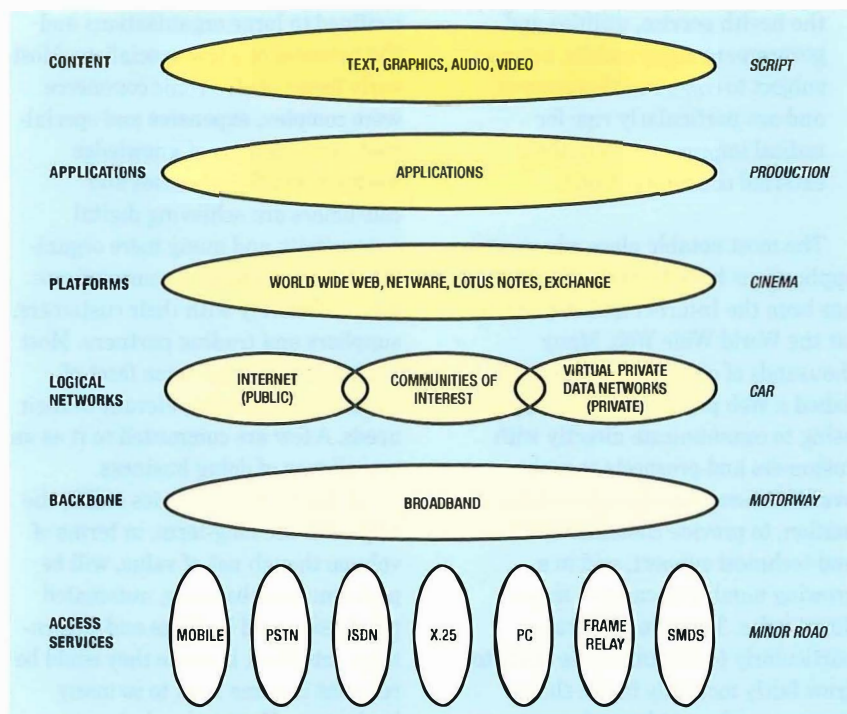
required to bring a service, over a data network, to the end consumer.

There is a hierarchy of service providers, as shown in Figure 1 and Table 1.

It is important to note that different types of service provider require very different skills and often different types of people. It is therefore unlikely that any single company will be able to offer a complete end-to-end solution. The information industry is one of alliances and partnerships.

Over time, as with increased competition, the cost of bandwidth declines and the cost of computer power and memory drops, the level of service offered by service providers will increase. (See Figure 2.)

Figure 1 – Hierarchy of service providers



Electronic commerce is just beginning for the normal consumer and therefore many of the methods, laws and relationships have still to be put in place.

Table 1 Service Providers

Content	Customer specific. Newscorp, Yellow Pages, Banking, Stock Exchange, etc.	<i>Script</i>
Applications	Customer specific. Microsoft, CompuServe	<i>Production</i>
Platforms	Internet service providers, telcos	<i>Cinema</i>
Networks	Internet service providers, telcos	<i>Car</i>
Backbone	Telcos	<i>Motorway</i>
Access	Telcos, cable, mobile	<i>Minor road</i>




Figure 2—Everything is decreasing

In particular, the applications will change over time from the predominately text and graphics-based network information services (NIS) of today to full video on demand (VoD) and virtual reality services during the next decade. During the same period, the method of interaction will progress from today's keyboard and mouse to include speech recognition and the like. (See Figure 3.)

Electronic Commerce

As an example, the process of service provision over a data network, or electronic commerce (EC), is compared to that of a visit to a cinema to demonstrate that there is little fundamentally new in on-line commercial services, just a change of technology. People en masse do not change at anything like the rate at which technology advances. Markets, and electronic commerce is a market,

have retained the same essential characteristics and requirements since they began 3000 years or so ago. They require a means of money exchange, an acknowledgement of property rights and other laws, police, a place to lay out the wares for browsing, a means of negotiating a purchase and guaranteeing delivery. They involve quality goods, shoddy goods and goods of dubious and illegal character, honest merchants and swindlers, genuine customers, fraudsters and thieves. Above all, as ever, caveat emptor.

Electronic commerce is just beginning for the normal consumer and therefore many of the methods, laws and relationships have still to be put in place. In the process, as with anything new, mistakes have been, and will be, made as the industry learns what is really required and how to facilitate or protect it. Each of the types of service provider described below has their contribution to make and it is upon the sum of their efforts that the success of electronic commerce rests.

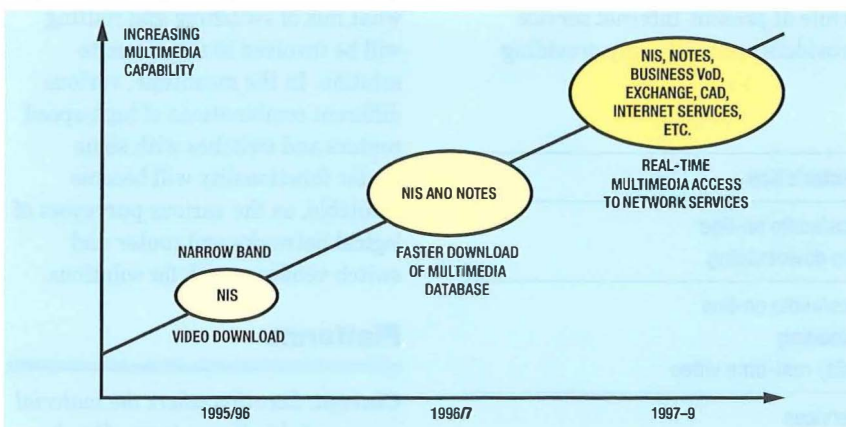
Access

Minor road: from home to the motorway and from the motorway to the cinema. Most often the cause of congestion, accident and delay. Normally the capacity of the minor road limits the speed of the whole process.

Access has historically been offered by the traditional telecommunications companies (telcos), but these are being increasingly joined by the cable television companies under pressure from the regulator. In addition, mobile telephony is becoming increasingly popular as a method of access to services by the travelling user.

The key technological change is the growing speed of modems, currently changing from 14.4 kbit/s to 28.8 kbit/s, and the penetration of the integrated services digital network (ISDN), giving 64 kbit/s–128 kbit/s over existing telephone lines. The next will be the advent of cost-

Figure 3—Evolution of multimedia capability



effective cable modems allowing higher bandwidth services to be offered over cable television services and later fibre, and new, higher speed, technologies for access over copper wires.

Today, access speed (see Table 2) is replacing backbone capacity as the limitation on what services can be offered to the user at an acceptable quality and cost and so determines the success of the whole endeavour.

The most interesting recent commercial change has been the emergence of virtual points of presence (VPoPs) services for dialled access from the main telcos. These allow nationwide coverage, at the cost of a local telephone call, without the requirement for the Internet service provider to have a physical presence in each local call area.

This means that local companies can now immediately become national in reach without requiring a large number of costly physical points of presence, thus reducing the capital limitations on growth of such companies. It also means there is no longer room in the market for as many small regional Internet service providers; they will all be competing on a national scale.

In the near future, therefore, a consolidation of the Internet service provider market can be expected.

Backbone

Motorway: the trunk, high-capacity, high-speed network. Usually safe and reliable but frequently with bottlenecks at national boundaries. For optimum total performance, motorway capacity

must be adequate otherwise congestion will progress back into the access network.

As with access, the backbone bandwidth used to be the sole province of the traditional telcos, but with liberalisation this is changing rapidly with over 100 providers now in the UK market. Naturally, this is causing a reduction in cost or, more commonly, an increase in the bandwidth provided at a particular cost; this is expected to continue. The key technological change to watch for is the increased availability of asynchronous transfer mode (ATM) services. In time, liberalisation of international links will lead to further cost reductions.

Logical Networks

Car: transports an item anywhere the minor roads and motorways go. Very dependent, for performance, on the underlying road system. If uncontrolled, congestion and other inefficiencies can result. Larger companies and stores often provide their own, private, controls. Pays a fixed subscription and a distance-related charge (fuel tax) for the use of the road system.

Logical networks, often referred to as *router networks*, have, until recently, been provided by companies for their own use or by universities, in the UK, in the case of the Internet. This is now changing with a proliferation of commercial Internet service providers (over 100 in the UK today) and with increasing outsourcing of companies' internal data networks. While at present Internet service providers are really only providing

Internet access, that is public networks, they can be expected to begin to compete seriously in the provision of private networks to industry as competition bites. Similarly, the telcos are offering services in this market.

The next change, once the necessary features have been added to the routers and security systems (firewalls), will be that, rather than each private network being designed and operated as a separate entity, the service providers will begin to offer virtual private data networks (VPDNs) which share the equipment, data links and operations and therefore most of the costs between different networks but provide the secure separation between different VPDNs that is missing from the public Internet.

These VPDNs will also allow the creation of sections of a network to be reserved for specific purposes such as Campus World. This is a service offered by BT to schools in which access to content is restricted to protect the children from the more dubious material on the Internet while allowing them access to the enormous quantity of material of genuine educational value.

Note the similarity with telephony, with privately owned and operated internal systems, connected to a public service, progressively changing to virtual private networks (VPNs) and Centrex, offered by the telcos.

Another aspect to watch is the evolution of networks as ATM services emerge. It is still unclear what mix of switching and routing will be involved in the ultimate solution. In the meantime, various different combinations of high-speed routers and switches with some router functionality will become available, as the various purveyors of logical networks and router and switch vendors reach for solutions.

Platforms

Cinema: Location where the material is presented to the customer. Can be

Table 2 Access speeds

Access speed	Service characteristics
64 kbit/s	<ul style="list-style-type: none"> ● text/graphics/audio on-line ● limited video downloading
128 kbit/s	<ul style="list-style-type: none"> ● text/graphics/audio on-line ● video downloading ● limited-quality real-time video
2 Mbit/s and above	<ul style="list-style-type: none"> ● real-time services

rough and cheap or plush and expensive. A number of different productions are likely to be playing serially and in parallel in the case of a multiplex. Customers often only choose which service to watch on arrival. While a particular film may have brought the customer to the site, other services are available and there is often significant interaction among customers themselves.

The provision of platforms as a service is only just emerging with many Internet service providers offering electronic mail, discussion groups and World Wide Web (WWW) space for rent by the megabyte, and a number of telecommunications companies beginning to offer services such as managed Lotus Notes. This is the area where the real activity is going to occur in the next year or two as the number of different types of platform and the facilities offered with these platforms increases.

Most of the commercial services are gradually moving from proprietary (to the service provider rather than the software supplier) platforms such as those listed in Table 3 to more open platforms, which as they become more Internet-based, also converge

Typical platform types are:

- World Wide Web,
- Lotus Notes,
- Microsoft Exchange,
- Novell Connect Services,
- audio servers,
- video servers,
- Internet telephony, and
- directory services.

The major development will be the emergence of true electronic commerce (EC). At present, most on-line commercial services are reliant on credit cards for payment. This is not really suitable for either business-to-

Table 3 Proprietary platforms

Microsoft	Microsoft Network, a service 'built in' to Windows 95, initially aimed at the Microsoft customer base
CompuServe	A self-contained service
Apple	E world focusing on information supply and discussion areas
Europe On Line	A consortium of publishers working on book, article and similar (originally paper based) information
BT, AT&T	Developing Lotus Notes and Internet and other information services
MCI	Shopping mall and information sources; Internet and WWW based

business transactions or for large numbers of very small transactions by consumers. The providers of closed services can provide their own guarantees of security and confidentiality, but for the more open Internet-based services, this is a major issue which is under investigation.

Customers do not wish to have to use different passwords for different services, nor do they wish to have to provide their credit card details for each transaction. What customers want is a single, secure, method of authenticating themselves and any merchant, an account on which all transactions are itemised regardless of service and a single point for any complaints. However, they do require access to services from many providers, potentially in many countries.

Some of the facilities that will need to be offered before true electronic commerce can really begin on a large scale are:

- common authentication,
- privacy,
- non-repudiation of transactions,
- protection of copyright,
- money exchange,
- billing, and
- audit trails.

As with the telecommunication companies and Lotus (IBM), alliances

and partnerships between platform service providers and software companies (for example, Microsoft, Netscape, Open Markets) and between software companies and providers of certain facilities (Microsoft/Visa, Netscape/Mastercard) can be expected. Increasingly, some of the larger software companies, such as Microsoft and IBM, are actually trying to become service providers in their own right.

The major problem will be the usual one with networks: getting the balance right between standards to facilitate interoperability and suitably rewarding and encouraging innovation. For example, at the time of writing, the method for charging to credit cards is an area of competing proposed solutions.

Applications

Production: process whereby the script is converted into the film often with distinctive characteristics. This is an artistic as well as a technical process. There are often a number of different competing productions of the same script. A single production will be sold through many different outlets.

Applications, in this context, are pieces of software that combine, in various ways, information from one or more different sources to create something that is of direct use to the end customer or, most importantly, for which customers are willing to pay. Such applications are normally extremely customer oriented and are often specific to a particular type of customer. An actual service will

usually consist of a number of different applications to cover the areas of interest of the target customers.

Content

Script: the original script, often difficult to appreciate raw. There are many different genres to cover the differing tastes of the potential audience.

There are a number of different business models for the relationship between the content providers, those who own the right to the content, and the various service providers.

The simplest is when the content provider owns the platform, the application and the content. This is most commonly seen with companies providing marketing content on the World Wide Web. In this case, the content provider does not wish to collect fees for access but rather to encourage the maximum number of people to view the content.

Another is when the content owner rents space on a platform provided by another service provider. Here the cost of the platform, its management and the communication link is shared with a number of other content owners, therefore reducing the cost to the individual application. This is most commonly seen with marketing content from non-technical enterprises.

Then there are the application service providers who will collect content from various owners, package it together to create a complete customer solution, perhaps with embedded advertisements and other marketing material, and then sell it on to the final customer. This is where CompuServe and America Online are today with their closed services although they also include the access, backbone and platform services as well.

Other models can be expected to emerge as this market develops and more and more existing and new players enter it.

Conclusion

The future of electronic commerce will bring open services where subscription to a single service provider at the same level as today's Internet service providers, but offering authentication, billing and privacy, will give access to a multitude of content and application service providers offered on a number of different platform types.

All the different types of service provider are crucial to this development. It is the skills of the different service providers in solving the many outstanding problems, and their willingness to cooperate as well as compete among themselves, that, while not perhaps determining whether electronic commerce ever arrives, will certainly determine when it becomes truly available on a large scale.

Biography



David Giddings
BT National Business
Communications

David Giddings graduated in Physics in 1971 and, after obtaining a Ph.D. in High Energy Particle Physics, he worked as researcher at the Rutherford Laboratory, UK, and at CERN, Switzerland. In 1983, he joined Digital Equipment Co. working on data networking, initially on X.25 and then on single and later multi-protocol router development. After leaving Digital in 1993, he joined BT and is currently the manager responsible for advanced services to business customers.

Interview with Keith Teare

The Web of Circumstance

Over the past year, many new opportunities to join the information age have opened up for the man and woman in the street. These have included new, easier Internet access via services such as Easynet and computer cafés, such as Cyberia, where customers can surf the Net over their cappuccinos. In this interview with Bonnie Ralph (the theme editor for this series of articles on 'The Information Industry and its Key Technologies'), Keith Teare of Easynet, Cyberia Cafe and Cyberia Records explains how one thing leads to another in the information industry, and how the occasional commercial revolution occurs on the way.

BR: What came first, Easynet or Cyberia? How are they connected and is there more to come?

KT: The decision to open Easynet happened on 29 July 1994. David Rowe and I, who knew each other, and had common business interests, because we were both into software development (in fact, databases) for large corporations, were asking each other 'How are we going to do multimedia?' I'd been trying to get the Woolwich interested—the housing market seemed to me an obvious place to start—and David had been considering other approaches. Then we said 'Well, Internet's already there. We ought to use it'.

At that point, David and I had both been Demon customers for three months, and we were not impressed with the user-friendliness of the Net. In fact, I'd written some software for my own use to make using Demon easier. So we decided we'd set out to be everything we thought Demon wasn't. Especially, we would set out to be user-friendly. That's why we called it *Easynet*,

because you didn't need a Ph.D. in Computer Science to use it.

David provided the funding from his previous business. And at that stage he funded the whole thing.

But then Eva Pascoe and my wife were thinking about getting people to actually try it. And they came up with Cyberia Cafe. Easynet opened on 9 September, and Cyberia on the 12 September—that's everything, leases, telecommunications links, computing, furniture, making it all happen, from a decision taken on 29 July.

BR: Were you involved in the Internet Society or anything?

KT: No. Not at all. In fact, we trod on some toes by mistake because we were going to call Cyberia *Cyber-Cafe* and then we discovered that someone had been running a bulletin board called *Cyber Cafe* for some time. We actually own the trade mark on both *Cyber Cafe* and *Cyberia*, but we didn't want to upset anyone. As it turns out, just fortuitously, *Cyberia's* a much better name.

Bonnie Ralph in conversation with Keith Teare





Cyberia screen

BR: How successful has Cyberia been in providing a way into Easynet?

KT: Very. We use People Training to provide our training material. Our training software has sold 20 000 copies and 6000 people have done our on-site training. It's called *TransCyberia* and we charge £25 for two and a half hours of non-intimidatory familiarisation. It is run by Eva Pascoe. She and my wife have become something of media stars—they've been in *Vogue* and *Cosmopolitan* and on TV.

BR: How did Cyberia Records happen?

KT: That's much more recent. Lynne Franks, the PR expert who does lots of work on branding, said that we must do something with the Cyberia brand. We'd paid nothing for it, but it was getting noticed everywhere. We really needed to exploit it. So we decided to do records. The idea is to create the equivalent of a 'Young Virgin'—although I think we're very different from Virgin really.

This all happened about the time that George Michael was having trouble with Sony. I decided we should design a copyright contract that George Michael, for example, would happily sign. So our lawyers, Allsop and Wilkinson, put together our 'New Concept Recording Contract'. In this contract, the record company relates to the recording artist purely as a contractor.

This contract does not require the recording artist to assign copyright for his work to Cyberia. He *licenses* Cyberia to produce and manufacture his creative work, and the licence has a set period. Profits, net of costs which are very clearly defined, are divided 50% to each party. There is

also provision for the recording artist to transfer elsewhere if he isn't satisfied with the way the record company handles distribution of his work. There are rules

for the calculation of a transfer fee, based on costs incurred by the recording company, within the limited definition I have already mentioned.

BR: How does Cyberia Records distribute the recordings?

KT: It uses both traditional means and electronic distribution. We use ordinary distributors, but we also use Ricky Adar's Cerberus Jukebox and we distribute directly from our own World Wide Web site, as well.

BR: How is it that HMV and Warner Brothers, and so on, need the kind of contract they use, and the share of profit they take, while you can operate comfortably on your New Concept Recording Contract?

KT: I suspect that HMV and the others don't really need the type of contractual arrangement they have. It was worked out to cope with a very broadly-based operation covering a very wide range of musical ownership. In that environment, it was necessary to make a huge investment to get the music to market. We operate by *not* providing things the musicians don't want. For example, most artists now have a large part of their own studios. We have a small studio here, although we use a larger one when we need to.

BR: What asset base do you really need?

KT: Well you just don't need a very large one these days. The big recording companies keep a database of privately-owned studios for their new artists. They don't run their own studios.

But the way things happen is that the recording artist will have a manager and he just wants a successful artist, so he can get his 10%. He is

perfectly happy to hand all the administration side over to the recording company. Once that is done, the artist really has no alternative but to sign the contract because he's dependent on the recording company.

BR: So, what is coming next? Have you things in mind to do?

KT: The digital world is a mirror of the real world. You can do things in the digital world, and they open up opportunities for you in the real world. It's very exciting, but the temptation is to take up every opportunity that comes your way. It's very hard to resist. For example, we have a fully developed digital payments system. But we decided we just couldn't launch it or use it, because it represented too much of a distraction.

BR: The human network of contacts seems very important to what you do.

KT: Yes it is. There's a very cooperative spirit in our area of the industry. We're competing yet working together. We've just set up an 'Internet Developers' Association'. Nearly everyone doing commercial work on the Internet has joined. We are aiming to shorten the period during which we have to cope with the perception of Internet service providers as 'cowboys'. We'll be deciding on standards and measurement systems, and advertising rates etc. within the Association.

BR: Talking of advertising, what do you think is going to happen to that in an information network-based environment?

KT: There are lots of different models of advertising emerging, and it isn't yet clear which will be successful. Grolsch has a nice interactive ad where you are 'teased' to click to the next page, and the next, and finally, there's a glass of Grolsch. It's very engaging. Guinness has a good one, where every time you click, a Guinness glass fills a little more. But the real success at present is the Wall Street Journal. It has a very strong, very rich, electronic journal which gets lots of hits (visits) because

it's very good. And down the side of each page there are six small static ads. And they are charging £200k a year for each of those! I was talking to BAA at Heathrow about a shop there and I was told 'You have to make £3k a year per square metre to be here.' So I showed him the WSJ and said, 'How about £200k per square inch?' He was a bit taken aback. But you do have to have the hits. The Cyberia site gets 2 million hits a week—that's beating the average TV programme and many channels.

BR: Do you think 'cyber-publishing' is the real future? Or will traditional publishing always be the main method?

KT: Well of course the easy way out of that question is to say they'll work side by side. And indeed, there is an element of truth in that. We *like* books and CDs. We like to see them on our shelves. When a friend comes round and he says he likes a particular piece of music, it's nice to be able to pull it off the shelf and play it. So *owning* a book or a CD is sexy. But one might speculate that it might be even more sexy to be able to call up *any* piece of music, at whim, from a vast music library. There are, after all, still people who insist that records are better than CDs, and so they are, if you're prepared to pay £3000 for the equipment to play them. But since you can get nearly as good from a CD on a piece of equipment costing £99.

And thinking about books, while I'm not sure about people wanting to choose the ending, you can see a taste developing for a stronger and more interactive engagement with the story among children. We like the passive book and passive TV, but the next generation may well feel differently. It's pretty clear that they *don't* want passive experience so much. So when I say I believe that books as we know them will remain,

† *The Easynet Book* by Keith Teare is published by Thomson Computer Press and includes a CD-ROM, to use alongside the text.

* BT's Internet access product.

it may be because I can't imagine anything better, not because there's nothing better to be imagined.

BR: So what is next on your agenda?

KT: Well a year ago, Cyberia was ahead of its time. And we're expanding. We opened a Cyberia Cafe in the Pompidou Centre in Paris, and we are due to open one in Tokyo in March. Actually, we hit a snag in the Pompidou Centre, because it's a government-owned building and the strike closed it down. But the strikers were interested in what we were up to, and we let them use one of our computers to issue their news bulletins, and so on. So they agreed we could open and, halfway through the strike, we did. The Pompidou Centre management weren't too keen at first, but they finally agreed.

The other new item on the agenda, also in March, is Channel Cyberia, which is what the Cyberia Web site is due to become. It will offer 24 hours-a-day, 7 days-a-week interactive programming, paid for by advertising revenue and programme sponsorship.

And, of course, there's my new *Easynet Book*†, which is intended to give straightforward, helpful advice on how to get the best out of Easynet.

But it is time to get into the next things.

We want to get into content. We want to *create* content. We've been putting the relationships in place to enable it. We have financial backing. We could make that element more important, but we have to be careful. We don't want to lose control.

Then there's our relationship with BT. We do want a global communications supplier. We recently became a reseller of MegaStream and we have a range of agreements with BT, including as a service provider over BT Net*. And we have a deal to train BT people on how to use the Internet. Mind you, we did hit a snag some time back because BT wasn't offering a virtual point-of-presence service. We had to go to Mercury. And there are *some* services which just will not take off unless and

until there's ISDN at a price that doesn't put off the home-based user.

We have other alliances in the pipeline that we'll need for our new projects. We plan to use our own programmers to develop the applications.

We already host a virtual job centre but we don't want it to be just a place where you *find* jobs. We want it to be possible to *do* the job, half the world away from your employer, and he can pay you through the virtual job centre. It's going to be a terrific help where people need more flexibility in what they do, or whom they employ, as well as for people who need to find jobs and are prepared to do them this way. We also have a virtual dating agency. It's better for people who are too reticent to use other methods of meeting people. These applications are free at present. I do believe that cyberspace is a less expensive medium than others for actually *creating* content.

We offer a product which consists of training to set up Internet content, and rented space on our server called *Power to the People*. But we hope to have a higher-quality authoring capability to offer when the Microsoft Blackbird product becomes available. It's an object-oriented authoring tool and should help with material like the interactive books we spoke about earlier.

So all in all, it's exciting, and we think it's likely to stay that way. As I say, when you do something in the digital world, just watch for the opportunities in the real one.

Biographies

Keith Teare, Easynet, Cyberia Care and Cyberia Records, has a degree in Politics and Sociology from the University of Kent, and a Ph.D in Sociology. He formed his own publishing company in 1981, becoming its Technical Director. He formed his own computer consultancy in 1985. He formed Easynet and Cyberia in 1994.

Bonnie Ralph is in BT Networks and Systems.

Julian Stubbs

The Fourth Information Revolution and its Impact on Some Public Information Services

The approaching information revolution, the IT information revolution, is only the last of four major information revolutions which have punctuated man's own evolution. Technology per se is not, in many senses, an important issue here – it is the way that technology is, or will be, utilised that will be key. Operational models of public information access will need to be redrawn to reflect this revolution and its attendant information superstructures, and this will be far harder than getting access to, or learning to live with, the technology. This article looks at information in the context of new technology as it might affect the public library and the museum documentation services.

The Information Revolutions

It can be argued that mankind has seen four major information revolutions.

The first, the development of **speech**, allowed man to communicate information to his fellows, and, eventually, to store that information for future use through folklore and oral tradition. Speech allowed man to stop living for the moment, and gave him a framework through which to develop such abstract ideas as religion. Speech was the medium through which man first began to create and use information.

The next big step was the development of **writing**, which allowed the storage of information in a semipermanent form not so easily capable of the distortions of an oral tradition. It also allowed immensely boring, but no doubt useful, information to be readily stored and accessed, such as detailed inventories—and it is no surprise that the eventual translation of Linear B, and much of the content of cuneiform writings, should turn out to be lists of things. It is arguable that it is writing (and its associated skill of the recording of, and eventual calculation with, numbers) that allowed the creation of (or at the least supported the swifter movement towards) developed economies, moving away from subsistence to wealth creating societies.

Reading and writing, if the developed civilisations of Greece and Rome are put to one side, were

primarily the remit of a scholarly, and frequently priestly, class, and it is not until the next major revolution, that of **printing**, that the possibilities of a universal access to stored information became practicable within numerically large societies. Printing not only led to more copies of existing works becoming available, but changed the economies of creating new material for dissemination, so that more ephemeral, or debatable, works could be created and circulated for discussion. It can be argued that this has led both to mass education and, through the ready sharing of knowledge, analysis and speculation among many people, to the increasingly rapid series of technological and scientific developments which we have seen since the Enlightenment.

The last revolution (if one views records, tapes and videos, broadcast radio and TV as a logical development of printing, though in some instances more ephemeral) is that of **information technology (IT)**, and it is a revolution because it allows, for the first time, direct mediation of information and interaction with it. The IT information revolution has been stimulated by (and even necessitated by) the end results of the printing revolution, which have not only led to vast stores of information, but have encouraged its continuing creation and updating, until the amounts of relevant information have snowballed to such an extent that they have become difficult to access or use effectively.

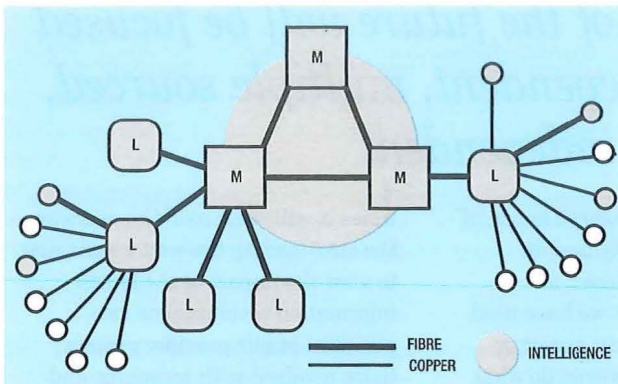


Figure 1—A typical network topology as a connectivity model

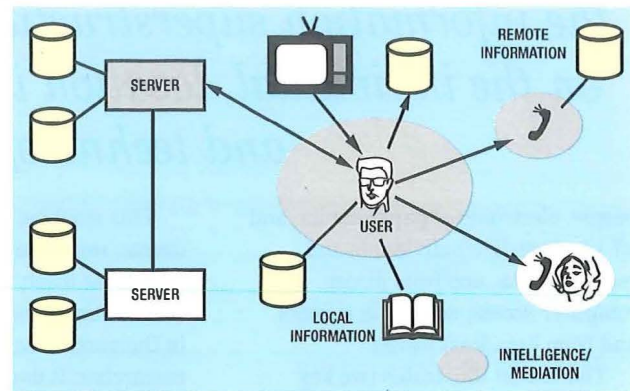


Figure 2—Information superstructure topology

What Differentiates the IT Information Revolution?

What, then, are the characteristics which mark the information technology effect?

One view is that it in some way describes the convergence of telecommunications and computing, and before too long arguments begin to revolve around a whole series of exciting acronyms and buzzwords, fibre, broadband, ADSL¹, T-PON², ISDN and B-ISDN³, SMDS⁴, digitisation, compression (MPEG1 and 2)⁵, client/server architecture, intelligent agents, multimedia, Internet and so on. These are the technical building blocks which allow the IT information revolution to be realised, but they are no more important than that. Almost all the technological developments needed for the IT information revolution are either already designed and working publicly, working in large-scale trials, or working in the laboratory. Over time they will be further refined, get quicker, smaller, cheaper, easier—but the Gutenberg press stage is already achieved and bettered—now we are on the polishing, not the initial carving out.

Another view is that what is key to all this is the coincidence of **content**, with all the richness that this offers, with **mediation**, broadly the software (rather than the technical hardware), including again, such exciting concepts as intelligent agents, browsers, cut and paste, groupware, trading, filters, morphing and so on. In the past, we have tended to treat information sources as separate from each other—there is the written word, the recorded word, direct information from other individuals, images. The introduction of intelligent software allows these to be brought together effectively, and

media translations to be made so that text can be read out loud by a machine, and a voice message turned into a fax. More importantly, perhaps, it allows machines to winnow the great mass of information to isolate the relevant and up-to-date. Almost as important, the information is delivered in ways in which it can immediately be reused within newly created documents (which might also include multimedia documents, films, animations, voice messages and so on), cutting down time to create and cost and often increasing quality. This gets closer, in my view, to the heart of the matter, but it still very much lists enablers and causes, rather than the real effects of the IT information revolution.

I take, I am afraid, a very much more simplistic view, which may involve everything I have mentioned above, but which can be summarised as: 'Getting the right information (and only the right information) to the right person, in the right format, at the right time.' Those who have been involved, as I have, with searching for information for a purpose, will know that, simple as this sounds, this is a step change which is not yet, in most cases, being effectively delivered.

Conceptual Models

Having reached a definition, I hope, of what the end game looks like, let me step back again to make a slight digression into conceptual models.

In the telecommunications business, and clearly that business is and has to be a player in the information superstructure, we have a very simple model of 'the network'. At the centre is the core network which links digital switches by huge (in terms of capacity) optical-fibre pipes. These then link

local switches, which in turn link to customers and customer sites, some already served by fibre links. (Actually, we almost never show customers in our diagrams, and only rarely show terminal equipment.) We demonstrate connectivity by showing how calls pass from access points at one side of the network through the central core and out to the other side (Figure 1). The switches (main and local) are now, in fact, huge and complex computers and there is a cloud of 'intelligence' which envelops them, as well as intelligence applied at some customer points.

As a further digression this diagram illustrates that the 'superhighway' is already in existence—what the arguments in fact are really about is the capacity of the slip-roads on to the superhighway—the access points. There is, in fact, no shortage of technology to support this access. Technology is not an issue, although cost (who will pay, and how much), applications (what it will be used for) and, to some extent, regulation are.

When we come to look at a schematic for an information superstructure, a different, and interesting, picture emerges (Figure 2) which reflects a very different 'mind-set' from the one which we habitually, as telecommunications people, use. 'Information' is not about connectivity or switches, nor does information have any obvious or predetermined method of storage or transmission. To write this article I talked to individuals face-to-face and on the telephone, exchanged e-mail, took published information from books, physical journals and from computer servers, some accessed over the Internet. The only commonality, and the natural centre for an information superstructure schematic, is the individual user, as information is sought on-site from

the information superstructure of the future will be focused on the individual, location independent, multiple sourced, and technology independent

people, electronic or paper media, and off-site from live individuals and recorded data, and from direct computer access, access via servers and from broadcast media.

This model illustrates two key things. The first is that information is available, and is delivered, via a multiplicity of different media. Many of these, most in fact, are mono-media, and it is people who internalise them into personal multimedia experiences, only to externalise them, as I am doing in writing this article, back into a mono-medium, in this case of print. In a few cases, true multimedia 'documents' (if that is any longer an appropriate word) are created, which include print, sound, moving and still pictures, and still and moving diagrams (animations); no doubt in some future a fuller multimedia experience will be possible, with scents and flavours synthesised, so that in reading, or perhaps experiencing, Proust we actually will taste the madeleine as he did. The multiplicity of source types is important, because it reflects the heterogeneous nature of 'information'—even where it is used to meet homogeneous needs.

The second interesting thing, I believe, is in the clouds of 'intelligence' associated with some of the information streams. It is the ways that this intelligence, which allows information to be found and to be manipulated and which will eventually allow information to find you through the somewhat nebulous offices of the much vaunted 'intelligent agents', which mark the IT information revolution as being a real revolution, and not simply a further development of the last.

Even more interestingly, the intelligence will not just support the user of the information but will also provide the necessary control and management for the information suppliers, building up profiles of customer use, providing charging mechanisms and ensuring that the information flows freely and effectively to where it is wanted.

This machine intelligence is not, of course, something really new, in concept at least; it replaces, or disintermediates, what we have used in the past, a live human expert or researcher. It does, however, do what the human researcher or information manager can no longer do—it copes with the extraordinary volume of information (both the sourced information being requested by the user and the metadata which surrounds the usage) which is now available, and presents it in ways which are of immediate use. Sitting at home with a computer, surfing the Internet, with a few relatively simple applications, and perhaps some stored images on CD-ROMs and a halfway decent colour printer, I can, in a few hours, print (or send electronically to up to 30 million Internet users) a piece of work which could have taken a dozen researchers, typesetters, photosetters and printers a month to produce only a few years ago. I can already use search routines which will hunt-out what I need (and much more usefully, discard the gigabits of information I don't need). Running a Web server, I can track who is using what information, and what information is not being used and could be discarded.

In summary, the information superstructure of the future will be focused on the individual, location independent, multiple sourced, and technology independent. Mediation will be the key for the customer, and service management for the supplier.

Impact on Public Information Services

The changes which have been enabled by this fourth information revolution, if not yet fully realised, will undoubtedly have a significant impact on all society, but initially this impact will be felt first in those areas where information can be seen as core to what they do. The providers of information will be the ones who respond to the revolution first, the users will follow, even though at

times it will seem as if the users are the ones leading the way. I want now to view the impact of the fourth information revolution on two potential public provider groups, those involved with museums and museum documentation, and those involved with public libraries.

Public Library of the Future

Why, then, should all this worry providers of library services? They already have books and magazines and journals, videos and cassettes and records and audio CDs, some already have access to the Internet or dial-up information services (more often within business or educational libraries). With enough money (of course always a key caveat) they can provide terminals for more 'readers'. Access to the superhighway, whatever that is, can be seen as just another information source, just like books, magazines or CD-ROMs—one position would be that this is a change in focus perhaps, but not in the nature of what a library is.

But what is a library? I would argue that a library is (certainly was) most like a warehouse—a book warehouse, carrying stock for a mass market of customers. Librarians' chief concerns can be seen as storekeepers' concerns—inventory control, location identification, storage conditions—to serve mass classes of customers—young children, schoolchildren, students, businesses, the elderly, fiction readers. What is so frequently requested it should be kept at the front of the warehouse? What so infrequently it can be kept at the back, the stacks? What needs to be specially housed (perhaps because of value or because of some fragility)? How many of which groups are to be served? What is the catchment area? How much resource is to be committed to serving which customers? What stock is not carried because demand for it is low, or because the customers for it, as a matter of policy, will not be served? How much of the stock location work can be placed on the

customer, so relieving the warehouse staff for more important work, such as cataloguing and recording?

Libraries (for traditional librarians) can be seen to contain artefacts, not, despite a romantic longing that this would be so, ideas. Many, perhaps most, librarians are primarily concerned with the containers which hold the information, not with the information itself. Librarians give customers access to the information (or entertainment) containers, directly, through cataloguing or through brief content summaries. The customers still have to locate and release the information from the containers themselves, nor can they be sure, without considerable personal research, that they have all, or the most relevant, or the most up-to-date containers to hand. While librarians frequently get to know their regular customers, their likes and needs, and many have some knowledge of the content of the containers, or their comparative value as sources (particularly in specialist libraries), in general librarians have neither the time, the training or the knowledge to act as information access experts for their customers as individuals.

The IT information revolution can change, and probably will change, this focus. In the future, it will be the information, not the source of the information, which forms one hub of the information superstructure, and the individual customer, not a class of group of customers, which forms the other. The overriding need will be for the two hubs to be joined effectively—for the customer to be put in touch with the information, delivered eventually through the medium of customer choice.

Where a human cannot hope to locate specific topics or pieces of information from within all published works, or to hold in his or her mind the needs of individuals he or she may never even have met before, a machine can. Individual profiles of needs and interests can be held on computers (or on smart-card chips on

new style library tickets?) and can immediately prompt offers to deliver regularly used information, or suggestions as to new classes of information which are now available. Programmes can be set to 'learn' from changed uses the changing interests of individuals. Most new information is created electronically, and could be 'held' in that way, rather than the traditional ways of libraries, and increasingly, important and relevant 'old' information is being, and will continue to be, digitised.

This does not come as any news to librarians, and already attempts, some of them less successful than others, have been made within the public library service to take on board this change of focus, often thwarted, or substantially watered down, by local funding influence.

However, libraries run within companies have been generally more adept (and resourced) at recognising, and implementing, solutions for individual user information needs. The BT business and marketing information library is called, quite intentionally, the *Information Resource Centre* (IRC), and its stock-in-trade is not volumes or publications, but information. While it provides some of the traditions of a library—a reading room, books on shelves, a catalogue—this meets the needs of only a few of its internal customers, who have the time, or the inclination, to undertake their own hands-on research. Most of its customers either contact the IRC by phone, fax or e-mail to request specific bits of information, relying on the IRC staff to know which is the best source for this; or pre-specify their interests, allowing electronic profiles to be made which will be delivered to them automatically and regularly. IRC staff in addition develop proactively, through contact with customers, new 'publications' which concentrate and group relevant information, thereby reducing the levels of ad hoc information enquiries and releasing staff to work on high-value-add research. The IRC has launched a World-Wide-Web

Server⁶ and is experimenting with groupware, which will allow customers direct access to source material with appropriate navigation tools to allow them to pinpoint the information they require—the IRC is already a very sophisticated user of a number of software tools, such as TOPIC.

The IRC is targeted with achieving high value for money from its expenditure on primary sources, which means ensuring that information is available in a timely manner to as wide a range of people within BT as possible, and that it is actually being used. Clearly the driving forces behind a library such as this is different from that of public libraries, though not all that different I would suggest, but the tools and capabilities being developed, and which are being developed by other commercial information suppliers, are as applicable in their use to public libraries.

The role of the IRC is to act as a window, which provides a user-friendly front end to people who are not information scientists, to the information they need. Some of that information is held in-house, on paper or electronically, but some is held by other libraries or information providers. The IRC is used by people within the company because it is cost-effective, because it is easy to use (and it intentionally offers a variety of entry strategies to the information to suit different ability levels and customer requirements), and because it delivers.

Libraries such as the IRC are creating new expectations among sophisticated information users, which will readily percolate to a more general public, about how information should be presented to users, and how much the drudgery of information location should be taken away from users. With reductions in the costs of computing and communications, and the real creation thereby of a true global village, there will be both commercial and public libraries, very possibly based outside the UK, which will start to offer these services in an affordable manner to a significant proportion of the general public in the UK.

The key aspects of the services offered by the IRC are that it is an information gateway (which can be used on-site or remotely), which has used the techniques of mass customisation to pre-prepare information digests for users, in order to concentrate its experts on high value-add work, with a focus on proactive information sourcing and digesting, rather than simply reactive response, aiming to automate as far as possible standard searches and information delivery to achieve a highly cost-effective service. Users can now gain information from the IRC without troubling information researchers at all, and can themselves set up regular deliveries of updated information, on a daily basis if required. (Other companies' systems, particularly those involved with trading, provide real-time updates of selected information.)

If UK public libraries do not themselves move towards this new model, they face the danger of increasing marginalisation, and just as the advent of printing and the reduction in cost of books led to the demise of the chained library, save as a curiosity for tourists, so will the book warehouses become curiosities in a world of information resource centres.

Museums of the Future

If we look at museums, and the services which surround museums, particularly documentation, we can see the future beginning to arrive. Already I can 'visit' museum and gallery sites, in the UK and abroad, view exhibitions, read commentary and critical analysis, download images, all on the Internet, or I can buy CD-ROMs which offer, normally, greater detail.

But what is a museum? I will not attempt to discuss the full philosophy surrounding museums here, but I would suggest that museums fulfil a number of roles. Firstly, and most obviously, they are storers (and conservers and restorers) of artefacts. They are warehouses of the past.

Secondly, they are (or can be) centres of scholarship, where, through analysis of artefacts, greater levels of understanding about past society, and hence often about present society, are achieved. Thirdly, they are centres of education, where that understanding can be transmitted to others. Fourthly, they are centres of entertainment, where individuals can be transported to different times, simply for fun, and where they can be brought into the presence of artefacts which carry the resonance of awe. To my ill-tutored eye, copies of the mask of Tutankhamun and the real thing are indistinguishable, but I still remember that sense of awe when I saw the real mask when it was on exhibition in London 20 years or more ago.

It is in the second and third of these roles that I see the IT information revolution playing its greatest part, although electronically shared information, particularly of techniques, will support the conservation role, and it may be that IT will allow 'virtual' restoration of objects so that the 'as new' look of an artefact can be demonstrated electronically, perhaps hologrammatically, without 'invasive' restoration on the actual item itself. Indeed, electronically scanned potsherds may even be 'reassembled' by computer.

As they are now, museums (which in their scholarship and education roles are warehouses of three-dimensional information objects and the textual commentaries that support them; that is, labels and guidebooks) suffer from two significant problems. They are rarely comprehensive in their nature—few museums carry a full range of known artefact types to illustrate a subject (and where they do, some of these are likely to be in storage and not available to the casual visitor)—and they, being physical by nature, are laid out in a physical manner to meet what they believe are the common needs of their mass of visitors. While, for any individual physical museum, this model is likely to continue, the virtual museums which are inherent

in the fourth information revolution will alter all this.

This does not come as any news to curators and museum documentors, and already attempts, some of them less successful than others, have been made within the museum service to take on board this change of focus, often thwarted, or substantially watered down, by funding influences.

If we look again at the BT example of the IRC (which, although BT does run a museum, may be a better exemplar), how is the learning we have gone through applicable to museums? First, let me be clear that it is the museum aspects of scholarship and education I am mainly addressing here. I may be in awe of the Internet working, but I am not as in awe of its content, which sits electronically on my computer, as I am of a real artefact in a museum. I can certainly surf museums on the Internet for entertainment, but it is of a different order to that I obtain from wandering the halls of the Ashmolean, or my local Horniman Museum in South London.

However, if I want to learn from a museum, or to exercise scholarship (should that have been my metier) in one, then it is the information, in the main, I am playing with. The information may be, in part, three dimensional, have tangible form, a particular size, but it is information, and most of it is capable of being digitised. Better, by handling digitised information, items in store are as accessible to me as items in galleries, and I can rearrange galleries at will. If an ethnographic museum keeps objects of the same culture together, I can group objects of similar function from multiple cultures, or if functional objects are displayed together, I can regroup them into cultures to look for commonalities of design or decoration. I can even, given cooperation, look at all the stored objects from many museums which I am interested in, before perhaps visiting the examples closest to me, and, via the Internet, booking access to them even if they are in storage.

Implications of the New Role for Public Institutions

I would expect this to impact, eventually, the public library and museum services in five main areas:

Proactive navigation and delivery systems

I expect the increased availability within libraries and museums of systems which will allow automatic sort and filter, initially of sources, eventually of information. If I want to know about the Korean economy post-1988 or decorative influences on metal work in pre-iron age cultures, I will key in those key words, and expect to be offered only information containers (and eventually information) which qualify, without having to go through a paper or electronic catalogue myself. I will expect in museums both graphic information about the subject, and a 'guide' to relevant exhibits in the museum. I will expect that convenience in every library or museum I visit, and preferably operated in similar ways, so that I can transfer my learning experiences. I will also expect, if I am doing a particular course of study, to key in (or click on) some generic description (NVQ level 3 in applied woodwork, Graeco-Roman influence on pre-Renaissance European art) and be pointed immediately to all the (automatically updated) required sources and appropriate texts. I will eventually expect text-based navigation systems through an increasingly voluminous digitised information base, and delivery of appropriate texts directly to the outlet, and in the medium, of my choice. I will not be averse to carrying with me some form of personal electronic profile which allows me to 'log in' to a museum or library I have not visited before, which will then inform me of the presence and location of exhibits or volumes which match my declared interests.

Location

I expect the future museum and library both to have, within the

foreseeable future, a real physical location, but also to be accessed remotely, from multiple locations, including schools, other 'public' places and home. Off-site access will probably require higher levels of user skill (and of course user investment), but this, in itself, will allow librarians and museum curators greater 'space' to concentrate their skills on the less well economically favoured, and the less skilled, while still being perceived as serving (and hence deserving financial support from) the wider community. Indeed as (and if) the public library of the future starts to look more like the internal business libraries of major companies, such as BT, they may well be able to provide smaller companies with similar levels of service, which they otherwise could not afford, and use this as a revenue earning opportunity.

Cooperation

I will expect, in 'visiting' (either directly or on-line) any museum or library to be offered access to information from other museums and libraries and from other (commercial) information sources—for example, university departments, commercial galleries, commercial information suppliers—for which I equally expect that I may be charged. I will expect museums to be pooling resources with other museums, and public libraries with other public libraries, and both with school, college and university departments and libraries to develop software and applications (such as profiles) which will then be shared (freely or on a commercial basis) with each other thus allowing maximisation of skills and creative input.

Internal use of IT

I will expect libraries to make full use of information technology to manage effectively their services, from data capture which allows profiles of usage and of users to be built, thus allowing better customisation of services, and better 'individual' attention, to the use of electronic ordering and stock control so that the drudgery of clerical work associated with acquisi-

tion and loan management can be off-loaded to machines, thus releasing people resource to more value-add roles. I will similarly expect museums to make full use of information technology to manage effectively their services, particularly by reviewing how individuals 'look' at artefacts, and which artefacts they look at, to reorganise their exhibits, and plan special exhibitions, as well as using electronic 'visiting' of stock in store to plan exhibit rotation.

New skills

The librarian, museum documentor and museum curator of the future will gain new skills which reflect an increased value-add role, both in terms of helping users through multiple information sources (and the systems that surround them) to arrive at the information they need, and in understanding differentiated needs of customers, linked to better understanding of information content so that proactive routines and information delivery systems can be developed. They will act as information concierges, like the concierges of the top hotels, knowing the best places to visit, making the arrangements for customers in languages they don't speak, giving the customers what they need before they have even asked for it, rather than the more traditional gatekeeper role where access is seen as limited and as a privilege.

Dangers of Non-Adoption—Public Libraries

If the funding authorities of the public library service do not eventually allow them to follow this path they will find that the growth of alternatives—home-base surfing of the Internet, commercial companies providing similar services—will begin to erode substantially the public-library information-user customer base and may lock public libraries effectively into a much more limited role (though still an important one) of meeting the entertainment needs of their audience by concentrating on the fiction/ music/

video end of the market. These are clearly important, and will continue to be so, but a substantial reduction of the information side of public libraries, while possibly being supportable within a more limited library service, will also lead to the further creation of a two nations effect, with an increasingly large group of 'information disenfranchised' growing within the community. Investment in an information-focused (although not exclusively so) public-library service is investment in the future of local communities.

Dangers of Non-Adoption—Museums

Similarly, if the funding authorities of particular museums also do not eventually allow them to follow this path, they will find that the growth of alternative access to other museums who do 'go digital', as numbers already have, will begin to erode substantially their visitor and user base and they, too, may find themselves restricted to a much more limited role of meeting only the local 'entertainment' needs of their audience and concentrating on the 'ace-caff' end of the market. This is also clearly important, and will continue to be so, but a substantial reduction in the proactive scholarship and education side of particular museums, while possibly again being supportable within a more limited museum service, will lead to the communities they directly serve becoming culturally disenfranchised. Although on-line access to other museum collections will be possible, using the local museum in an information-rich way will not. As for libraries, investment in information-focused museum services is investment in the future of both local and national communities.

Don't Panic

This will not (and need not) happen overnight—the UK is still not a computing- or information-focused society. A recent study⁷ shows that only 25% of UK households have a computer (as opposed to games

machine), and that only a further 11% are prepared to indicate that they plan to buy one. 'Accessing the Internet' does not even 'register', according to a summary of this report in the *Independent*⁸, 'on the list of most popular PC usages, where word processing and games dominate'. Growth rates in Internet users in the UK have been spectacular only because the base is so low, and many who access do so from work or educational sites, not from home. In the same *Independent* article, Jamie Muir of Packard Bell UK is quoted as saying 'If they can get on to the Internet, what are they going to use it for? They don't know where to go, there's no index, it's something you have to be very smart to use. For the average consumer, until the interface becomes much simpler and there are good indexing tools available, the Internet won't realise its potential.'

The museum service and public libraries (with their funders) do have time to make the changes that I have suggested; the window of opportunity is not even starting to close, but those changes will take time and investment, both in physical resources and in skills enhancement, and museums and libraries will not only not be the only players in the game looking for these resources, but are also likely to be competing even with each other.

References

- 1 Asynchronous digital subscriber loop—a method of carrying a multi-megabits per second (Mbit/s) down-link and a voice channel up-link on a normal copper twisted pair to support video-on-demand (VoD) services. It now works effectively at 2 Mbit/s over a 2 km radius.
- 2 Passive optical networks for telecommunications—a method of feeding in and carrying multiple channels on fibre as part of an access network, thus allowing a single fibre to service multiple customers.
- 3 Integrated switched digital networks and broadband ISDN—protocols to carry digital channels on copper, coaxial or fibre access systems. Basic-rate ISDN (ISDN2) puts two 64 kbit/s channels and one 16 kbit/s control channel (2B+D) on a standard copper pair.
- 4 Switched multi-megabit data services—a method of transporting high-speed, high-bandwidth data traffic, particularly between local area networks (LANs).
- 5 Compression algorithm standard allowing high quantities of data (such as is required for TV or film transmission) to be transported over 2 Mbit/s links.
- 6 The IRC Web sit (for internal BT use only) is at <http://turing.irc.bt.co/index.html>.
- 7 IDC. The UK Home Computer Market. April 1995, IDC No. UK Q020B.
- 8 NEWSOME, CLARE. Not caught up in the Net. *Independent*, 10.04.95.

Biography



Julian Stubbs
BT Group Business
Management

After Julian Stubbs graduated from Balliol College, Oxford, where he read

Modern History and took a Post Graduate Certificate of Education, he worked for three years in advertising as an account executive before joining, in February 1975, what was then the Post Office Telecommunications Division as a Telecommunications Traffic Superintendent. Since privatisation he has worked in various marketing and strategic planning roles, most recently leading a team of analysts and information scientists in National Business Communications Marketing. He is also secretary of the cross-Divisional Business Opportunities Council. In January 1996, he took up the appointment of programme manager for the Third Party Billing Programme, in Group Business Management. He is a member of the Chartered Institute of Marketing and holds their Diploma and a Diploma in Management Studies.

Nigel Hickson

Digital Signatures: A Solution for Security in the Information Society

The information society offers significant benefits to business, but the resolution of issues relating to the security of information on the new information highways is crucial. Digital signatures are a key enabler. The Department of Trade and Industry (DTI) has been playing a leading role in facilitating debate on the issues involved.

Introduction

It must have been much simpler in the 'old days' is an expression that is often heard said. And for those who find it difficult to grapple with increasingly complex technologies it is something which is often heartfelt. For it is because of rapidly developing technology that we find ourselves in the situation where we even need to contemplate digital signatures. If the copying machine could have been said to have put the coffin lid on the concept of the original signature, electronic commerce is probably responsible for driving in the nails.

This article explains why the Department of Trade and Industry (DTI) is interested in digital signatures and gives the results of a workshop it held in 1995; outlines ongoing developments (both domestic and international); discusses the possibility of regulation; and considers possible ways ahead.

For those new to this subject, some background in the concept of digital signatures and their potential business use is given in Appendix 1.

Why is the DTI Bothered?

The main objective of the DTI is to help British business compete successfully—in the UK, Europe and worldwide. Industrial competitiveness is key to the success of British (and of course European) business. One way in which the DTI can help is by assisting organisations to strengthen their innovation base. Within such a process, the Depart-

ment believes that the effective and efficient use of information technology systems is very important to most businesses today. Computers and networks are key tools to help people do their jobs more efficiently. They are so important that we need to be able to rely on them and have confidence in them. The role of good information security practice is to give business this confidence and allow the use of information systems to maximum effect. This applies whether the members of an organisation are simply using an individual PC or are avid users of the Internet.

The forthcoming information society raises the stakes. It is clear that if UK and European businesses are going to be able to compete in the global market place of the future, they will need to take advantage of the new technological opportunities proffered. In doing this, however, they will need to have confidence that the information and data they are committing to the new information highways are going to remain secure and intact. At the G7 Conference on the Information Society in Brussels last year, Al Gore noted that the lack of security on networks could effectively be a 'showstopper'. The theme was also taken up in the Eurobit† submission which, *inter alia*, recommended that, for electronic commerce to become a reality, '...worldwide agreement be reached on the legal

† Eurobit-ITI-JEIDA paper on Global Information Infrastructure (27/1/95).

aspects of electronic transactions in the GII, and on the use of electronic signatures for the resolution of disputes'.

But it is not just the Brussels conference, or others like it, which demonstrates the requirement for digital signature services. Almost every day, we read about new business services on the Internet, many of which involve users in ordering, or even paying for, goods. A recent survey in the US indicated that, by 2005, 10% of global business will be transacted on the Internet. To enable such business to develop, it will, in our view, be vital that appropriate security is put in place. Users may well be willing (and evidence suggests they are) to order products where payment is off-line, but there is much anecdotal evidence that they will not be happy to commit their credit card details to the Net until adequate security is implemented. Looking further ahead, the advent of 'digital cash' on the Internet will also demand strong security. Users will want to be satisfied that their transactions with a bank, or a supplier of products or services, cannot be modified. They may also want a guarantee of confidentiality: do we want others to know what type of magazines we purchase or the on-line videos we order? A spokesman for Digicash BV announced at the EITC Conference in Brussels on 27 November 1995 that two banks, in the US and Sweden, were now using their electronic cash payment system over the Internet. Known as *E-Cash*, the Digicash system incorporates strong cryptography† between both first and second parties (that is, the purchaser and the supplier) and the bank.

DTI Workshop

Purpose and format

Receptive to the perceived need for digital signatures, the DTI, with help

† 120 bit or more cryptographic algorithms

from Fischer International, convened a workshop in September 1995 with interested parties from the business world. Its objective was to determine the extent to which digital signature technology, standards and trusted third party (TTP) services, for management support of these signatures, were required for business applications.

The DTI, in holding this workshop, hoped to solicit cross-sectorial views from business on requirements and standards, and what actions should be taken by UK government, standards makers and industry; and what messages should be taken to the European Union.

There have already been debates and questions on the issue of standards including the recent question of whether there is a requirement to standardise on a single digital signature technique; for example, the digital signature algorithm (DSA), which is the subject of standardisation in the US (see below). In addition, there is the issue on what standards may be required for a TTP infrastructure to support the use of digital signatures. These and other related questions were addressed by the workshop with the aim of defining a way forward for the development and implementation of standards and specifications, both at the European and international levels. A fuller description of the relevant standards issues is given in Appendix 2.

Over 80 different organisations and businesses attended the workshop, held at the new DTI conference suite in Victoria Street, London. After an opening address from the DTI, and Ted Humphries (from Fischer International), two panels of speakers were convened: one representing the user community and the other IT vendors. Discussion was lively and constructive; many different views and opinions were put forward.

Results

The main conclusions from both panels were, however, as follows:

- The business solution for digital signatures needs to be international, supported by government, recognised in law, cost-effective, long lasting and to fit any established business infrastructure and the emerging global infrastructure.
- Digital signatures are required by business for customer and supplier relationships, for internal management and management of joint ventures.
- There is merit in having just one signature scheme, but, in practice, there may be a need to live with two (or more) and any infrastructure must be able to cope with this. In general, support was greatest for just one scheme.
- There is much concern about intellectual property rights (IPR) and export controls; this is something Government needs to address.
- There is a need for a signature scheme and supporting infrastructure that business can trust, have confidence in and which performs well.
- There is a need for more checks and balances to support the practical use of signatures across a wide range of business processes.
- There is a need to support the concept of a 'witness' signature as in the manual process.
- There is a need to establish an international infrastructure which business will have trust in.
- Is there a need to apply the same controls to digital signatures as with hand-written signatures?
- Although some believed that technology for confidentiality should be dealt with concurrently with that of integrity, others felt

the users want digital signatures (although they are not entirely sure in which form) while the vendors want to supply them (but they are not sure what the users want)

that the signatures could be dealt with first as they did not see an immediate requirement for confidentiality.

- There need to be rigorous procedures for licensing and managing TTP services.
- For some industries, a higher risk model may be required (for example, in safety applications), although users in general thought it would be difficult, and costly, to introduce different 'trust' level systems.
- Legal issues need to be addressed (both in identifying the pertinent issues and how these might be addressed).
- Several procedural issues were identified relating to the handling of keys. In particular, the replacement of keys on stored records that have expired or become invalid were identified as issues.
- Is there a need for a 'two person' mechanism; that is, someone witnessing someone else signing a document electronically?
- It should be possible for people to sign as an entity; for example, finance director, as well as an individual, although there must be an auditable record of who acted as finance director when the signature was made.
- It would be desirable to have a single standard algorithm and key length. In reality, however, it was recognised that more than one algorithm might be used across the whole of the business sector.
- It would be useful for the DTI to facilitate the development of a guide on the environment and management of digital signatures, possibly along the same lines as BS 7799.

- Consideration needs to be given to exception condition handling; for example, loss of locally stored secret keys, loss of TTP (or certification authority), implications on disaster recovery processes, etc.
- There is a need to extend current legislation to bring the use of digital signatures fully within the law. In the meantime, business needs to sign bilateral agreements on the rules for trading electronically.
- There is a future need for a broader set of trust services beyond that provided by digital signatures and supporting TTPs.

In summary, *users* wanted:

- usable solutions (simple and manageable);
- trusted digital signature scheme with supporting infrastructure;
- global and exportable implementations;
- reliable and cost-effective solutions;
- interworking standards, guidelines and codes of practice; and
- business confidence and usability (simple and manageable);

while *suppliers* wanted:

- support for technology
 - infrastructure developments,
 - export control relaxations (RSA), and
 - legal aspects sorted out; and
- promotion of standards as necessary
 - signature algorithm, TTP infrastructure and application protocol interfaces (APIs).

The above may, when first read, seem rather confusing and perhaps even contradictory. However, the results are repeated in full because they reflect the real feeling of industry that something needs to be done. It is almost as if there were a vicious circle on this issue: the users want digital signatures (although they are not entirely sure in which form) while the vendors want to supply them (but they are not sure what the users want). It will take action by all parties (see below) before the circle is truly joined. Although many of the conclusions above may not be surprising, there was at least one issue which was reasonably new and important. This concerned the level of 'trust' that would be required of any TTP service offering digital signature services. While most users would, as a matter of course, accept a certain level of business risk, those that wanted to use electronic signatures to back up safety-critical systems (for example, in the aviation industry) felt they could not. In other words, in any introduction of TTP services on a national, or even an EU scale, it must be recognised that certain industry sectors will require a greater level of assurance (with consequences for cost) than others.

Recommendations

From all these deliberations, and an analysis of the conclusions reached above, the following five recommendations were derived:

Recommendation 1: Guidelines and codes of practice

An industry working party should be established to consider and develop standards, guidelines and best practice on digital signature schemes and TTP infrastructures (along the lines of the UK Code of Practice group). The objective is to take a business lead in a much needed area of security.

Recommendation 2: Promotion and awareness

A user-group activity should be established to promote a better

trusted third parties (TTPs) should be able to provide the conduit whereby businesses that do not know each other will have the confidence to transact business

awareness of the requirements for, and application of, digital signature techniques and TTPs.

Recommendation 3: Legal, export and IPR issues

Government should address what legal measures are required to bring forward the use of digital signature services on public networks; the study would also address any necessary changes to export control regulations.

Recommendation 4: European activities

The UK should reflect the various user requirements identified as a result of this workshop (and any requirements resulting from Recommendations 1, 2 and 3) in any related European activities, either via industry involvement in European projects or through government activities in SOG-IS†.

Recommendation 5: International activities

The various activities suggested by these recommendations and the results they produce should also, where appropriate, be reflected in other international fora such as ISO/IEC and OECD.

The DTI has accepted these recommendations as a basis for further work, although action on them will, to a great extent, depend on resources.

Who Should Do the Work

Clearly, then, there is a role to be played to provide the necessary digital signature services so that the information society will bring with it the advantages we are all looking for. The next question then is **how** will such services be created: is it a job for Government, for industry, or for the European Commission? It will be of no

† Senior EU Officials Group on Information Security.

* RSA—Rivest, Shamir and Adleman Algorithm

surprise that the view of Her Majesty's Government (HMG) is that it is for the market to provide and run the services which are required. Government, however, does have a role, particularly in facilitating the appropriate regulatory and legal environment in which these services can operate. Indeed, Government probably has a crucial role to negotiate with the Commission, and other international bodies, so that the services can operate on an EU-wide basis.

The DTI has, for some time, been urging the European Commission to initiate work concerning both algorithm standardisation (that is, on RSA* or DSA) and on the infrastructure that would need to be established to support the use of such algorithms. HMG has generally supported the adoption of the DSA as a standard because of its (apparent) royalty-free use, and the fact that products implementing it are not subject to export control. RSA, as a competitor to DSA, is not technologically inferior, but it does suffer from IPR restrictions, and it is subject to export control because of its dual use in encryption applications. Given the lack of a standard, businesses throughout the EU have tended to use the algorithm most convenient for them. Therefore, both algorithms have been used in integrity applications in member States. The Commission, has, unfortunately, exploited this difficulty over standardisation on algorithms and has not been willing to tackle the more complex task of standardising the infrastructure needed to support them.

Because of this impasse on the European standardisation front, and the concern which the DTI felt to uphold the momentum for standards, the DTI held the workshop described above to consider the **real need** for digital signatures and what business considered as priorities in the standards debate. As a result of this, and in particular the recommendation for further action on the standards front at an EU level, the DTI was able to go back to the Commission with a clear

mandate for further work. Unfortunately, the focus of the information security work is now purely on the promulgation of the second Infosec decision (described below), and it is therefore unlikely that standardisation activities with respect to digital signatures will be addressed until that new decision is agreed.

TTPs and Licensing

From what has been said above, it will be clear that the DTI considers that TTPs may be able to play a significant part in delivering integrity services to business at large. By being responsible for the register of the public cryptography keys of their clients, TTPs should be able to provide the conduit whereby businesses that do not know each other will have the confidence to transact business. This model can then be expanded onto a national infrastructure (a network of TTPs in the same jurisdiction) or one throughout the EU. In either case, the TTPs of the respective clients would be able to guarantee, where appropriate, the integrity of the use of the respective public keys. Such an infrastructure could solve some of the confidence problems identified above, and form a major part of the superhighways on which the new multimedia services for the information society will be delivered.

To do so, however, it is clear—from the consultations the DTI has held with industry—that TTPs would need to have the complete trust of their clients. Without such trust the first and second parties concerned would just not deal with them. To develop such trust, it would probably be appropriate to introduce some form of licensing regime under which TTPs would only be 'approved' if they held certain competencies. In addition to being considered honest, TTPs may have to satisfy certain fiduciary and financial independence criteria. Clearly, a client would require some assurance that its TTP was not suddenly going to cease trading. Finally, the TTP would need to show

that it complied with BS 7799 (the standard on information security management) and that its IT resources were evaluated and certified under ITSEC† to the appropriate assurance level.

Although the licensing regime may perhaps be implemented under Government control, that does not imply it has a role in deciding who should be TTPs. As in the current licensing regime for telecommunication services, applicants who satisfy the licensing criteria would be deemed to have qualified. Whether they then traded successfully as a TTP would depend on their attractiveness to the client base. In addition to licensing, approved TTPs may also be offered other benefits by the state. For example, it may be possible that licensed TTPs would be formally authorised to verify the authenticity of 'signed' documents. In other words, contract law would, perhaps, recognise the validity of electronic documents if they were signed by the public key of a TTP's client.

The nature of businesses that became TTPs would, as indicated above, be a function of both the licensing regime and market acceptance. Banks, utilities (such as BT) or non-commercial institutions (such as the IEE) may be obvious candidates. The TTP, in addition to the provision of digital signatures, may well wish to offer other security services to their clients. These could perhaps include time stamping (where the time of a message is recorded), file retrieval and confidentiality services. The latter two services may well be offered in conjunction with *key-escrow* arrangements (see below).

Legal Issues

Apart for the legislation that might, at some stage, be required to license

† The IT Security Evaluation Scheme (ITSEC) is managed by DTI/CESG and enables the security assurance of IT products and systems to be evaluated.

TTPs, it was clear from the workshop that business requires confidence that 'electronic documents', when duly signed, have the same legal standing as corresponding paper ones. This would be particularly pertinent where digital signatures were used in applications which would have some legal effect; for example, funds transfer, contracts or performance bonds.

Although a number of studies have been carried out into some of the relevant issues at the national, European and the international level, there has so far been little real consensus on exactly what may be required. In the various implementations of digital signature that have so far taken place, both within the UK and Europe, the respective parties have included the recognition of the 'signature' in their working contracts. In other words, relevant documents duly signed, under the procedures and methodologies approved by the contracting parties, have been accepted as genuine. The Bolero project, explained below and in Appendix 3, relied on multiple contracts between the various partners to overcome the legal difficulties of 'signing' documents in several different jurisdictions. Such cooperation between partners may be sufficient in limited applications but would probably not be appropriate to the general application of digital signature technology in the TTP-type arrangements described above. One of the main advantages of an EU-wide network of TTPs, offering integrity services, is that the clients do **not** have to have definite contractual relationships with each other before corresponding. The 'trust' is dealt with by the third party. Some form of legislative framework will probably be needed where electronically signed documents have a legal entity of their own.

In recognition of these difficulties, the Law Commission recently made a paper available to the DTI which highlighted a number of legal issues, concerning electronic commerce, and

the advent of the information society, which warranted further discussion. As a result of this, Ministers have endorsed the establishment of a working party led by the IT Law Unit of Queen Mary and Westfield College at the University of London. Its task will be to establish the business needs for law reform to facilitate electronic commerce and to make recommendations for legislative changes where necessary. The working group, which includes various industrial representatives, is working closely with the Law Commission and the DTI and hopes to produce a draft report by next summer.

In addition to the above, the European Commission is also looking closely at the legal issues connected with electronic commerce and the information society: no doubt some proposals will come from Brussels in due course.

EU Dimension

A new Infosec decision

It was noted above that, for commercial reasons, any provision of digital signature services would have to be offered on an EU basis rather than in a single member State. Therefore some form of coordination at EU level will be necessary to set up the infrastructure required. Fortunately (as noted above) the possible adoption of a new Infosec decision by the European Commission may allow the necessary work to be taken forward to facilitate this. A recent draft of the decision contemplates, *inter alia*, the establishment of European trust services (ETSS) across the EU. Such services, incorporating a network of TTPs, would offer a variety of security services including both integrity (digital signature) and confidentiality. An action plan (the programme of work put in train by the decision) could trial the necessary technology on an EU-wide basis and facilitate coordination of the licensing arrangements—both for integrity and confi-

dentality services—which would need to be put in place by member States. A good start has already been made with respect to integrity through the various trial digital signature projects carried out under the first Infosec decision†. These proved that the technology (whether using RSA or DSS as the cryptography standard) worked and gave confidence to both IT users and vendors.

One of the most interesting projects under the former Infosec decision was one which considered the communications which have to take place between the various parties when freight is ordered, delivered to a port, shipped by sea, delivered to a customer and finally paid for. This project, known as *Bolero*, replaced all the paper transactions, in the above processes, with electronic ones using digital signatures to give confidence to the various commercial actors. A more detailed description is given at Appendix 3.

Confidentiality Concerns and Key-Escrow

The particular difficulties concerning the provision of encryption services for the confidentiality of information and data are now discussed. Whereas strong cryptography for the provision of integrity does not interfere with a member State's legal right to intercept (and understand) certain types of communication, the same cannot be said about the use of encryption for securing confidentiality. In the latter case, data and messages encrypted—perhaps by criminals or terrorists—would not be able to be read by the authorities if strong algorithms were used. Thus there is a real problem: on the one hand, business must be provided with good-quality encryption to protect their commercial information, and to encourage them to use information highways. While on the

other, the use of the same by criminals and terrorists must be protected to protect citizens and the security of the state. In other words, a 'balance' has to be found to equate both requirements. This dilemma, which is certainly not unique to the UK, has exercised a good many minds in recent years. A senior interdepartmental committee has been set up in the UK to consider possible solutions and work has also taken place within the EU Commission and in the US. The universal solution which most of these studies have identified—possibly the only one which might preserve the 'balance'—is the use of public key cryptography in a TTP framework where the latter hold the encryption keys of their clients. Such an approach—often referred to as *key-escrow*—does provide the 'balance'. Business and commerce acquire access to good-quality encryption (perhaps Data Encryption Standard (DES) or other public-domain algorithms) in return for depositing their private encryption keys with their TTP. The national security concerns of the authorities are also satisfied through legal access (under the process of law) to the keys held by the TTP and hence to the messages of their clients.

This solution also fits well with the discussion above on the requirements for a TTP architecture. TTPs, with the adoption of the appropriate key-escrow arrangements, would be able to offer a full range of security services on (hopefully) a pan-European basis.

Decisions, however, on whether and how such arrangements—for confidentiality—should be introduced are still under discussion within HMG. The view of Ministers, therefore, may radically affect how both integrity and confidentiality services are introduced into the European market place.

US Developments

Developments in the US have, not surprisingly, been somewhat more rapid than they have in the EU. Electronic commerce is, in many

areas, already well developed and the appropriate security mechanisms have followed. As in Europe, substantial use has been made of products incorporating the RSA algorithm for end-to-end integrity between business partners and within a single organisation. Recently, however, greater attention has been paid to the DSA* which the US authorities have adopted for their so-called *Digital Signature Standard* (DSS). It is this standard, rather than one based on RSA, which is being cited in federal procurement programmes and for general public sector communications.

There are, however, possible legal problems associated with the use of the DSA algorithm by commercial organisations both within the US and further afield. Various claims have been made concerning the 'ownership' of the algorithm to the effect that its use may be restricted by still active patents. The US authorities dispute this, however, and in an announcement on 10 October 1994 the Department of Commerce publicly stated that '.....[DSA] is available for use without any written permission from or any payment of royalties to the US Government'. Last year, in their desire to promote the use of the standard in the commercial sector, they took a further step by publicly stating that they would indemnify commercial users against any patent infringement actions against them with respect to their use of the DSS for communications with Federal Agencies or Government Departments. Such an indemnity does not, however, apply to the use of the DSS by commercial organisations outside of the US and so, it could be said, the debate goes on. HMG has recently asked the Commission for a ruling on the use of the DSS in Europe; no reply has yet been forthcoming.

Way Forward

In conclusion, it is important to reflect on the **real** needs of industry and commerce. While the Govern-

† 92/242/EEC.

* DSA (US patent: 5231,688) FIPS 186.

ment has only a minor role to play in the whole creation of the information society, it is an important one. If it somehow can create the environment in which the important industrial players can function, then it will have done some good. For although information security should not, and cannot be, at the top of the information society agenda, it must be recognised as a factor that could make an enormous difference in the rate of take up of the new opportunities that will be on offer. For it may only take a few well publicised breaches (such as the copying of credit card numbers or the interception of private political negotiations) to sour the appetite for the 'new' revolution. In all of this, the successful implementation of digital signatures services could play a vital role: it is up to the developing partnership between industry and Government to ensure it does.

Appendix 1—The Digital Signature Concept

Digital signatures are based on the use of what is known as *public-key encryption* techniques. The technique employs the use of two keys, one that is kept secret and the other that can be made public. Each user is issued with a secret/public-key pair. The two keys of this pair are mathematically related in such a way that knowing the public key (which many would) does **not** enable the secret key to be determined without having knowledge of other secret attributes that were used in the generation of the pair.

This technique provides a method for producing what is commonly called *digital signatures*. The method in its simplest form involves the user encoding a message with his secret key applied to the digital signature algorithm. Anyone then receiving the message and the signature, and having access to the public key, can verify this signature. Thus this process provides a secure method of verifying who actually signed the information

since only the owner of the secret key could have produced this signature as the owner is the only one with access to this key. This method also provides a means for checking the integrity of the information signed by the owner of the secret key.

An important issue with regard to the use and application of digital signatures is the management aspect. Each public key has a certificate associated with it which is a signed version of the public key. This process is introduced to ensure the integrity of the public key. One of the important issues then is the management and distribution of public keys and their certificates. This is one area of development which requires some form of infrastructure, this being sometimes referred to as a *trusted third party* (TTP) infrastructure.

Digital signatures have emerged over the last ten years as one of the important elements in the provision of security within open systems. They are of commercial significance and apply across the whole spectrum of electronic commerce, trade and business. Their use can be applied to documents in electronic form including invoices, orders, contracts and other legal documents, medical records, copyrighted material and many more business-related information.

Business applications

Digital signatures can be used for several different applications. These applications include:

- protecting and verifying the integrity of a document, time and legal stamps (notarise);
- verifying the origin and content of electronic funds transfer messages conveying authority or approval for payment;
- notifying proof of delivery for electronic mail and electronic document interchange (EDI) messages;
- official statutory returns involving a statement of declaration that the information given is correct and complete;
- conveying authority and providing evidence that an electronic document was signed by an authorised signer (for example, authorisation by a healthcare professional); and
- verifying the origin and the integrity of safety critical information being transmitted for control purposes.

Appendix 2—Types of Standards

Three issues relate to the so-called *standardisation* of digital signatures:

- the algorithm and scheme for generating and verifying the signature;
- an infrastructure for managing the various security parameters that are necessary for using digital signatures, including the secret/public-key pairs and the public-key certificates; and
- interface standards for applications software to communicate with the range of available computer systems.

Standards for techniques

Over the past 2–3 years, there has been much debate on whether a standard is needed and, if so, which algorithm should be standardised. This debate has taken place within standards bodies in Europe and North America, and at the international level in user and vendor meetings (for example, within IBAG) and at the European Commission.

Currently, the two main candidates for digital signature algorithm standardisation are RSA and DSA. From an implementation point of

view there is no difference between them as both can be easily implemented in software, firmware or hardware. The real difference appears to be which is more politically acceptable and technically applicable for worldwide use. Such acceptance concerns export restrictions (for example, DSA can be freely exported while RSA can't), IPR (that is, patents and licences) and functionality.

There are several standardisation options:

- do nothing and let the market adopt whatever it wants;
- standardise on one algorithm (for example, DSA or RSA); or
- standardise on two or more algorithms (for example, DSA and RSA).

There is also the question whether it should be a European or an international accepted standard. The latter of course makes more sense for international trade and global interworking.

Standards for infrastructure

Whatever algorithm is used by business, there is an urgent need to develop an infrastructure to support the use and management of signatures. Such an infrastructure is required for all business transactions and communications between organisations which might use digital signatures.

Various standards, specifications, technologies and services might be used to implement such an infrastructure. They might be used for a group of organisations operating within a sector, in the sense of a closed system, or by organisations operating in an open system. These standards need to be able to deal with **both** situations, and with a variety of signature algorithms; for example, in the case where DSA, RSA and perhaps other algorithms are being used by business.

Appendix 3—Bolero

Introduction and process

The Bolero project, funded 50% by the European Commission under its Infosec project, and 50% from commercial sources, is a pilot scheme to establish an electronic equivalent of the negotiable Bill of Lading. This is not the first attempt to solve the problem: many groups in the international trading arena have looked at the problem in the past but to date no working solution has been found

One of the main building blocks in this solution is to define the negotiable Bill Of Lading as more than one electronic message. This approach had already been adopted by the Maritime International group of lawyers. Instead of just one message being exchanged between the parties (the carrier and the shipper), a record of the Bill with a unique identification number is set up on a central registry which is a neutral party to the process. The record of the Bill is initiated by the carrier sending a message to the shipper, who confirms acceptance of becoming the first holder of the Bill.

The shipper may then want to transfer the message on to a bank, another trader in the chain or the final consignee. The shipper sends a message to the registry requesting a transfer of holdership to the proposed new holder. This message is then transferred to the proposed new holder, who can review the Bill and accept or decline the transfer. The shipper is then informed of the outcome. These messages together achieve the transfer of 'holdership' of the Bill of Lading thus **replicating** that function of a negotiable Bill which has so far been most difficult to achieve by electronic means.

Security

This in itself does not, however, achieve the full function of electronic negotiability. There must be some way in which the electronic Bill can

be proved to be unique. And **here** is where the concept of the TTP and digital signatures is introduced. In fact, within the Bolero system, a number of TTPs are used. The services could, theoretically, all be sourced from the same service provider or TTP, but in splitting them up between different service providers, an extra level of security is achieved, since to defraud the system more than just one single party would have to be involved.

The registry on which the record of each Bill of Lading is kept is a neutral TTP who is the 'holder' of the rights to the Bill at any one time, thus achieving the same uniqueness as an original paper Bill, since all the users of the system will be willing to accept that the information in the registry is correct. Anyone authorised to do so can check who that holder is at any one time. Also there needs to be a method of ensuring that a message has not been tampered with at any time.

These authenticity, non-repudiation and integrity functions are of major concern, and an adequate level of security must be provided in the system. This is achieved through the use of digital signatures using public and private key pairs. An asymmetric cryptosystem is used whereby each user has a public key and a private key. The public key is distributed widely to all users, while the private key is kept on a smart card held only by the user and with a personal identification number (PIN) code for access. When creating a message, the user authenticates it by calculating a hash value of it and signs it using his or her secret key. This digital signature is then appended to the document. The recipient uses the public key to decrypt the digital signature back to the hash value. This is then compared with the hash value calculated by the recipient on the electronic document using the same hash function. If the two values are the same then the message **does** come from who it purports to. If it is not, the message has been tampered with.

The process described above thus achieves the following security:

- secure identification of the individual responsible for sending the message;
- protection against attempts to tamper with the message contents after signature; and
- efficient administration of security based on a star-shape of one-to-one relationships with the central registry.

Registration and Certification

In order to use the Bolero service, each trading partner must register with the TTP service known as the *registration authority* (RA). The RA has two functions, the first of which is the registration of the organisation as a whole. This involves the RA in checking the credentials of an organisation and either accepting or rejecting it. The RA then enrolls the individuals within the organisation who are to hold the keys on behalf of the organisation. The RA will issue a smart card for creating key pairs to the user. A copy of the public key is sent by the RA to the certification authority (CA) for certification purposes. The CA ensures that key pairs are assigned to the registered name of only one user. The CA assumes that the user's credentials are in order because of the trust they place in the RA. Then the CA produces a public key certificate by signing the user's public key with its own certification key (private key) together with the time-frame for which the key is valid.

Development of Bolero


The initial Bolero pilot project had a total budget of nearly ECU 4.5 million, part of which was funded by the EC Infosec Program (DGXIII), with further funding through commercial partners. Starting in April 1994, the project involved eight trading chains from Europe to the US and Hong Kong, with a total of 26

pilot users. Although activity in the pilot concentrated on certain geographical areas, the benefits of such a system could be achieved on a global basis; thus attracting widespread interest in the project. The Bolero pilot project is based on international standards such as X.400 communication standard, X.500 directory standard and EDIFACT messaging.

Conclusion

The pilot ran from June–August 1995 with a group of around 20 actual users of the negotiable Bill of Lading including exporters, shipping companies, freight forwarders and banks. The Bolero project proved the technical concepts of digital signatures and the use of multiple trusted third-party services under the Infosec program. This project is thought to be the first real implementation of electronic commerce with organisations from different industries trading together on a global basis. The technical approach has had to build in a very high level of security, which has necessitated a considerable amount of work. This work has, however, thrown up new innovative ways of dealing securely and legally with negotiable documents electronically.

Biography



Nigel Hickson
Department of Trade
and Industry

Nigel Hickson obtained a B.Sc. in Electronic and Electrical Engineering from City University. After graduating, he went to work for Thorn EMI in its Enfield factory where he was responsible for technical quality on the company's first teletext colour televisions. He joined the DTI in 1982 and for the next nine years worked in a section responsible for export controls on high-technology goods. After a spell on financial services regulation in the DTI and then in HM Treasury, he joined the Telecommunications Division in 1994 to head a team working on commercial information security and information society standards. He is a member of the Institution of Electrical Engineers and is heavily involved in Scouting.

David Wilcox

Community Networks

Community networks which provide access to network systems, local information services and interactive communication are over 10 years old in North America, and now number several hundred there alone. Could similar models offer significant social and economic benefits for the UK? This article examines the technology they use, what we mean by community, and the nature of the networks which may be appropriate. These networks are as much social as technological.

Introduction

Over the next year, we are likely to hear much more about local authorities and voluntary bodies using the Internet and other local systems to provide public information services.

More Internet-access kiosks will open in public libraries, and cybercafes will be started in community centres. Systems which began with teletext listings of council services will migrate to the World Wide Web and expand to become 'virtual cities' of entertainment, shopping and information about every aspect of local life.

Voluntary bodies and community groups who previously limited their IT activities to word processing and databases will join in bids to Government 'challenge' funds, sponsors and the EU to develop multimedia projects, community networks or even telematic regions.

Bulletin board systems run by enthusiasts in their back bedrooms, carrying esoteric discussion of shareware programs, will become the electronic equivalent of community newspapers.

Champions of community networking hope that these various projects together will enhance local democracy, provide new education opportunities, and help bind together communities by communicating across boundaries which divide us in an increasingly specialised world.

However, before we join in hyping the possible benefits the information superhighway may bring at local level, we should examine North American developments, and consider the nature of 'community' and 'networking' in the UK.

I suggest we need to consider three issues with complex linkages: the technology, the information, and the connection of computer-based networks to 'real world' networks and activities.

North American Developments

North American community networks are varied, but have some common features which provide a framework for similar systems in the UK.

Technically, they help computer users in a locality connect with each other, mainly using modems and telephone lines. They may do this through one computer, the server, dedicated to acting as a hub in a closed system, or by building their network on the server of a commercial Internet service provider. Some provide that access free, covering costs through grants, donations and earnings from elsewhere; others charge subscriptions.

Users connect from home, work or public terminal, perhaps in a classroom or library. The connections available on different networks range from a couple of telephone lines to fibre-optic links. These links, perhaps between key sites like schools and training centres, could offer two-way videoconferencing, audio and data networking. Others are experimenting with wireless radio networks.

Through these connections users can exchange messages privately, conduct public discussions, exchange computer files and access centrally held information. Their exchanges may be limited to the local system, or they may have gateways to the wider Internet and other systems.

The information the networks offer to network users may range from timetables to job opportunities, electronic library catalogues, restaurant listings, tourist attractions and news reports to material from community groups.

The electronic networks depend on local organisations and real-world networks for their support in funding, information provision and volunteer

This new medium offers immense potential for helping people address many of the challenges to their individual success and the vitality of their communities.

help—local government, libraries, universities, voluntary bodies, sponsors. In turn, the networks will train local groups in using the technology and help them develop new projects.

The general form of organisation of a community network is a non-profit-distributing company, which may consequently gain tax advantages. Funding and support come from a mix of grants, donations, sponsorship, subscriptions and volunteer help.

The Morino Institute, one of the champions of community networks in the US, has published a directory of public access networks, which it splits into two sections: firstly Free-Nets, community networks and civic networks aiming to serve communities as a whole; and secondly special-focus networks. The latter is split into further categories including economic development, government information, education, community service, health and education.

Morino draws a distinction between relatively passive broadcasting and interactive communication where information users are also providers. It argues that:

'This new medium offers immense potential for helping people address many of the challenges to their individual success and the vitality of their communities. Those who have experienced the richness of interactive communications understand its ability to empower individuals, inspire collaboration, facilitate learning and enhance our patterns of access to people and information.

'On the grand scale, interactive communications is already connecting millions of individuals around the world in unrestrained dialogue and helping them to reach vast resources of knowledge and information. Closer to home, it is helping local communities energize citizen participation, reinvent institutions, provide outreach services and spur economic development. Perhaps the greatest opportunity is that it may provide a vehicle for bringing

together groups of people in collaborative efforts to solve the interconnected social problems afflicting those communities.'

Further details about the Morino Institute, other sources, and community networks themselves, are given at the end of the article.

Some Communication Basics

Before considering how the mix of technology, information handling and real-world connection may operate in the UK, I would like to ensure we are on common ground by dealing with the basic functions of interactive communication and the systems they use. In the process, I will declare a few personal dispositions which colour the later discussion.

One of the great difficulties in this field is the lack of a common language and framework of understanding between (social) community developers and (technical) network developers, so I make no apology for a simple approach which aims to close the gap a little by trying to make this article accessible to both.

The following functions are common to most systems.

E-mail, mailing lists and conferences.

I have found general accord that the most useful function of electronic systems for general users is e-mail—the one-to-one transmission of messages which might otherwise go by fax or post.

Two forms of public e-mail provide additional benefits. Mailing lists allow members of an interest group to subscribe to a central e-mail address, and then receive any message posted by any other subscriber.

Conferences, also called *newsgroups* or *forums*, perform the same function except messages stay at the central address and users must take the trouble to access them periodically rather than find them automatically piling up in their mailbox each day.

Personally I prefer mailing lists, but over-enthusiasm in subscribing to them gives you the equivalent of an overflowing intray every day, with the constant temptation to check what's just come in and to join in the discussion.

Conferences can work well, but I believe that their likely value in community networks depends upon the software available. Older systems use text-based menus which are frustrating for those used to the more friendly graphical interfaces of Macintosh or Windows.

File transfer

File transfer theoretically enables system users to send anything on their computer to that of any other subscriber to the system: whether text file, magazine layout, spread sheet or graphic. In practice, the ease of transfer depends on the nature of the system and agreement of common standards.

Files sent attached to Internet e-mail have to be encoded and decoded, and this can be troublesome. Unless users agree simple matters like which word processing format, anything more than a simple text file must go through tedious processes of translation.

Other means of retrieving files from distant computers tend to baffle novice users.

These difficulties apparently disappear when using World Wide Web (WWW), software which presents the user with attractive pages of text, graphics and hypertext 'hot links' to other pages which may be on the same computer or anywhere else on the Internet. Just click on a link and you can 'surf' across the world to further information, download a file, e-mail a response or join an associated conference.

WWW appears to be the application which makes everything work easily, but there are problems even there. All pages must be designed and prepared using HTML—HyperText Markup Language. While this is becoming easier, it means that WWW

is principally a one-to-many publishing medium, compared with, for example, many-to-many publishing through a mailing list.

And although WWW can be read as text only, its main benefits require a computer capable of running Windows, or a Mac. Links can be slow, and combined with the temptation to just try the next connection to find what you really want, can consume a lot of time.

Systems and platforms

The pros and cons and possibilities of different applications and approaches come into sharper relief when considering the different systems or platforms on which community networks may run.

Community networks in North America started on simple bulletin board systems (BBSs) as early as 1980, when David Hughes set up Old Colorado City Communications in Colorado Springs to help people become more active and involved in local government.

These early systems required their users to deal with text menus to work their way through to discussion and file areas, and would limit e-mail addresses to others on the same system. Later BBSs would connect together, so that users could send messages across the country, and the world, by Fidonet.

Today a community network using a BBS might use First Class software which can run on equipment costing under £2000, and a software licence costing a few hundred pounds. Subscribers are greeted with the equivalent of a new Mac or Windows desktop (unless they are running DOS), and a range of e-mail, conferencing and file transfer options.

Costs will, of course, increase substantially with the numbers of users, storage and connection requirements, but the look and feel of the system is much the same whatever the scale. It has some similarities with commercial systems like CompuServe, that have for some years offered subscribers a friendly

graphical interface, with e-mail, conference areas, specialist databases and a host of information services from timetables and weather reports, to news services and film reviews.

The BBS-based system could expand its services to include e-mail out to and in from the Internet. It might offer both direct dial-in and access via the Internet for long-distance users.

Many community networks have seized the opportunities offered by WWW, and created community networks on systems run by universities or Internet service providers. These systems are excellent for publishing information attractively, but demand that all core material is designed and coded using HTML. There may be associated discussion groups, but in my view these are not as flexible as those running on First Class systems.

No doubt, in a few years, personal computers and television sets will be one, and conferencing and WWW software will converge, but we aren't there yet. Even so, I think that the constraints on community networks lie more with information handling and their real-world relevance than with technical limitations.

Organising the Information

I suspect that good community networks need a mix of skills drawn from journalism, librarianship and facilitation. News keeps users coming back to check what's happening. Well-structured information on screen with signposts elsewhere is the core of the system. Conferences are likely to need moderators to keep people on track and ensure an acceptable signal-to-noise ratio.

Community networkers should consider how much central work they can put in to develop and maintain a Web site, and how far they want to encourage users to be information providers under central guidance.

Whatever the software and the system, experience suggests that a great deal of effort and skill is needed

to create an information-rich environment and keep it updated.

The Social Context

At one level it has never been easier for the UK home or business user to enjoy the benefits of interactive communication. While there are no developed community networks on the North American model, he or she can open an account with CompuServe which will provide a wealth of information plus Internet access at about £6 a month plus use charges, or simply gain access to the Internet for £10–15 a month plus only local call charges.

Once on either system, he or she will have international e-mail, an ability to transfer files, and access to the WWW. A look at the North American community networks may prompt a socially-conscious user to think about developing something similar, locally. Here I would offer the enthusiast a few cautions, which lead to where I think the main issues for community networks lie—in the real, not the virtual, world. Here, first, are the cautions:

- The level of computer use and connectivity in the UK is much lower here than in North America, and getting on-line is a big jump for most people. The potential user base in most localities is still very low.
- Little will be achieved by creating a WWW site called *Anytown Community Network*, with some basic information, and hoping enthusiastic volunteers will e-mail you with pages of fascinating material. Ask the editor of a community or parish newspaper how easy it is to get good copy.
- Even creating mailing lists or conferences around local topics and hoping people will send basic messages may not work. People simply 'lurk' on the lists, uninterested or under-confident about joining in.

- Managing even a modest BBS can be time-consuming, maintaining a Web site even more so.
- Key public sector partners for a community network—local council, libraries, colleges, training agencies—may not even have e-mail. They probably won't understand what you are talking about.
- Community and voluntary sector interests are unlikely to see on-line working as a priority, unless there are very obvious benefits.
- Potential private sector supporters will be more likely to offer help in kind than funding.

Behind these observations lie some lessons and possible guidelines for community networking.

Networking the Real-World Community

Community networks are significantly different from public or private ventures. They rely on a wide range of different interests to commit technical support, information, and 'animation' voluntarily.

The collaboration and commitment needed to build them does not come easily, particularly when the different interests probably haven't worked together before and are new to the technology.

This means a lot of early effort to build up teams, assess technical needs, develop information structures, plan training and support. Later this work may be done by paid staff, but in the early days there may be no funding to cover the work. Start-up costs cannot easily be repaid from later funding or earnings.

The more 'commercial' a system tries to make itself, the less voluntary effort it may attract.

However, there is a positive side to these difficulties. In reality, there is no one 'community'. There are communities of locality and of interest. Every place has a myriad of

links of family and friendship, of clubs for sport and recreation, support for health and learning, campaigning for the environment and citizens rights.

These communities need to communicate better within their interest groups, and between themselves. They need to share experience with others nationally and internationally, and increasingly they need to form partnerships to bid for funds from Government and Europe and demonstrate continuing collaboration.

They are the human infrastructure which already uses the technology of telephone and fax, as well as print and face-to-face meetings, to maintain the web of relationships which makes up what we call *community*.

So for the new community networker a good starting place is the existing networks which want to enhance their activity. They have

Organisations and networks need customised solutions which go beyond a simple Internet connection.

information and they have members who could become network users. They will do your marketing and provide content. They understand voluntary effort, fundraising and non-profit-distributing companies.

Once a critical mass of users develops, there will be scope for lively informal networking between individual users, finding others with similar interests, trading information and maybe holding the electronic equivalent of parties or car boot sales.

This is in addition to the more conventional information providers in local government, libraries, schools and colleges and other public agencies who will soon begin to feel some political pressure to join in.

A UK Model?

I wouldn't like to propose an off-the-shelf UK model for community networking. We start from a different

place than the Americans or Canadians. We have better Internet connectivity, but higher local call charges. Our traditions of volunteering and civic responsibility are different. We have better public service broadcasting and, perhaps, a better press. There may be less perceived need for local electronic networks.

On the other hand there is growing interest in getting on-line, and a realisation that, while the Internet may provide good connections, it doesn't necessarily provide useful information. There is too much, which is too diverse and too difficult to find.

Organisations and networks need customised solutions which go beyond a simple Internet connection.

UK community networks may well be networks of networks and projects, as much as information providers. They may be teams of technical experts, information brokers, trainers

and facilitators who ensure there is an appropriate platform, and help people use the new medium.

Their work will reinforce our sense of real-world place, as well as create new virtual communities on-line, because they will be making face-to-face connections which might not otherwise happen. I have found e-mail has increased the people I know across professional and other boundaries, and I have now met quite a few. Community networking could help get us away from the computer screen into new real worlds, as well as create virtual ones.

David Greenop, writing in the October 1995 issue of this journal, touched on many of the social issues raised by the new technology. He remarked: 'Above all else in bringing about the information society is the necessity of partnership between the builders of the information infrastructure and the representatives of society.'

He added: 'The challenges and the opportunities of new technologies must be made apparent to all members of society and its institutions, and a healthy debate must be encouraged.'

Much of this debate will take place through normal media and political channels. However, community networks—broadly defined—could provide the telecommunications industry with an unparalleled opportunity to test different approaches and contact a wide range of individual and organisational users not readily accessible through normal public and private sector channels.

Even if we discount the hype, community networks could demonstrate the social benefits of connecting computers at a time when people are becoming fearful of an increasingly machine-dependant society.

The only way to find out what may or may not work is to try it. We need a number of local pilot projects which share their experience nationally and internationally, and which draw on expertise in telecommunications, computing, information management and community development.

Further Information

The Morino Institute is dedicated to opening the doors of opportunity—economic, civic, health, and education—and empowering people to improve their lives and communities in the communications age.
<http://www.morino.org/>

The National Public Tele-computing Network is the parent body for Free-Net systems worldwide.
<http://www.nptn.org/>

Computer Professionals for Social Responsibility has a reference section on community networking. <http://www.cpsr.org/dox/community.nets.html>

The Center for Civic Networking is a non-profit organisation dedicated to applying information infrastructure to the broad public good.
<http://civic.net/ccn.html>

Community networking documents and resources can be found at the International National Federation of Library Associations and Institutions.

<http://www.nlc-bnc.ca/ifla/services/commun.htm>

Among the useful documents are Community Computer Networks: Building Electronic Greenbelts by Steve Cisler of Network Outreach, Apple Computer, Inc, and his Can We Keep Community Networks Running?

As I was writing this article, a lively discussion started on the (mainly North American) Communit mailing list about the idea of an International Association for Community Networks, and the nature of the model. Discussion ranged over many of the points raised in this article, particularly whether networks should focus on providing access, publishing information, encouraging discussion, and building real community.

To join the mailing list, send a message to listproc@list.uvm.edu with no subject, and message content subscribe communit yourfirstname yourlastname

You can then post to the list by addressing messages to communit@list.uvm.edu.

To see archive files, send a message index communit to listproc@list.uvm.edu.

I am also happy to respond personally to enquiries, preferably by e-mail.

Finding networks

Free-Nets and community networks are listed by Peter Scott.
<http://duke.usask.ca/~scott/free.html>

Two WWW-based UK community networks are:

Coventry Community Network.
<http://www.ecosaurus.co.uk/coventry/>

Capital Net-Cardiff.
<http://info.cf.ac.uk/ccin/homepage.html>

Books

The Virtual Community, by Howard Rheingold, Minerva, 1994, £6.95, provides an optimistic view of the benefits of the technology.

Silicon Snake Oil, by Clifford Stoll, Macmillan 1995, £9.95, offers sceptical second thoughts on the information highway.

Biography



David Wilcox
Partnerships for
Tomorrow

David Wilcox spent 12 years as a journalist, working mainly for the *London Evening Standard*, before specialising as a consultant and trainer in the development of partnership organisations and techniques for community participation. He now plans to take the lessons of journalism and partnership building on-line. After visiting the United States to research community networks earlier this year he set up the Partnerships for Tomorrow network for people interested in exploring how new media technologies can benefit communities. He recently ran the Communities Online conference at BT Centre, and is now developing the Community Regeneration Network with a number of national voluntary bodies.

David Wilcox:
Partnerships for Tomorrow,
13 Pelham Square,
Brighton BN1 4ET, UK.
Tel: +44 (0) 1273 677377
Fax: +44 (0) 1273 677379
E-mail: dwilcox@pavilion.co.uk

Mike Whybray and Mark Shackleton

Image Processing for Telecommunications in the 21st Century

Image processing in telecommunications has, so far, concentrated on compression to allow still or moving picture transmission over narrow-bandwidth circuits. With the provision of high-bit-rate channels, this need will diminish though not disappear. Instead, as the electronic transfer of pictorial information becomes commonplace, there will be a demand for 'intelligent' image processing, which is responsive to the content.

No Future for Compression?

In telecommunications, the main driving force for still- and moving-image processing has been compression—squeezing the same picture quality into fewer bits. Because transmission channels had limited capacity, there was an economic benefit. With optical fibre bringing the promise of many orders of magnitude increase in capacity, will we still need compression in the 21st century? The answer is 'Yes' for three reasons.

Firstly, although transmission capacity may be very high, some source data still has to be stored. Considering the number of hours of film and video already in archives, and the expected growth in capacity of known storage technologies, there will still be a very significant benefit from data compression.

Secondly, optical-fibre cable will not be everywhere, nor serve every purpose. This is particularly the case for portable or mobile systems, which will still rely on radio (or even free-field optical) transmission, with the finite overall capacity shared between users.

Thirdly, history shows that what at first seems to be an unimaginable increase in transmission capacity soon gets used up by increasingly voracious services. Copper pairs have given way to coaxial cable, waveguides, multi-mode fibre, mono-mode fibre, and wavelength-division multiplexed fibre. Video will be a major capacity driver here, as new camera, display, and processing

technologies evolve through high-definition to super-high-definition television (double the bandwidth), stereoscopic viewing (two views at once), holographic television (view from any angle), and beyond. This is the technologists' view of the future!

An alternative view

Programme distributors such as satellite TV companies see the future differently. To them, digital TV is not a means of delivering higher quality, but a means of cramming more channels into a given bandwidth! Reasons for this are the vast investment in old TV and film footage, exploitation of the existing receiver investment, the cost of larger higher-resolution displays, and the perceived benefit of more programmes as opposed to higher visual quality. It may not be until fibre is used to distribute the majority of broadcast TV, and the airwaves are used for the more appropriate purpose of mobile services, that the potential quality improvements to TV pictures will be realised. In this sector at least, image compression very definitely has a future into the next century.

Finding the Bits You Want

Providing more and more channels of even standard-quality TV brings its own problems. With more than 100 channels to choose from, how do you decide what to watch? Channel-hopping at a rate of a few seconds per channel will take many minutes to cycle round all the channels, and those few seconds will not be enough




Figure 1—A set of 'eigenfaces'

to identify what many programmes are about, unless you recognise a key character for example. The situation is even worse when faced with millions of hours of video in a database! Although there will still be high demand for conventional programmes, where the viewer simply watches a broadcast, increasingly people will control their viewing. Control will range from the simple level of choosing when a programme starts (rather than having to organise one's life around it or remembering to set the video recorder), to selecting subject areas, the level of censorship applied to a programme (for family viewing say), or the categories of news items of interest.

The solution to this information overload is likely to be a combination of smart video-searching and browsing tools, with well-structured, annotated and cross-referenced databases of video material, and artificial intelligence 'agents'.

Annotation

Perhaps the two most obvious ways to enable users to find the information they want from a large video database are to structure it hierarchically, and to provide an index based on keywords. Text attached as annotations to individual frames, sections, or complete programmes can then be used to select items using conven-

tional text searching methods. However, someone has to annotate the material first. This is not a trivial task, and is confounded by the subjective nature of deciding how to describe a particular segment of video. For example, in a shot of two people in a car, does the make and colour of the car matter? Is the time of day significant? Is the mountainous scenery behind important? Or is the fact that the two people are Bogart and Bacall the only relevant information? The answer of course depends on who is searching and why.

A more open-ended approach is to search through video material by means of the image content itself. For example, if you want to find a red car, you would show the system an example, and the system would search through its stored video for similar objects, and present you with a choice of possible matches. Or if you wanted Bogart and Bacall together, you would show the system examples of each and it would search for matches to their faces occurring in the same frame of video. Although this sounds far-fetched, in fact this can already be achieved to some extent.

Face matching

Techniques have already been developed to find an example face in a database of several thousand pre-stored face images, with correct

matches being achieved over 95% of the time. Such matching technology will ultimately allow automatic labelling of video with the names of the actors and actresses, and the sections where they are on-screen, making it possible for a viewer to fast-forward to the next appearance of a particular star². In the shorter term, there are obvious commercial applications of the technology in police databases of suspects, and many security-related areas such as secure site entry, screening for known terrorists at passport control, and the reduction of the fraudulent use of credit cards and bank ATM cards.

The face matching process is very efficient as it transforms a full face image into a very compact code consisting of only a few numbers, rather than using more traditional techniques which match every one of hundreds of thousands of pixel intensities of the test face against every face image in the database being searched. The short code is derived by taking a small set of so-called *eigenfaces* (see Figure 1) carefully constructed to represent faces, and working out how much of each eigenface is required to re-synthesise the test face. This process gives a single weighting number for each eigenface, and the resulting set of numbers provides the code describing the face. These codes are pre-computed for a database of faces, and can be compared extremely quickly, allowing matching of thousands or even millions of faces in real-time. The image matching process is reasonably tolerant of changes in facial expression and hairstyle, and can be combined with traditional textual database entries to permit powerful searches such as 'show me people who look like this, but who now have a beard'. Using these techniques in a police suspect database for example, a witness need only look at a small subset of similar-looking candidates matching the specified criteria rather than leaf through large photo albums of potential suspects.

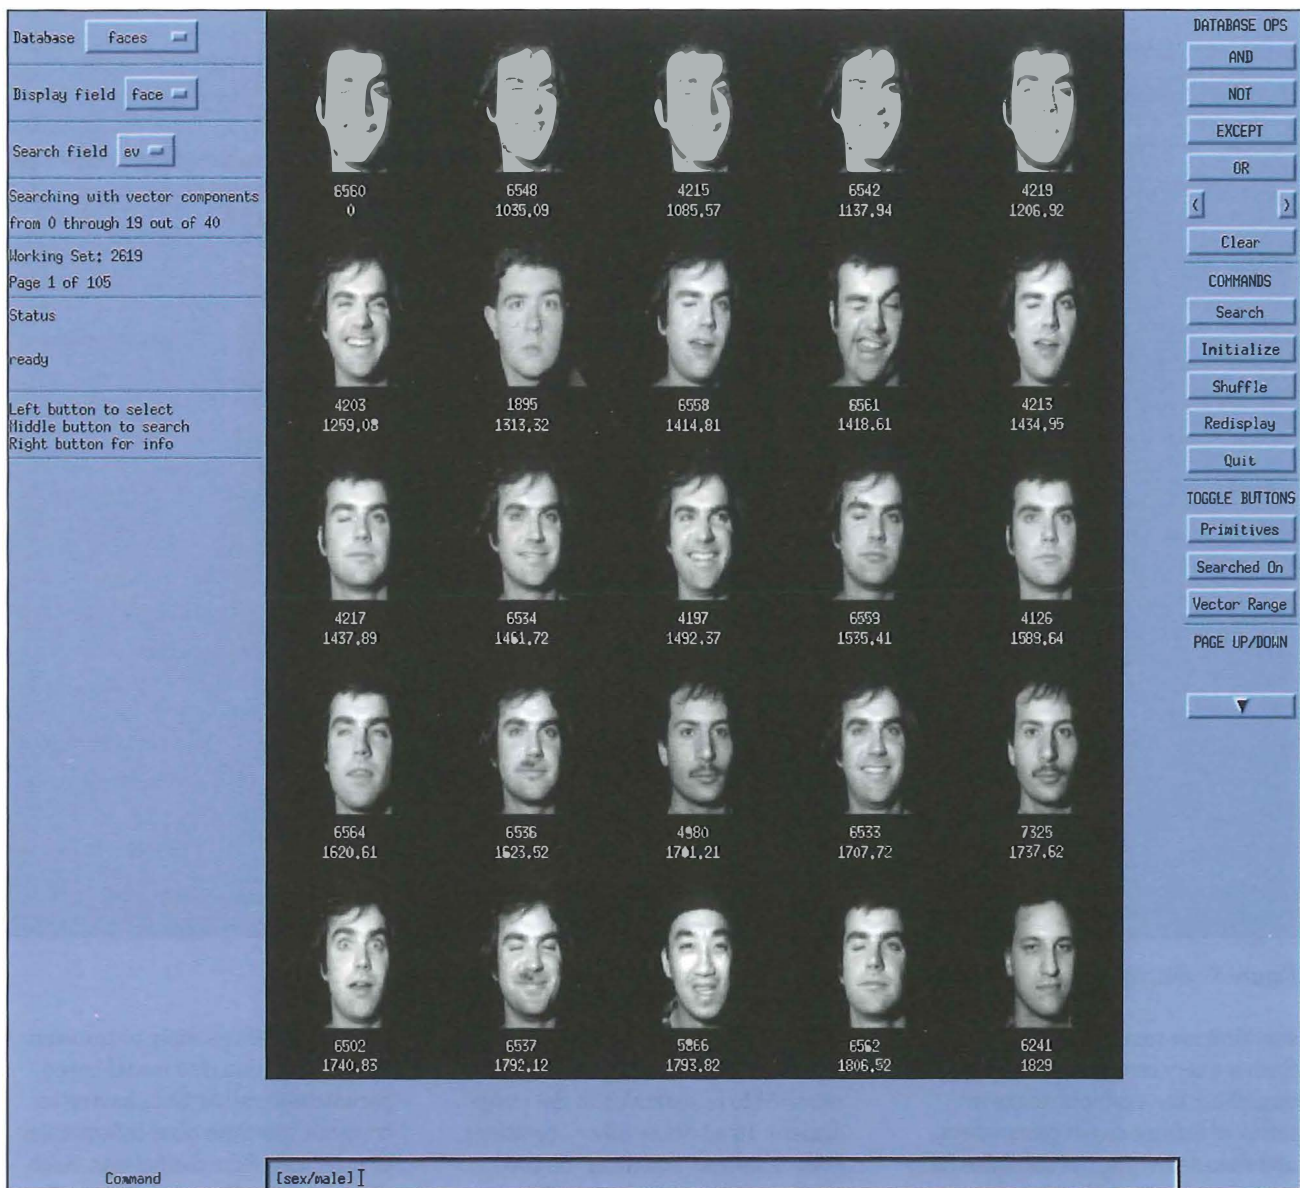


Figure 2—Top matches after a face database search

Figure 2 shows an example search, where the face in the top left corner was used as the original image to search a database of 2619 faces, containing several photographs each of many hundreds of different people, of mixed age and ethnic origin. The screen shows the resulting matches ordered in decreasing closeness of match scanning across the screen and then down row by row. Most of the faces are different photographs of the same person, but here and there other people appear. This illustrates an important point about face

matching generally—it can never guarantee to be 100% correct. However, it can bring close matches, including the desired match, up to the top few percent of a database, thereby greatly reducing the number of faces that a human operator actually has to look at. This makes many tasks possible that would otherwise be infeasible.

Generalised image matching

Faces are a very specific example of objects one might wish to find in databases; however, other properties

are more difficult to describe, and are typically known in computer vision research as *textures*. Again, techniques have been developed to allow textures to be described efficiently as a small set of numbers. For example, descriptions of image regions corresponding to grass, water or buildings can be derived and classified allowing a user to search for similar properties in other images. Sometimes a user might find it difficult to describe exactly what is sought but might prefer to point to an image, or highlight a region of an image, and

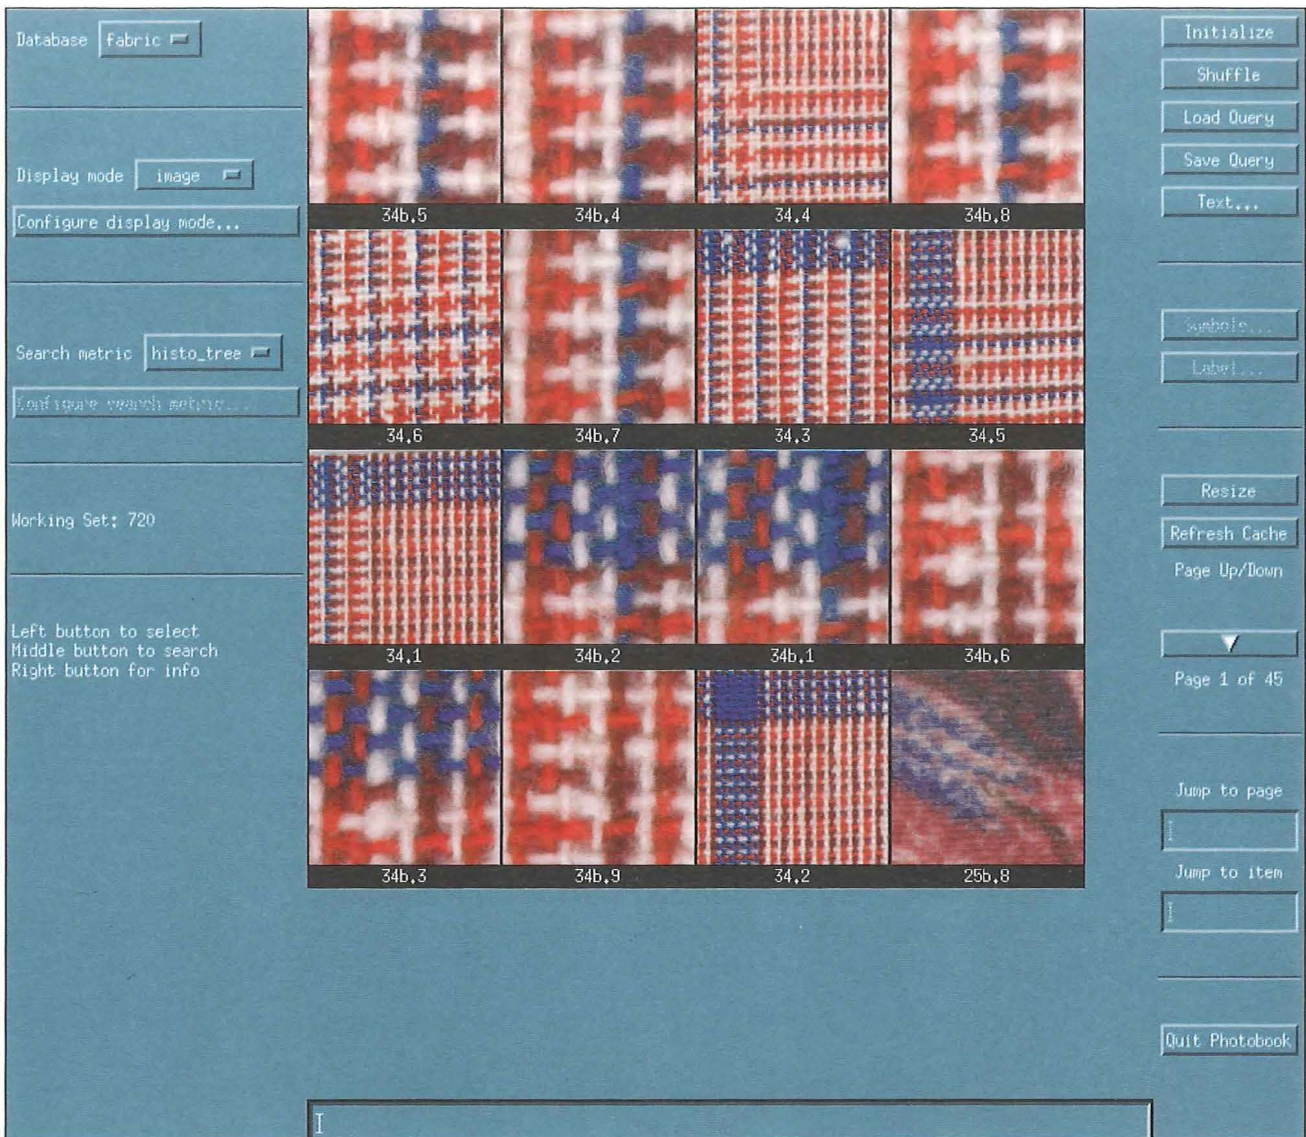


Figure 3—Search for matching fabrics

say, 'find me something like this'. Such a query can be answered by describing the example image in terms of texture model parameters, and then searching the database for images containing similar pre-computed descriptions.

An application of texture-based matching which has recently been explored is automatic matching of fabrics for the fashion industry. A database of fabrics, which can be automatically matched using various criteria such as overall colour composition or fabric 'roughness' (see Figure 3), has been created. Matching could be incorporated into a database holding details of suppliers, and stock information to allow customers wishing to source a particular type of cloth to search for suppliers of similar fabrics. In essence, these automatic image content-based matching techniques are allowing the tradi-

tional alphanumeric database search tools such as searching, sorting and retrieval to be extended to the image domain. In addition, other operations such as inexact matching (no two objects in an image are usually identical) and browsing are required.

The recent developments in intelligent image content analysis have an important knock-on impact on image coding for compression. Increasingly, rather than simply transmitting visual information, people are going to want to store, search, browse and retrieve that information.

Searching image content typically requires describing images in terms of the objects an image contains such as faces, cars and trees as well as regions of texture such as grass or water. If an image can be effectively described in terms of these objects, regions and texture attributes then it

becomes feasible simply to transmit that information (the model-based parameters) rather than having to transmit complete pixel information or even waveform coefficients. Such *object based* coding algorithms offer the hope of far greater compression than existing coding standards, as well as a better match to the requirements for content-based searching and manipulation^{3,4}.

It may take several decades before computer understanding of images can achieve the accuracy of a human being, but in the shorter-term computers will be able to 'power assist' a human. An example is to speed up the process of manual annotation of video by parsing the video to detect *keyframes*—places where there is a significant change in scene content such as a cut, violent motion, or a new character enters a scene. Rather than view the entire

Figure 4—Segmenting a background from a person...and replacing it

video sequence in real time, the annotator has only to skip through the keyframes and annotate those containing the vast majority of useful information.

Looking at People


In much video material, the objects of most interest to humans are other humans. We tend to fix our gaze mainly on the faces of other people if we are involved in a conversation with them. People therefore have a special place in the field of image processing, and many applications can be envisioned based on developments in this specialised domain.

Smart videophones

Today, videophones are the equivalent of two empty bean tins and a piece of string. They simply transmit the image received by a camera at one end to a display at the other. The result is often sadly disappointing to people weaned on expertly composed broadcast television images of newsreaders, panel discussions and the like. Imagine a videophone that also operates as cameraman, lighting designer and director, keeping you reasonably central despite your shifting position, correcting for poor lighting, and substituting an alternative background to replace your office clutter.

The first is already possible, using techniques similar to face-recognition processing to find accurately the position and scale of a face or faces in an image, and hence select and scale an appropriate window from the output of a camera. Having found the face, a better judgement of camera settings to achieve a well-exposed image can be made, avoiding the problems of back-lighting for example. Some correction for excessive side-lighting will also be possible.

Finally, people could be segmented from their surroundings without resorting to the chroma-keying used in broadcast television. This final stage will have several benefits. It will afford greater privacy to videophone



users as it is only themselves, not their surroundings, which are exposed to view. The users will be able to choose their new background (see Figure 4): a corporate logo, a picture of their city, or start a new craze in video wallpaper—perhaps a new outlet for computer screen savers.

Only a few years ago such complex processing would have been impractical, but by the 21st century even desktop and portable computers will have sufficient processing power to tackle such tasks, as well as implementing the basic image compression work required in a videophone.

Continuous presence multipoint

To widen the use of videophone and videoconference services, multipoint control units (MCUs) which allow more than two video terminals to be linked together are already in service. All participating terminals connect into the MCU in a star configuration to create a distributed meeting. Existing MCUs operate in 'switched' mode, however, meaning that only one incoming video stream at a time is available to any given terminal. In essence, people at one terminal can only see those at one other terminal at

As well as enabling better person-to-person communication, image processing will increasingly improve the human-to-computer interface.

a time. The chosen one is selected dynamically either by direct control of a designated chairperson or, more commonly, automatically according to who is shouting loudest! Both methods have obvious disadvantages, and many nuances of a normal meeting are lost.

An answer is *continuous presence multipoint*, where sound and pictures from all other locations are available at every terminal. There is significant market demand for this, as most first-time users of multipoint conferencing assume this is what they will get, and are disappointed when they discover that the current system is a switched one. In its basic form, continuous presence multipoint simply requires the MCU to transmit all incoming video streams out to all terminals. However, with N participating terminals, this would require an MCU-to-terminal channel capacity $(N-1)$ times that of the terminal-to-MCU link. As in practice most channels are symmetrical, the video bit rate must be squeezed down by a factor of $(N-1)$.

The brute force approach to this is to decode the compressed bit streams back into images in the MCU, then re-encode them at the lower bit rate. Unfortunately, this can introduce additional coding artefacts and a significant delay. New methods have recently been proposed⁵ which allow the bit rate to be reduced by a process of re-quantisation, with the use of error feedback to avoid distortion accumulating in an uncontrolled manner. The result is a bit stream which has been further compressed, but with little extra delay, and a quality only slightly lower than if the full compression had been performed in one pass.

Another step will be to apply the automatic segmentation of people from backgrounds previously mentioned, to remove people from their real surroundings, and merge the images together into some form of virtual meeting environment. This could range in complexity from using the same background colour for everyone, through arranging people round a synthetic meeting table, up to

a full three-dimensional virtual reality environment.

Smart interfaces

As well as enabling better person-to-person communication, image processing will increasingly improve the human-to-computer interface.

The goal is to enable humans to interact with computers in a natural manner. Speech is an important part of this process, but people also use hand and body gestures, eye contact, gaze awareness and facial expressions.

Techniques which use optical flow analysis to monitor small movements of parts of the face can now detect the difference between a fake smile and a real one—the eyes crinkle as well⁶.

Similarly, simple gestures can be recognised and used to replace some of the pointing and selection functions of the computer mouse⁷. An interesting application is sign language interpretation, where very complex hand gestures need to be correctly analysed and translated into alternative modalities such as text or speech. Also, facial expressions are used to modify the meanings of many signs, and so form an integral part of the language. Currently, recognition of about 40 signs has been demonstrated⁸, but within a decade it is likely that unconstrained vocabularies will be possible. The main group to benefit from this will be profoundly deaf people (about 50 000 in the UK), most of whom use sign language as their preferred means of communication, but who currently have to use text-based systems to communicate over the telephone network. Video-phones will increasingly allow them to converse with each other by sign language, but unconstrained sign language recognition will enable them to converse freely with anyone on the network, by whatever method they choose.

As well as providing new input modes, image processing can also provide new output modes. One possibility is that of synthetic people who are animated by a computer to provide a focus of attention. When the

computer speaks, it animates the mouth, facial expressions body position, stance and gesticulations. Since this form of computer output can now be driven by unconstrained text⁹, it is very flexible, and does not suffer from the disadvantages of using recorded video segments of a real human such as a fixed set of responses, and the requirement to store large amounts of video. By using pictures of real people to provide the basic 'texture map', it is possible to produce results which casual observers do not even notice are actually synthetic.

Finally, an increasingly important application of image processing is virtual reality, where graphics techniques are used to render images of three-dimensional worlds that users can interact with. Increasing processing power and more advanced texture mapping, modelling and rendering techniques are transforming these worlds from jerky, blocky places to satisfyingly 'analogue' ones. These worlds can offer ways to navigate through information by analogy to the physical world (for example, creating a virtual office with familiar items such as file cabinets and telephones), and by inventing entirely new ways of visualising data such as *emotional icons*¹⁰. Techniques such as automatic segmentation, and synthesis of realistic images of people will enable people to meet and converse with each other in these worlds, both for work, socially and at play.

Conclusion

In the 21st century the need for compression for transmission will gracefully wane to be replaced by compression for image storage. This will see new functional requirements such as the ability to deconstruct images into their component objects, and to search on image content. Object-based coding algorithms may address these issues, as well as providing higher compression. Image analysis and understanding will be used to search or browse video material

by content, rather than relying on attached textual annotations.

There will be a particular focus on analysis of images containing people, resulting in computers able to see and respond to human gestures and facial expressions, and recognise individuals. This will have applications in many areas, particularly more natural user interfaces, virtual reality systems, and security.

Achieving the above vision will rely on significant but foreseeable improvements in image processing algorithms, and a parallel increase in the processing power available in desktop and portable computers. Of course, the more interesting developments are bound to be those that are not foreseen!

References

- 1 PENTLAND, A. P., PICARD, R. W., DAVENPORT, G., and WELSH, W. J. The BT/MIT Project on Advanced Image Tools for Telecommunications: An Overview. *Image'Com 93*, Bordeaux, France, March 23–25 1993, pp. 397–400.
- 2 PENTLAND, A. P., PICARD, R. W., and SCARLOFF, S. Photobook: Tools for Content-based Manipulation of Image Databases. *SPIE Conf. Storage and Retrieval of Image and Video Databases II*, No. 2185, February 6–10, 1994, San Jose, CA.
- 3 MUSMANN, H. G., HOTTER, M. and OSTERMANN, J. Object-oriented analysis-synthesis coding of moving images. *Signal Processing: Image Communication*, 1(2), Oct. 1989, pp. 117–138.
- 4 WELSH, W. J., SEARBY, S., and WAITE, J. B. Model-based image coding. *Br. Telecom Tech. J.*, July 1990, 8(3), pp. 94.
- 5 MORRISON, D. G., NILSSON, M. E., and GHANBARI, M. Reduction of the bit-rate of compressed video whilst in its coded form. *Sixth International Workshop on Packet Video '94*, Portland, Oregon, Sept. 26–27, 1994.
- 6 ESSA, I. A., and PENTLAND, A. P. A vision system for observing and extracting facial action parameters. *IEEE CVPR 1994 Conference*, Seattle, Washington, June 1994.
- 7 DARREL, T. J., and PENTLAND, A. P. Space-time gestures. *IEEE Conference on Computer Vision and Pattern Recognition*, 1993, pp. 335–340.
- 8 STARNER, T. E. Visual recognition of American sign language using hidden Markov models. M.Sc. thesis, Feb. 1995, MIT.
- 9 WELSH, W. J., SIMONS, A. D., HUTCHINSON, R. A., and SEARBY, S. Synthetic face generation for enhancing a user interface. *Image'Com 90*, Bordeaux, France, Nov. 19–21 1990, pp. 177–182.
- 10 FISHER, KIM, TAYLOR-HENDRY, ROB, LINTON, DAVE, and COOPER, MARTIN. Non-Verbal Guidance for Cyberspace Explorers. *Br. Telecommun. Eng.*, July 1995, 14, p. 131.


Biographies



Mike Whybray
BT Networks and
Systems

Mike Whybray gained a degree from the University of Birmingham in 1977, joining BT Laboratories to work on reliability physics, studying water vapour contamination in semiconductor packages, failure modes of magnetic bubble memories, and CMOS latchup problems. In 1983, he gained an M.Sc. at the University of Essex, and moved to the Visual Telecommunications Division of BT Laboratories, where he became leader

of a group whose activities included developing videophones for deaf people, writing software for DSP-based videophones, and building demonstrators of advanced image processing applications such as intelligent surveillance and synthetic 'talking head' displays. He currently leads the Image Coding and Processing group within Advanced Applications and Technology, developing new image-coding algorithms and standards, face-recognition technology, and image analysis and understanding systems. Mike is a Chartered Engineer and a Member of the IEE.



Mark Shackleton
BT Networks and
Systems

Mark Shackleton graduated with a B.Sc. in Computer Science from the University of Sheffield in 1986. He then worked for three years for a manufacturer of flight simulators developing algorithms and software to generate realistic three-dimensional terrain imagery for aircraft cockpit displays. In 1989, he joined the Image Coding and Processing group at BT moving from computer graphics generation into the field of image processing and computer vision. Areas of research have included automatic face recognition and content-based retrieval of images and video, using computer vision techniques. As part of a collaborative research initiative, he has spent about a year on secondment working alongside researchers at the Massachusetts Institute of Technology Media Laboratory investigating image interpretation algorithms and their application.

The Virtual Science Park

The virtual science park developed by the University of Leeds is one of the first to be constructed and tried in a real-world situation. It is the first step in the creation of a virtual workplace and market for the 21st century.

Introduction

The University of Leeds has constructed a virtual working system (VWS) to enhance the university's ability to interact with industry and develop applied research and workplace learning activities. Many UK universities have joined forces with local authorities and financial backers to invest in physical science parks to bring industrial and academic partners together. However, in the case of Leeds, a feasibility study¹ showed this to be inappropriate for the following reasons:

- the university is close to the city centre without adjacent green field sites so the collocation of a science park and the university is impossible;
- there are already considerable numbers of science parks in the UK and Europe in attractive areas suggesting that it would be difficult to persuade key companies to move to a northern city; and
- there is only a small high-technology industrial base in the surrounding region.

The technical challenge was thus to develop a system that delivers many of the services of a conventional science park within an integrated electronic environment. The *Virtual Science Park* (VSP) provides services such as consultancy and consortium building; access to specialist knowledge (databases and expertise); delivery of workplace learning; and technical support. The current prototype VSP has been trialled across the campus to demonstrate

The authors are from the School of Computer Studies and School of Geography, University of Leeds

feasibility and explain the concept to potential users. Tests between Doncaster Business Link offices and the Leeds Environment Centre (LEC) have provided environmental consultancy between small/medium-sized enterprises (SMEs) and environmental experts. PC videoconferencing, using the ISDN-2 based ScreenCall (BT VC8000) conferencing package, was used to provide collaborative working tools for audio, video and data conferencing. This article describes some of the design detail and practical experience with this prototype VSP.

VSP Services

The main purpose of the VSP is to enhance the university's ability to interact with industry and deliver educational services. It achieves this by characterising organisations in terms of facilities, resources, expertise of employees and its current and completed contracts or work projects. This enables it to support a number of services:

- consultancy,
- access to specialised databases and experts,
- support,
- delivery of workplace learning, and
- assembly of focused consortia.

In order to illustrate a typical service, consider the following situation: a funding agency invites bids to address problems posed by a particular environmental hazard. The environmental agency (EA) needs to identify the people with the required expertise and to assemble a bid. The VSP provides the EA with access to the necessary expertise from aca-

demographic departments and industry. Once the right people have been identified and the contract set, the VSP provides the necessary communication tools to undertake the consultancy.

VSP Information Model

The VSP is designed to support interaction between information providers and those seeking information within the context of industrial-academic research activities². Information providers must be represented as accurately and concisely as possible. A structured information model has been defined which ensures that all the VSP tools consistently manipulate the same information. Experience has demonstrated that the model needs to be person- rather than organisation-centred. Where companies, for commercial reasons, do not wish to offer direct access to staff, the VSP model is centred on a group of people with particular knowledge and skills. Interaction then takes place between a 'gate-keeper' for the group and the VSP user.

There have been many suggestions as to how the knowledge and skills of an organisation might be represented³, although there is no consensus on how to achieve this. In fact, there are some fundamental differences between skill and knowledge. Skill implies know-how and the ability to make something happen, while knowledge implies understanding and knowing what something means. Faulkner⁴ suggests a three-way distinction between knowing as understanding, knowing as holding information and knowing as holding skills. However, in characterising the personal knowledge and skills of the employees of VSP member organisations, it proved impractical to maintain such distinctions. In some cases, it is very difficult for individuals to identify and distinguish between their knowledge of particular fields. This difficulty was overcome through the use of a tree-like hierar-

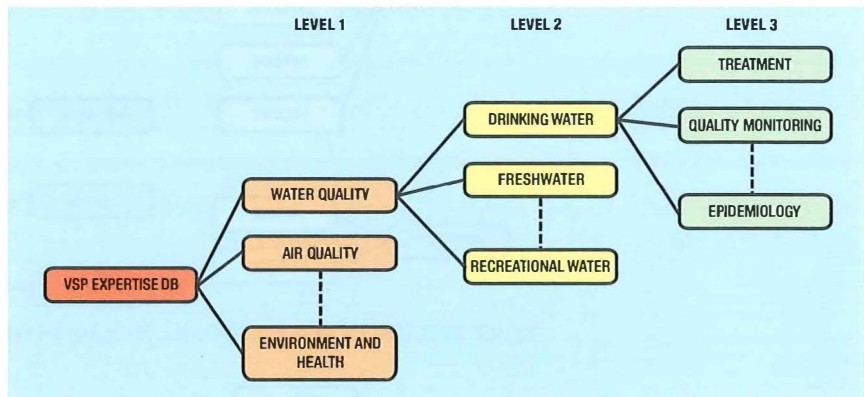


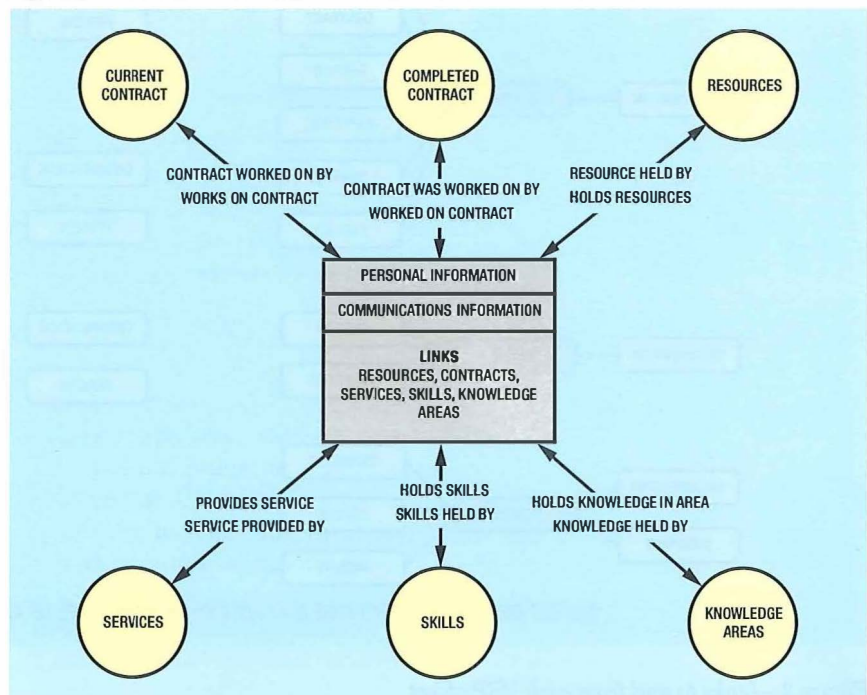
Figure 1 – Hierarchically-structured expertise database

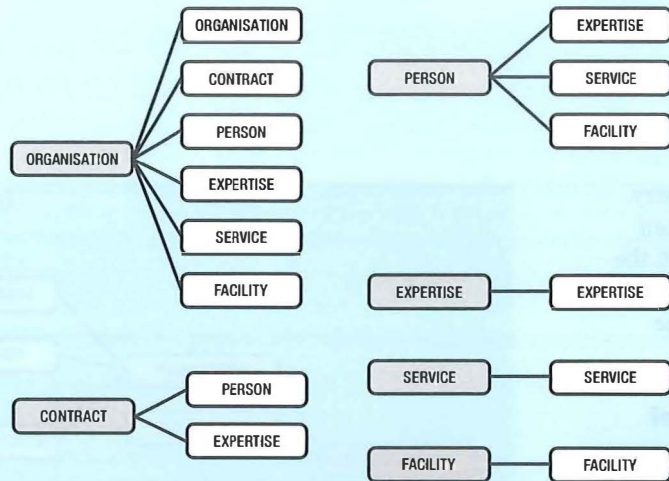
chically-structured expertise database. Figure 1 shows an example from the environmental section of this database where level 1 represents a broad subject field or major area of activity. Each level 1 may be subsequently expanded to the second level, resulting in a set of more narrowly defined areas of activity. Each level 2 expertise may also be expanded, until a leaf object comprising a specific expertise is found. During the realisation of the information provided by the pilot organisations in the VSP, three levels of expertise have proved sufficient. However, if a user has expertise that

does not fit into any level of the database, a new classification may be added.

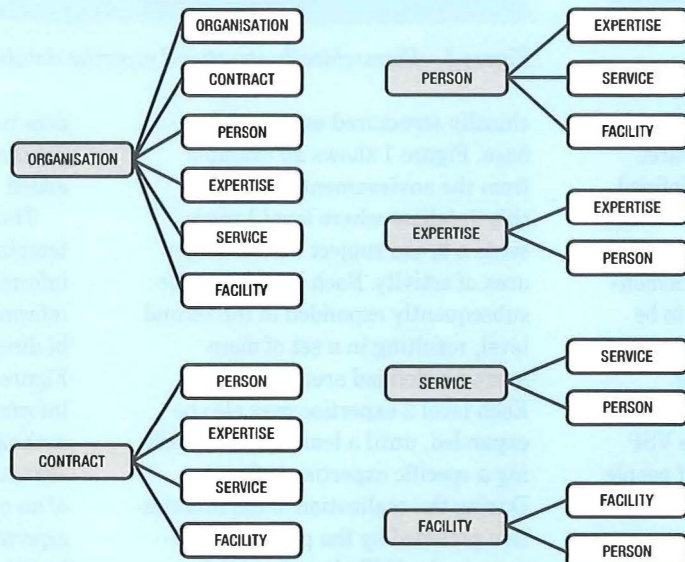
The information model employs templates to store and present information in the VSP database. The information is cross-referenced using bi-directional links as shown in Figure 2. These provide additional information and can be used when making a query. For example, searching for people who are members of an organisation that provides expertise in the area of environmental health is a powerful facility providing more flexibility and expressiveness than simple membership relations.

Figure 2 – Person-centred information model used in the VSP

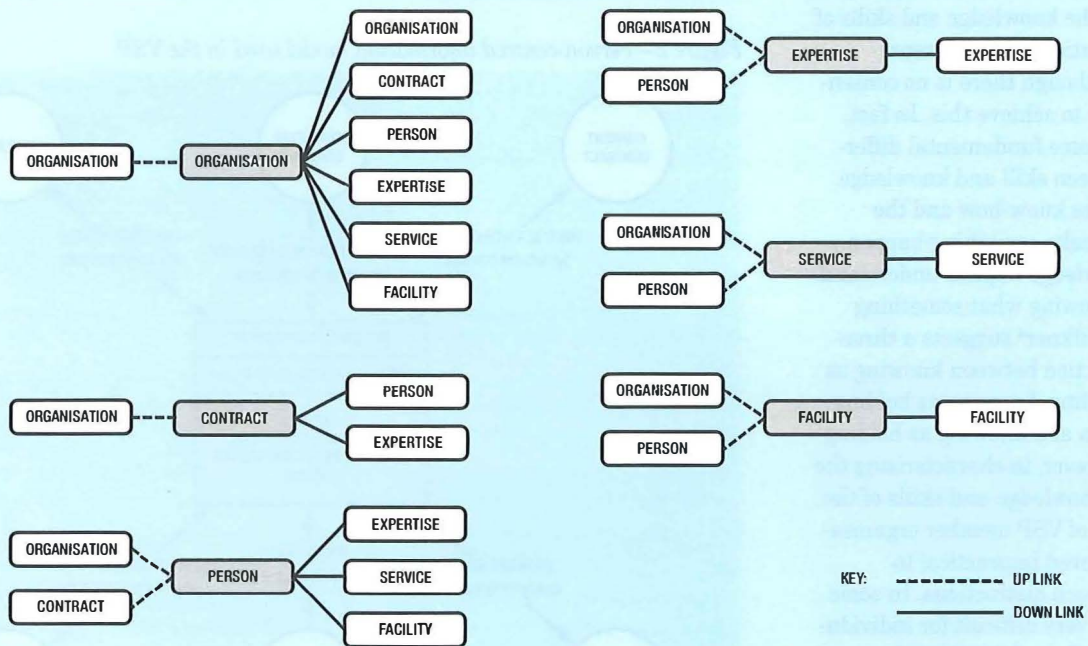




3(a) VSP DIRECTORY ENTRY TYPES AVAILABLE FROM A ONE-LEVEL DOWN SEARCH



3(b) VSP DIRECTORY ENTRY TYPES AVAILABLE FROM AN ALL LEVELS DOWN SEARCH



3(c) VSP DIRECTORY ENTRY TYPES AVAILABLE FROM A ONE LEVEL UP AND DOWN LINKS SEARCH

Figure 3—Links to and from each VSP object

Projects, resources, services, skills and knowledge areas are all VSP groups. The links between a person and the relevant information in other groups are automatically created based on the knowledge of the group type.

The VSP navigation tool allows information to be searched for in three ways:

1. *One level down*: searches will traverse 'down' links from the current object to the depth of one level.
2. *All levels down*: a recursive search from the current object via all children. The links to and from each VSP object can be seen in Figures 3(a), 3(b) and 3(c)
3. *One level up and down*: searches traverse all links to a depth of one level from the current object, whether they be to parents or children

A typical VSP user's entry is shown in Figure 4. The circles represent VSP objects linked to the Person object. This is stored in tabular form. The symbol 'UoL' denotes the 'University of Leeds' which is the enterprise to which the user belongs. The fields 'Name', 'Aims', 'email' and 'tel' are used to specify specific information about the user. The field 'member_of_dept' is linked to another named VSP object which in this case is 'vsp.UoL.Computer_Studies'. Similarly, the VSP object 'works_on_project' is a link to another VSP object called 'vsp.UoL.VSP_Development'. Finally, links to the knowledge areas that 'Richard Drew' holds are included in the entry.

VSP Toolbase

The VSP provides four main tools:

- *editor*, which allows the content and structure of the VSP database to be modified;

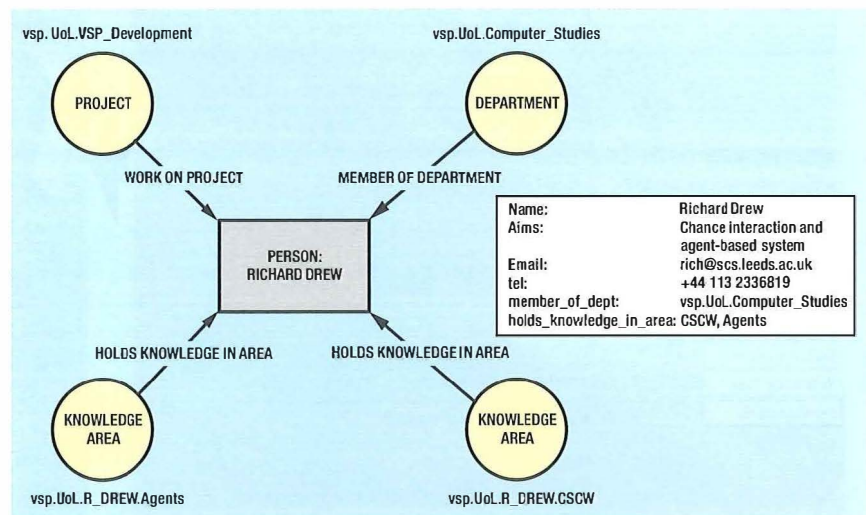


Figure 4—Example of a VSP user's entry

- *navigator*, which provides graphical navigation of the VSP database;
- *browser*, which provides facilities for viewing information stored in the VSP database; and
- *communicator*, which allows users to be contacted using a number of collaborative tools.

The use of these tools is illustrated in Figure 5 and described in more detail in sections that follow. The images are from the prototype VSP. When users first enter the VSP through an organisation, they record their personal information and skill details they are offering. This is entered using the *editor* which is also used to create the structure of the VSP database by allowing relationships between VSP objects to be specified. The *navigator* tool can then be used to search through the full set of VSP objects. Levels in the hierarchy can be expanded to show child entries based on a specific object type (for example, 'show all people') or by relationship type, (for example, 'show all people who manage another object'). Individual entries can be viewed using the *browser*. If, after navigating and browsing, someone is discovered who offers a service or has information of interest, the *communicator* can be used to establish a link.

Navigator

To navigate a large information space such as the VSP, it is necessary to be aware of your current location. In the case of the VSP, the information space is structured using the hierarchy for each of the VSP members. The navigator (Figure 5) allows the organisational structure to be viewed graphically by displaying each group type as a node in a tree that correspond to the VSP objects in the information model (see Figures 1 and 2). The lines represent the links between the objects. For example, a project within the VSP typically has several people involved in it, and is represented as a node with several 'child' nodes (people) linked by 'member of' type links.

The interface to the navigator is divided into four sections:

1. *Object selection buttons* these are colour-coded to represent different object group types.
2. *Navigation window* displays the current of several possible views of the organisational hierarchy.
3. *Function buttons* control display and search options.
4. *Search Facility* allows the user to search for a group or groups in the information hierarchy.

1. Edit personal information using the VSP editor. Include expertise, qualifications, communications, information, etc.

2. View the information in the VSP browser

3. Browse through the information using the VSP navigator. Locate useful expertise and services, and the providers

4. Contact other VSP users via the VSP communicator. Select the people and collaborative tools that you wish to use and hit 'Dial'.

Figure 5—Example of the VSP toolbase used in the Proof of Concept

Browser

In the prototype VSP, a modified World Wide Web (WWW) browser, Mosaic⁵, has been used with the information pages consistently structured within the hypertext system. This contrasts with the present use of the WWW, where there is often no consistent representation for information, no guarantee that links point to anything, or even what the description says it points to. The VSP's information model enables indexes to be automatically created using the object group types, (that is, index of clubs, people, projects). These are accessible from a top-level VSP 'home page', and from each information page allowing the user to return to the indexes as required.

Communicator

The communicator is the tool used to support interaction between members of the VSP. Users indicate their wish to communicate by selecting the relevant groups in the VSP. If a user has made his or her picture available, the communicator then retrieves and displays it. For all available users, the communicator displays the available set of synchronous tools with which they can be contacted. Unavailable users can still be selected, but they can be contacted only by using asynchronous tools, such as e-mail. The synchronous tools currently supported in the VSP are shared text and image editor, white board, audio and video telephone, and a three-dimensional collaborative tool that allows multiple users to interact with three-dimensional models. A later version of the VSP also supports teleconference facilities.

Editor

The editor has two functions. Firstly, for information entry into the VSP database—with each VSP object having an associated template (for example, person template, see Figure 4). This has a number of fields, some of which are designated 'must have', which must be filled in for the entry to be considered valid.

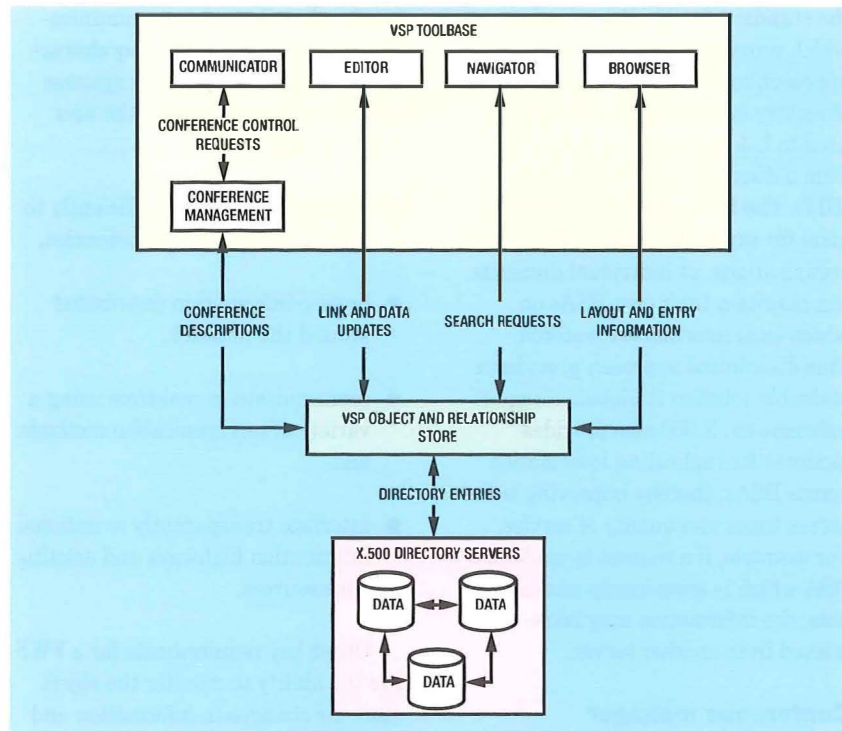


Figure 6—Overview of the organisational structure of the VSP system

Other optional 'may have' fields may also be specified; these provide additional non-essential information for an entry. For example, 'must have' entries for a person include name and a description. 'May have' fields include qualifications. To create a new VSP entry, the user selects the group template: a form with default fields which must be filled in. The second function of the editor is to create the links between VSP objects that:

- provide the structure for services and organisations (that is, the membership links); and
- provide the relationship links, (for example, person *provides* service, person *manages* project).

Implementation

The overall structure of the VSP is shown in Figure 6. At the top level are the interactive desktop tools providing access to the VSP object and relationship store which underpins the VSP system. Below this level

are the distributed directory services provided by X.500 directory servers and the network-independent data-transport layer (which provides transparent access to asynchronous transfer mode (ATM), integrated services digital network (ISDN) and Ethernet services).

VSP object store

The VSP object and relationship store underpins the whole system. It manages the information using the VSP information model and can be viewed as a distributed, persistent store of named objects together with their relationships. The VSP object store is an abstract layer that provides application programmer interfaces (APIs) for interacting with information stored in the VSP. The implementation of the underlying database requires careful thought. The database must be scalable to avoid bottlenecks when performing searches or information retrieval, provide adequate security measures and should use standards for interoperability. Given these requirements, an ideal candidate is X.500,

the standard for directory services which provides a distributed approach to information storage. Directory system agents (DSAs) are used to hold information linked to form a directory information tree (DIT). The DSAs provide a mechanism for supporting distributed organisations, as individual divisions can maintain their own DSAs on which local information is stored. This distributed approach provides a scalable solution to global storage of information. X.500 also provides facilities for replicating information across DSAs, thereby improving both access times and quality of service. For example, if a request is made to a DSA which is temporarily unavailable, the information may be retrieved from another server.

Conference manager

The conference manager is responsible for finding, running and connecting the collaborative applications, such as shared whiteboards, used by the communicator desktop tool. A conference manager is created for each active VSP user (that is, when the user first logs onto the VSP) and the network location, and details of the user are recorded in the VSP object store. When the VSP user wishes to make a call, the conference manager associated with the caller looks up the location and contacts the corresponding conference manager to determine the set of tools available to both users. Any whiteboards or other communication tools required are also enabled and connected. Collaborative tool information is stored in the VSP object store, enabling it to be updated remotely.

Virtual Working Systems

The conceptual framework for the VSP has wider applications since it provides a virtual working system (VWS) which allows people within organisations to interact using person-to-person communication. It is virtual in the sense that the linkage between people within the environ-

ment only exists when communication is taking place. The key characteristics of virtual working systems are that they will provide the user with the ability to:

- navigate quickly and efficiently to identify skills and competencies,
- browse information distributed around the network,
- communicate in real-time using a variety of communication methods, and
- interface transparently to national information highways and existing data sources.

Other key requirements for a VWS are the ability to monitor the object store for changes in information and user state, and to act on events generated by them. For example, when users with whom you wish to communicate log onto a VWS, it would be useful if they were informed in some way. If some new information becomes available in an area of interest to you, then providing some notification is useful by preventing the need constantly to perform searches⁶.

This VSP development is one of the first to be constructed and tried in a real world situation. It is a first step in the creation of a virtual workplace and market for the 21st century that will allow global interaction. We believe that such environments will have a major impact on the world of work and play, and of course telecommunications, on which they rely.

Acknowledgements

The authors would especially like to thank the Vice Chancellor, Professor A. G. Wilson for his incisive leadership and support of the project. Chris Cole, BISS Ltd. has contributed to the conceptual framework of the VSP. Special thanks go to the local DTI office (John Baragwanath and Glenys

Timmons) for supporting the project at a critical stage. In addition, Leeds Environment Centre (Professor David Kay) and the Regional Technology Network (Dick Clark and Kevin Riley) are making important contributions to the VSP as potential users. Andy Walker has supplied project management skills and made major contributions to the user interface design. The project is supported by DTI grant SU/010/901.

References

- 1 LEIGH, C. M. Science Park System for Leeds, Regional Research Observatory, University of Leeds. 1991.
- 2 CURSON, J. M. Information Architecture of the VSP, Internal Report, University of Leeds, June 1994.
- 3 KOGUT, B., and ZANDER, U. Knowledge of the firm, combinative capabilities and the replication of technology. *Organisation Science*, Aug. 1992, 3(3).
- 4 FAULKNER, W. and SENKER, J. Conceptualising knowledge used in innovation: a second look at the science-technology distinction and industrial innovation, Edinburgh PICT Working Paper No.43, Edinburgh: Research Centre for Social Sciences, University of Edinburgh. 1993.
- 5 BERNERS-LEE, T. J., CAILLIAU, R., GROFF, J-F., and POLLERMANN, B., CERN, World Wide Web: The Information Universe, published in *Electronic Networking: Research, Applications and Policy*, 2(1), pp. 52-58, Spring 1992, Meckler Publishing, Westport, CT, USA.
- 6 DREW, R. S. System Architecture for Information Location and Monitoring Using User Interests and Agents, Internal Report, University of Leeds, School of Computer Studies. 1994.

Biographies



Peter M. Dew
University of Leeds

Professor Peter Dew is Head of the Division of Computer Science and Deputy Director of the

Keyworth Institute. His research interests cover interactive solid modelling, scalable systems and algorithms (he has three EPSRC grants in this area) and virtual working systems for applications in manufacturing (in collaboration with the Keyworth Institute) and education (he is Technical Director of the University's Virtual Science Park project). In addition, he directs the Distributed Computing Research Environment (DCRE) project which provides a high performance, collaborative computing environment to key researchers in Leeds. The environment which has substantial external funding is built around a state-of-the-art asynchronous transfer mode (ATM) local area network.



Richard S Drew
University of Leeds

Richard Drew is a postgraduate student with the School of Computer Studies at the University of

Leeds. He graduated in Computer Science at the University of Leeds in 1992. He worked on the Leeds University Virtual Science Park Proof of Concept, and has continued work on further VSP projects, including a WWW gateway to the VSP project in 1993. In 1995, he started an EPSRC sponsored Ph.D. at Leeds University School of Computer Studies in the area of virtual working systems. His work looks at providing an architecture for supporting chance interaction and information location and retrieval using agents.



Christine Leigh
University of Leeds

Christine Leigh is a social scientist with the School of Geography at the University of Leeds. She has

long-standing research interests in the theoretical and empirical analysis of restructuring processes in manufacturing and service industries, including research into the application of new information technologies in industry in support of innovation and technology transfer. Her research into issues of regional economic development, particularly as they relate to traditional industrial regions, has included collaborative work with a network of European academics in Lille, Dortmund, Paris and Liège. Her work is widely published in academic journals and books. She is Director of the Yorkshire and Humberside Regional Research Observatory, a distributed network of 14 research units in universities across the Yorkshire and Humberside region. She is co-director of the Virtual Working Environment Laboratory in the University of Leeds and is directing the detailed planning of the VSP for the university, some of which is being supported by ERDF funding through the DTI. She is Managing Director of the Shadow Board of the VSP Company, which is chaired by the Vice-Chancellor of the university.



Dr David Morris
University of Leeds

David Morris works in the School of Computer Studies at the University of Leeds, where he is

jointly funded by Silicon Graphics UK Limited. He has a Ph.D. in Parallel Computing for 3D Computer Graphics from the University of Leeds, and worked in industry for four years before taking up his current post. He was responsible for the technical vision and development behind the Virtual Science Park Proof of Concept demonstration systems, and is currently working in the VSP technical development team. His current research interests include distributed directory systems on high-speed wide area networks, intelligent agents for Internet browsing, searching and communication, and innovative user interfaces for Internet browsers.



Jayne Curson
University of Leeds

Jayne Curson is a postgraduate student with the School of Geography at the University of

Leeds. She graduated in Business and Economic Studies at the University of Leeds in 1989. She then gained an M.A. in Economics in 1991, also at Leeds. In 1993, she started a Ph.D. in the School of Geography at Leeds investigating the use of information and communication technologies for the inward transfer of scientific and technical information. This involves looking at the organisational impact on advanced technology firms and the implications for regional economic development. She has also worked on the VSP project, working on the design and population of the VSP information architecture.

Alan Jackson

Transport Network Architecture

The ever increasing functional complexity of the world telecommunications transport network is becoming seriously difficult to understand and manage within the current equipment-oriented perspective. A new function-oriented model is needed to bring clarity to the organisation and deployment of transport functions in the network. This article describes such a model, as set out in the new ITU Recommendation G.805, and its value to network designers and future operations support systems.

Introduction

The first article in this series on Architecture, 'Building on the Architecture Framework'¹, described the various kinds of architecture that can be found in the telecommunications business and showed how they complement each other. The broadest architecture is a business model of the fundamental processes which a company undertakes. A *logical* architecture describes in more detail the implementation-independent abstract functions that need to be carried out; to system designers, these are the input 'requirements'. Then there are *physical* and *data* architectures which show what hardware and software are fitted together to implement the required functions. This article is largely about the logical domain, and its relationship with the physical domain.

The second article in this series, 'An Architectural Framework for Networks'², described network architecture in general. It showed how network functions can be categorised into *transport*, *intelligence*, *applications* and *management* domains. In this context, transport is the function of moving information from place to place, including the actual switching of information; intelligence is the function of deciding where and when information should be switched; applications are the users of the transport platform; and management is the function of maintaining the health of the network and organising changes to its size and structure. These four domains interact closely with each other, but each has its own distinct architecture. This article falls firmly within the transport domain.

Transport network architecture is usually described in physical terms, typically with boxes representing switches, connected together by lines representing transmission, and the whole sprinkled with interworking units. This is fine for describing what equipment makes up a network, but poor at describing what the equipment does. As the functions of networks get ever more complex, and as more functions get packed into ever smaller equipment, there is an increasing need for a generic architecture in which transport functions can be modelled, analysed and discussed. There have been several specialised functional architectures before now, but none has proved suitable for generic application to all kinds of transport. This article describes some of the problems now looking for an architectural solution and a new internationally standardised generic model for transport network functions that goes a long way toward providing a solution.

The Problems

Network architecture until now has focused heavily on the physical disposition of equipment and their interconnections, with insufficient attention paid to the functions going on inside the many boxes. This is the understandable result of the past largely one-function-per-box approach. Equipment 'boxes' identify the main functions, but hide most of the auxiliary functions or relegate them to interworking units. Transmission 'lines' typically name their bit rate, but hide virtually all the complexity of the protocol stack they carry. A poor view of the functions leads to poor functional organisation,

with duplication and unnecessary complication, which leads in turn to more and bigger boxes, with their expensive overhead of supporting metalwork, cabling and building. The physical approach to design tends to start from the presumed hardware solution, and the functional organisation thus becomes a consequence of the design process, rather than the initial requirement. What is needed is more understanding and analysis of the functional requirements in the first place, to enable more functional integration within less and smaller equipment, and this needs in turn a new 'language' through which telecommunications engineers can discuss transport network functions clearly and unambiguously.

International standards for switching and transmission have, in the past, naturally echoed the physical approach, which made it difficult to purchase more functionally integrated equipment, and still does for 'older' technology. Most newer standards are now more functionally oriented, but the familiarity of the physical, coupled with new technologies usually being stimulated by new physical components, still makes functional standards hard to write.

Network management becomes ever more difficult as networks and the services on them become ever more complex. Each new technology, new manufacturer, new network, seems to require yet another management system. Many of these systems cannot communicate with each other, even though they increasingly need to, or only with difficulty and expensively. There is not even agreement on how to represent the date or a customer's name and address. They are very equipment-oriented, which is fine for locating faulty equipment, but poor for understanding equipment-independent network concepts such as circuits and paths, especially across multi-operator, multi-manufacturer, multi-platform networks. A common functional language for transport networks of all flavours is

desperately needed that can lead to generic transport management processes and software that can be reused and can interwork with the minimum of effort.

In the past, most countries had only one network operator, with a handover at national boundaries. Having handed over a call or private circuit, an operator's responsibility would then end. The concept of managing a worldwide end-to-end circuit had not been born. Nowadays, there are multiple networks within countries, and circuits and paths can even start in one network, leave it, and then re-enter it elsewhere. Customers expect the whole network to work just as well as if a single operator was in charge of all of it. This means that operators now need a means by which they can simultaneously model their own networks in detail and other operators' networks abstractly. Such a model would describe interconnect in terms of network functions offered, independently of particular equipment deployed in particular places to suit the current technology at the time.

There has been a tendency to manage whole groups of equipment (switches, multiplexers, line systems) according to their 'customers'; for example, the plain ordinary telephone service, or the narrowband digital leased line network. In consequence, bulk bandwidth transport is operated by several different management systems (and thus several management systems have to know how to manage it) and cannot easily be reassigned according to changes in demand. It would be more economical if such common network features could be managed as a single entity in common for all their customers, with service level agreements covering performance, resilience, etc.

The New Transport Functional Architecture

Between 1988 and 1992, experts from leading operators and manufacturers around the world worked intensively

within CCITT Study Group XVIII (now ITU-T SG 13) to develop a functional architecture addressing all these concerns. Although it was initially required for the synchronous digital hierarchy (SDH), it was intended to be generic and applicable to other network technologies, and so far it has proved capable of describing many kinds of transport network. The original Recommendation, containing both generic and SDH-specific material, is ITU-T Recommendation G.803³, although this is now being revised to remove the generic material which has now been put in Recommendation G.805⁴, approved by the ITU in November 1995. SG 13 has also recently approved a version specific to asynchronous transfer mode (ATM) and others will follow.

The model provides two compatible views. A top-down view describes networks from a whole-network perspective and can be used to hide details when they are not relevant or not available. In this view, two techniques are used to divide the complexity of the world telecommunications network into more manageable portions, called *layering* and *partitioning*, which are described below. A bottom-up view describes the elementary functions which can be used to model the transport functions of actual equipment, described under 'Architectural Components'. Some of these components extend up into the top-down view, thus linking them together.

Although the top-down view will not seem particularly strange, the bottom-up view almost certainly will. However, it should be remembered that a major objective of this model is to solve network management problems, which these days are too complex to be done by humans alone, so the primary customer of the model is the computer. Computers work best with large numbers of very simple objects, thus the model has few components, but many replications of them. Although network designers will need to learn this new language, its small number of components means that it is a lot simpler than

Figure 1—Dividing a network into layers


learning a standard computer language.

Layering

Layering means the horizontal division of the world telecommunications network into layers, each layer being responsible for a distinct kind of *characteristic information*, namely information of a certain characteristic rate and/or format. Familiar examples of layers include 64 kbit/s; frame relay; 2, 34 and 140 Mbit/s plesiochronous digital hierarchy (PDH) paths; the various SDH virtual containers; ATM virtual channels and virtual paths; various physical section layers including PDH intra-office sections and the SDH STM-N section. Although all these layers have been standardised, it is also possible to have proprietary layers—for example, the PDH line system sections and the passive optical network section—since even though the networking is confined to one manufacturer, the equipment implementing them still produces management information.

Layers are supported one on another, with higher layers supported by lower layers. Figure 1 shows a selection of layers in groups and an example of one layer supporting another. Higher layers are closer to users' applications and lower layers are closer to the physical media. Where one layer supports another, the higher is called the *client layer* and the lower the *server layer*. While there are many layers, only a small number will be present at any one point in the network. The process of getting information from one layer to another is called *adaptation* and is described under 'Architectural Components'.

Most client layers can be supported on a choice of server layers and most server layers can support a number of different client layers. It is even possible in some cases, though usually undesirable, for a layer to be used more than once in the same stack; for example, for a (high bandwidth) ATM virtual channel to support an SDH virtual container



which supports another (lower bandwidth) ATM virtual channel. The transparent support of one layer by another is often called *encapsulation* in the data network community.

The layers closer to the users' applications are collectively called *circuit layers*; circuits are carried, usually in multiplex, on path layers; paths are carried on section layers, usually medium-specific. (Note that *higher-order* multiplex signals are in *lower layers*.) Although this distinction may appear fairly firm in the ITU-T Recommendations, it is arbitrary so far as the model itself is concerned because the same rules apply to all layers. For example, circuits are normally thought of as being *switched* and paths as being *cross-connected*, but there is no difference in basic functional terms (the origin of the switching commands and the response time of switches is not directly modelled). Optical signals ostensibly belong in

section layers because they are medium-specific, but multiple-wavelength multiplexing and optical switching look more like path layer functions.

There is a relationship between transport layers and the layers of the Open Systems Interconnection (OSI) seven-layer protocol stack: each transport layer is effectively an OSI stack in itself, with the OSI application layer representing the transport client layer above and the OSI link layer representing the transport server layer below. The OSI physical layer is only seen in the appropriate transport section layer.

Figure 2 illustrates how a particular multi-layer network can be simply represented and how various kinds of services can be shown to be supported on the various layer networks. The same representation can also be used to illustrate the stack of layers at a particular mid-network interface.

Figure 2—A possible layer stack for a future network

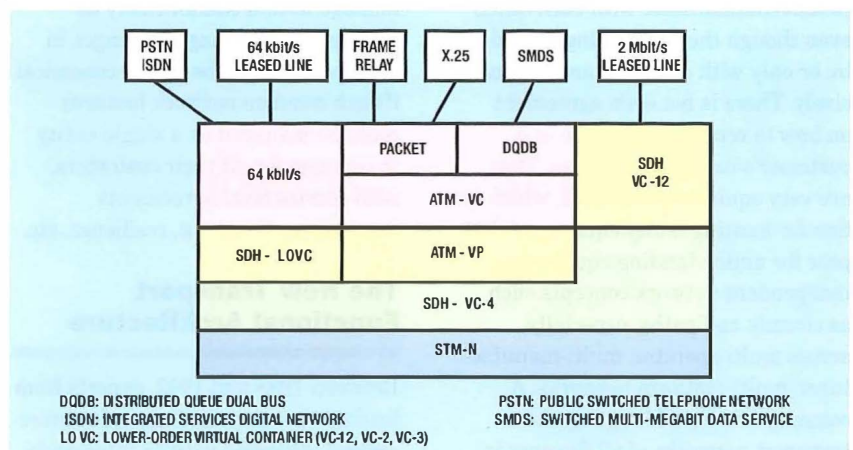



Figure 3—Partitioning within a single network layer

These are the layers of most interest to network operators, but above these could be considered higher layers; for example, an ideas layer, carried by either a words layer or a pictures layer. The words layer can be carried by a speech layer or a writing layer. The speech layer can be carried on a general analogue waveform layer which can be carried by various circuit layers. Below the section layers, further layers can also be considered such as copper or fibre layers, a cable layer, a duct layer and a road layer. To show how powerful the model is, it is equally possible to model an alternative set of layers between words and roads, via a writing layer, a paper layer, an envelope layer, a sack layer, and a van layer. Also note that while the generation and interpretation of signalling and management messages is not modelled, the transport of these messages between equipment can certainly be modelled.

The criterion for whether or not to define a particular layer network is whether it is useful; for example, could it be networked, does it generate any management information? It is possible to expand a single layer into two or more separate layers when new capabilities need to be modelled, or to combine two or more layers back into a single layer when a theoretical capability is never used in practice.

Any one item of equipment will participate in a number of network layers. For example, a public switched telephone network (PSTN) switch whose major transport function lies in the 64 kbit/s layer also participates in the PDH 2 Mbit/s frame layer (G.704) and various physical layers. A few kinds of equipment may major in more than one layer; for example, an ATM switch may well switch both virtual channels (VCs) and virtual paths (VPs).

Although each layer has its own distinct characteristic information (bandwidth, bit rate, format), it will be shown that each layer can be functionally modelled using the same basic components, with layer differences being described more by



flavours of data than by new suites of software. Thus, to a large extent, the task of managing all the layers can be reduced to a simple recursion of managing a single layer.

Partitioning

Each layer network is composed of many fixed point-to-point links which, in most layers, can be connected together (and disconnected) at points of flexibility such as switches, concentrators, cross-connects and distribution frames. Each layer network can thus be modelled by *subnetworks* representing these flexibility points, with *links* between the subnetworks and *access links* between the subnetworks and the logical layer boundary (that is, where client layer signals are adapted into and out of the layer). This is a very detailed (or low level) partition of a world layer network.

In order to hide some of the detail, a group of low-level subnetworks and the links between them can be represented as a single higher level of subnetwork (in the same layer). The low-level subnetworks which have links to the outside world sit topologically round the inside edge of the higher-level subnetwork and their links to the outside become simultaneously links from the higher subnetwork; any low-level transit

subnetworks without links to the outside sit topologically in the middle of the higher-level subnetwork.

A group of higher-level subnetworks and the links between them can then be represented as a single even higher-level subnetwork which hides even more detail, and so on, recursively, until they have all been combined into one world subnetwork with access links out to the layer boundary. This recursion is illustrated in Figure 3.

A fairly high level of subnetwork could represent an entire country. An intermediate level of subnetwork could represent an administrative domain; for example, that part of a network belonging to a particular operator. A fairly low level of subnetwork could represent, for example, a local exchange with all its associated concentrators. At the bottom level of partition where subnetworks correspond to actual physical switches, cross-connects, etc., there is no point dividing these further because they are inside equipment. They are called *matrices*, but they are still modelled as subnetworks. Using this model, intelligence and management entities can simultaneously view their own subnetworks in detail, nearby networks in less detail and distant networks as amorphous blobs.

Figure 4—The components of the G.805 model

Architectural Components

There are 12 architectural components in the G.803/G.805 model, in four groups, as illustrated in Figure 4. The topological components have already been described. The transport processing functions and transport entities are defined by the functions they perform on information. They are, so to speak, the basic building bricks that actually ‘do’ something and are of most use when describing what equipment does. The reference points identify where the individual building blocks connect together.

A set of all these components exists in each network layer. Figure 5 shows a generic model fragment containing each of the components, using standard drawing conventions. Although most subnetworks and links contain many connections, there is only room to show one here.

The adaptation function adapts the characteristic information of a client layer network into a form suitable for transport in the layer network of interest, and back again at the far end. The two halves are called the *adaptation source function* (client to server) and *adaptation sink function* (server to client). Examples of adaptation processes, which may occur singly or in combination, include coding, modulating, rate changing, aligning, justifying, mapping, encapsulating and multiplexing. The results of most adaptations are the payloads of frame and packet structures, though adaptations onto media result in volts and optical powers and suchlike. The output of the source function and the input to the sink function are access points to the layer network. An access point is only a unique name for that individual point in the total network; it doesn’t ‘do’ anything. An association between these two access points, across the layer network, is called a *trail*.

Trail is a generic word for circuit, path or section and was chosen as one of the few ‘pathlike’ words not already meaning something else. The trail is

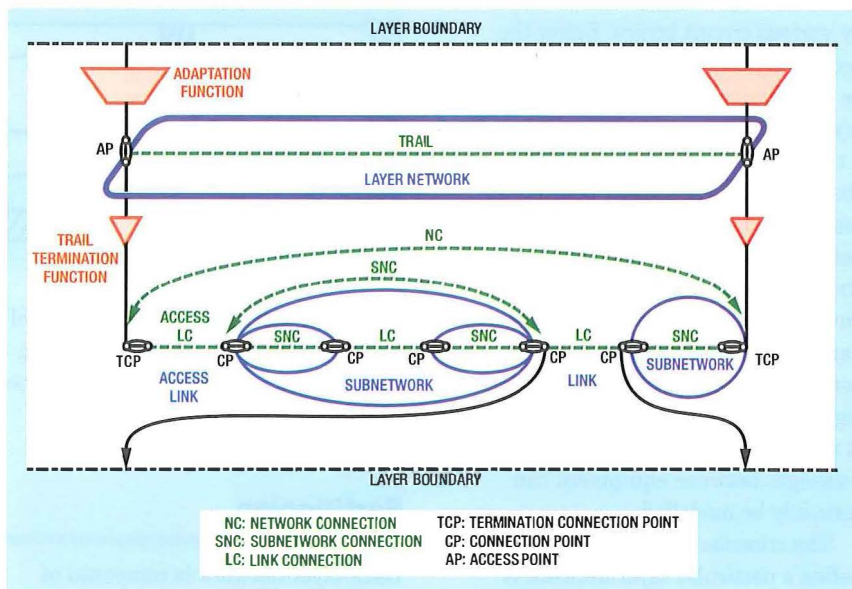
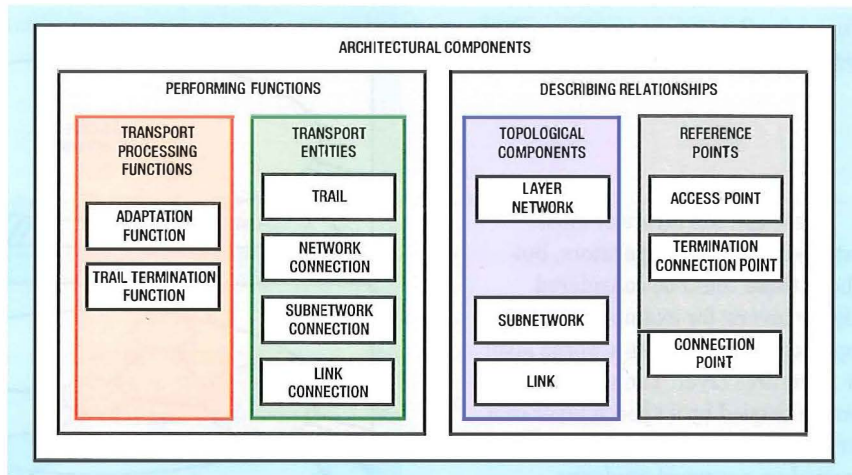


Figure 5—An illustrative fragment of a network using G.805 components

an indivisible, end-to-end transport entity. In most layers, a trail represents a fixed bandwidth between the two ends, but in packet and cell layers, when nothing is being sent, there is only a logical association. Depending on layer, trails may be set up for minutes (as in calls), months (as in leased lines), or years (as in inter-equipment sections).

The trail termination function measures and/or maintains the integrity of information transfer on a trail across a network. Typically, additional known information is added to the adapted information at the start of the trail and it is monitored at the far end of the trail; where degradation is likely, it can provide correction; where degradation is rare, it provides detection. Trail termination is also split into source (generation end) and sink (monitoring end). Even when present, the function cannot detect every error or degrada-

tion. Trail termination must be carried in some kind of ‘overhead’ channel; for example, a digital frame structure, redundant line code, carrier modulation. Some layers, by choice or mistake, have no overhead channel, so trail termination is not possible, but for consistency it is modelled as a null function. The input to the source function and the output from the sink function are the same *access points* mentioned above. The output of the source and the input to the sink are termination connection points. These are likewise simply unique names for individual points. An association between these two termination connection points, across the network is called a *network connection*.

The network connection is also an end-to-end transport entity, but, unlike a trail, it is divisible. As a network can be partitioned into smaller subnetworks with links

between them, so the network connection across it can be broken into corresponding subnetwork connections with link connections between them. Thus an end-to-end network connection consists of a series of alternating link connections and subnetwork connections; the subnetwork connections 'connect' the link connections together into a long chain. At each join is a connection point.

Finally, since at the equipment level, the subnetwork connections are inside switches, it is only the link connections between switches that need to be adapted onto server layers and finally onto cables. Thus the input to an adaptation source function and the output of an adaptation sink function is at a connection point (or a termination connection point in the case of an access link).

Any one layer carries many different kinds of client information; for example, the 64 kbit/s layer carries A-law encoded speech, μ -law encoded speech, and data of many kinds. However, the model does not show such 'flavours' within a layer. If a trail were set up between different kinds of adaptation, then nonsense would result in the client layer(s); the transport network relies on the network intelligence to prevent it. Interworking between different kinds of signal in one layer or between one layer and another is shown as de-adaptation of one signal up to some common client layer (which may need to be created for the purpose) and a different re-adaptation back down to the other.

The Relationship of the Functional to the Physical

The adaptation and termination functions must be implemented inside equipment and so must the lowest level of subnetwork, the matrix. The lowest levels of subnetwork connection are also therefore inside the equipment. Getting from a subnetwork in one item of equipment to a subnetwork in

another requires a link connection, the ends of which must be inside the equipment. Thus all the reference points are also inside the equipment.

Any item of equipment has a number of physical interfaces and some 'essential' functions—for example, a switch or multiplex—in one or more higher layers. It is therefore modelled as a stack of adaptations and terminations from the physical medium layer of each interface port up to the 'essential' layers where they are switched, multiplexed, etc., by matrix subnetworks and adaptations.

All the layers of the model are functional and all are implemented in some physical way. Physical fibres, for example, are implementations of the function of guiding optical signals. The higher layers are specified in abstract, media-independent terms, but they are physically implemented, by pushing electrons backwards and forwards in copper or launching photons into fibres; the electrons and photons are indisputably physical, just much faster and lighter than the so-called physical layers which are simply the layers that are big enough to trip over.

Nothing in the model implies anything about how these functions are implemented inside an item of equipment, so, for example, it is a mistake to assume that layer adaptations and terminations are necessarily done one after the other, in a series of modules. In the PDH, this is the only practical way of extracting 2 Mbit/s frames from a 140 Mbit/s frame, because of the way PDH justification works, but in the SDH, implementation short cuts are possible; there may well not exist physically separated streams of VC-3 bytes, VC-2 bytes, etc.

How signals are carried inside equipment is also entirely the implementer's business. When an equipment de-adapts from an incoming physical interface, the result is a digital frame structure or an analogue signal. Such frames or signals may be carried in whatever way the

manufacturer chooses, be it electrical, optical, acoustical.

The functional model can lead to unexpected effects. In the model for ATM, there is a virtual channel (VC) layer and a virtual path (VP) layer. Neither the VC nor VP layer contains complete cells. The VC-to-VP adaptation consists almost entirely of multiplexing and adding the VC identifier bits to distinguish one VC from another. The adaptation of VPs onto a server layer (SDH or PDH, for example) consists of multiplexing VPs together, adding the VP identifier bits to identify one VP from another, and adding the remainder of the cell header (including the HEADER ERROR CHECK field). Thus only in the server layers are complete cells seen. It is only cell payloads that travel transparently across an ATM network.

How to Model a Specific Technology

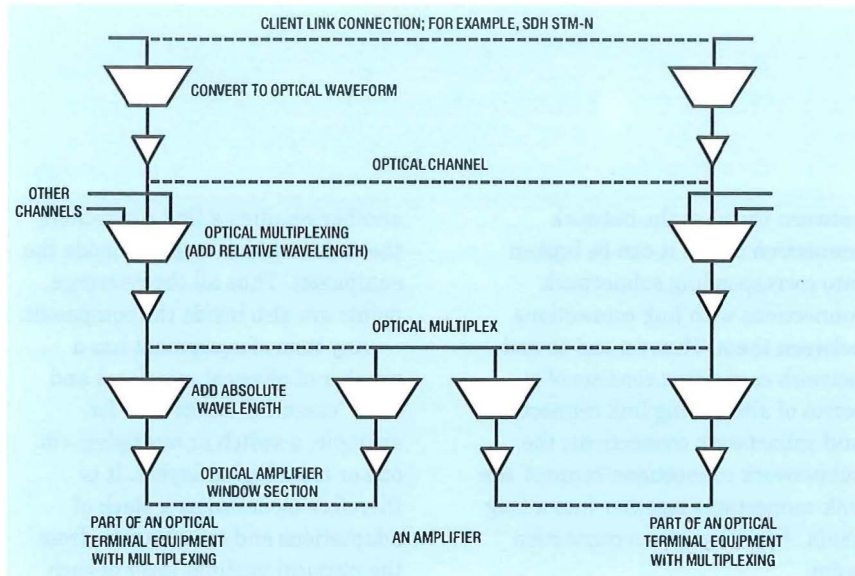
The G.803 model is SDH-specific; the G.805 model contains the generic principles, which can be applied to other technologies. When setting out to model another technology, the first essential is to ignore the actual equipment. This is so important as to be worth repeating: the first essential is to ignore the actual equipment. It is vital to think in terms of the underlying fundamental functions; and in the case of a new technology, also what they might do in future, not what they are limited to on day one. The second essential is to identify how many different kinds of characteristic information will be transported around the network—what 'networked' entities will there be. The third essential is to work out the minimum characteristic necessary to describe each layer; the layer is usually more general-purpose than the initial application being considered. After this, a common difficulty is deciding whether a function is adaptation or termination—adaptation moves information into the layer from a client layer and is thus client-specific; termination measures or

Figure 6—Modelling multiple wavelength optical transmission

actively ensures integrity of layer information, independently of client.

The example of multiple-wavelength optical transmission is described below and illustrated in Figure 6, to show how to approach new technologies. This particular model is currently under debate in standards fora, since it is a genuinely open question precisely how many layers are needed in a specific model and not always obvious what goes in them.

In multiple-wavelength optical transmission, a layer is needed to represent the signal carried by each wavelength; the name *optical channel* is gaining acceptance. Such a signal will sooner or later be 'networked', even if, today, multiple wavelengths are only used within self-contained transmission systems. The underlying function of the layer is to transport any client signal within the limits of the channel bandwidth and signal-to-noise ratio. Its characteristic information is thus any optical signal within the information capacity of the channel. The wavelength assigned to an optical channel is **not** the business of this layer. The layer below is currently named an *optical multiplex* layer, in which a number of optical channels have been multiplexed together. Even this bundle of channels might at some point in the future be 'networked'. The characteristic information of this layer is thus all the optical signal(s) which can fit within the greater information capacity. The channels within it still do not need an absolute wavelength, but they do need wavelength relative to each other. A third layer, below, is currently called an *optical amplifier window* layer. This is admittedly an equipment-oriented name, but the characteristic information is the total optical signal which can be transmitted along a fibre equipped with optical amplifiers which only amplify within a particular optical wavelength 'window'. In this layer, the absolute wavelength is assigned to the client layer signal and the signal applied to the optical fibre.



These three layers are telescoped into a single layer for existing single-channel systems.

Amplification and dispersion accommodation are both termination processes since they maintain signal integrity, independently of the signal content (that is, client layer). Timing recovery is part of the adaptation between an optical channel and its digital client, since the optical signal (even a perfect one) must still be measured against a threshold at the right instants of time in order to recover the client's digital sequence.


Conclusions

The transport functional architecture described above provides a language and framework for analysing functions more precisely than is usually done, currently, thus making it more likely that functional duplication and complication will be spotted before implementation. It is already beginning to be used to guide new international standards. It is both simple and powerful enough to be a foundation for new management system software. It provides a means for describing services offered by one network to another, independently of physical implementation. It also forms a foundation for more enterprise-oriented operation of multi-layer networks. In short, from a quiet beginning it is destined to become a major factor in the future design and operation of telecommunication networks.

References

- 1 STEWART, H. Building on the Architecture Framework. *Br. Telecommun. Eng.*, Oct. 1994, **13**, p. 186.
- 2 WRIGHT, T. C. An Architectural Framework for Networks. *Br. Telecommun. Eng.*, July 1995, **14**, p. 141.
- 3 ITU-T Recommendation G.803, Mar. 1993.
- 4 ITU-T Recommendation G.805, Nov. 1995. Until copies are available, the text can be found in Report COM 13-R 41-E with amendments in Report COM 13-R 47-E.

Biography



Alan Jackson
BT Networks and Systems

Alan Jackson graduated from Birmingham University with a degree in

electronic engineering in 1971 and joined BT to work on procurement specifications and prototype testing of FDM line systems and multiplexers, and later on, analogue and optical field test equipment. More recently he has worked on transmission network element management interfaces. He now provides guidance on architecture frameworks to BT's platform evolution planners and supports the generation of commercial direction for BT's delegates to network standards fora.

Usability Engineering of Speech-Based Services

Speech technology offers intuitive information services accessible from any telephone worldwide. However, only user-centred design techniques can deliver interfaces that live up to this promise. BT has compiled a style guide that includes guidelines and methods based on its experience of speech interface design. These have been applied in user trials and field evaluations throughout the development of Call Minder—a network-based telephone call answering service. Improvements were made to the usability of the Call Minder service as a result.

Introduction

The perceived quality of speech services offered by a telecommunications company, and hence the perceived quality of that company is, to a large extent, derived from customers' experiences when using or encountering network-based services¹. The design of such services has a major impact on ease of use and so customers were involved to steer the design process towards optimal usability². This approach was used throughout the design of Call Minder—a network-based telephone call answering service—to demonstrate the effectiveness of the application of usability design and evaluation principles.

Multidisciplinary teams of software, systems, marketing and human factors people are needed to design successful interactive services³. The human factors contribution

includes requirements capture, design, prototyping and evaluation supported by research into user psychology. BT has compiled checklists, guidelines and methods to support the design of usable services⁴. This provides designers and engineers with guidance in the following areas:

- **requirements capture** checklists to help specify user requirements;
- **dialogue structure and message composition** design guidelines for the format and content of speech messages;
- **keypad guidelines** for mapping functions to the telephone keypad (see Panel 1); and
- **evaluation guidelines** for evaluating services at all stages of development.

Panel 1—The DTMF Interface

(Extract from BT's Voice Applications Style Guide)

The following rules are provided to assist in the assignment of keys to functions and tasks that typify a DTMF interface:

- Whenever possible, actions should be available from single rather than multiple key presses, as some callers may be unable to complete the key presses in the allotted time span and the system may interpret them as key ahead.
- Menu options should be numbered to correspond with the numbers on the DTMF keypad rather than named with mnemonics.
- Key assignments should be consistent across all applications.
- Menu options offered by order of frequency, i.e. the most commonly used first, or logical order where this exists, e.g. record a message before sending it.

Figure 1 demonstrates how information from different sources is utilised in the design of speech services.

Designing Usability into Call Minder

The Call Minder answering service generates voice prompts and interprets spoken or keyed responses. There are two main dialogues:

- **Caller Dialogue** When someone calls and no-one answers or the phone is engaged, Call Minder answers the call and prompts the caller to leave a message. Panel 2 shows a typical interaction with Call Minder.
- **Customer Dialogue** Call Minder customers can listen to their messages from their own telephone, or, after keying in a personal identity number (PIN), from any other telephone. The service informs customers how many messages have been recorded and allows them to be retrieved. During the interaction, customers can also change various options including the greeting message played to callers, the number of rings before Call Minder answers the telephone and the customer's PIN.

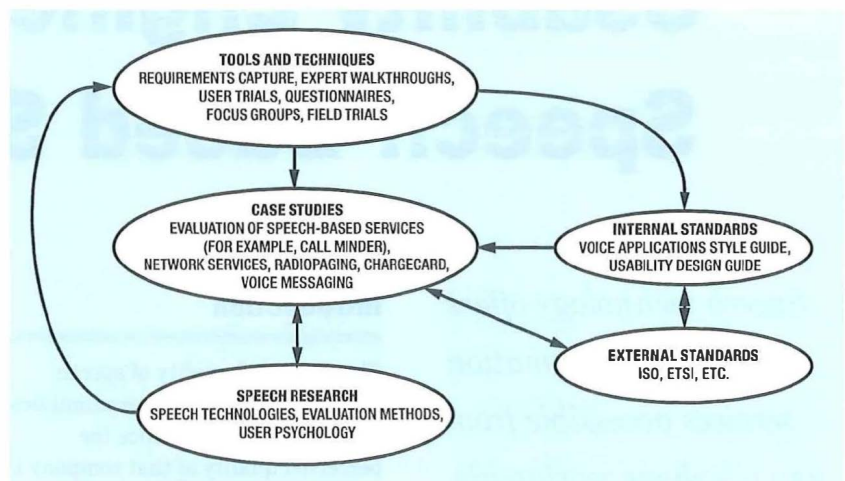


Figure 1—Optimising usability in developing speech-based services

Call Minder's success depends on dialogues that customers can use effortlessly. BT Human Factors involvement in the development of the service can be seen in Figure 2. The tools and techniques used are discussed below.

Service concept and design

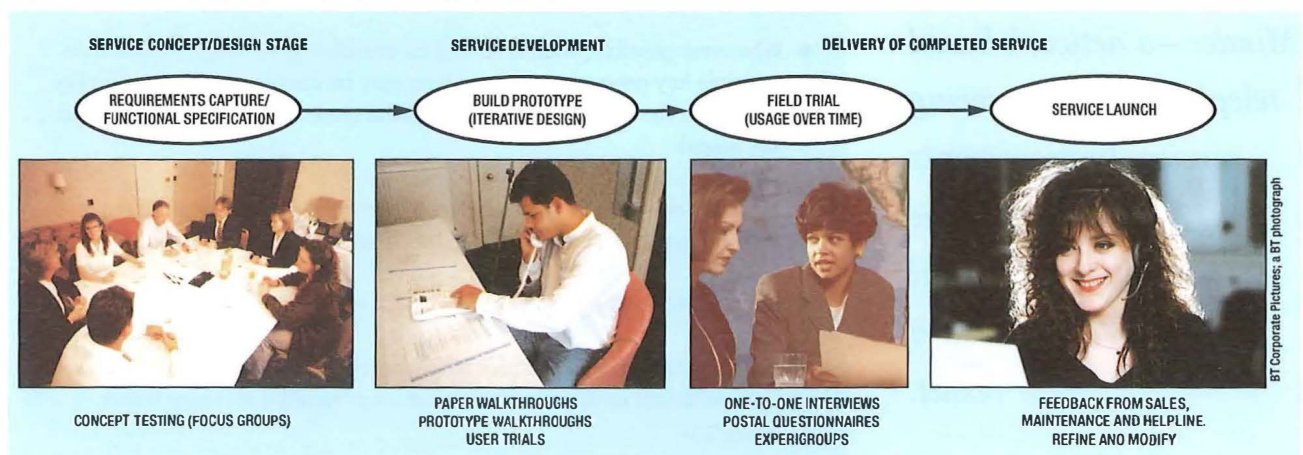
As soon as market research confirmed the basic service concept, preliminary dialogue designs were constructed following principles such as those below:

- Provide adequate feedback so that users feel in control of the interaction, know where they are, what they can do next, and how to correct errors.

- Minimise complexity so that users' memory is not strained. They should not have to remember information from one part of the interaction to the next.
- Speak the users' language and provide clear and easy-to-understand prompts.

The specification of the service was improved by conducting a series of paper walk-throughs. These provide early visibility of the service, but there are important aspects that cannot be assessed using this technique. They do not fully address the impact of essential components such as time-outs, concatenation, and intonation of the recorded messages.

Figure 2—Stages in the usability engineering of Call Minder



This is where a working prototype provides an advantage.

Service development

A dialogue prototyping tool was developed to build rapid prototypes and simulate the service. This allowed objective and subjective usability data to be gathered in laboratory trials with users.

The results from these trials allowed many improvements to be made. Early trials showed that voice prompts did not always elicit the expected responses. For example the prompt:

'Would you like to leave a message?'

elicited responses such as,

'um...yes I would' or *'Thank you, yes'* or *'Hello Gerry,*

rather than with the expected 'yes' or 'no'. As a result, the caller dialogue was radically changed by eliminating 'YES/NO' questions and introducing open-ended responses.

Another benefit from carrying out these trials was that the timing data collected allowed appropriate time-out durations to be specified. Early prototypes had time-outs that were so short that hesitations triggered the next stage in the dialogue. Not surprisingly, users found this very frustrating!

When all the improvements identified in the laboratory had been incorporated, extensive field evaluation was undertaken to test Call Minder in a real-life environment to investigate usage over time.

Field trials

The aims of the field trials were to refine and establish:

- the robustness/effectiveness of the technology,
- the processes needed to deliver the service to market, and
- data on usage, customer attitude and usability.

Panel 2—Example Call Minder dialogue

Good afternoon. There is no reply from that number at present, but I can take a message for you. Can I have your name please?

'Charanjit Sidhu.'

Who are you calling?

'Gerry Coyle.'


Please leave your message and telephone number after the tone.

'Hello Gerry, just demonstrating the service.'

I recorded your message as: 'Hello Gerry, just demonstrating the service.'

You now have a choice. If you are happy with your message, please hang up now. If you would like to replace it with a new message, please hold on. Your new message will overwrite the old one. Please leave your new message after the tone.

'Hi Gerry, I'm demonstrating the overwrite facility.'



The field evaluation involved over 200 customers for one year. Usability was investigated by means of questionnaires and interviews, and objective data from the service platform. The results indicated that customers were very satisfied with the Call Minder service. However, the data showed that there was scope for improvement in the following areas:

- **Recognition performance/speed of operation** The high level of background noise in some environments reduced recognition performance leading to time-consuming error-correction dialogues. As a result, some customers found the service too slow. Fast-track dialogues and keying options were thus introduced.
- **Attitude towards caller dialogue** Customers reported that many callers were hanging up on encountering the service owing to the unconventional format† of

the caller dialogue. They were still attempting to leave their name and message in response to the request for the name. The need to investigate caller attitude towards the service became imperative, as it represented a key element of the service. As a result, customers were provided with an option to remove the question which asks for the caller's name.

Subjective and Objective and Benchmarking

Both subjective and objective measures were taken throughout the course of the trials which enabled different techniques to be benchmarked.

Objective measures were obtained from the service platform and by observing users. These provided data on performance and usage, such as

† The caller dialogue requests the name of the caller and the message separately.

the number of times and at what point in the dialogue users hang up, recognition performance and responses to voice prompts. They allowed problem areas within the dialogue to be identified; for example, confusing voice prompts which cause users to hang up could be pinpointed.

Subjective assessment through interviews and questionnaires allowed customer satisfaction to be monitored. Questionnaires measured key dimensions relevant to speech services including speed of use, level of concentration needed, reaction to the voice and clarity of messages. The results from questionnaires were represented on attitude profiles (see Figure 3), which allowed easy comparison between different versions of a service. The questionnaire was validated and refined during the development of Call Minder and is now part of the BT Human Factors tool-kit for designing future speech services.

Conclusions

Increasingly, usability is becoming the key differentiator in steering the


design process towards producing services that customers will find easy to use and effective. The Call Minder project demonstrates the importance of designing usability into speech services. Users were involved throughout the development and identified numerous improvements to the service. Implementing such improvements in future speech services will yield enhanced customer satisfaction, increased return on investment and significant competitive advantage.

References

- 1 ROSE, KEVIN, and HUGHES, PAT. The BT Speech Applications Platform. *Br. Telecommun. Eng.* Jan. 1995, **13**, p. 304
- 2 COOPER, M. B. Human Factor Aspects of Voice Input/Output. *Speech Technology*, March/April 1987, pp. 82-86.
- 3 ATYEO, M., and GREEN, R. User Friendly Weapons for the Competitive Fight. *Br. Telecommun. Eng.*, Oct. 1994, **13**, p. 201.

- 4 Speech Applications Division. Voice Applications Style Guide: Design Manual, British Telecommunications (7.7).

Biographies



Charanjit K. Sidhu
BT Networks and Systems

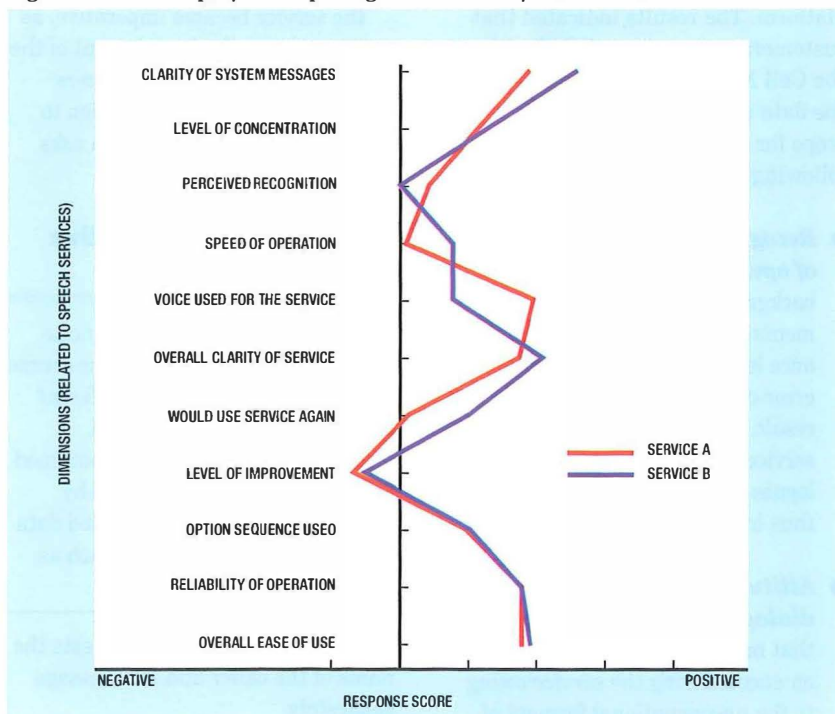
Charanjit Sidhu graduated from Hatfield Polytechnic with a B.Sc. in Psychology and went on to gain an M.Sc. in Computer Science before joining BT Human Factors in 1991. She has been involved in the application and promotion of usability design and evaluation, providing human factors consultancy on a wide range of project domains including multimedia, graphic user interface design, network services and speech-based services. Charanjit coordinated the Human Factors contribution to the design and development of the Call Minder service.



Gerry Coyle
BT Networks and Systems

Gerry Coyle joined BT Human Factors in 1992 after graduating with a B.Sc. in Psychology from the University of Ulster and an M.Sc. in Applied Psychology from Cranfield University. His main areas of technical interest are multimedia evaluation, speech-based services design and customer satisfaction measurement. While at BT, he has worked on a variety of projects including Call Minder, and currently, BT's Interactive TV service.

Figure 3 - Attitude profile comparing two versions of Call Minder



Access Network Support Systems—An Overview

Modernisation of BT's access network and its operation is vital to reduce costs, enhance services and improve the quality of service to customers. It is key to ensuring that BT keeps its leading role as a provider of telecommunications in the UK. Remote control and automation of the provision and restoration processes from customer reception are critical to meet the expectations of customers and control costs. This relies upon modern access transmission technologies that may be managed and configured remotely.

Introduction

A special issue of the *Journal*, published in April 1991, was devoted entirely to the access network. Since then, BT has carried out a major review of its plans to take account of the increase in competition, the consequences of regulation and continuing improvements in technology. The major areas recognised as vital to the long-term success of the company were:

- the control and reduction of costs through improved operating efficiency;
- improved quality of service as well as a major new focus on private circuit network performance and major system failures;
- the provision of new services and features that increase the value for money from BT's network services;
- the flexibility and speed of response to customers' needs, particularly a reduction in time to provide service; and
- the need to develop, test and launch new services in much shorter timescales.

This article describes how access network support systems (ANSS) will accommodate the delivery of new and existing services over a range of emerging modernised and advanced access technologies.

A high-level overview of the automated systems needed to support the modernisation of the

access network provision and maintenance structure is shown in Figure 1.

The Access Network

The access network is the link between the serving telephone exchange and the customer's premises; that is, between the main distribution frame (MDF) up to, and including, the network terminating equipment (NTE). This network serves more than 26 million customers with over 30 million circuits. The access network represents a third of BT's total assets and currently accounts for 80 per cent of total network failures. This results in customer complaints and massive fault clearance costs. Improvements in the management and operation of the copper access network to reduce the fault rate are being sought, through Breakout initiatives on field effectiveness and quality (FEQ) and the access operations unit (AOU) pilot. These were part of the Breakout Genesis Integrated pilot in the North Midlands. The fault volume reduction (FVR) programme is also helping to reduce the fault rate.

Investment in fibre, radio and copper-based transmission systems is now being undertaken in modernising the access network. These will support new services and reduce the inherent cost associated with the existing copper bearer.

Existing Legacy Problem

Currently, in most cases, the link between the modern exchange and the customer is a pair of copper wires that passes through a number of flexibility points. When a fault

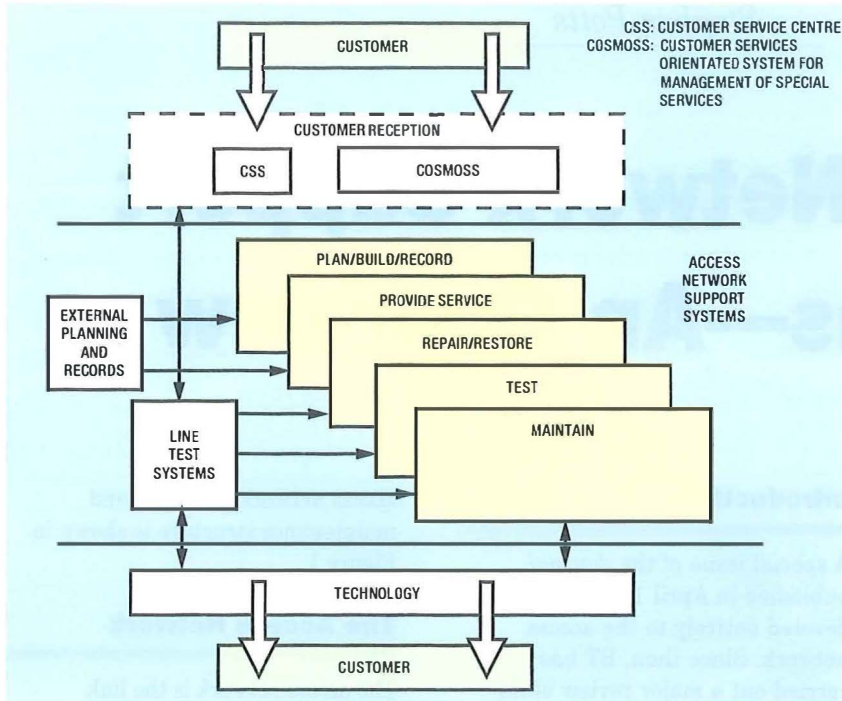


Figure 1—High-level overview of the automated support systems

occurs, the customer is relied upon to report the problem. The fault is then localised and repaired. When the access network needs to be expanded, a team of planners use hand-produced drawings to perform a survey of available capacity. In the absence of automated records, only a limited monitoring of cable capacity and available duct space is performed. This means that there are few early warning mechanisms for identifying any problems relating to the supply of additional circuits or new cables.

The predominant use of copper technology together with a high degree of manual construction, circuit provision and repair processes means that management of the access network is still extremely labour intensive. As more services are introduced and new technologies evolve, it is imperative that a highly automated network is provided. Furthermore, the introduction of new services and technologies, together with the emergence of many more other licensed operators (OLOs), is making the job of planning and operating the access network an increasingly complex challenge.

Overall Strategy

The access network is controlled through the following processes:

- Plan/Build/Record.
- Provide Service.
- Repair/Restore.
- Test.
- Maintain.

The 'Plan/Build/Record' process aims to install access network capacity in anticipation of customer demand and record its existence. This will ensure that service is provided with a minimum of delay. Rapid, reactive build must also be supported to meet unforeseen customer demand.

The 'Provide Service' and 'Repair/Restore' processes need to be fully automated to ensure demands are met within timescales agreed with the customer. These need to be supported by an automated 'Test' process that indicates that the service requested is operational and functioning correctly.

The 'Maintain' process needs to be able to drive either the 'Plan/Build' process for planned remedial work or the 'Repair/Restore' process if a fault affecting a customer is detected.

Access Network Management Systems

The existing access network support systems (for example, parts of CSS,

LTS) were built and designed to support a purely copper-based network, where one pair of wires carried one communication channel to the customer. The main aim of the access network support systems (ANSS) programme is to enhance and supplement these systems to support similar processes using fibre and radio, as well as, enhanced copper technology.

Advances in access systems to support the 'Plan/Build/Record' processes will enable monitoring of stabilised copper networks down to the level of the distribution point, of which there are more than three million. This will include the ability to highlight equipment that is approaching full capacity. These enhancements are the foundation for cost-effective, just-in-time provision of network capacity. Early versions of this access capacity planning system (ACPS) have already been delivered to support this process. The available data will be used to design network routing and to monitor the capacities of the different bearers.

An important part of the 'Plan/Build/Record' process is to enhance the existing copper-based systems, so that they are capable of handling the enhanced copper services and technologies as well as those provided over fibre and radio.

Increasingly, customers expect to be provided with service quickly and to an agreed timescale. One of the early thrusts in systems support for the access network has been to increase the level of automation in the provision-of-service area to support these expectations, irrespective of the technology being used. This will give customer receptionists the ability to agree a mutually acceptable date for provision of service with the customer. Indeed, in some cases, receptionists will be able to switch services on immediately. The combination of private services and public switched telephone network access delivery technologies and support systems will also see

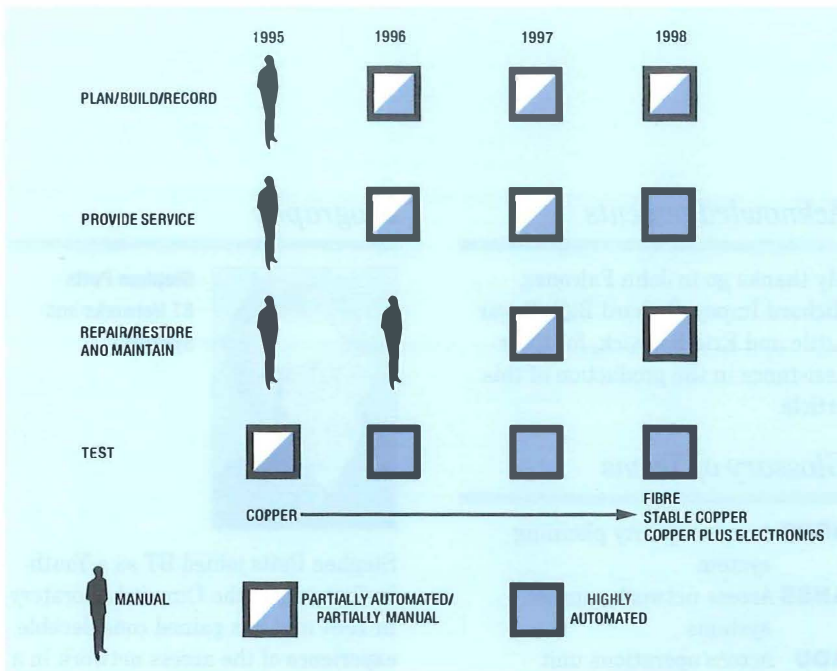


Figure 2—The programme of system and technology deliveries

improvements being made to the provision time of private services.

The introduction of new technology, which has the inbuilt capability to monitor its own ability to support service correctly, will automatically produce status reports. On the detection of a fault, the appropriate actions will be taken automatically. This allows faults to be reported and acted upon with minimal delay. This watchdog facility will also identify any maintenance work that needs to be performed and should lead to a reduction in customer detected and reported faults.

Access network support systems will analyse all the fault reports they

receive with the aim of identifying a single point of failure in the network. This fault correlation allows the cause of the failure to be detected and automatically logged. Faults will generally be highlighted before they are noticed by the customer. Future systems will be able to determine the availability of alternative routes or equipment, thus allowing automatic reconfiguration of the link around the faulty element. This facility will only be possible in situations where an alternative network path exists. The faulty item will then be placed in a repair queue for repair at some convenient time. The use of a

proactive, rather than a reactive, approach will improve the quality of service and reduce costs.

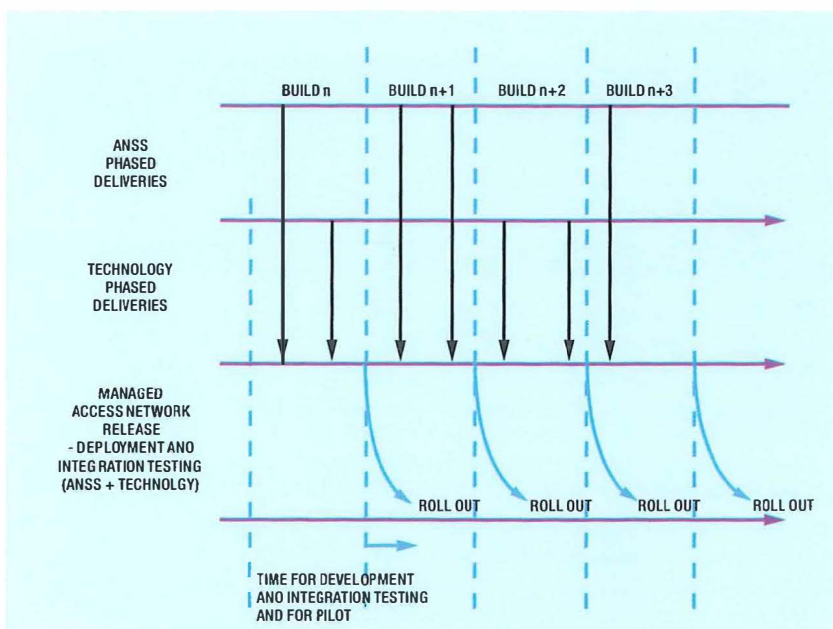
If a customer does have cause to complain about the service, the customer service centre (CSC) staff will be notified of any known network faults affecting that customer. They will also be able to carry out tests, if required.

Current Situation

The aim of the access network support systems programme is to help bring about radical improvements to services and much greater automation of the access network. The strategy being used to meet this goal is to deliver the new and improved systems in phases. These phases will allow small but qualitative advances in automation to be provided to the customer facing duties, with minimal disruption to their current processes. This strategy is demonstrated in Figure 2, which shows how system deliveries are phased to align with technology deliveries to ensure an integrated roll-out of fully tested services and technologies. The phased deliveries are referred to as *BUILDS*, which enable the contents of each delivery to be managed.

The system developments are linked to the technology programmes. It is only when technology and systems are combined that any benefits are achieved. The current thrust is to support the TPN and one-per-customer radio trials, together with the technologies that will deliver such exciting new services as video-on-demand and home shopping as well as giving customers a more reliable service and better value for money. The combination of ANSS builds and technology deliveries into managed access network release (MANR) is illustrated in Figure 3.

Figure 3—The linking of system developments and technology programmes



The Next Steps

Further developments to access systems will occur over the next few

years to increase the remote management of the network by CSCs. The increase in ANSS automation will allow BT to reduce its costs substantially and remain competitive.

Facilities will be provided to allow services to be switched on (or off) by CSC within seconds of the customer request. Eventually, customers may have the ability to configure their own equipment to meet their individual requirements immediately.

New systems will allow the introduction of new network services and technologies while streamlining existing business processes and work practices. They will also provide a vehicle for the fast and effective deployment of these new facilities.

Completion of access systems deployment is key to BT having one of the most advanced, reliable and cost-effective access networks in the world.

Conclusion

The further automation of the business processes is required to increase efficiency and quality of service and to reduce costs. This is being achieved by combining the process, systems and technology programmes into phased releases. These releases ensure that a fully integrated and validated product is delivered into the operational environment.

This article provides an overview of the ANSS programme. Further articles are planned, providing more detail on how ANSS will support the various business processes.

Acknowledgements

My thanks go to John Falconer, Richard Impey, Richard Ball, Roger Little and Eric Bonwick, for their assistance in the production of this article.

Glossary of Terms

- ACPS** Access capacity planning system
- ANSS** Access network support systems
- AOU** Access operations unit
- COSMOSS** Customer services orientated system for management of special services
- CSC** Customer service centre
- CSS** Customer service system
- EPR** External planning and records
- FEQ** Field effectiveness and quality
- FVR** Fault volume reduction
- LTS** Line test systems
- MDF** Main distribution frame
- NTE** Network terminating equipment
- OLO** Other licensed operator
- TPON** Telephony over passive optical fibre

Biography



Stephen Potts
BT Networks and
Systems

Stephen Potts joined BT as a Youth-in-Training in the Circuit Laboratory in 1967 and has gained considerable experience of the access network in a variety of development and management appointments. He has been the Access Systems Programme Manager for the past 18 months managing the full project life-cycle of the access network support systems projects. He is a Chartered Engineer and a Member of the British Computing Society.

Access Networks Evolution and Preparation for Implementation: EURESCOM Project 306

This article describes the achievements, from BT's perspective, of the EURESCOM (European Institute for Research and Strategic Studies in Telecommunications) Project 306 – 'Access Networks Evolution and Preparation for Implementation', which investigated fibre in the loop in non-green-field situations. The need for collaboration is explained and how the project has worked towards achieving common agreement on a broad range of access-network issues. The implications of this work on further collaboration are explored.

Introduction

The access-network infrastructure is based on dedicated copper pairs for analogue telephony with a reach of typically 3 km for 90 per cent of lines. While transmission, terminals and switching have evolved through several generations, copper pair technology has not changed significantly for over a century. Although new services have become available on copper, these could be much improved if the capacity of fibre is made available. Upgrading the access network is very complex because of the vast choice of technology options and the number of business activities that are affected by such changes.

The telecommunications market demands service flexibility coupled with high reliability which the copper network struggles to provide. Passive optical networks (PONs) using optical splitters in the network to share costs were recognised 11 years ago as an economical way forward for fibre in the loop (FITL)[†].

There is a consensus among public network operators (PNOs) that fibre provides a more future-proof solution in terms of service capability, responding to new service demand, network reliability and lowered costs. In areas with an established telecommunications infrastructure, optical fibre could provide a solution to rehabilitate the obsolescent copper infrastructure or cope with provision when there is a rapid growth in service demand. But areas with an established infrastructure may need

different strategies and equipment than the green-field situation.

There is a need to investigate and further clarify the application of FITL systems in urban, suburban and rural areas in order to be able to supply PNOs with knowledge of technical solutions. This article describes the results of EURESCOM† Project 306 – 'Access Networks Evolution and Preparation for Implementation', which investigated FITL in non-green-field situations. The project started in August 1993 and finished in July 1995.


The Need for Collaboration

The size of investment required to upgrade the access network is far greater than the cost of the core switching network and is, at present, inhibitive. Before any widespread development of new technology in the access network can take place, there must be:

- Agreement on the economic benefits, with respect to cost savings or additional revenues.
- Consensus on the system and network requirements. What interfaces are needed? What

[†]EURESCOM was founded in 1991. It now includes 24 PNOs as shareholders and is the only organisation representing PNOs alone. Its main function is to perform collaborative research and development in telecommunications through projects carried out by the shareholders.

Figure 1 – Relationship between EURESCOM Technical Advisory (ETA) and European Technical Standards (ETS)



capacity is needed in each direction of transmission?

- An understanding of where and how the technology will be deployed. What is the first application going to be? What is the range, coverage area and penetration over the system life?
- Technical feasibility of the practical solution.

With common agreement on a broad range of access network issues, this results in:

- a common set of requirements which are clearly stated to industry;
- speeded up development and roll-out of new technology;
- suppliers encouraged to invest because a wider market is potentially available;
- lower equipment costs;
- pooled expertise of PNOs;
- trade of experience and expertise between PNOs, increasing the total understanding in area;
- development of key standards by coordinated activities; and
- investment in technical strategy.

The benefits of collaboration are clear for the PNOs and their suppliers, in terms of reducing the risk of the investment required and increasing the potential size of the market. Project 306 is developing the areas in need of common agreement to realise the advantages of collaboration.

Project 306

BT was involved in Project 306 task areas where it was considered that the tasks would benefit most from establishing a common strategy.

Objectives

Project 306's objectives are to establish an understanding of PNOs' common requirements for FITL in areas with an existing telecommunications infrastructure, to reduce product diversity by common specifications, to specify operation and maintenance (O&M) requirements and procedures and to provide a common European PNO input to the European Technical Standards Institute (ETSI) and the International Telecommunications Union (ITU). The common requirements for narrowband FITL systems is captured in a EURESCOM technical advisory (ETA).

Structure and BT's involvement

Project 306 is divided into seven tasks, BT's involvement has been in:

- task 1—development of a technical advisory (TA) for narrowband FITL systems to meet the requirements of all participating PNOs using experience from implementation;
- task 2—strategies for FITL introduction;
- task 4—practical implementation issues; and
- task 7—operations and maintenance.

Technical approach

After a verification of the existing standards and an inventory of experiences, strategies were developed for the introduction of FITL in existing network situations which address the different technology solutions for the customer drop. The tool developed in RACE (Research into Advanced Communication in Europe) Project R2087 TITAN (tool for introduction scenarios and techno-economic studies for the access network)² was used for techno-economic evaluations. The developed strategies were aimed at enabling

future upgrading for broadband services. After setting this framework, practical issues were dealt with; for example, cabinet design, powering, O&M, etc.

Deliverables

The project aimed to produce three documents.

- A EURESCOM technical advisory 1 on narrowband FITL equipment for passive optical networks. (This provides a set of requirements for telecommunications terminal equipment to be used on PONs.)
- Strategic recommendations and guidelines for the efficient introduction and management of FITL networks, including planning for possible upgrades to broadband.
- A summary of the conclusions of the project for the operation of FITL systems, including practical implementation considerations.

Development of a TA for FITL Systems

TAs³ are relatively unknown in Europe, in the context of telecommunications. In the USA, they have served to focus the development of telecommunications equipment for the US market. Within Europe, much of this activity occurs through direct contact with industry, RACE, and activities in ETSI.

The aim of this TA, entitled 'EURESCOM Technical Advisory for a Fibre-in-the-Loop TPON System', is to focus the market and stimulate volume by bringing together the requirements of several separate PNOs in a single document. It may also accelerate the deployment technology as PNOs gain understanding through the sharing of information and confidence in the technology. Figure 1 shows the relationship between European technical stand-

Table 1 : Service penetrations for the cases examined in task 2

Service Evolution Case 1: Obsolete existing infrastructure, high subscriber density, no competition

Penetration of services (%)					
Year	POTS	ISDN	CATV	ASB	SSB
1995	30	0.0	0.0	0.0	0.00
1996	40	0.5	0.0	0.0	0.00
1997	50	0.8	5.0	1.5	0.00
1998	60	1.5	10.0	4.5	0.20
1999	70	2.5	20.0	7.5	0.50
2000	80	3.5	30.0	10.0	0.75
2001	90	4.0	33.0	12.5	0.95
2002	100	4.5	35.0	14.0	1.10
2003	100	5.0	40.0	16.0	1.20
2004	100	6.0	45.0	18.0	1.35
2005	100	6.5	50.0	20.0	1.50

Service Evolution Case 4: Modern existing infrastructure, high subscriber density, PNO not allowed to provide CATV

Penetration of services (%)					
Year	POTS	ISDN	CATV	ASB	SSB
1995	90	1.0	0	0.0	0.0
1996	89	1.5	0	0.2	0.0
1997	88	2.0	0	0.4	0.0
1998	87	3.0	0	0.8	0.0
1999	86	4.0	0	1.5	0.0
2000	85	5.0	0	3.0	0.0
2001	85	7.0	0	5.0	0.1
2002	85	9.0	0	7.0	0.3
2003	85	11.0	0	8.0	0.6
2004	85	13.0	0	9.0	1.2
2005	85	15.0	0	10.0	2.0

ards and EURESCOM technical advisory.

The term TA is used in EURESCOM to mean a framework requirements specification for a technical system. TAs are distributed to supply industries as drafts and comments are requested. Once agreed, TAs communicate the harmonised requirements, from the operators' point of view, to manufacturers. For their procurement, operators are invited to use the TA in their individual specifications. The TA started as a consolidation of previous TAs jointly prepared by BT and Deutsche Telekom, but has been influenced by subsequent work in ETSI.

The TA requirements refer to two basic network elements of the PON system: the optical line termination (OLT) and optical network units (ONU) of different capacity. They range from type 0 with two lines (2 × 64 kbit/s equivalent capacity) to type 5, a 128 line unit.

Service Evolution Case 2: Modern existing infrastructure, high subscriber density, low level of competition

Penetration of services (%)					
Year	POTS	ISDN	CATV	ASB	SSB
1995	99	1	0	0.0	0.5
1996	98	2	0	0.2	1.0
1997	97	3	0	0.4	2.0
1998	96	5	0	0.6	3.0
1999	94	6	0	0.8	4.0
2000	93	7	10	1.0	5.0
2001	91	9	14	2.0	6.0
2002	88	11	18	4.0	7.0
2003	85	14	22	6.0	8.0
2004	82	17	26	8.0	9.0
2005	80	20	30	10.0	10.0

Service Evolution Case 5: Incomplete existing infrastructure, medium subscriber density

Penetration of services (%)					
Year	POTS	CATV	ISDN	ASB	SSB
1995	75.00	0.0	0.00	0.00	0.00
1996	76.96	1.0	0.02	0.02	0.02
1997	78.82	2.3	0.04	0.04	0.03
1998	80.57	4.4	0.08	0.08	0.06
1999	82.21	6.8	0.16	0.17	0.11
2000	83.76	8.9	0.32	0.34	0.20
2001	85.20	12.2	0.63	0.70	0.36
2002	86.54	17.0	1.26	1.42	0.66
2003	87.78	23.2	2.51	2.89	1.20
2004	88.93	32.6	5.01	5.89	2.19
2005	90.00	50.0	10.00	12.00	4.00

General requirements for FITL broadband PON systems

The uncertainty in the strategies of PNOs with respect to broadband services and in market take-up means that it is too early to develop a broadband TA. An internal project 'General Requirements for Fibre in the Loop Broadband PON (BPON) Systems' has been created that gives some early thoughts of the European PNOs on the general requirements of BPON systems for the delivery of broadband and narrowband services to residential and small business users. It is not intended as a TA or a technical standard, but it could be the basis for future standardisation work. The focus is on the architectural, transmission, performance and management requirements for BPON systems. But it is not clear whether BPOs are the infrastructure that PNOs will use for the transport of broadband services in the access network.

Service Evolution Case 3: Modern existing infrastructure, medium subscriber density, competitive environment

Penetration of services (%)					
Year	POTS	ISDN	CATV	ASB	SSB
1995	90	5	0.0	0.0	0.0
1996	86	6	0.0	0.0	0.0
1997	82	6	0.0	0.0	0.0
1998	78	7	0.0	0.0	0.0
1999	74	8	0.0	0.0	0.0
2000	70	10	1.0	1.0	1.0
2001	66	12	1.6	1.6	1.6
2002	62	15	3.0	3.0	3.0
2003	58	18	4.5	4.5	4.5
2004	54	21	6.7	6.7	6.7
2005	50	25	10.0	10.0	10.0

Strategies for FITL Introduction (Task 2)

Europe exhibits a wide diversity of existing infrastructure, short- and medium-term needs, cultural and social acceptance of new services and economic capability to afford them. The five hypothetical cases below represent situations common among several European operators:

Case 1: Obsolete copper pair infrastructure, high subscriber density, no competition.

Case 2: Modern copper pair infrastructure, high subscriber density, low level of competition.

Case 3: Modern copper pair infrastructure, medium subscriber density, competitive environment.

Case 4: Modern copper pair infrastructure, high subscriber density, PNO not allowed to provide CATV.

Case 5: Incomplete copper pair infrastructure, medium subscriber density, no competition.

Five network transport services were considered, plain old telephone service (POTS), integrated digital subscriber network (ISDN), asymmetric switched broadband (ASB) (for example, video-on-demand), symmetric switched broadband (SSB) (for example, local area network interconnect), and community access television (CATV)—analogue TV distribution. Table 1 lists the service penetrations for each case generated from within the project. To carry these services, several technical solutions were proposed for each case, as shown in Table 2.

Table 2 Technical solutions modelled for each case

Case	Solution title	Technical solutions considered
1	1. Full copper pair network, no CATV	Concentrators for old POTS, ADSL for ASB, HDSL for SSB, HDSL to a RSM for new POTS and ISDN.
	2. Fibre-based multiplexer, no CATV	Concentrators for old POTS, ADSL for ASB, fibre based RSM for SSB, new POTS and ISDN.
	3. PON for new customers	Concentrators for old POTS, HFC for CATV, ADSL for ASB, TPON for SSB(FTTH), new POTS(FTTK) and ISDN(FTTK).
	4. PON for all customers	TPON for POTS(FTTK), ISDN(FTTK) and SSB(FTTH), HFC for CATV and ASB.
	5. ATM PON	TPON(FTTK) for POTS and ISDN, ATM PON(FTTK) for ASB and SSB, HFC for CATV
2	1. Copper network, regenerator and ADSL	Concentrators for POTS and ISDN, SDH regenerators for SSB, ADSL for ASB, HFC for CATV.
	2. Copper network, regenerator and CATV	Concentrators for POTS and ISDN, SDH regenerators for SSB, HFC for CATV and ASB.
	3. Copper network, HDSL, ADSL	Concentrators for POTS and ISDN, HDSL for SSB, ADSL for ASB, HFC for CATV
	4. Copper network, ADSL, optical point-to-point SSB.	Concentrators for POTS and ISDN, fibre feed to kerb for SSB, ADSL for ASB, HFC for CATV.
	5. PON, HDSL, ADSL	TPON(FTTK) for POTS and ISDN, HDSL for SSB, ADSL for ASB, HFC for CATV.
	6. Copper network, ATM PON for ASB and SSB	Concentrators for POTS and ISDN, ATM PON(FTTK) for SSB and ASB, HFC for CATV
3	1. Enhanced copper network	Concentrators for POTS and ISDN, HDSL for SSB, ADSL for ASB, HFC for CATV
	2. PON	FTTK (TPON) for POTS and ISDN, FTTH/FTTK (TPON) for SSB, HFC for CATV and ASB
	3. ATM PON, 8 subs/ONU	FTTK (ATM PON) for POTS, ISDN, SSB and ASB, HFC for CATV
	4. ATM PON, 32 subs/ONU	FTTK (ATM PON) for POTS, ISDN, SSB and ASB, HFC for CATV
	5. Advanced FTTH	FTTH for POTS, ISDN, SSB, ASB and digital broadcast TV
	6. Hybrid fibre and coaxial cable	HFC for POTS, ISDN, SSB, ASB and CATV
4	1. Enhanced copper network, no CATV	Concentrator for POTS and ISDN, ADSL for ASB, HDSL for SSB
	2. Fibre-based multiplexer, no CATV	Point to point multiplexer to cabinet for POTS and ISDN, ADSL for ASB, HDSL for SSB
	3. PON, no CATV	TPON for POTS (FTTK), ISDN(FTTK) and SSB(FTTH), ADSL for ASB
	4. HFC, no CATV	HFC for POTS, ISDN, SSB and ASB
	5. ATM PON, no CATV	ATM PON (FTTK) for POTS, ISDN, SSB and ASB.
5	1. Enhanced copper network, no CATV	Concentrators for POTS and ISDN, HDSL for SSB, ADSL for ASB
	2. HDSL, CATV also for ASB	Concentrators for POTS and ISDN, HDSL for SSB, HFC for ASB and CATV
	3. PON for new customers	Concentrators for old POTS and old ISDN, TPON(FTTH) for SSB, HFC for ASB and CATV, TPON(FTTK) for new POTS and new ISDN
	4. PON for all customers	TPON(FTTK) for POTS and ISDN, TPON(FTTH) for SSB, HFC for ASB and CATV
	5. ATM PON	TPON(FTTK) for POTS and ISDN, ATM PON(FTTH) for SSB, HFC for ASB and CATV

ADSL	Asymmetric digital subscriber loop	HDSL	High-speed digital subscriber loop	POTS	Plain old telephone service
ASB	Asymmetric switched broadband	HFC	Hybrid fibre and coaxial cable	RSM	Remote subscriber multiplexer
ATM	Asynchronous transfer mode	ISDN	Integrated services digital network	SSB	Symmetric switched broadband
FTTH	Fibre to the home	ONU	Optical network unit	TPON	Telephony over PON
FTTK	Fibre to the kerb	PON	Passive optical network		

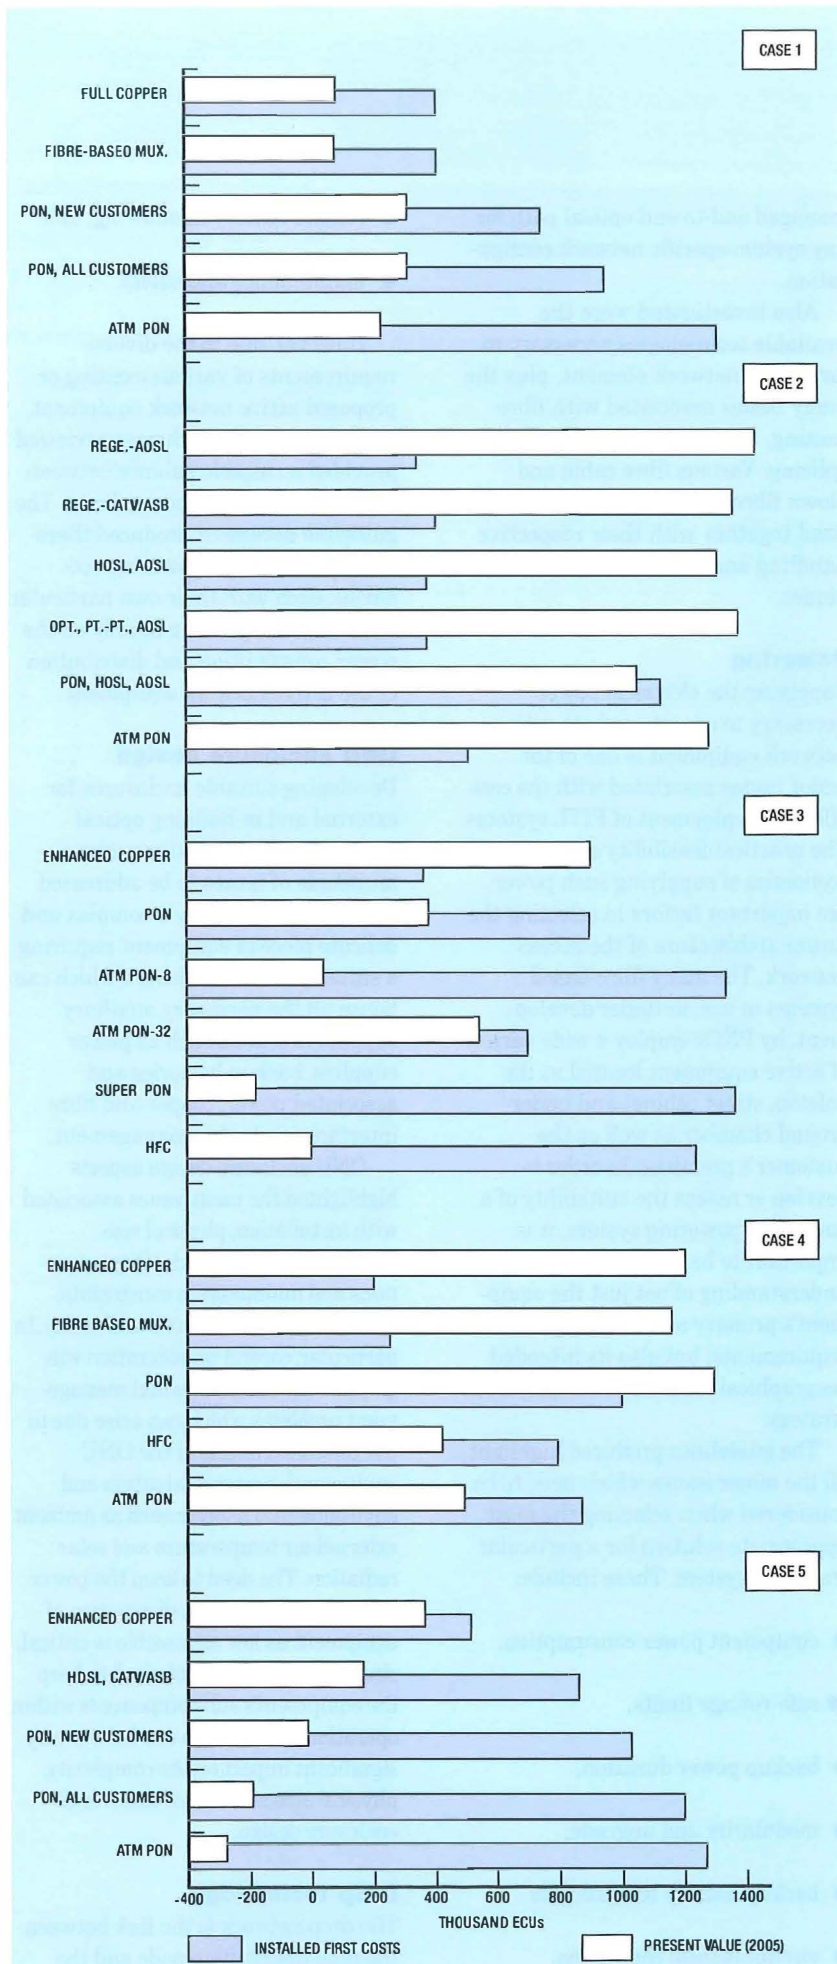


Figure 2—Installed first costs and present values of the solutions evaluated

Figure 2 shows the installed first cost and present value of the different solutions for an area serving 1024 customers. These results were generated using the TITAN techno-economic access network modelling tool. Upgrading an existing access network for broadband services represents a substantial PNO investment in all cases. The results between cases are not directly comparable, because the cases correspond to completely different area types, service patterns and technical solutions. The results are based on the following constant annual tariff assumptions: 160 ecu for POTS, 530 ecu for N-ISDN, 250 ecu for CATV, 600 ecu for ASB and 1000 ecu for SSB. In cases where the existing access network is obsolescent, and in technical solutions where the whole access network needs to be rebuilt, the pay-back periods are remarkably high in the order of 10–12 years.

In developed areas, where a good existing copper network is available, introduction of broadband services can be done cost effectively by using enhanced copper solutions. This conservative strategy is advisable if the demand for broadband services remains low. This assumes that enhanced copper technologies work reliably in practise and current account expenditure calculations in the cost model are relatively simplistic.

In Eastern European countries and some other developing areas, the demand for new POTS lines is one of the main driving forces for access network evolution. Building a fixed network to satisfy this demand for new POTS lines and the rising demand for broadband services require high investments from the PNO. The pay-back period of these investments is far too long to make business profitable. Therefore, the PNOs could consider alternative access solutions such as radio access to satisfy the urgent need for POTS.

FITL introduction strategies depend strongly on penetration of

various services, the set of services, tariffs and dwelling distributions. With different sets of services, service penetrations and tariffs, an access network solution that now seems uneconomical can turn out profitable. Similarly, fibre-based access solutions, which are uneconomical for single houses, can be cost effective for apartment houses. Extensive sensitivity analyses of all these aspects is highly recommended, as well as including more detailed current account expenditure modelling.

Practical Implementation Issues

In most countries, the existing local distribution network for customer access mainly consists of copper pairs. In many such countries however, substantial experience has been gathered during recent progress in the application of optical plant technology. The objective of reviewing these experiences is to form a consensus of opinion from which guidelines could be produced concerning network architecture, powering, cabinet design and drop technologies.

Network architecture

In order to begin evaluating the optical performance and practical feasibility of various network infrastructure configurations, a top-level reference model for the optical access infrastructure is required to develop an in-depth appreciation of this fundamental information. The physical reference model is system independent and highlights the physical position and functional requirements of all the network elements required to establish any specific transport system (point-to-point, PON, FTTK or FTTH) between the local exchange and customer premises equipment.

The tasks' guideline document contains a portfolio of network elements, each with a clear functional description. These can be used in a modular fashion with the flexibility required to create a completely

managed end-to-end optical path for any system-specific network configuration.

Also investigated were the available technologies necessary to form each network element, plus the many issues associated with fibre routing, termination, storage and splicing. Various fibre cable and blown fibre technologies are summarised together with their respective handling and installation techniques.

Powering

Supplying the electrical power necessary to operate remote active network equipment is one of the major issues associated with the cost effective deployment of FITL systems. The practical feasibility as well as the economics of supplying such power are important factors in selecting the future architecture of the access network. The many fibre-based systems in use, or under development, by PNOs employ a wide variety of active equipment located at the poletop, street cabinet and underground chamber as well as the customer's premises. In order to develop or assess the suitability of a particular powering system, it is important to have an in-depth understanding of not just the equipment's primary and backup power requirements, but also its intended geographical location and deployment strategy.

The guidelines produced highlight all the major issues which need to be considered when selecting the most appropriate solution for a particular transport system. These include:

- equipment power consumption,
- safe voltage limits,
- backup power duration,
- modularity and upgrade,
- backup battery technologies,
- environmental conditions,

- remote battery monitoring, and
- maintenance and safety.

However, due to the diverse requirements of various existing or proposed active network equipment, no single powering schemes reviewed provided a suitable balance between reliability, maintenance and cost. The guideline document produced therefore describes four powering scenarios, each with their own particular advantages depending mainly on the power consumption and distribution of the active network equipment.

ONU enclosure design

Developing suitable enclosures for external and in-building optical network units (ONUs) requires a multitude of issues to be addressed. The ONU is generally a complex and delicate piece of equipment requiring a suitably robust enclosure which can house all the necessary auxiliary support functions such as power supplies, backup batteries and associated power, copper and fibre interfaces and cable management.

ONU enclosure design aspects highlighted the main issues associated with installation, physical size, modularity and upgradability, operations and maintenance constraints, environmental conditions and safety. In particular, careful consideration was given to the severe thermal management problems which can arise due to the combined effects of the ONU equipment's power dissipation and environmental aspects such as ambient external air temperature and solar radiation. The need to keep the power consumption and heat dissipation of equipment as low as possible is critical, since the technology required to keep the equipment's sub-components within operational tolerances can have a very significant impact on the complexity, physical size and cost of the overall enclosure design.

Drop technology

The drop network is the link between the final distribution node and the

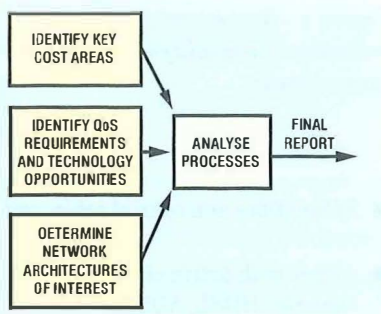


Figure 3—Approach to O&M task; the four main work areas

customer's premises and may consist of a variety of interfaces depending on the particular services supported and the associated transport technology. Generally, interconnection between the customer and the network is realised with two basic techniques: cable or radio. All the basic cable types such as balanced twisted pair, coaxial, optical and hybrid are reviewed, together with their specific transmission characteristics. Installation methods and principles for both underground and overhead cables are summarised with comparative costings for cables, civil engineering, installation and future operations and maintenance activities.

Operations and Maintenance

Operations and maintenance (O&M) is a key area to address when introducing new technology into the network. New revenue can be generated through the new services that the technology enables, but for this to be turned into real profit, the O&M costs associated with the technology must be acceptable. PNOs are thus faced with a major problem: how can the O&M cost implications of a new technology be determined, and hence the commercial viability of a new technology be assessed? A second, related, problem is the need to specify O&M process support functionality in the new technology early in its development life cycle.

The importance of O&M has been recognised within EURESCOM Project 306, with a discrete task being set up to consider the O&M aspects of new access network technology.

Objective of the O&M task

The objective of the O&M task within Project 306 was to make recommendations in the following areas:

- the key O&M cost areas that new technology needs to address,
- opportunities for improving key business processes the new technology presents, and
- the O&M functionality that should be included in any new technology.

Approach

The starting point for any work on O&M must be the dual business imperatives of cost and quality-of-service (QoS). Invariably, the objective for any PNO in a competitive environment is to strike the optimum balance between the two, such that customers will select them as the supplier that delivers the right quality of service at the right price.

The O&M task was broken down into four key work areas reflecting these business imperatives (Figure 3). 'Identify key cost areas' and 'Identify QoS requirements and technology opportunities' tackled the cost and quality issues. 'Determine network architectures of interest' put forward five possible network scenarios. These three work areas summarised the information required for the fourth work area, 'Analyse processes', in which the final recommendations were produced. The remainder of this discussion on the O&M task looks at the four work areas in detail.

Identify key cost areas

What are the areas?

In order to focus the limited resources available onto the most important areas, it was necessary to identify which of the O&M areas represented the major sources of cost to PNOs. Before this could be done, an agreement was needed on what the relevant O&M areas are. Using telecommunications managed network (TMN) recommendations as a starting point, the following four key areas were identified:

- *Administration management* This is the management of all data-

bases connected with the access network, from records of the network infrastructure to information regarding which customers have which services available.

- *Provision management* This is effecting new service connections and cancelling existing connections.
- *Performance management* This is measuring network performance, and includes the cost of collecting and analysing the relevant performance data.
- *Fault management* This is dealing with fault reports, and includes handling the initial fault report, localising the problem, all activities that take place to resolve the problem and repair the network, and notifying the customer that the repair has taken place. Also included in this area is preventative maintenance.

Obtaining the cost information

A detailed questionnaire was compiled which requested cost-per-line estimates for the four areas identified previously. In addition, the questionnaire asked for cost estimates to be broken down into urban, suburban and rural geographical environments, and by network elements (main distribution frames, primary connection points, main cables, etc.).

All PNOs partaking in the O&M task attempted to complete this questionnaire by contacting relevant people within their organisations. There were a number of barriers to the success of this exercise, the most significant ones being:

- much of the information requested was commercially sensitive, especially at a time when (for BT at least) access network cost information was a key issue for the regulator,
- many PNOs did not have financial systems in place that enabled




Figure 4—Brainstorm for cost reduction and quality-of-service improvement

them to provide the relevant information, and

- there was, inevitably, only an approximate relationship between any one PNO's interpretation of the four functional areas and another's. This made cost comparisons between PNOs difficult.

Despite these shortcomings, the overall aim was achieved, which was to highlight those areas where the most-significant costs are incurred.

Significant O&M cost areas

The output from this work area was a detailed document which broke down O&M costs by functional area, geographical environment and network element. Because of the shortcomings identified, it was difficult to draw hard conclusions from this deliverable. A few general conclusions could however be drawn:

- By far the most significant cost areas are provision and fault management. The cost of these two is broadly similar.
- Spend on performance management was small in all cases. It is worth posing the question: could the expenditure on fault management, and hence the total overall cost, be reduced by increasing the expenditure on performance management?
- Fault costs rise very rapidly as one moves to more rural areas. Provisioning also rises but less rapidly. Perhaps different O&M strategies should be pursued in different geographical areas?

Identify QoS requirements and technology opportunities

The next step was to consider the technology used in various network elements (existing and new) and

assess what functionality could be incorporated into these elements to improve the QoS delivered to customers and/or reduce operational costs. The cost of providing the functionality was not considered at this point. The aim was to identify anything technically feasible that could potentially give benefits.

What is meant by QoS?

Quality of service is an area where detailed work has been done within standards organisations such as ETSI. This standards' work was used to produce a list of QoS parameters which could be influenced by access network O&M. There are, in fact, very few parameters which are not influenced in some way by a PNO's access O&M processes. This is well illustrated by the fact that there were over 50 parameters in the final list. These were grouped into the following general categories:

- service support performance,
- service accessibility performance,
- transmission performance,
- dependability (includes availability, reliability and maintainability),
- maintenance support functions, and
- security.

Scope of technologies considered

The following access technologies were included in the study:

- passive analogue copper,
- point-to-multipoint radio,
- dropwire replacement radio,
- point-to-point fibre,
- PON,

- hybrid fibre and coaxial cable, and
- copper with active electronics (for example, HDSL, ADSL).

Determining the impact of new O&M functionality

At this point, the scope of the work was expanded to include all of the following O&M areas:

- database administration,
- network design and build,
- network configuration (primarily to support provision and cessation of service),
- fault management,
- performance management, and
- security management.

Having identified the QoS parameters, the technologies and the O&M areas of interest, it was then possible to brainstorm each combination to identify opportunities for cost reduction and QoS improvement as shown in Figure 4.

Key conclusions from QoS opportunity analysis

A detailed document was produced which recorded, for each of the O&M areas considered, general principles which should be followed (independent of technology type) and the QoS parameters on which these principles would have the most impact. Following on from these general principles, technology-specific opportunities were identified which could be used to support them.

It is difficult to highlight any one of the principles or opportunities that were arrived at as being more or less significant than any other. The exercise was undoubtedly successful, however, in that a wide range of opportunities were identified and recorded in a systematic way for use in subsequent analyses. For example, if a PNO was looking to improve a

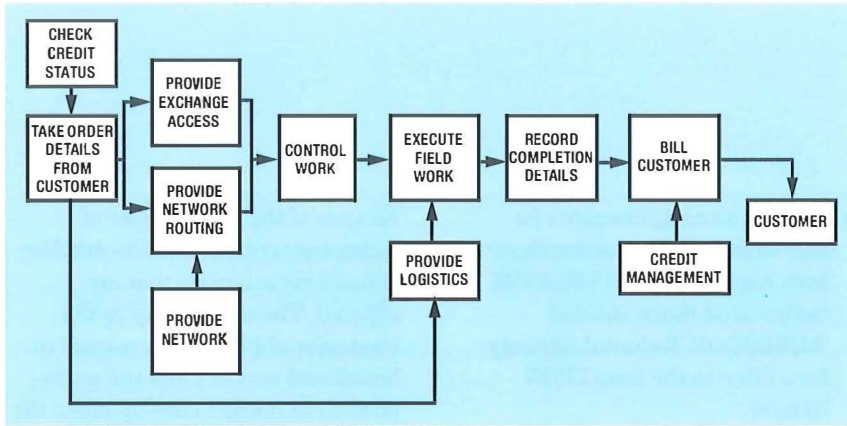


Figure 5—A schematic provision process diagram

particular QoS parameter, then it would be easy to identify the technology opportunities for doing that from the document produced.

Determine network architectures of interest

The work described so far had been done without reference to any specific network architectures. Before moving on to the final stage where processes were considered, it was necessary to make some assumptions about the networks on which the processes would be used.

The two key network characteristics that have an impact on O&M are technology and topology (that is, ring, point-to-point or branched). From the 27 possible network architectures given in Table 1, five were chosen which covered a sufficient range of both characteristics.

The implications of having a multi-technology access network, where, for example, the same workforce might be used to effect the repair process for customers served by differing network technologies, were also considered. This included sending a questionnaire to the PNOs taking part in the O&M task to find out how the O&M of multi-technology networks is currently implemented.

Key conclusions from network architectures work

- No PNOs are experienced in operating multi-technology networks.
- All PNOs recognise that a key issue is how the workforce should be organised: groups of relatively few people highly skilled in a specific technology, or many people with basic skills for dealing with all technologies.

Another key issue is the integration of management systems in a multi-technology environment. This is believed to be crucial in delivering acceptable O&M costs. Adherence to the ITU-T TMN recommendations is a key enabler to this integration.

Analyse processes

This final work area brought together all the previous work and used it to produce recommendations for O&M processes together with any resultant O&M requirements for the network technology.

Which processes matter?

The first challenge was to identify which business processes should be considered. Initially, all major business processes within a typical PNO were identified. This list was then reduced to focus on processes where the access network technology has a significant impact. The resultant list became:

- remove network fault (including repair faulty element);
- restore customer service;
- provide customer service;
- build the network;
- control the inventory;
- allocate people to tasks;
- provide customer premises equipment (CPE);
- maintain the network (including monitor network condition and performance); and
- plan the network.

In addition to the above list, it was recognised that it must be possible to measure the performance (and ultimately cost) of all business processes. This may have an impact on the functionality of the technology.

Process descriptions

The next step was to draw up descriptions of each of the above processes as they are today.

An important principle behind this work was that each process should be broken down into activities which were small enough to fit in only one of the TMN O&M functional areas. This then allowed easy cross-correlation between the processes and the technology opportunities identified earlier.

Another key principle was the use of the 'black box' concept. Any process will contain some activities which are not influenced significantly by the access network technology. For example, the 'provide customer service' process includes taking customer order details and checking customer credit status, neither of which make use of the access network technology. These activities were considered as black boxes—we were interested in the interfaces to them but not interested in what happens inside. By drawing black boxes around activities in this way, it is possible quickly to focus on those activities where new technology can have an impact. An example process diagram is shown in Figure 5.

Identifying the process opportunities

The final step was to consider the activities identified on the process descriptions for each of the five network solutions.

Key conclusions

It became apparent early on that, in order for the output to be of value, both a detailed list of processes and a simple but effective methodology for describing them was required. Both these prerequisites have been met and the level of understanding of the

process design problem has been substantially increased as a result.

There now exists a structured method for determining process improvements that could be made and, with the information now in place, it is highly likely that process recommendations of significant value will be produced once the work is complete.

Overall conclusions

The nature of O&M, being such a huge area to consider, meant that inevitably the output from task 7 was of a broad nature. To illustrate this, it is worth summarising the extent of the parameters that must be considered: over 50 QoS parameters, 25 network scenarios, six O&M functional areas and 9 key business processes.

The breadth of the work undertaken makes it difficult to summarise the findings in a short article such as this one. It is, however, fair to say that a great deal has been achieved, resulting in a large number of recommendations of how the O&M of new technologies should be addressed. The task has also succeeded in producing a well-organised framework to help PNOs do this.

The real value of task 7 will only be known when each of the many recommendations produced are taken up and costed by PNOs and suppliers of the technology. One thing is certain, without the kind of thinking on O&M issues that has taken place in this task, cost-effective deployment of new technology will not be possible.

Summary of Achievements

The Project 306 objectives were to establish an understanding of PNOs' common requirements for fibre in the access network in areas with an existing telecommunications infrastructure, to reduce product diversity by common specifications, to specify O&M requirements and procedures and provide a common European PNO input to ETSI and the ITU.

The objectives have been achieved as discussed below:

- The common requirements for narrowband FITL systems have been captured in a EURESCOM technical advisory, entitled 'EURESCOM Technical Advisory for a Fibre in the Loop TPN System'.
- An internal project result 'General Requirements for Fibre-in-the-Loop BPN Systems' has been created. It presented some early thoughts of the European PNOs on the general requirements of BPN systems for the delivery of broadband and narrowband services to residential and small business users.
- Access network techno-economic modelling has shown that FITL-introduction strategies depend strongly on penetration of various services, the set of services, tariffs and dwelling distributions. Further work is required in this area, especially with respect to sensitivity modelling of key variables in the access network.
- The powering requirements of remote active equipment have been established and four possible powering systems identified.
- The generic functional requirements for an external ONU enclosure design have been defined and the issues of installation and maintenance addressed.
- A greater understanding of O&M costs, and the impact of O&M on QoS, has been obtained.
- A number of recommendations on how the O&M of new technologies should be addressed have been produced and a well documented framework to help PNOs do this.

Implications For Future Collaboration

The problem of how to upgrade the access network is very complex

because of the vast number of technology options and the number of business activities that are affected. The uncertainty in the strategies of PNOs with respect to broadband services and the uncertainties in market take-up make the decision on how to upgrade the current ubiquitous copper pair access network very difficult and very specific to a particular PNO. The focus of the cost modelling activity within this project has been on technology options. However, operations and maintenance are equally important and need to be fully accounted for in cost modelling, but this may be best handled within the PNO.

The upgrade to broadband could result in a divergence in the PNOs' networks and processes. Collaboration has the potential to de-risk this upgrade, but the goals, objectives and areas of study have to be carefully chosen to ensure that consensus will be of value to all participants.

BT has proposed a project (pre-study) within EURESCOM entitled: 'Co-operation in broadband capable infrastructure following de-regulation in the European telecommunications market'. The aims of this project are to establish the value of cooperation in broadband infrastructure in a deregulated market and pave the way for rapid agreement on broadband requirements at the appropriate time.

Acknowledgements

This document is based on results achieved in a EURESCOM project; this does not imply that it reflects the common technical position of all the EURESCOM shareholders/parties. The authors gratefully acknowledge the support of EURESCOM for carrying out this work. Thanks are also indebted to all the participants of Project 306 tasks 1,2, 4 and 7, and the project manager Wsewolod Warzanskyj since this article reflects the results of their work.

References

- 1 STERN, J. R. Optical wideband subscriber loops and local area networks in the UK Proc. ICC, Amsterdam, 1984.
- 2 ZAGANARIS, A., *et. al.* Methodology for risk assessment and techno-economic evaluation of optical access network, ECOC '94, 25-29 September 1994.
- 3 FAULKNER, D.W. and FITCH, P. Purpose and Benefits of EURESCOM Technical Advisory for a Fibre in the Loop TPN System. EFOC&N '94, Heidelberg, Germany, 21-24 June 1994.

Glossary

- ADSL** Asymmetric digital subscriber loop
- ASB** Asymmetric switched broadband
- ATM** Asynchronous transfer mode
- BPON** Broadband passive optical network
- CATV** Cable television
- ETA** EURESCOM technical advisory
- ETS** European technical standard
- ETSI** European Technical Standards Institute
- EURESCOM**
European Institute for Research and Strategic Studies
- FITL** Fibre in the loop
- FTTH** Fibre to the home
- FTTK** Fibre to the kerb
- HDSL** High-speed digital subscriber loop
- HDTV** High-definition television
- HFC** Hybrid fibre and coaxial cable
- ISDN** Integrated service digital network
- ITU** International Telecommunications Union
- O&M** Operations and maintenance
- ONU** Optical network unit
- PNO** Public network operator
- PON** Passive optical network

- POTS** Plain old telephone service
- QoS** Quality-of-service
- RACE** Research into advanced communications in Europe
- RSM** Remote subscriber multiplexer
- SSB** Symmetric switched broadband
- TA** Technical advisory
- TMN** Telecommunications managed network
- TPON** Telephony over PON

Biographies



Alan Quayle
BT Networks and Systems

Alan Quayle graduated from York University with an M.Eng. in Electronic Systems Engineering. Since joining BT, he has worked on high-speed opto-electronic integrated circuit design, represented BT in RACE EURESCOM, developed future TV/HDTV broadcast networks, investigated compressed multimedia transmission/networking and developed video-on-demand systems. Currently, he is developing advance fibre access networks and domestic customer networks.



Simon Fisher
BT Networks and Systems

Simon Fisher graduated from Liverpool University with a degree in Electronic Engineering in 1986. He worked in the design and manufacture of data transmission and multiplexer equipment, before joining BT in 1989. His work in BT Laboratories Access Transport Engineering Centre has been on the development and specification of hardware and software for copper and fibre optic systems. He is currently specialising

in the evolution, practical implementation and planning of BT's future access network.



Jerry Trigger
BT Networks and Systems

Jerry Trigger joined BT in 1985 after graduating from the University of

Manchester Institute of Science and Technology with a B.Sc. in Electrical Engineering. He worked initially on fault location techniques for the copper access network, with a particular interest in new opportunities presented by digital transmission systems. From this, his work broadened into considering the complete repair process and how line-testing technology might be used to reduce its cost. Jerry is currently responsible for a team designing new operations and maintenance processes and systems, with the dual objectives of enabling cost-effective use of the new network technology, such as TPN, and reducing the cost of operating the current network. He has been involved in EURESCOM Project 306, task 7 since it began in April 1994.



Dave Faulkner
BT Networks and Systems

Dave Faulkner has been involved in optical fibre local access networks at BT Laboratories since 1983. From 1986 to 1990, he carried out pioneering work on both narrowband (TPON) and broadband (BPON) systems. More recently, he has been involved in access network restructuring projects and the standardisation of FITL. His recent publications have shown how TDM over PONs offers the prospects of significant cost savings and service opportunities for both narrowband and broadband services.

A View of a Future Intelligent Network

This article presents a view of a future intelligent network (IN) as developed within the EURESCOM forum. The future IN consists of three components, a set of service creation environments, a set of reusable service constituents and a network resource model. These three components are shown to be sufficient to overcome current IN limitations.

Introduction

European telecommunications operators find themselves in an increasingly competitive telecommunications market. Witness the UK, where BT faces competition from cable companies, other licensed operators, such as Mercury Telecommunications, and mobile telephone companies. Each company strives to offer services that give them a competitive edge.

The introduction of new services in the past often required the re-programming of every switch in the network. Consequently, major European telecommunication companies, with their large networks, faced delays, sometimes of the order of years, in getting new services to market. As a result, new service offerings often missed their marketing window of opportunity.

In 1992 the ITU-T published a set of standards for an intelligent network (IN)¹. The essence of the IN is the removal of the service logic from the switches to an independent computer platform. This allows rapid service delivery, timed to meet market needs. The IN provides telecommunication companies one means of obtaining a competitive edge. Some IN implementations have already occurred. In the UK, customers can subscribe to some IN-like services, an example being the 0800 numbers, although they are not, as yet, implemented using IN.

Although the IN standards have made long strides in removing a major barrier to rapid service delivery, they still contained significant drawbacks. Firstly, services are modifications of a basic telephone-call model, limiting services to

telephony. Secondly, the reusable service components, the so-called service-independent building blocks (SIBs), are too big in functionality, limiting their reusability aspects. Finally, service creators still require detailed network knowledge since the standards do not specify how to map the SIBs to network-related functionalities.

This article describes work carried out under EURESCOM² project P103 'Evolution of the Intelligent Network', set up to address the problems described above. The aim of the project was to define a service-driven IN, not rigidly bound by present implementations and network technology. EURESCOM is a company, based in Germany, with the European telecommunication operators as shareholders. It aims to initiate, coordinate and supervise research and study projects in the field of telecommunications on behalf of its shareholders. Project P103 was one such project.

This article presents results of project P103. It does not, necessarily, represent BT's policy or view of the future IN or the service creation environment.

Overview of a Future IN

The fundamental concept of the IN is the removal of service functionality from the network switches. This allows service and network technologies to evolve independently from each other, enabling the coexistence of existing and future services and more rapid service creation.

Figure 1 shows a conceptual view of a future IN as developed within P103. The future IN has the service and the network domains separated.

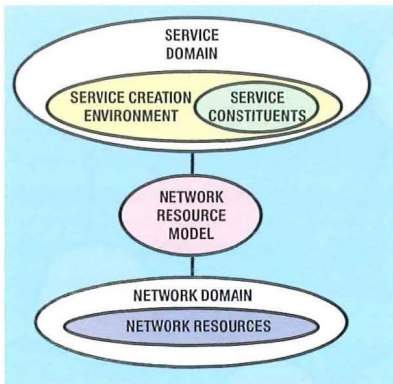


Figure 1 – Overview of a future Intelligent Network

The interface linking the two domains allows them to evolve independently using new technology as it becomes available. Current INs allow limited flexibility in incorporating new technology since the service and network domains are tightly coupled—services are built on the existing network functionalities, not network functionalities required to support the services.

The service domain has two components, the service creation environment, and a library of service constituents. The service domain allows rapid service creation without detailed network knowledge, permitting services to be built to satisfy market needs and meet customer expectations. The use of service constituents, rules for using them and a particular service constituent representing the network, overcomes the need for detailed network knowledge. Thus services are designed top-down—the network supports the functionalities of the services.

Since P103 was a collaborative and pre-competitive project, no attempts were made to construct a service-creation environment tool. However, the project did produce a generic service-creation environment model that the project participants can take and tailor to their specific needs.

Enabling Rapid Service Creation

This section describes the service-creation environment, how it enables rapid service creation and how it is modelled.

People have different interests and expectations from telecommunication




Figure 2 – Model of a service life cycle

services. For example, the customer's interest lies in using the service, the service provider wants to sell the service to customers, and the network operator provides the network. A service creator needs to take into account all these interests. The incorporation of these differing needs, when creating services, is termed the service-creation environment (SCE). The SCE creates services from pre-defined service constituents, ready for deployment on a target network. It allows the creation of new service constituents whenever services require functionality not covered by existing service constituents.

Building an SCE model starts from an understanding of how services progress from ill-formed ideas to being installed onto a network and ready for customers' use. The service life cycle provides an understanding of how services progress from initial ideas to final removal from the network. The service life cycle also contains a description of the people involved in service development. Project P103 has derived a description for a service life cycle, suitable for use within the telecommunications industry, from published descrip-

tions^{3,4,5}. Figure 2 shows the P103 service life cycle model. The service life cycle has three separate phases, creation, deployment and execution. Each phase contains a number of segments, called activities, containing detailed descriptions of the activities occurring within that phase. For example, part of the requirements-capture is to derive and structure the requirements for a specific service.

The next stage in building a model of the SCE is to develop an understanding of how the service life cycle impacts service creation. As shown in Figure 2, service creation results from taking a path through the service creation activities. This service creation path is unique to a particular company, as it reflects the technical and administrative culture within it. Hence the SCE should support an organisation's particular service-creation path.

A modelling technique applied to the service-creation phase of the service life cycle, that can link the activities within the phase, yields the SCE. The SCE will then contain both the processes of service creation and the people influencing that process. The modelling technique used, as

Figure 3—Role diagram for the requirements analysis activity

developed within the project, has three parts. Firstly, the application of an object oriented method (OORAM⁶—object oriented role analysis method) gives a set of role models. A role model is made up of a set of interacting objects. Within the role model each object has a defined *role* which describes the objects responsibility within the model. Secondly the resulting role model's dynamic behaviour is modelled using message sequence charts (MSCs)⁷. Finally, the application of SDL-92⁸ gives detailed specifications of the role model. These steps can be applied recursively.

Figure 3⁹ shows an example of a role model for the requirements analysis activity. It is important to keep in mind that a role describes what an object does, not how it does it. The following is a brief textual description for the roles in Figure 3.

Service analyser

A role that analyses the service requirements with respect to detail, completeness and consistency. When instructed, the role can prioritise and, or, generalise the requirements.

Requirements database handler

This role stores the requirements. Only the client is allowed to request modifications to the requirements set.

Network requirements analyser

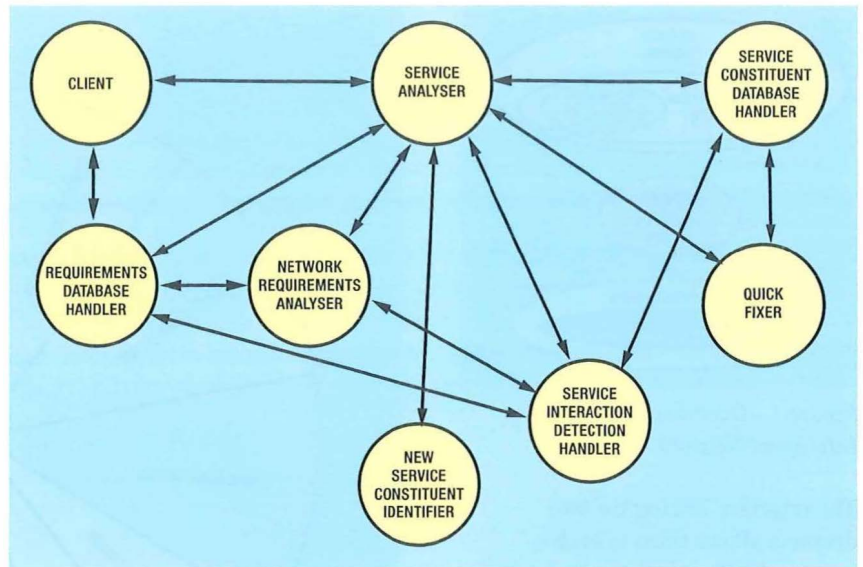
This role analyses the requirements with regard to the distribution platform the service will execute within.

Service interaction detection handler

A role that detects potential interaction of the proposed service with installed services. The role offers solutions by suggesting modifications to the service-requirements set.

Service constituent database handler

This role has information on service constituents in a form usable in the requirements-analysis phase.



New service constituent identifier

A role that generates a set of requirements for new service constituent(s).

Quick fixer

This role suggests modifications to the service requirements in order that existing service constituents can be used.

Client

A role that initiates and receives the results of requirements analysis.

Similarly, role models for the other service creation activities are derived. These role models represent some of the components to build an SCE. The combination of the role models, through a process called synthesis⁶, gives an SCE role model—other SCE models are possible from the same role models. A concrete SCE results from assigning persons, software, or machine, to play the roles in the role model. For example, a person plays the role of the Service Analyser in Figure 3 while a requirements analysis tool contains the other roles. In this way a service creation environment incorporates, in a common setting, the service creation process and the people influencing that process. Figure 4 shows an example of a possible future distributed SCE. The objects and roles which together make the SCE are identified in bold and normal text, respectively. The service provider is the person offering services. He makes use of people with the required expertise, identified by their roles, to create services to customer needs and perhaps also to

shorten time scales. Services are installed on the service provider's chosen network. In many cases the service provider and the network provider will be the same. The interested reader is referred to reference 10 that shows a role model for a service creation environment.

Service Constituents

The SCE uses a set of pre-defined and reusable service constituents for rapid service creation. A service constituent is defined as a role model, since it is a natural unit of modularity within the object oriented method employed within P103. Re-use of service constituents occurs in all service-creation aspects, from requirements analysis to designing and implementing the service. To this extent, the role models have additional information that encourages their reuse. This additional information includes the set of requirements they fulfil, their specifications, design and implementation details. By having this information, for example, it is possible to identify the service constituents required to design the service. Project P103 identified 13 service constituents that were needed to construct models of a multimedia conferencing service and a telephony-based card-calling service. A model for service-constituent storage can be found in reference 10.

Network Independence

At some stage, the service-creation process introduces the network

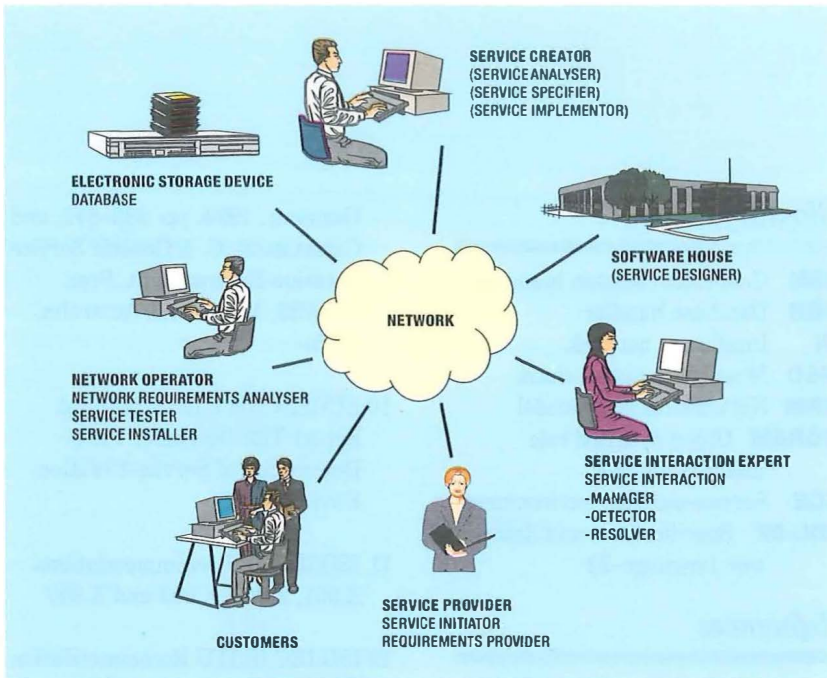


Figure 4—Simplified view of the role diagram for the service-creation environment

asynchronous transfer mode virtual connection or a physical bearer.

Using the Future IN

Consider a customer who requests a service that allows three, or more, people to hold a conference. The customer also requests that the service supports the use of a multimedia terminal, capable of supporting video and voice, as well as a telephone.

The SCE takes the above customer requirements and captures a more complete set of service requirements. The requirements are then analysed and specified. At this stage, consultation with the customer ensures the needs of his service. The next stage of service creation identifies the required service constituents to fulfil the service functionality. In this example, the required set of service constituents already exist, hence removing the need to design additional service constituents.

Next, the service is designed using the design tools and techniques applicable to the particular SCE. Using the NRM, the design stage of service creation extends the service design by incorporating the network related aspects. A simplified view of the designed service¹⁴ in which three users are connected to a network via terminals is shown in Figure 5. Users A and C have multimedia terminals while user B has a telephone. Each terminal is modelled as an object with an interface containing input and outputs for each data type (video and audio). The shaded grey region represents the underlying architecture and technology that bind the objects together. Any subsequent connectivity requests made by the service can be handled through the set of functions supplied by the network resource model¹⁵. The CSM object—connection session manager—supplies the view of the real network required by services.

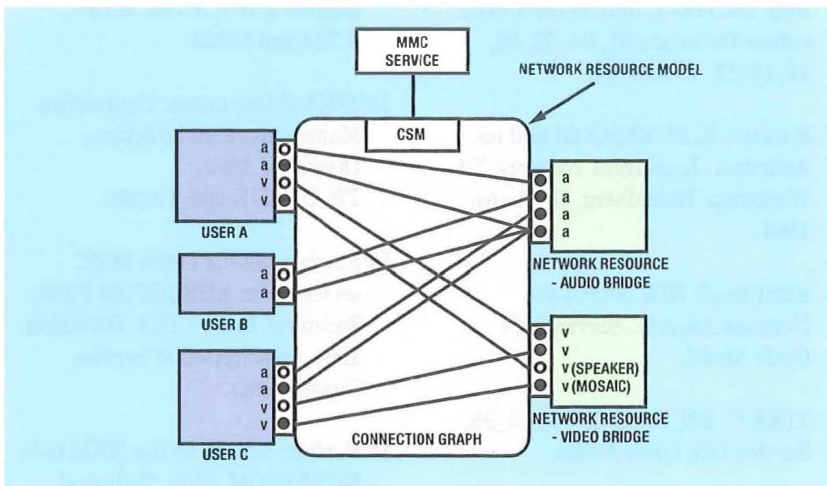


Figure 5—The role of CSM in presenting the network view to services

complexities into the service. To assist rapid service creation, service designers require a view of the network that hides its full complexity. P103 tackled this issue by specifying a special service constituent that allows services to use network resources and connections without the need for detailed network knowledge. Figure 1 illustrates this concept by showing the connection of the service domain to the network domain via the network resource model (NRM).

Using the NRM offers two main advantages. First, it provides services with a technology independent view of the switching and transmission infrastructure. Secondly, it offers a specified interface between services and the network, compliant with both existing and developing standards.

The NRM developed within P103 supports connections within a B-ISDN, since this type of network was anticipated to be capable of supporting advanced IN services. Within the NRM services view network resources, for example an audio bridge, as a set of abstractions—known as computational objects— with defined interfaces. The service needs these resources for its proper execution. Once objects have been defined, the service requests network connections between the objects. The service can then manipulate these connections as it requires. Note that a network connection can be viewed at various levels of abstraction. At the highest layer it is a set of objects and interconnections, and at the lowest layer it could be an

Conclusions

This paper describes a future intelligent network composed of a generic SCE model, network resource model and a set of service constituents. The three components are sufficient for the IN to achieve its main objectives, namely, a means to rapidly create and deploy services to meet market needs.

These components remove limitations of current IN implementations. Specifically, the SCE and the service constituents allow rapid service creation from a set of a re-usable software blocks within a setting supporting the service-creation aspect of the service life cycle. The service constituents allow the IN to cater for services other than just limited to modifications of a basic telephone call model. A particular service constituent, the NRM, provides service designers with a model of the network supporting the service in sufficient depth and complexity for them to incorporate into the service design. This aspect allows top-down service creation and allows the porting of the service logic to more than one network architecture.

Current INs are tightly coupled to the telephone switches. The future IN reduces this coupling. By modifying current intelligent networks to incorporate the ideas on the SCE, NRM and the service constituents, it is possible to extend current IN to cater for advanced services. Furthermore, this would also enable the IN to achieve its purpose, a means to rapidly create and deploy services.

Glossary

- CSM** Connection session manager
- DBH** Database handler
- IN** Intelligent network
- MSC** Message sequence chart
- NRM** Network resource model
- OORAM** Object oriented role analysis method
- SCE** Service-creation environment
- SDL-92** Specifications and description language-92

References

- 1 Q.12xx series of recommendations from the ITU-T, March 1992. (The values for xx are 01, 04, 05, 08, 11, 13-15, 18 and 90.)
- 2 KATZEFF, K. EURESCOM and its Activities. Intelligent Network '94 Workshop, Heidelberg, Germany, 1994.
- 3 ETSI Draft ETR NA601-09 Network Aspects. Service Life Cycle Model.
- 4 TINA-C, EN_B.MDC.006_1.3_93. Service Life Cycle Model.
- 5 EURESCOM P103, ST1.4.1 Report, 1992. Service Life Cycle.
- 6 REENSKAUG *et al.* Seamless Support for the Creation and Maintenance of Object Oriented Systems. *Journal of Object Oriented Programming* 1992, 6(6), p 27.
- 7 CCITT Recommendations Z.120.
- 8 FÆRGEMAND, O. and OLSEN, A. Introduction to SDL-92. *Computer Networks and ISDN Systems* 1994, 26, pp. 1143-1167.
9. For requirements capture and specifications role models, see, respectively, MUDHAR, P. Service-Creation Environment for a Future Intelligent Network. Second International Conference on Intelligence in Broadband Services and Networks, Aachen, Germany, 1994, pp. 333-342, and CAPELLMANN, C. A Generic Service Creation Environment, Proc. TINA'95, Melbourne, Australia, 1995.
- 10 EURESCOM P103, Technical Report TR3 December 1994. Description of Service-Creation Environment.
- 11 ISO/IEC/ITU recommendations X.901, X.902, X.903 and X.905
- 12 ISO/IEC IS/ITU Recommendation X.700, 1991. See also recommendations X.701, X.720, X.721, X.722 and X.723.
- 13 TINA-C Document. Connection Management Architecture. December 1993, TB_C_2.JJB.005_1.0_93.
- 14 Further details of the MMC service is in EURESCOM P103, Technical Report TR4, December 1994. Description of Service Constituents.
- 15 Further details on the NRM is in EURESCOM P103, Technical Report TR7, December 1994. Network Resource Model.

Biographies



Parminder Mudhar
BT Networks and
Systems

Parminder Mudhar graduated with a first class BSc (Hon.) in Physics from

the University of Lancaster, England, in 1982. In 1986 he gained a PhD from the University of St. Andrews, Scotland. He joined BT Laboratories in 1987 after a brief spell in academia. Since 1987 he has researched into photoluminescence properties of novel semiconductors at low temperatures, properties of semiconductor lasers at low temperatures and optical spatial switches for use in optical networks. In the last three years he has been involved in the EURESCOM project P103. He is now working on future advanced services.



Peter Glen
BT Networks and
Systems

Peter Glen graduated from the University of London in 1973, and obtained an

MSc in Telecommunications from Essex University in 1979. He joined BT Laboratories in 1973 where he worked initially on data transmission, covering transmission over the PSTN, ISDN local access and BT's Switched Star Cable TV system. He then moved into the area of broadband transmission. In 1988 he joined the RACE BUNI project carrying out studies on the specification of the broadband user/network interface. This was followed by the RACE MAGIC project developing signalling protocols for B-ISDN. In 1993 he joined the EURESCOM project P103 where he led the BT involvement in this project. He is currently working on high level design for the Concert network.



Carla Capellmann
Deutsche Telekom

Carla Capellmann studied Computer Science at the University

of Koblenz. After finishing her studies in November 1991, she joined Deutsche Telekom's research group, Functional Aspects of Networks, which deals with Intelligent Networks, Telecommunication Management Network aspects, and Formal Description Techniques. Her main interest lies in software engineering, especially methods for analysis and specification of telecommunication systems and services. In the last three years she has been involved in the EURESCOM project P103.



Ronald Janmaat
KPN Research

Ronald Janmaat received his M.Sc. degree in electrical engineering

from Twente University of Technology, Netherlands, in 1990. As a member of the technical staff at PTT Research he has worked on several projects in the context of Intelligent Networks, including service interaction, service specification and IN protocols. In 1994 Ronald acted as task leader in the EURESCOM project P103. He is currently managing projects in the areas of B-ISDN and TINA and Service Control at PTT Research.



Jørgen Nørgaard
Tele Danmark
Research

Jørgen Nørgaard has a MSc in computer science from

Aarhus University, Denmark. From 1989 to 1991 he worked for Nokia Telecommunications with object orientated (OO) techniques and applications in telecommunications software, using the notation of the emerging standard at the time, SDL-92. From Nokia he rejoined the computer science department at Aarhus University as a research assistant, working on co-designing and developing an open, distributed hypermedia, system on top of an OO database system and the X window system under UNIX. In 1992 he joined Tele Danmark Research. He has worked within the software engineering department at Tele Danmark Research on the RACE Cassiopeia and the EURESCOM P103 projects. In the EURESCOM project he was task leader for the activities on service creation and methodology work.

In this new regular column in the Journal, Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, gives some of his views on topical issues in telecommunications. In his first contribution, he describes how e-mail has enabled him to be a more effective and efficient manager and mentor.

The Great Socket Hunt— From 12 Days to 12 Hours by Peter Cochrane

During my early career, I worked in a government institution organised and run on classic Civil Service lines: rigid, over formalised and bureaucratic. The primary means of communication was paper, usually in triplicate, with the time needed to move letters and notes between departments typically 12 or more days. Why 12 days? For senior managers, simply dictating a letter, then having it typed, corrected and signed off could take two or three days. For junior managers, who had to use the typing pool, this could extend to over five days during busy periods. The internal and external mail system would then take a further two days, or more. If the original letter was dictated on a Monday morning, this meant that it was posted late on Friday, arrived the following Monday, and then the process began again in reverse.

Today many organisations achieve delays of less than three days, but this still seems excessive in a world of rapid industrial restructuring, company downsizing and market transformation. To say the least, the telecommunications industry is in the vanguard of this technically driven change process, which is being intensified by deregulation and competition. Getting results fast and minimising the time to market are now the name of the game. So three years ago I decided to move my entire 660 strong department onto electronic working. I did this with the dual promise that I would destroy all internally generated paper and

respond to any electronic communications within 12 hours. As I am blessed with an understanding family, a robust constitution that requires little sleep, and a determination to see just how far electronic working could go, I took this as a 24 hour, 365-day-of-the-year obligation on my part. Of course, people outside my organisation still write me letters and I likewise respond, but if their letter contains an e-mail address, I reply electronically.

So how has this all progressed? Statistically, the change from my old Civil Service days has been dramatic. My average response time to any communication is now about three and a half hours. I have replaced over half of my external paper mail with electronic communications, 98% of which—both internal and external—is completed within 12 hours. All my managers are on-line, and have access to laptop computers enabling them to communicate from all points of the globe. During the working day, the vast majority respond to communications well inside my 12 hour deadline. Their desire, like mine, was to find ways to work smarter—not longer and harder. Commercial pressures had already seen workloads increase and the working day expand. We had to rise to the challenge and find ways to become more effective—better communication was just one obvious ingredient. Overall, we have seen our operational performance improve dramatically with the adoption of electronic working.

Inevitably, however, statistics tell only the least interesting part of the story. For myself, the greatest benefit I have received from the 35–60 e-mail messages I process each day has been



a dramatic increase in the time I spend engaged in the oldest means of communication: talking. The number of letters I now write has fallen from an average of 12 a day to less than five a week. I also make fewer telephone calls—and most of those I do make are made on the hoof, from pocket or car telephone. True, I send far more e-mail than I used to send paper letters, but I spend less time doing it because it can be less formal, more terse and to the point.

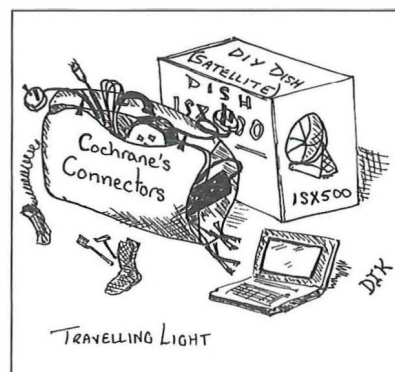
In reply to a page-long business case to buy equipment from David, my response would typically be: **D = GO. P**. In response to a group message requiring several diverse actions by Bill, Mike, Dave and Anne, my response—sent as a single e-mail copied to all four—might be of the form:

B = OK, do it, but take care. P
M+A = I have no idea, but Roger might. P

D = I suggest you buy one and try it. P

A = Can I have a full copy to read? P

E-mail does not suffer from decades of formalised ritual, it is new and experimental. You can do what you like; discover what works and what doesn't in your organisation. That is the real power. For me it not only cut the time I spend on formal communication, but it has also put me into direct communication with a wider range of people. With a widened span of contacts, I have a better foundation from which to tackle the most important part of my job: a dialogue with my customers and



colleagues devoted to understanding and directing BT's research effort. With less time spent on formal communication, there is more time to walk the floor and talk to people. As a result, I am now a more effective and efficient manager and mentor.

With more effective communication we were able to reduce the managerial hierarchy from four to two layers. It was also possible to empower everybody in the organisation to respond more quickly to the needs of customers and colleagues. Answers to questions, approval and agreement are only a message away. So decision-making can be immediate and cerebral reassurance is always on-line!

So what are the problems? Well e-mail can be too easy and too popular. Messages get copied to everyone. In some organisations, people receive over 300 messages a day, just sorting through them is impossible. You have to repel boarders if you are to avoid electronic overload, which requires both self-

discipline and management to discourage unnecessary communications. Living in a faster-paced world dictates the discipline of keeping up with communication, and adjusting to a business life that becomes more chaotic, less ordered and more opportunistic. Also, everybody has to remember that, even if you can reply quickly to a message, you don't always want to. Some replies need careful consideration if feelings are not to be hurt, or complex situations made worse.

And finally there is the problem of sockets. To try to honour with my 12-hour promise, I travel with screwdrivers, crocodile clips, a set of international connectors, and a nose for sockets. As an engineer, it has been fascinating to discover the number of different socket types and communications technologies used around the world—and as an executive it is immensely frustrating. For example, the same RG11 connector is used in Europe and the USA, but

with a different pin-out. The most recent development designed to thwart my efforts, is the installation of arbitration units in some North American hotels. These add long delays while some mysterious negotiation decides which carrier and circuit will carry my international call. From the point of view of my computer, this is just as bad as 30 years ago when international calls were all connected manually.

Perhaps in a few years, digital mobile radio will totally eliminate my socket hunt? I now have a GSM cellular telephone connected to my laptop, and within Europe I can roam from country to country. But the majority of my overseas travel sees me in North America, where the mobile radio standards are different. I don't expect to see a dramatic improvement in my performance as responding to 98% of messages in less than 12 hours is about as fast as I am likely to get for a while. Long live my screwdrivers and crocodile clips!

book review

Successful Marketing Strategy for High-Tech Firms

by *Eric Viardot*

The first thing that hits you when you read this book is that the author has obviously done his homework in terms of research and in gathering information from various professionals in very reputable high-tech companies, each chapter ending with a comprehensive reference list.

The book starts with various definitions of what marketing means and why successful high-tech companies do not necessarily have the best products but definitely the best marketing strategy. Coverage then includes corporate and marketing strategy including the need for a company to have a mission statement, the importance of defining and choosing strategic segments and the need to understand what the marketing plans for high-tech

products are, including situation analysis, target markets and monitoring procedures.

The book goes on to explain why you not only need to know and understand the markets you are in, including the needs/wants of customers, but also need to be able to analyse competitors and their activities. The next chapter looks at markets including defining, evaluating and selecting segments.

The second half of the book concentrates on product strategy. First it looks at the pricing of high-tech products, which includes understanding various pricing techniques and the basic economics of supply and demand. The remaining chapters cover the importance of a communication strategy (how you set a budget, corporate advertising, public relations etc.), distributing and selling high-tech products and the position of marketing within high-tech companies. The book concludes

with two useful appendices, one on the key success factors of a marketing department in a high-tech company and the other about how to prepare a market plan.

The book is well laid out with easy-to-use reference sections, but it tries to cover too much and as a result the content is thin in places. Hence, the title of the book is rather overambitious. However, for those people who are new to marketing and are working in a high-tech company, I would recommend this book as a useful starting point, and the references will certainly help by providing more detailed information to support the basic principles outlined.

*Published by
Artech House Books
£39.00. xiv + 184 pp.
ISBN 0-89006-770-8*

*Reviewed by Jenny Hawksworth
and Don Powell*

BT and MCI Provide Concert Services to Meet Grand Metropolitan's Global Communication Needs

Grand Metropolitan (Grand Met), one of the world's largest consumer goods companies, has chosen Concert's global communications services to meet their international voice and data networking requirements

Concert is the joint venture company set up by BT and MCI to provide seamless, consistent communication products and services to meet the needs of international business customers. Bill Brant, Group Information Services Director of Grand Met, said that the services offered through Concert put MCI and BT ahead of its competitors in providing genuinely global solutions.

Concert Virtual Network Service, a virtual private voice network, will be deployed in North America, Ireland, France and the United Kingdom. Phase two will expand this network into additional European countries as well as the Asia-Pacific region.

Concert Global Managed Data Services—low- and high-speed packet-switched and frame-relay services—will be available in additional countries including, Japan, Spain, Italy, Germany, Belgium, Holland, Portugal, Singapore and Australia.

Grand Met will use the networks to interconnect its global food and drinks businesses, which include Pillsbury, Burger King, Paddington, Haagen-Daz, Pearle Vision, Heublein and others. The seamless networks will allow Grand Met to manage better their disparate global businesses, enabling global stock control and financial analysis as well as establishing consistency of brand marketing around the world.

Michael Rowny, Executive Vice President of Ventures and Alliances for MCI, said: 'Concert was created in order to service the special needs of multinational corporations such as Grand Met.'

Bruce Bond, BT's Director of National Business Communications,

said that, originally, Grand Met was going to take a regional approach, using several large suppliers. But Concert changed all that. BT and MCI became the obvious choice for Grand Met because they were able to offer a global supplier framework that was a better fit for their needs than any other company.

Concert was formed in 1993 by MCI and BT to provide a wide variety of advanced global communication services to leading companies worldwide. MCI and BT, the first of the large carrier alliances to secure all regulatory and other approvals, began marketing its services last June. Under the terms of the joint venture, MCI and its distributors market Concert services in the Americas, while BT and its distributors market Concert services in the rest of the world.

New World Speed Record Set on Fibre Superhighway

Scientists and engineers at BT Laboratories have used photonic technology to shatter telecommunications world speed records on the 'information superhighway'.

They have designed a new type of telecommunications network based entirely on the transmission of laser beams through optical fibre, using devices that can almost completely eliminate conventional electronic circuits used in today's telephone exchanges.

The advances mean that the content of 50 full sets of the Encyclopaedia Britannica could be transmitted along a single optical fibre in a future telecommunications system in just one second—and that is utilising just one wavelength.

Key elements for this super-net of the twenty-first century are new photonic components invented and built at the laboratories. The new components include an 'optical transistor' that can process optical information at rates of 80 Gbit/s, and a high-speed optical switch many times smaller than existing designs. These are the basic building blocks used to construct the world's first all-optical telephone exchange that can read

information from a packet of data at a rate of 100 Gbit/s, decode the identity of the intended recipient and switch it onto the right path to its destination.

At the same time, BT Laboratories have demonstrated multi-wavelength transmission along single fibres, each optical fibre carrying several lightwaves, with each one capable of data rates of 100 Gbit/s. The scientists predict that 10 wavelengths are a practical proposition—a staggering data transmission rate of 500 'Britannicas' a second.

The Future of British Education on Trial in Bristol

BT and ICL have announced the launch of a joint trial of a world-leading interactive on-line education system in Bristol. It offers the chance to peer into education's future, in which open access to knowledge will supplement the more traditional resources, spurring a revolution in teaching practices.

The six-month trial is expected to go live at 11 schools in the Witherwood area of Bristol in the near future. Its objectives are to assess educational benefits as well as the practical applications of the technology. It is hoped that the system will encourage pupils to explore their curriculum and lighten teachers' administrative load.

The heart of the system is an education on-line network, which will complement teaching by allowing schools to draw on extensive educational resources, integrated learning systems and multimedia communications. It will be supported by the internationally-renowned School of Education at the University of Exeter.

The trial will involve more than 100 teaching staff and 2000 pupils in both primary and secondary schools. The educational benefits will be evaluated by the National Council of Educational Technology (NCET).

David Oliver, General Manager of ICL Lifelong Learning, commented that the need for a revolution in the use of technology in schools had already arrived. What was needed was to demonstrate a solution which

allowed schools to afford and manage the appropriate level of technology. It was essential that all pupils in the UK had access to IT at an early age to avoid creating a society of 'haves and have nots'.

The aim of the trial is to see how an educational on-line network can work in practice and be used beneficially by teachers and pupils.

The personal computers and printers will be supplied by ICL. These will be linked together within schools on local area networks. External wide-bandwidth links will be made to educational resources and, via BT's CampusWorld service, to selected areas of the Internet. Video enhanced desktop conferencing equipment and links will be part of the package, allowing dialogue between schools and remote teaching support by the University of Exeter School of Education.

Bid for Telecommunications Licence in the Netherlands

BT has announced that it has been chosen by NS Telecom, the telecommunications division of Nederlands Spoorwegen, the national state-owned railway company, as its partner in its bid for a national telecommunications licence in the Netherlands.

On obtaining a licence, the new company—Telecom 2—will address the market for data, voice and video services. It will also act as a distributor for Concert—BT and MCI's global networking company.

NS Telecom's national network currently extends to most major towns and cities in the Netherlands. On obtaining a licence, the companies intend to make a substantial investment in developing the network. Initially, Telecom 2 will address the business market; however, after January 1998 when the switched voice network is liberalised, it will investigate addressing the residential market.

Mike Grabiner, BT's European Director, speaking at the opening of BT's Netherlands global customer service centre in Amsterdam, said that BT saw Europe as its home

market and had already concluded alliances in four other European countries, as well as a number of distributorships. BT viewed the Netherlands as a strategically important country and was delighted to partner NS Telecom.

Syntegra Helps Prevent Social Security Fraud

Syntegra, the systems integration business of BT, has won a strategic three year contract to manage the Department of Social Security (DSS) Benefits Agency's new Generalised Matching Service (GMS). Worth £1.4M, the new system is of key importance to the DSS's highly publicised strategy to reduce fraud within the benefits system.

Syntegra's GMS enables the Benefits Agency to compare over 20 million separate DSS records, highlighting fiddles and multiple claims. Already one of the largest systems of its kind in Europe, the GMS was developed after extensive conceptual trials by the Benefits Agency and the DSS's Information Technology Services Agency. It will ultimately become the world's largest when the whole nation's national insurance records come on stream.

The Parliamentary Under-Secretary of State said that the new hi-tech system meant that any two income support claims could be compared within an hour, rather than ploughing through the old manual checking system. Even records from different benefit groups could be cross-checked within two hours.

BT Heralds Next Generation of Videoconferencing Systems

Natural meetings usually feature information exchange together with the conversation itself. In addition, senior people rarely want to spend time learning how to operate technology. These are the driving influences behind BT's latest videoconferencing room system called the VS3.

Adrian Butcher, General Manager of BT Visual Services, said

that the new system offered much improved ease of use and the ability to share data electronically, not just by showing documents under a camera.

The VS3 is the second step in BT's strategy to adopt a solutions-based, not product-based, approach to the market.

Based on the new VC2400 video codec, the VS3 is a comprehensive room videoconferencing system which will replace BT's successful VC5000 series. The VS3 has a new graphical user interface for simplicity and flexibility together with a host of performance enhancements. In addition, the VS3 will be BT's first group system to include compliance with the T.120 series of data interworking standards.

The VS3 and VC2400 offer a host of new features to their users. Improved video performance is provided by using pre/post processing, CIF and QCIF video. It offers a higher level of video compression for the transmission of video, high-quality audio and data over low-rate digital links. It can dynamically alter its parameters to maintain compatibility with other vendor's systems—including picture format, audio coding, data rates and frame rates. It includes JPEG still-image graphics transfer for higher resolution still-image transfer.

The VS3 Dual System features two video monitors—29 inch as standard but with the ability to support up to 35 inch monitors. The VC2400 codec is housed in a cabinet which is custom-designed to blend in with the office environment. The whole system, including the two high-resolution cameras, is controlled by a graphics tablet with an electronic pen or hand-held controller.

A document scanner allows the presentation of scanned images into the conference and the VS3 can store up to 10 images. In addition, VS3's document camera will display high-resolution graphics for discussion. As new functions and peripherals are added to the system, extra icons can easily be added for continued ease of use. Additional peripherals include printers, hand-held cameras and video recorders.

OFTEL Publishes Market Information

Don Cruickshank, Director General of Telecommunications, has issued a new edition of OFTEL's ground-breaking publication 'Market Information' giving facts and figures on the UK telecommunications market. All the figures are collected directly from the operators—over 30 fixed link and the four cellular network operators submitted information. This second edition includes figures for the 1994/5 financial year and more detailed information on most market segments.

The Director General said: 'This publication is an important part of OFTEL's drive towards openness in the telecommunications market through the publication of high-quality information. With reliable information, both consumers and those within the industry can make better informed decisions and get a better deal.

'The data in the new edition cover £12.9 billion (about two-thirds) of the UK market which is currently worth about £19 billion per annum. BT's overall share is just under three-quarters and 82 per cent of the market segment covered in this publication. For the first time there are quarterly figures giving a better idea of trends, separate information on the cable operators and a wider range of figures on the cellular networks. In future we will extend the range of data collected further to include 'value-added' services, and we will begin to publish a selection of key indicators quarterly.'

Proposals for Competition in Directory Enquiries

Don Cruickshank, Director General of Telecommunications, has published proposals to boost competition in directory enquiries and other directory-based services. The proposals are set out in a consultative document—'The Use of Directory Information'.

The Director General said: 'I am concerned that eleven years after the privatisation, BT is still the virtual monopoly supplier of directory products and services. It has been put to me that this is damaging to the interests of consumers and is inhibit-

ing the growth of competition in telecommunications generally. I am therefore consulting on proposals which will encourage competition in the running of directory information databases from which services and directory products are derived.

'However, I am determined that we must maintain the public's confidence that directory information will not be misused. Customers have a right to expect that levels of privacy of information will be safeguarded and my proposals allow for strict control of any organisation having access to sensitive directory information. I particularly look forward to receiving the views of consumers and representative consumer bodies on whether the proposed restrictions, along with data protection law, will provide sufficient protection of consumers' interests.'

European Commission Telecommunications Directory Guidelines

Telephone users would be guaranteed access to a directory of subscribers, under a communication setting out guidelines on the directories market put forward in October. Directories form an important sector, representing 7.5% of the European Union (EU) advertising market, and directory services play a key role in telecommunications services. To the extent that they are both a product and a service, directories are subject to a range of provisions under EU law. As well as the universal directory requirement, the guidelines include provisions to protect privacy and intellectual property rights, conditions for access and marketing, the promotion of new technologies and the abolition of the exclusive directory rights that exist in some Member states.

C&W's Cable Ship

Cable & Wireless Marine (CWM) has launched *Cable Innovator*, a radical new approach in cable ship technology, which will provide customers with an improved level of service.

Cable Innovator, as the largest purpose-built cable ship in the world,

is fully stern operating—the most radical new design for 30 years.

The most significant advantages of a fully stern-working cable-laying ship are that it can travel faster to the site as it is not slowed down by conventional bow sheaves, it can operate in weather that traditional cable ships cannot withstand, and it provides optimum protection from severe weather conditions for all cable handling operations.

'As the worldwide demand for cable capacity increases, so does the customer's need for greater efficiencies in cable management. When we started thinking about designing a new ship, we went out and talked to customers about what they wanted. *Cable Innovator* is the result.' said Cable & Wireless Marine (CWM) Chairman, Edward Astle.

European Commission Satellite Proposals

In November, the European Commission adopted proposals on satellite personal communications services (PCS). The PCS sector, which uses low-earth orbit satellites, is expected to expand greatly and offers significant opportunities for European industry. The Commission's paper seeks to ensure that licensing arrangements use the limited frequency resource efficiently and that Europe adopts a coherent approach. As a first step, the Commission is to publish a call for information, addressed to those planning to provide PCS.

Opening Up Cable TV Networks

Interactive telecommunications services, including education packages and teleshopping, should become more widely available as a result of a Directive adopted by the European Commission. Cable TV networks are able to carry these and other multimedia services throughout the European Union from 1 January 1996. The Directive requires Member States to lift restrictions on the use of cable networks, in order to encourage investment and foster pilot projects and new initiatives in this field.

Existing national regulations in some countries still limit cable networks to carrying simple, one-way TV broadcasting, thus preventing operators from offering any of the new switched (interactive) multimedia services. The new Directive amends the 1990 telecommunications services Directive (90/388) and has been adopted after consultation on the draft text.

OFTEL on International Direct-Dialled Calls

Don Cruickshank, Director General of Telecommunications, has published a consultative document on the future basis of charges for interconnection for the international direct-dialled calls (IDD). These are the charges paid by other UK telecommunications operators to BT and Mercury for the handling of international calls which originate on other operators' networks.

Introducing the document, the Director General said: 'Customers in the UK, as elsewhere, have been paying over the odds for their international telephone calls. The benefits of competition in infrastructure have not yet reached this part of the telecommunications facilities. That is a decision for government.

'However, some progress can be made next year, and the question is whether I should moved closer to a fair deal for customers via an interconnection determination from April 1996 onwards on the basis of net accounting. This method of charging would enable other operators to share with BT and Mercury the profits from international telephony and would encourage greater competition in this market.'

Removal of RPI+2%

The Director General of Telecommunications, Don Cruickshank, has announced that he is proceeding to statutory consultation on proposed modifications to BT's licence to remove the current constraint limiting any increase of BT's exchange line rental charges to RPI+2% a year.

Introducing the document setting out and explaining the proposed modifications, the Director General

said: 'I see the removal of the RPI+2% line rental constraint as a deregulatory step which will be good for both customers and the competitive framework in the UK. I believe the change will be broadly commercially neutral for BT within the present price control period ending in 1997.

'The RPI+2% constraint originally protected customers, in the absence of competition, from significant rises in rental charges. Now, however, the domestic market is seeing substantial competition. Further competition will come when radio-based operators launch their services.

'I feel, therefore, that it is now an appropriate time to lift the constraint. This would give BT freedom, if they chose, to adopt a much more imaginative approach to the way they price services, particularly for residential customers. I believe BT would use this freedom to introduce a family of price packages. Customers would choose which package suited them best. BT has told me that it would expect to structure any such package so that no customer, for a given level of usage, would see any real increase in their bill. Most would see reductions.

'I do, however, need to be satisfied on a number of issues before the launch of any packages:

'Customers would need enough time to understand what was being proposed. I would expect BT to have a major media and television campaign to promote customer awareness of any shift to price packages, advise customers of the details of the prices on offer and explain how to decide what would be the most advantageous package for them.

'The overall RPI-7.5% price cap will continue and BT will still have to make overall price reductions of around £400 million per annum under the current controls.

'Customers eligible for the Light User Scheme (approximately 20 per cent of all residential customers whose bills are the lowest) will be guaranteed no real increase in their bills for the same usage.

'Another indirect benefit of removing the RPI+2% restraint is the demise of the complex Access Deficit Contribution regime under which

competing operators pay BT contributions to the accounting deficit it reports on its access business. This has made the UK regulatory regime very complex and uncertain and has placed disproportionate discretion with the regulator. OFTEL will be glad to see the back of this—it has not proved a helpful regulatory measure.'

RACAL Proposed Acquisition of BRT

Racal Electronics plc has announced it has entered into a conditional agreement with the British Railways Board to acquire the immediate holding company of BR Telecommunications (BRT) for which Racal will pay £132.75 million. In view of the size of the acquisition, it is conditional upon the approval of Racal shareholders at an extraordinary general meeting.

BRT's business principally consists of the provision of telecommunications services to the railway and employs approximately 2800 people, based at some 180 locations nationwide.

BRT will become a separate subsidiary company within Racal Network Services group, managed on a fully integrated basis with the other businesses.

France Telecom and Deutsche Telekom Receive US Approval

The US Federal Communications Commission (FCC) approved the global alliance among France Telecom, Deutsche Telekom AG and Sprint Corporation on 15 December 1995. France Telecom and Deutsche Telekom are thereby permitted to invest, on a 50/50 basis, a total of approximately \$4 billion for a combined equity stake of 20 per cent in Sprint, the third-largest US long-distance carrier.

The FCC approval came on the same day that the European Commission (EC) published in the Official Journal of the European Communities its decision to authorise the creation of the France Telecom/Deutsche Telekom Atlas joint venture and its global alliance with Sprint.

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.

© 1995: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Price: Single copy—companies, universities, libraries and other bodies £9.25 (£10 including postage UK, £11.25 including postage overseas); private individuals (paid by personal cheque) £5.00 (£5.75 including postage UK, £7.00 including postage overseas). Annual subscriptions (including postage)—companies etc. £40.00 UK, £45.00 overseas; private individuals (paid by personal cheque) £23.00 UK, £28.00 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: +44 171 356 8050. Fax: +44 171 356 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: +44 171 356 8050. Fax: +44 171 356 7942.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured 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/95 \$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.

Centre	Local Secretary	Membership Secretary
Aberdeen	Steve Murray (01224) 571726	
East Midlands	Dave Bostrom (01908) 656215	Ian Bethell (0115) 947 8587
Lancs & Cumbria	Trevor Hughes (01706) 40840	
London	Jim Skinner (0171) 843 8552	Peter Martin (0171) 492 6138
Manchester	David White (0161) 600 2468	Mike Ku (0161) 833 2669
Martlesham Heath	Peter Shaw (01473) 224488	Nisha Moudgill (01473) 224801
North Downs	Nick Smith (01227) 474823	
North East	Richard Jackson (0191) 496 0251	Frank Balloch (0191) 261 3226
North Wales & Marches	Peter Clay (01743) 274301	
Northern Home Counties	Terry Birdseye (01702) 373330	Stephen Graham (01753) 564179
Northern Ireland	Allan McDougall (01232) 215301	Brenda Sharvin (01232) 320092
Paris	Delphine Lalire +33 1 46 67 27 89	Terry Ferniot +33 1 46 67 27 54
Scotland East	Graham Neilson (0131) 668 5878	Keith Doig (0131) 668 5315
Sevenside	Mike Stenning (0117) 920 5691	Dean Shephard (0117) 920 6457
South Downs	Graham Lindsay (01273) 222633	
South Wales	Peter Coleman (01222) 514524	Keith Voyce (01222) 555075
Staffordshire	Ray Piper (01785) 813483	
West Midlands	Richard Harwood (01527) 550252	Paul Jervis (0121) 230 4181
West Scotland	Ian McReady (0141) 220 7118	Jill Maxwell (0141) 220 5837
Westward	Rob Rand (01392) 267114	Chris Gould (01392) 267113
Yorks & Lincs	Steve King (0113) 237 8124	Paul Horncastle (0113) 237 6542

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.

Zone Committee	Contact	Telephone No.
London	Terry McCullough	(01707) 601151
Midlands	Mark Broughton	(0116) 242 3811
North East	Geoff Jenkinson	(0113) 237 8163
Northern Home Counties	Graham Lovell	(01473) 646056
Northern Ireland	Brian Walker	(01232) 232413
North West	Phil Holt	(0161) 339 1971
Scotland	Raymond MacMillan	(01463) 665913
Southern Home Counties	Martyn Harvey	(01273) 568388
South West	Roger Garland	(0117) 920 5485
Wales	Ferg Brown	(01952) 249886

IBTE International HelpLine — +44 171 356 8008



GLOBECOM 96 *London*

GLOBECOM '96

CALL FOR PAPERS



IEEE GLOBAL TELECOMMUNICATIONS CONFERENCE

"COMMUNICATIONS: THE KEY TO GLOBAL PROSPERITY"

18-22 November 1996 • Queen Elizabeth II Conference Centre • Westminster, London, UK



IEEE
COMMUNICATIONS
SOCIETY

Sponsored by the Institute of Electrical and Electronics Engineers, Inc. (IEEE) Communications Society and IEE

General Chair

Dr Alan W Rudge
BT - Deputy Group
Managing Director

Executive Committee Chair

Dr Tom Rowbotham
BT

Vice Chair

Dr Bob Blake
BT Laboratories

Technical Committee Chair

Professor Hamid Aghvami
Kings College London

Finance Chair

Dr Gerry White
BT Laboratories

Local Arrangements Chair

Professor Chris Todd
University College London

Secretary

Mr Geoff Richman
BT Laboratories

Executive Committee Office

Mrs Vikki Pollard
BT Laboratories

Publications Chair

Professor Trevor Clarkson
Kings College London

Publicity Chair

Dr Michael Miller
BT Laboratories

You are invited to submit original papers addressing topics of interest for presentation at the conference and inclusion in the conference proceedings. Submissions should be in-depth, technical papers describing recent telecommunications research and development results.

TECHNICAL SUBJECT CATEGORIES

- Network Protocols and Architectures
- Software Systems, Technologies, and Tools for Network Platforms, Services and Interworking
- PCS, PCN and Mobile Systems and Networks
- Radio, Satellite, and Infrared Systems
- Modulation, Coding and Synchronisation
- Modelling and Simulation Techniques
- Network Management: Testing and Modelling
- Security and Privacy
- Video and Audio Compression and Coding
- Signal Processing for Transmission or Recording
- VLSI for Communications
- ATM Systems and Networks
- LANS, MANs and WANs
- Testbeds and Trials of Networks and Services
- Multimedia and Broadband Services, Applications and Development Tools
- Optical Systems and Networks
- Buses and Channels for High Speed Interconnection Within Digital Systems
- Routing and Control of Congestion, Admission and Flow
- Intelligent Networks and Network Operations
- Quality Assurance Management
- Switching, Transmission and Signalling
- Network Operations and Management & Distributed Network and Systems Management
- Personal Communications Services Management

SUBMISSION GUIDELINES:

The title page of your submission must include

- 1) the name, complete return address, email, telephone and fax numbers of the author to whom all correspondence will be sent;
 - 2) a 100-word abstract;
 - 3) the designation of the technical subject category (see listing) to which the paper is most closely related.
- All other pages should be marked with the title of the paper and the name of the first author.

IMPORTANT DATES:

Submissions due: March 1, 1996

Acceptance Notices mailed: July 5, 1996

Camera ready papers due: August 23, 1996

SUBMISSIONS TO:

Send five double-spaced copies of the manuscript (limited to 3000 words) in English to:

*Professor Hamid Aghvami
Technical Programme Committee Chair, Globecom'96
Kings College London
Department of Electronic and Electrical Engineering
Strand, London WC2R 2LS, UK*

Tel: +44 171 873 2898

Fax: +44 171 836 2664

Email: h.aghvami@kcl.ac.uk

GLOBECOM'96 GENERAL INFORMATION

Contact Vikki Pollard, BT Laboratories.

Email: pollarv@btlip10.bt.co.uk

Tel: +44 1473 644799 • Fax: +44 1473 647438

WWW URL <http://www.labs.bt.com/profsoc/globecom/>

ATTENTION STUDENT

AUTHORS / PRESENTERS:

You may qualify for a Student Travel Grant to defray some of your costs for presenting a paper at this Conference. IEEE Communications Society will award a limited number of Student Travel Grants to full-time engineering students. To qualify you must:

- ✓ Be a member of IEEE (or have applied for membership);
- ✓ Live outside of Europe, Africa & the Middle East (Region 8);
- ✓ Be an author who will present an accepted paper at GLOBECOM;
- ✓ Not have received a similar grant in the past 24 months; and
- ✓ Provide all of the information requested on the Student Travel Grant Application you will receive in your Author Kit, once your paper is accepted.

Grants consist of a check for \$800 (U.S.) plus a copy of the *Conference Record*. IEEE Communications Society Student Travel Grant Applications for this conference must be submitted by 23 August 1996.

Contents

Guest Editorial: Towards the Information Age	273
Convergence—Synthesising a New Industry	274
Electronic Commerce	281
Service Providers for Electronic Commerce	286
The Web of Circumstance	291
The Fourth Information Revolution and its Impact on Some Public Information Services	294
Digital Signatures: A Solution for Security in the Information Society	301
Community Networks	310
Image Processing for Telecommunications in the 21st Century	315
The Virtual Science Park	322
Transport Network Architecture	330
Usability Engineering of Speech-Based Services	337
Access Network Support Systems—An Overview	341
Access Networks Evolution and Preparation for Implementation—EURESCOM Project 306	345
A View of a Future Intelligent Network	356

