

VOLUME 17 PART 3 OCTOBER 1998

BRITISH TELECOMMUNICATIONS ENGINEERING



Included in this Issue

*Customer Service—The Key to
Satisfaction and Loyalty*

Future Switched Network Strategy

Powering the Internet



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 17 ■ PART 3 ■ OCTOBER 1998

Guest Editorial: Customer Service Delivery	113
Ian Smith	
<i>Customer Service</i>	
Customer Service—The Key to Satisfaction and Loyalty	114
Megan Brown	
<i>Structured Process Improvement</i>	
Strategic Resource Planning for Optimum Service Quality	120
Gail Lochtie, Andy Noble and Mark Gilbert	
Implementing Software Tools for Planning Automation	126
David Butler	
<i>Optical Technologies</i>	
Micromachining Future Networks	135
Graeme Maxwell	
<i>Telecommunications in the 21st Century</i>	
Broadband Multicast on BT's Futures Testbed	139
Chris Gibbings, Mark Barrett, Margarida Correia, Uma Kulkarni and Kevin Smith	
Engendering Trust in Electronic Commerce	150
Tim Rea and Peter Skevington	
Whither Video?—Television and Telepresence	158
Alwyn Lewis and Graham Cosier	
The Future World of Telecommunications, 2005? 163	
Chris Russell, Maff Holladay and Pete Barnsley	
Future Switched Network Strategy—Profiting from an Uncertain Future	172
Roger Ward	
World-Class Network Power—Could That Ever Mean Us?	180
Sarah Bennett, Bernie Johnson and Peter Howell	
Powering the Internet	
Datacommunications Equipment in Telecommunications Facilities:	185
The Need for a DC Powering Option	
John Parsons, John Åkerlund, Chris Riddleberger, Bo Lindemark, Dan McMenamin, Seiichi Moroyama, Katsuichi Yotsomoto and Jacques Poulin	
A People and Information Finder for Virtual Construction Project Teams	198
Alistair Duke	
37th European Telecommunications Congress	204
Tim Jeanes	
Journal Awards for 1997/98	207

Cover Picture

BT Corporate Picture Library;
a BT photograph

Regular Features

pcochrane@btlabs	209
Field Focus 210	Book Reviews 210
	Telecom Focus 212
	Information 216

Board of Editors

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

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

Assistant Editor Richard Kemp, B.SC., C.ENG., M.I.E.E. Roger Blake, C.ENG., M.I.MECH.E. Jason Chilton Martin Evans, B.SC., PH.D.
David Greenop, B.SC., M.SC. Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E, M.I.O.M. Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT.
Keith Ward, C.ENG., F.I.E.E., M.I.MGT. David White, M.I.P.D., M.B.I.M., L.I.A.A.

Theme Editors

Customer Service Mark Jakeman **Structured Process Improvement** Keith Beales **Optical Technologies** David Smith
Telecommunications in the 21st Century Steve Sim

Ian Smith

Customer Service Delivery



Thankfully, the vast majority of service we deliver goes like clockwork and within target.

On 1 October 1998, BT created one of the largest customer service organisations in Europe by bringing together call reception for residential and business customers, operator services and field engineering for the access network. As Managing Director of this newly created division I will have within my responsibility the end-to-end service delivery processes for provision and repair as well as responsibility for managing over a billion customer contacts per year. Anyone familiar with the world of customer and field service will soon recognise the benefit

of managing the full *cycle of service* as one process through one organisation.

Managing processes as cycles means starting with customer need and moving through a series of *moments of truth* until completing the process with confirmation that the customer need was satisfied. The learning experience from such instantaneous feedback contributes to reduction in error and process improvement in such a way that both customer and employee involvement and satisfaction can be considerably enhanced. By developing a choice of robust yet flexible processes the business can offer what feels like customised service delivered through standard, low-cost tasks.

The key to getting process development right, however, is always to look at your delivery mechanisms through the eyes of your customers. From that first contact, where *active listening* is the key skill required, through *keeping customers informed* of progress and a determination to *meet your promise* to deliver the result and concluding with *confirmation of completeness* and *customer satisfaction*, the customer's perception of the process is the essential measure. In an ideal world, these simple processes meet the expectations of customers and can, indeed, exceed their expectations, given the performance of service companies in the UK at present. However, on those occasions when all does not go to plan the system needs to be as slick at *recovery from failure* since one of the peculiarities of service is that customers can be more loyal to companies after failure than before, if they are dealt with correctly.

Accepting that not everything will go right on the day is one of the first rules of customer service and by anticipating all eventualities, and standardising the way in which you respond, service providers can

significantly reduce their *cost of failure* and prevent the disruption caused by panic! *Mandatory escalation* of problems, based on time or criticality, ensures failing to deliver to expectation is visible within your company in proportion to the disruption within a customer's household or business.

Thankfully, the vast majority of service we deliver goes like clockwork and within target. However, even after completing the delivery process, be it flawless or problematic, a *review of performance against objectives* is most advisable for any company committed to continuous improvement. Yet the process of selecting objectives is in itself fundamentally important to the success of the organisation, and this is where I return to how I started. By seeing your business through the eyes of your customers you are less inclined to use internal measures as objectives and more inclined to seek out what is truly important to customers, and use their perception of how you perform as the driver of reward, recognition and improvement.

The opportunity to manage end-to-end processes is not often available to people in large companies. So you can imagine how excited I am by the challenge of creating, modifying and delivering the provision and repair processes for BT in support of our ambition to be recognised as a world-class customer service provider.

Customer Service is the theme for a new series of articles featured in *British Telecommunications Engineering*, which begins on page 114 of this edition with the article 'Customer Service—The Key to Satisfaction and Loyalty.' I hope you will find these articles interesting and informative.

Ian Smith

Managing Director
BT UK Customer Service

Megan Brown

Customer Service—The Key to Satisfaction and Loyalty

Customer service delivery and its impact on customer satisfaction has led to the drive for service excellence for BT's business customers. This article considers the wealth of information on BT's performance and what is being done to influence customers' perceptions, to improve performance, and why.

Introduction

Competition has intensified and created greater choice; it has educated customers to expect greater value—expectations have risen and customers are both more demanding and less tolerant.

Product innovation and competitive response are the norm in today's business environment. As all companies strive to keep one step ahead of their competitors it has become the human factor, the relationships, that provides the key differentiator. Within BT, over time, the products and service offerings have changed and improved to meet customer requirements. It is now the service element (delivered by BT's people) that makes the difference and moves customers from being merely satisfied to very satisfied, or indeed completely satisfied, which in turn generates loyalty.

Satisfaction and Loyalty

In considering what constitutes loyalty, Pearson¹, writing about brands, suggests 'loyal customers stay longer and buy more'. He argues that 'loyalty means the propensity of customers to behave in the face of competition and choice so as to maximise lifetime value. A loyal customer has a high propensity to choose the brand, irrespective of competition because of the unique values associated with the brand and the customer's continuing relationship with the company that produces the brand.'

Interestingly, Pearson goes on to discuss the concepts of customer satisfaction and loyalty. The evidence in most industries is that satisfaction predicts future purchase only at the

very top of the scale—only the very satisfied have a high propensity to repeat purchase. World-class customer service is a vital part of BT's vision and recent research both within and external to BT has confirmed that superior customer service is a prerequisite of customer loyalty, leading to revenue flows.

Key Drivers

Ongoing research and measurement of BT business customer satisfaction and its key drivers is conducted. The key objectives are to provide an overall satisfaction score, to understand the key drivers of satisfaction and loyalty and to link quality of service, image, price and value for money in an integrated framework. A sample of some 2000 customers continues to be surveyed each month. The survey asks customers to evaluate BT from a number of different perspectives that include their experiences of sales and service as well as their image and perceptions of BT.

To determine the relative contribution or 'drive' of each of these areas on customer satisfaction and loyalty a model was constructed based on a derived rating of importance of each of these areas. The key drivers model uses average scores where satisfaction ratings are given out of 10.

The key findings are illustrated in Figure 1.

For all business customers, of the four major drivers of satisfaction (quality of service (QoS), image, value and price), QoS has the most significant impact on customer satisfaction. Figure 1 also indicates that overall QoS is dominated by repair and provision—hence BT's particular focus in delivery in these areas.

Figure 1—Key drivers of customer satisfaction

An interesting finding from the research was that customer satisfaction does not have a big direct impact on customer loyalty at all levels. Customer satisfaction impacts loyalty only when customers reach the 'very satisfied' level. Table 1 shows that moving customers from the satisfied to very satisfied categories has a definite impact on loyalty (31–37% change).

The overall conclusion taken from this research is that customer satisfaction is a reliable reflection of

Table 1 Satisfaction: Impact on Loyalty

	Fairly Loyal	Very Loyal
Satisfied Customers (7+)	36%	9%
Very Satisfied Customers (9+)	73%	40%

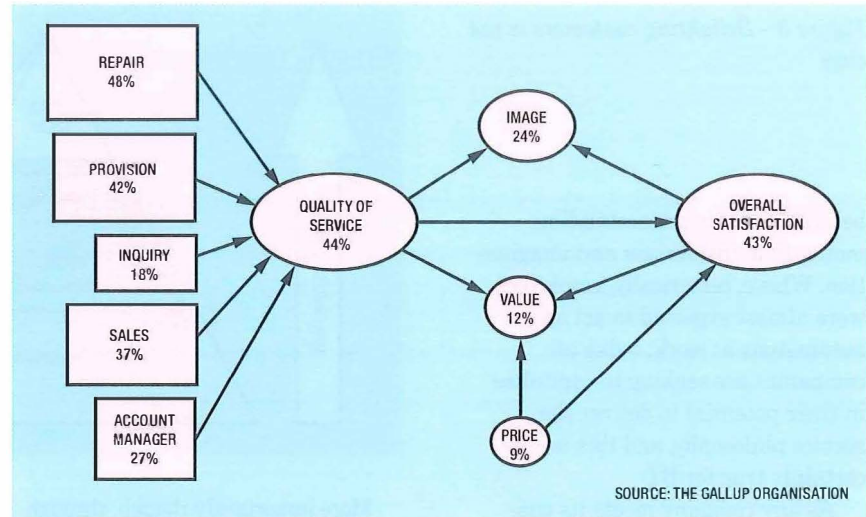
the perceived QoS provided by BT and that it impacts loyalty when customers become very satisfied, and when the level of service exceeds expectations. Mere satisfaction alone is not sufficient to guarantee customer loyalty. The most important influencer of satisfaction is the QoS, contributing some 60% of the total influence on satisfaction, and which has an indirect, but significant, influence on loyalty.

Revenue Implications

More recent research has also suggested that the level of customer satisfaction with BT not only has a direct impact on customer loyalty but also has a direct impact on customers' propensity to spend with BT. This relationship is shown in Figure 2, which illustrates the effect on incremental revenue at each satisfaction level.

'The absolute fundamental is to make money out of satisfying customers'

Sir John Egan, Jaguar.




Thus as customers move from satisfied to very satisfied the percentage change in the revenue opportunity for BT increases. Conversely customers' revenue streams into BT decrease dramatically at the dissatisfied level.

The key points are:

- Satisfying customers is no longer enough, averaging 7 out of 10 will result in a loss in revenue.
- By making half of customers 'very satisfied', BT could potentially increase its revenue by more than 10% over the next 5 years.
- If all BT customers were 'very satisfied' the potential increase in revenue over the next five years is almost 20%.

At present around four out of 10 customers are 'very satisfied' overall with BT. Because it is recognised that

Figure 2—Customer satisfaction the key



true customer delight is so important BT is now focusing key operational measures—for example, repair, provision, call handling—on completely satisfied (10 out of 10). The game is being raised because it is a commercial necessity.

Every contact between a company and a customer is an experience, and the effect of every positive experience is to increase loyalty. In marketing, advertising and customer service are usually considered as separate activities. They are in reality both aspects of building the brand by enhancing customer relationships, and every customer contact influences loyalty. Building excellent customer relationships is vital to loyalty because they enable the company, through dialogue, to identify and anticipate customer needs before the competition.

Mitchell² writing for the Chartered Institute of Marketing (CIM) suggests more and more companies are recognising that brand promises are being delivered by people rather than products. Thus BT must strive to ensure that customer service is not just skin deep and that all BT people understand this concept so that they may empathise and deliver the quality of service its customers' demand.

Marketing's roots lie in product marketing: its ability to identify changing and emerging needs, to translate this into a better product or solution and then to tell the world about it. Increasingly the trend is towards service as the differentiator. The focus for competitive advantage is shifting away from the product and its presentation towards the people

Figure 3—Delighting customers is not easy

behind it—their understanding, motivation, inspiration and imagination. Where, historically, employees were almost expected to act as automatons at work, today all companies are seeking to capitalise in their potential to deliver the service philosophy, and this is certainly true for BT.

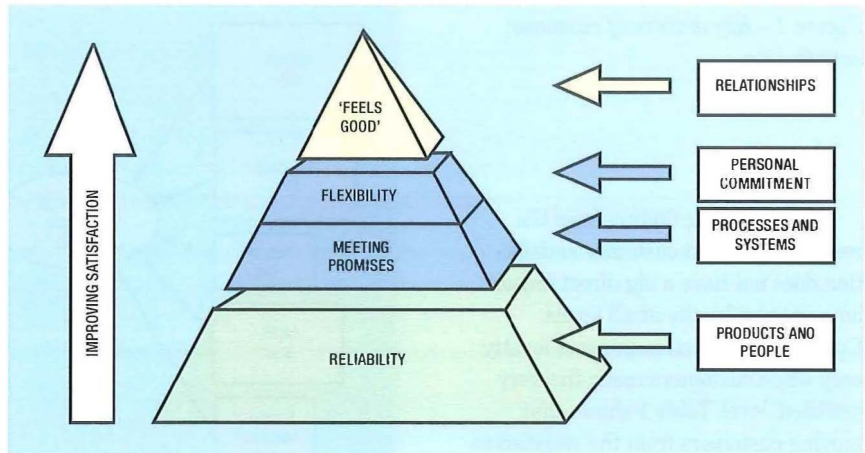
As any company meets its customer requirements so their requirements and expectations increase and what was once innovative and exciting becomes commonplace and expected. For a continued competitive edge, that can really make a difference to successful business results, the basic product offerings are not enough, nor is simply increasing reliability and meeting promises to customers.

The important differentiator comes from the feel-good factors—the trust and empathy portrayed by the brand values and brought to life by peoples' behaviours. This ensures the customer experience is not purely at a business level but is also pitched at the human level. The end result is to delight customers by creating positive memorable experiences by ensuring the key drivers are addressed at all times. (See Figure 3.)

Benchmarking

A determined focus on quality of service is a priority this year—not only because of what customers are saying they expect but also because feedback from the wider marketplace is indicating that BT must continue to improve. The TMA Report for 1998 shows that the most important consideration when choosing a telecommunications supplier is quality of service, as shown in Figure 4.

Benchmarking can be a valuable tool for any business that consistently aims to be the best. It provides a performance comparison of how good a business is in relation to its competitors and other industries, indicating areas where improvement is possible (or even vital).



More importantly though, through examining processes, benchmarking enables a business to understand how other companies achieve those levels of performance, whether it be in the area of quality of service, profit or customer satisfaction. By benchmarking across industry types and countries BT can identify those companies it would wish to emulate, to learn from and to make improvements.

Overall BT has a clear picture of how well it performs in the key customer service areas, and who is best in the chosen areas. It is noteworthy that in some cases current benchmark performance falls short of customer expectations. The key points from service delivery comparisons are outlined below.

Customer satisfaction

In terms of overall satisfaction BT is in the top half of UK service companies, but is well adrift of the best performance in the US service market, although, like for like, comparison is difficult owing to

different customer expectations in the two marketplaces.

Satisfaction with quality of service

BT again scores well but is behind best-in-class for UK service companies in general. Most aspects of its performance are consistent with those of the US telcos but fall short of the benchmark telecommunications company's performance.

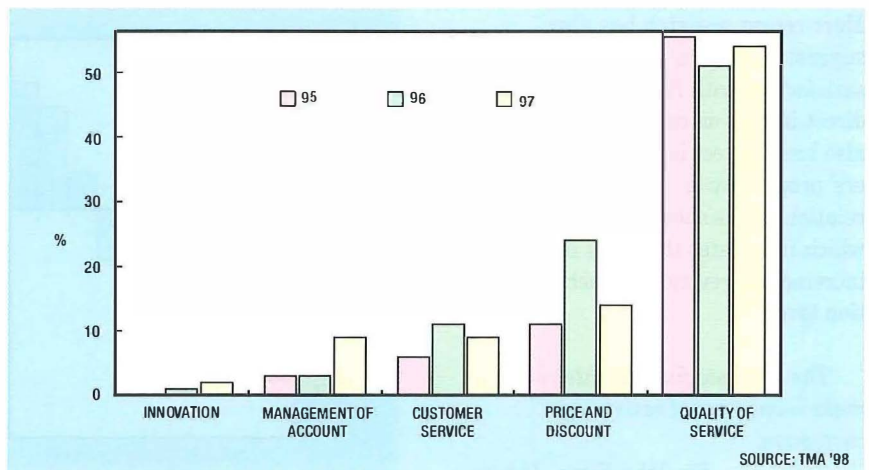
Call handling

BT has the best call-answer time threshold of 90% in 15 seconds compared to the North American service industry standard of 80% in 20 seconds.

Key practices from North America are to minimise time spent handling the order while maximising the quality of information gathered to avoid subsequent problems. The focus is on a few indicators of performance.

Recent call centre benchmarking studies show that centres with high employee satisfaction and retention,

Figure 4—Most important considerations when deciding upon a telecommunications supplier



plus high customer satisfaction ratings, trust their employees to work effectively. Hence they are measured and bonused on call quality rather than percentage calls answered (PCA) in target time, time on line and sales value. To reduce mistakes in order entry, rules-based systems and scripting are being increasingly used in North America. Best-in-class companies only introduce technology where it is acceptable to customers and it enhances their experience.

Provision

In the UK, BT is close to the benchmark performance in delivering service on the agreed date for both PSTN and private services. However performance is well short of the world best of 100% completion in three days achieved by minimising intervention once the order is placed and by pre-provision. Indications from North America are that orders not requiring a visit can exceed 90% although it is unclear whether this refers to dial-tone provision only.

Other best practices from North America are that only one field engineer is employed (70% of provisions) and 95% of field work is completed on the first visit. It should be noted that despite these practices, BT outperforms all such companies in meeting customer agreed dates, but probably less efficiently (see below the change of emphasis to customer required date).

Satisfaction with provision among BT's customers is at UK benchmark level. BT employ some best practices (a single records/production system, using customer satisfaction to improve performance and as a coaching tool, temporary resource for surges in demand, automatic work allocation to field via laptops), but needs to do more in terms of developing more flexible interfaces (for example, enabling direct order placing from Internet/third parties/home) and ensuring that service is delivered on time.

Repair

BT is at the benchmark level for PSTN satisfaction with repair although there is no easy comparison with speed of repair as companies report against varying target times and have different working hours. The world best attainment is 85% in two hours (real time). Notable effects arise from high use of re-provision to speed restoration.

BT's fault rate is currently 10 times the world best. The access network improvement programme (ANIP) for routining customer lines claims up to 50% reduction in reported faults and will impact here.

Complaint handling

In terms of speed of resolution BT is near the top of UK telcos and is steadily improving. Scottish Telecom is perceived as benchmark best in class. Only two UK companies

reported satisfaction with complaint handling in the comparable performers indicators (CPIs) survey, BT being a close second.

Service Improvements

The blitz on failure and other initiatives in BT's drive for service excellence have proved that determined focus and excellent teamwork can yield handsome results. This year's targets for business customers are focused much less on internal operational statistics and strongly on key customer-driven measures.

These include fixing today's fault today on repair, and delivery to customer required by date for provision. BT is concentrating on delivering what customers want when they want it and not when it suits BT and standard lead times. By focusing on what customers clearly value and is critical to them (reliability, keeping promises, empathy and flexibility), a marked improvement in satisfaction can be achieved. (Figures 5 and 6.)

Service as a Route to Market

Service people are a very effective way of keeping in touch with customers. BT has thousands of people making millions of customer contacts each year, both face to face and over the telephone. Not all of these contacts are positive of course—customers are reporting faults or making com-

Figure 5 – Customer required by date (CRD): percentage scoring 10/10

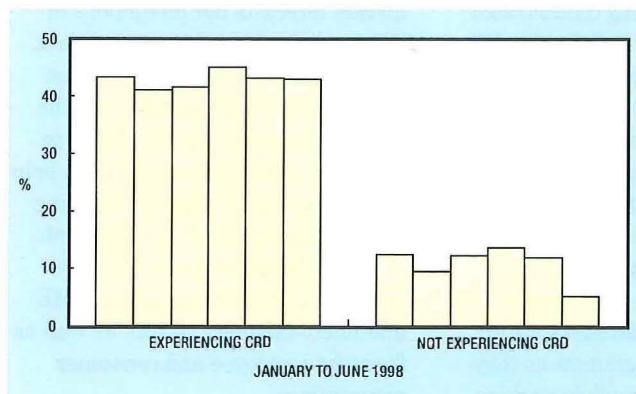
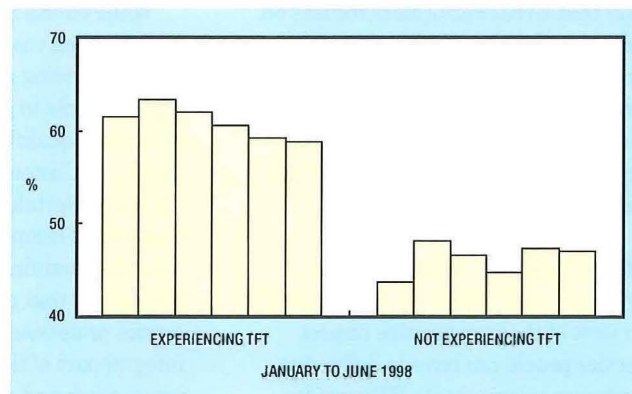


Figure 6 – Today's fault today (TFT): percentage scoring 10/10



plaints—however, the trust customers have in the service people means they frequently ask for advice. Achieving service excellence is not easy but the relationships people build with customers are a vital part of ensuring the company understands and responds to their needs and expectations. Through discussions and guidance BT people are able to build on these relationships and enhance customers' satisfaction with BT and are able to spot and realise revenue opportunities.

The more BT looks after its customers and meets their expectations the greater the effect on the bottom-line. Recognising this, and the potential revenue streams that result when service people are encouraged to look for new opportunities, BT has incentive programmes to enhance and support their efforts.

For example, *Great Performers* was launched in April 97 to reward individuals and teams for both revenue generation and customer satisfaction. It was built on experience and participants' feedback from similar programmes over recent years. From the early days when rewards were given simply for supplying a sales lead to a current focus supporting strategic marketing campaigns the programme has generated an excitement and commitment to deliver among service people. The programme pulls together a variety of opportunities to earn rewards and makes full use of technology, not only in communications, but also in enabling people to redeem their rewards via the Intranet.

But it is not only revenue generation that *Great Performers* focuses on. Because of the link between customers who are very satisfied, their loyalty to BT and the positive effect on our revenues, a large element of the reward pot goes to reward excellence in customer satisfaction.

Investing in Our People

In view of the very positive impact service people can have in delivering customer requirements BT provides

personal and professional development and training opportunities for all levels of responsibility within the organisation. Realising the importance of people development, encompassing enhanced ability, understanding and application, BT has tried to equip everyone with a broader perspective for the challenges ahead.

Senior manager modular programmes have encompassed such diverse elements as business strategy, building high-performance teams and the impact of technology on society. These managers then take the tools and techniques they have learned and cascade and apply them within their own teams. The modular format of the programmes ensures they can be used and adapted to suit individual and unit needs, with elements being included in team meetings and events. The same format is also used within working teams to improve their communications and customer focus skills in a supportive, cooperative environment where everyone's skills are used to optimum benefit. This approach assumes that people who work together towards a common objective deliver a more efficient, effective and customer-focused service.

BT's *Working Together Makes Perfect Sense* programme was specifically developed to create an atmosphere conducive to improving communications between all the links in the value chain, enabling people to focus on issues which create barriers to delivering world-class service, and empowering them to take action to resolve their own and customers' problems.

While efforts are being concentrated on increasing customer satisfaction, BT is also providing an opportunity for service people to gain a nationally recognised qualification, the NVQ Level 3 in Customer Service. Assessment is undertaken against various elements of competence and the resulting qualification provides the recognition that people are indeed service professionals. Managers are an integral part of the programme as they act as coach and mentor while working

on their own personal development with the opportunity to achieve NVQ assessor and verifier awards.

All of these programmes are aimed at allowing BT people to achieve their potential, linking their personal development to business needs, relevant to today's changing markets. These aims were recognised during BT's recently achieved Investors in People award.

Service Vision

Knowing why BT must continue to meet, and indeed exceed, customer expectations ensures that focus remains firmly on delivering excellent customer service. But BT's strategy must ensure that if it is to achieve its corporate goals the company must continue to plan ahead as well as improve service delivery today. The creation of the Service Vision helps shape BT's service strategy for the future. The work has led to an understanding of the forces shaping the nature of the marketplace in 2001 in order to create feasible scenarios of business customer requirements. From that understanding, implications for the UK's super-telco of the future have been identified and explored. A picture incorporating people, systems and process has been built to ensure that BT is best equipped to be that super-telco.

Experience shows that the past is a dangerous predictor of the future hence the scenarios were built rather than simply forecasting on historical trends. The scenarios are neither market forecasts nor predictions of what individual businesses will look like in 2001 but are stories built on solid evidence and research into the ways the business world is likely to operate in 2001—stories that can help BT recognise and adapt to changing aspects of the present environment.

Three business scenarios were constructed for the corporate, SME and micro-business sectors as well as from the employee and consumer perspectives.



Figure 7—The service vision

Using the scenarios and top-level assumptions about the telecommunications market context up until 2001, implications for the super-telco were developed and categorised into vision, organisational structure, systems and processes, service offering, customer service strategy, values and attitudes, skills and competencies (see Figure 7).

Based on the implications, guiding principles for the development of a service strategy were developed and clustered in the following groups:

- what the customer will want from the super-telco;
- what customers will want from the super-telco's people; and
- customer requirements from corporate, SME and micro-business perspectives:
 - values and attitudes,
 - capabilities,
 - service characteristics, and
 - the human resource environment

The guiding principles are now being translated into deliverables, which will help to guide the development of our service strategy and plans.

World-Class Service

The specification for world-class service, produced as a component of the Service Vision work sets out a view of the behaviours that should characterise BT's service offering for business customers in the year 2001. It also aims to provide the overall framework for optimising the creativity and enthusiasm of BT

people and for shaping the behaviour of the organisation to build long-term competitive advantage.

World-class service is an ability to understand customers and deliver precisely what they want each and every time. This is more than mere service. BT is a reflection of its people and what they say and do; therefore BT's reputation and the behaviour of its people cannot be separated. BT people have an opportunity to enhance this reputation every time they make contact with a customer—face-to-face, visually, electronically, in writing, in-bound or out-bound.

'A barrage of tiny positives overwhelm the customer and the competition.'

Tom Peters—About the Customer

The behaviour of BT service people is affected by many factors which include:

- training and on-the-job coaching;
- individual performance measurement and reward;
- recruitment criteria addressing personality as well as technical skills and knowledge; and
- the values and signals demonstrated by BT policies and management style and practice.

Pen-pictures of world-class service in 2001 have been built from five different perspectives: the customer, the customer service professional, the team member, the team manager and the senior manager. These pen-pictures provide a view of the desired characteristics of BT's business customer service people in the year 2001. They are not specific to a job function or part of the organisation.

The task is to convert the 'vision' into actions or programmes that will make a lasting difference. Although the key drivers of customer satisfaction are understood, the business world is constantly changing and the

challenge for BT will be to ensure it continues to delight customers now and into the future with their increasing demands and expectations.

References

- 1 PEARSON, STEWART. Building Brands Directly. Macmillan Business, 1996.
- 2 MITCHELL, ALAN. Marketing Business, May 1997, CIM.

Benchmarking studies:

- Service Delivery Benchmarking Study (1996) with North American utilities.
- Customer Satisfaction Measures Benchmarking (1997) with UK and US service companies.
- Oftel CPI Results (July–December 1997) with UK TelCos.
- COMPOST (1997) with UK TelCos.

Acknowledgements

The author acknowledges the help of Colin Pavey, Will Sewell, Kevin Thomas and Carol Reeve in the production of this article.

Biography



Megan Brown
BT UK Markets

Megan Brown is Manager, Customer Service Development. She is responsible for

programmes which change people's behaviours so that their contribution to customer satisfaction and to customer service profitability is maximised. Previous roles were in marketing, including both internal marketing and customer event strategy. Megan's early career encompassed exchange modernisation and planning. She has a degree in Economics, Diploma in Marketing, an Advanced Professional Diploma in World Class Marketing and is a member of the Chartered Institute of Marketing.

Strategic Resource Planning for Optimum Service Quality

With increasing competitive and regulatory pressures to raise quality and reduce costs, there is a growing need to ensure that BT's field operations are optimal. This article describes how the logistics of managing the field resources to achieve optimal working are being supported by state-of-the-art mathematical and computer tools. These techniques make it possible to quantify the trade-off between cost and quality when many variables are involved.

Introduction

The systematic application of quantitative analysis to the deployment of resources was first used in earnest by the military during the development and operational implementation of radar¹. More recently, operational research has been applied successfully to a range of civilian problems, and many large industrial organisations now maintain an operational research department. One area of operational research which has received a large amount of attention is the scheduling of resources in order to deliver a prescribed quality of service within a cost constraint.

Within BT, delivering a quality service, in terms of meeting customers' expectations on provisioning and repair of basic telephony, is key to its operations. In order to provide a flexible and comprehensive service to customers, BT's operational units are organised into small- and medium-size work teams. These teams have sufficient autonomy to solve day-to-day problems with minimal reference to a central control. Although the use of small teams encourages the local optimisation of resources, it does not necessarily lead to a company-wide optimum strategy. Thus, there is a continual need to ensure that an optimal cross-divisional strategy exists. One manifestation of this requirement is the need to understand the relationship between prescribed quality-of-service (QoS) targets, product volumes and available resources. By identifying the nature of these relationships it is possible to quantify the trade-off between QoS and resourcing costs.

Simple methods of predicting future resourcing levels are based on current productivity and service targets. The new resourcing levels are then adjusted to take account of productivity improvements, which are likely to occur due to the introduction of improved working procedures and new technology. This method of prediction relies heavily on average statistics, which are derived from data on the time taken to carry out individual tasks and the number of tasks carried out. Although this data provides an accurate indication of historical performance, it does not identify the interdependencies between the type of work, the productivity and task volumes. Thus, it is very difficult to predict accurately the effect on costs of changing the QoS targets of any individual product. An example of this would be if, in response to customer demand, an enhanced level of service was introduced, for which the customer is charged a premium. The volume of work remains the same but more people may be required to ensure that the targets for the new service are met without impacting on the performance in the other areas. The current methodology does not identify this increase in resourcing level and hence the cost that will be involved.

This article describes how the logistics of managing the field resources to achieve optimal working are being supported by mathematical and computer tools. These techniques make it possible to quantify the trade-off between cost and quality when many variables are involved. By understanding the nature of these trade-offs it is possible to identify a company-wide

operating environment which maximises QoS while minimising costs.

Balance of Cost and Quality of Service

The first step in understanding the relationship between cost and quality of service is defining what is meant by each of these terms. For the purposes of this article only human resourcing costs are considered. Thus, minimising cost implies maximising productivity by increased resource utilisation. Examples of factors to consider are the required size of the workforce, the efficient allocation of jobs and ensuring the workforce is sufficiently trained to enable each member to efficiently carry out the work required.

Quality of service is more difficult to define. Customer surveys suggest that factors such as rapid response to problems, quality of workmanship and keeping appointments play an important role in the perception of quality of service. However, in order to measure quality of service it is necessary to capture these factors in a set of targets and metrics. An example of a metric may be, 'How many jobs of type J were completed within X hours of being received?' The target for jobs meeting this criterion is a given percentage of the total number of jobs of type J. The time period X is often termed the *commitment time*. A number of factors, such as the size of the workforce, the time taken to complete a job and the priority given to different types of job, will influence the ability to meet a given target. Therefore, it is necessary to understand the relationship between these factors before it is possible to quantify the probability of meeting a given target or metric.

Meeting the required quality of service can be thought of in terms of the process illustrated in Figure 1. A job is allocated to an engineer who will travel to the job, complete the

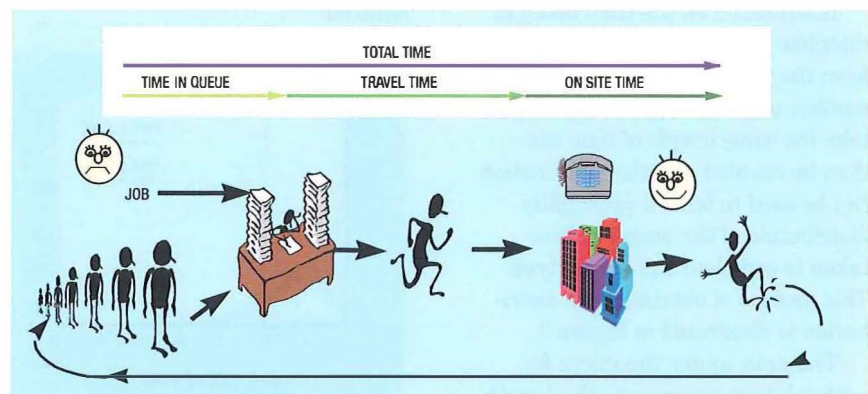
work and then rejoin the resource pool ready to be allocated the next piece of work. The total time taken to complete the work can be divided into three components: the time the job waits in a queue to be allocated to an engineer; the time taken for the engineer to travel to the job; and the time taken by the engineer to carry out the work at the site of the job—the 'on site' time. There are a number of different types of job, which require the same workforce resources to service them.

The allocation of each job to an engineer is determined by a number of constraints, such as the skills required to complete the work, the size of the workforce and the relative priority of jobs. As an example, consider the three stages of the process described in Figure 1. The first stage is the length of time a job is in the queue waiting for an engineer to be assigned to it. This time will depend on the availability of an engineer with the correct skills to service the job and whether it has a high or low priority. The second stage is the time to travel to the job, and this will depend on the relative geographical locations of the job and the engineer(s). The third stage is the time taken to carry out the work which will depend on the difficulty of the work and the skill of the engineer. Therefore, the total time taken to carry out each individual job is both job and engineer dependent. The requirement of the job allocation algorithm is to ensure jobs

are allocated to workforce members in such a way as to minimise the total time across all jobs.

Within BT, work is allocated with the aid of an optimisation algorithm embedded in a software work allocation system known as *Work Manager*². This software allocates jobs to engineers throughout the day and takes account of new jobs continually joining the queue of work. To provide a work schedule, the algorithm requires a set of inputs such as the number and skill of the available engineers, the priority of different types of jobs and the time by which each job has to be completed. The values given to each of these inputs can have significant impact on the cost of providing a given level of service, and therefore it is important the impact of changing any of these parameters is well understood. As an example, consider the case of multiskilling engineers. Enabling every engineer to undertake all types of job will increase the probability of an engineer being available to carry out a given job when required. However, the cost of training and providing the required equipment to enable every engineer to carry out all work is a costly process. Hence, a compromise of multiskilling a percentage of engineers is regarded as providing the optimum solution. The topic of multiskilling will be dealt with in more depth in a future article in this *Journal*.

Figure 1—Process used to complete work






Figure 2—Illustration of the job allocation schedule

Modelling the Problem

One method of identifying the optimum strategy and quantifying input values for the Work Manager algorithm would be to simulate all possible scenarios and determine the most efficient working method. However, there are a large number of parameters which can affect the job allocation strategy and it is impractical to systematically search such a large space in order to obtain the optimum solution. An alternative approach is based on statistical analysis and uses representative data to identify key relationships between the parameters. In this way the size of the search space can be dramatically reduced. To obtain the statistics, consider the allocation process. Figure 2 illustrates the allocation of jobs to engineers as a function of time. The letters (A) to (E) represent different engineers, while the numbers (1) to (6) represent six different types of jobs. The location and length of the coloured columns indicate when the job was allocated and the total time it took to complete, respectively. For clarity, the colour of the column is also used to indicate the type of job being allocated.

Information on the time taken to complete each job can be obtained from the job allocation schedule. The number of jobs of a given type which take the same length of time can then be counted and this information can be used to build a probability distribution of the length of time taken to complete a given job type. This method of obtaining the distribution is illustrated in Figure 3.

The area under the curve for each job type represents the length




Figure 3—Building the probability density function

of time taken to complete all the jobs of that type. Also indicated in Figure 3 is a red vertical dashed line, which represents the length of time by which the job of that type should be completed; for example, the commitment time. Therefore, by calculating the area under the curve that lies to the left of this red dashed line, it is possible to obtain the percentage of jobs, which were completed within the commitment time.

The above approach is appropriate to non-appointed work where the probability densities for each job type are continuous functions of the elapsed time. In some instances it is more appropriate to carry out the work by making a prior appointment with the customer. Under these circumstances it is not meaningful to include the time that a job waits in the queue as part of the measure of quality of service, and a more appropriate measure is to ensure that the appointments are kept. Thus, in the case of appointed work the distribution is binary in nature and takes a value which depends




Figure 4—Binary probability distribution for appointed work

discretely on the number of successfully met appointments. This is illustrated in Figure 4.

Coupling of distributions

For each job type there will be a separate probability density function. These functions will be coupled due to the fact that all the work must be serviced by the same workforce. This coupling is illustrated in Figure 5, where the different coloured curves represent the density functions for two separate non-appointed job types. The left-hand illustration in Figure 5

Figure 5—Illustration of the coupling of the probability density functions due to all job types sharing the same work force. The left-hand graph shows the situation where the commitment time for both jobs is the same. The right-hand graph illustrates the situation where the commitment time for one type of job is reduced

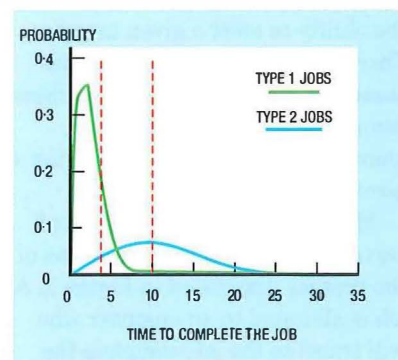

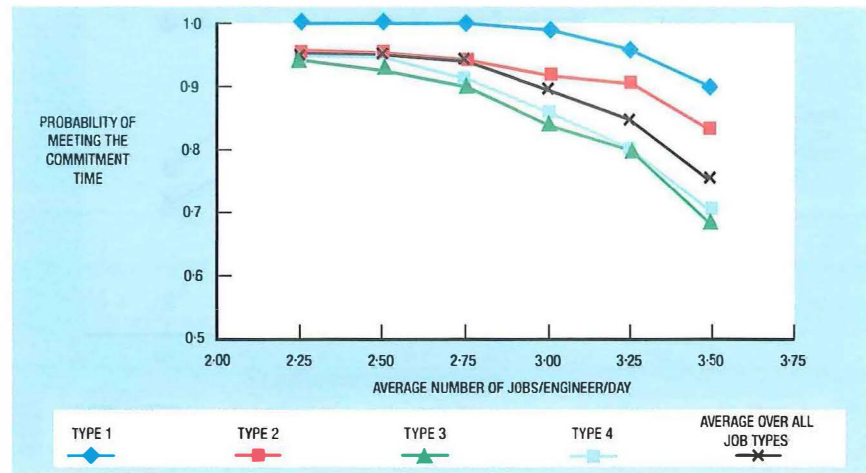


Figure 6—Impact on the probability of meeting the commitment time as a function of increasing workload

represents the situation when the commitment times for job types 1 and 2 are the same, as is the percentage of jobs which is required to meet this commitment time. If the situation is now changed so that the commitment time for job type 1 is reduced, then, if the same percentage of jobs are to meet the commitment time, the area under the curve of the density function which is to the left of the commitment time must remain constant. However, assuming there is little or no spare capacity within the work force, then the total number of jobs of type 1 and 2 that are completed within a given time will remain constant. Therefore, if there is a requirement to decrease the commitment time for one particular type of job, then there will be a corresponding increase in the elapsed time for the remaining job types. This situation is shown in the right-hand illustration of Figure 5 and can be explained qualitatively as follows. If it is assumed that the on-site and travel times remain constant, then in order to ensure sufficient jobs of type 1 are completed within the shorter commitment time, it will be necessary to reduce the time that each type 1 job spends in the queue. This could be achieved by increasing the priority of type 1 jobs. However, if type 1 jobs are given preference to type 2 jobs, the queuing time of type 2 jobs will increase and hence the total elapsed time for type 2 jobs will increase. This implies fewer jobs will be completed within the commitment time.

The coupling of the probability density functions is not only observed when the commitment times and job priorities are altered, but also when other strategic parameters are changed. The impact of changing these parameters can be investigated individually and the results used to determine the nature and strength of the coupling. In this way the parameters which are strongly coupled can be identified. Any change in these strongly coupled parameters is likely to have



a major impact on the ability to meet the quality of service targets.

Example results

Using the probability density functions described above it is possible to determine the statistics of a given situation. However, the complexity of the interrelationships of the coupling parameters makes it difficult to formulate the nature of the coupling in an analytical manner. Therefore, as an initial attempt to identify key relationships between parameters, representative data has been collected. These sets of data represent typical distributions of job type, duration and location, together with engineer attendance patterns and skills. As an example of the modelling technique, consider the problem of identifying the number of engineers required to service a given workload. The work consists of a mixture of four different job types. Moreover, each of these categories can be carried out either by appointment or within a predefined commitment time. The particular job mix and the geographical location of the jobs were chosen as representative of a suburban area containing a typical market town. The workload for a period of several weeks was constructed, with the daily variation of workload being chosen to be representative of that observed in measured data. The total number of jobs within the data set was varied to investigate the quality of service as a function of workload. While varying the total workload, care was taken to ensure the daily variability of workload and the ratio of different jobs type were kept constant. The pool of engineers available to service

the work was a mixture of single skilled workers, who can only service either repair or provision work, and multi-skilled engineers who can tackle both types of work. In this example, the number of engineers and their skill base remained constant.

Using the data described above, jobs were allocated to engineers to provide a schedule of work. The resulting schedule was then analysed to determine the percentage of each job type that met the quality of service criteria.

Figure 6 illustrates the number of jobs that meet the commitment time as a function of total workload for each of the four job types. All jobs have been given the same priority in the work stack; however, work of job type 4 has been given a longer commitment time. The workload has been varied from 2.25 to 3.5 jobs/engineer/day which is a typical range of workloads. As can be seen from Figure 6, the percentage of jobs which meet the commitment time decreases as the workload increases. For an average workload of 2.75 jobs/engineer/day, the probability of meeting the commitment time is between 0.99 and 0.92. Hence, although all job types are treated as having equal priority, the probability of meeting the commitment times is not the same. One cause of this difference is the very peaky distribution of report times of the work. When the work force is working near capacity, high fluctuations in workload will lead to more jobs failing than if the workload is constant for the same total number of jobs. Therefore, if the arrival times of one job type are clustered,

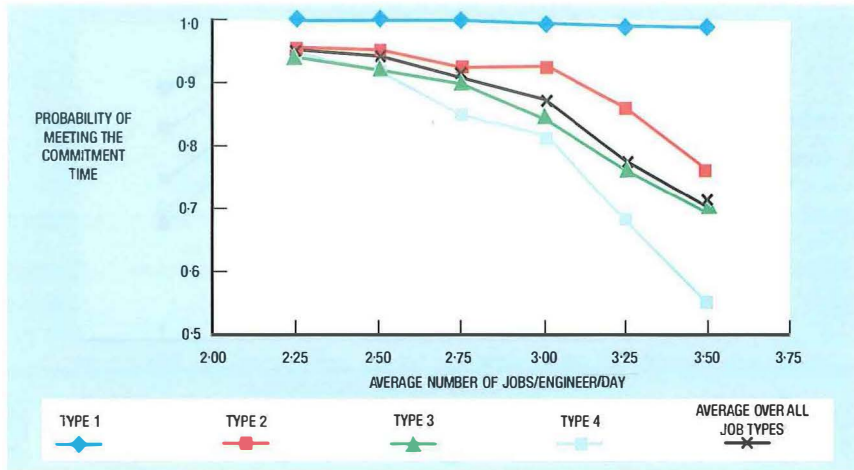


Figure 7—Impact on the probability of meeting the commitment time as a function of increasing workload with job given different priorities

this job type has a higher probability of failing the commitment than other job types with a more even distribution of arrival times. Another cause of jobs of one type failing more frequently than those of another is the different length in commitment time. Jobs with a longer commitment time are more likely to meet this time limit since there is a longer time period over which to spread the peaks in the workload.

Prioritising job types

One method of overcoming the problem of different commitment times is to increase the priority of jobs with short commitment times. By giving a job a high priority it is given precedence over other jobs in the queue when allocating that job to an available engineer. To investigate the impact of prioritising work, the same data was investigated with jobs given priorities. Job type 1 was given the highest priority, followed by type 2 then 3 and finally job type 4. The results of this investigation are illustrated in Figure 7.

It can be seen from the figure that introducing a job priority has little impact during periods of lower workload. However, as the workload increases, the probability of meeting the commitment for job type 1 remains close to 1. Unfortunately, if engineers are servicing high priority jobs, they are not available to provide the resource required for other work. Therefore, in high-workload situations, the probability of meeting the commitment time for lower-priority work is much reduced. In this example, the probability of meeting the commitment for job

type 4, which has the lowest priority, was reduced from 0.72 to 0.57. Figures 6 and 7 also indicate the average probability of meeting the commitment as measured by the proportion of total jobs. In the case when no job priority is imposed, the average probability of meeting the commitment varies between 0.96 and 0.76 for low and high workloads, respectively. However, when job types are given different priorities, the average probability of meeting the commitment varies between 0.95 and 0.70. Thus, when in high-workload situations, adding priorities to different work types to ensure high priority jobs are more likely to meet the commitment time will lead to a reduction in the average performance across all job types.

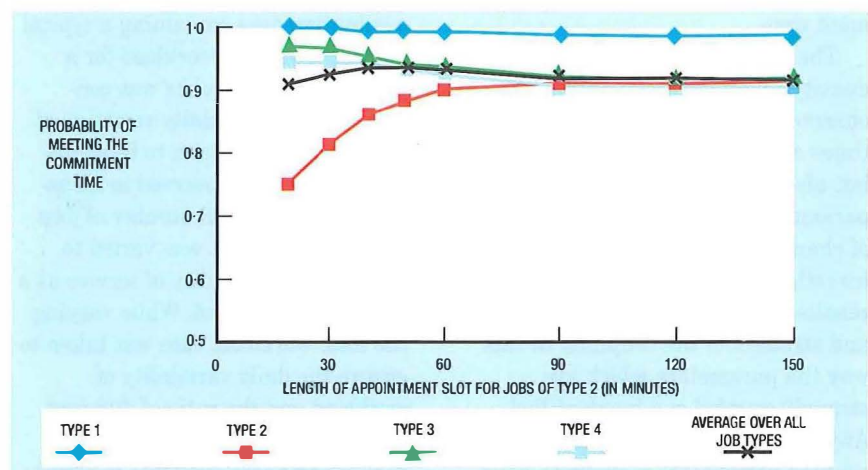
Length of appointment slots

As another example of coupling between jobs, consider the case of appointed work. As mentioned above, work can be done either with

or without making a prior appointment. In cases where appointments are made, the metric used to identify a successful appointment is whether an engineer arrives within the first made appointment period. Given the uncertainty in the duration of any given job and the variability of travel time, a key factor in determining whether it is possible for an engineer to keep a given appointment is the length of the time window in which he/she has to arrive. The length of the time window is known as the *length of the appointment slot*. From the customer's perspective, the smaller the appointment slot the more convenient the appointment, especially if the customer has to remain in one place in order to give access to the building. Thus, the problem is to quantify the length of appointment slot which is short enough to minimise the length of time customers have to wait, while being long enough to ensure that the engineer has a sufficiently high probability of arriving within the required time slot.

Figure 8 illustrates the example where the appointment slot for jobs of type 2 has been varied between 20 minutes and 2.5 hours. The workload in this example has been kept constant at 2.75 jobs/engineer/day. However, jobs have been given

Figure 8—Probability of meeting the commitment as a function of appointment slot



different priorities, with job type 1 given the highest priority, followed by type 2 then 3 and finally job type 4.

In Figure 8 it can be seen that the probability of the engineers arriving within the appointment window is very much reduced when the length of the appointment slot is reduced below 60 minutes. It is interesting to note that reducing the appointment window below 60 minutes appears to improve the probability of meeting the commitment time of other job types. This is due to the fact that if engineers cannot arrive within an appointment slot they are not allocated the job. This allows another job of a different type to be allocated to the engineer, which in turn increases the probability of other job types being completed within their commitment time. However, the magnitude of the increased probability in meeting other job types is not sufficient to offset the decrease in ability to meet appointment slots of less than 60 minutes. Thus, the overall performance of the system is reduced. This is indicated by the decrease in the average probability as illustrated in Figure 8.

Conclusions

This article has shown how it is possible to identify the key relationships between resourcing costs in terms of workload per engineer and the ability to provide a given quality of service. A methodology has been developed which reduces the size of the problem space that has to be numerically searched in order to quantify these key relationships. Two examples have been given of how the methodology can be used to determine an optimum operating strategy for the provision of a given level of service.

References

- 1 KIRBY, M.; and CAPEY, R. The Air Defence of Great Britain, 1920–1940: an operational research perspective. *Journal of Opera-*

tional Research Society, 48, 1997, pp. 555–568.

- 2 SIMPSON, J.; NOBLE, A.; EGAN, B. F.; MORTON, D.; RICHARDS, T.; and BURSTIN, M. Experience in applying OR techniques to the solution of practical resource management problems. *BT Technol. J.*, 13(1) Jan. 1995.

Biographies



Gail Lochtie
Networks and
Information Services,
BT UK

Gail Lochtie graduated from Edinburgh University in 1984 with an honours degree in Physics. On joining BT, she worked in the radio division studying the effects of tropospheric propagation on communications systems. During this time she gained an M.Sc. and Ph.D. in Mathematics. In 1992, she joined the system research division where she worked on the application of self-organising systems to optimisation problems in mobile radio. In 1997, Gail moved to the Process Engineering group where she currently leads a small team working on resource and capacity planning.



Mark Gilbert
Networks and
Information Services,
BT UK

Mark Gilbert joined BT as a trainee technician improver at BT Laboratories in 1980 working on the assessment of gallium arsenide (GaAs) for use in high-speed digital electronics. Over the following six years he

became involved in a team developing metallisation schemes for both digital and analogue GaAs-based electronic circuits, with particular emphasis on application to millimetre wave radio systems. The role developed to include the development of novel electronic and optoelectronic devices in a range of III–V and silicon-based semiconductors, including DTI and European-based collaborations. In 1993, he received a first class honours degree from the Open University. During 1995/6, he represented BT in a EURESCOM project looking at the integration of B-ISDN and IN signalling protocols, with particular emphasis on its application to mobile communications. During the following year he became part of a team looking at error correction schemes for VDSL (very-high rate digital subscriber line). He is presently employed in the Process Engineering group within Applied Research and Technologies, looking at various aspects of the application of production management techniques and modelling within BT.



Andy Noble
Networks and
Information Services,
BT UK

Andy Noble graduated from the University of Newcastle-upon-Tyne in 1976 with a B.Sc. in Statistics. Since joining BT he has worked in a variety of areas including manpower planning, depreciation studies, performance measurement and market research. He also spent time as a support manager for the roll-out of BT's Total Quality Management programme. He is now head of a team of management scientists working in BT's Networks and Information Services division.

David Butler

Implementing Software Tools for Planning Automation

This article describes the implementation of a BT software planning tool called BECS (budgetary estimate calculation system), from initial specification to product support, highlighting essential lessons about how to bring about successful software development. The BECS development, culminating in the roll-out of the tool within BT, has significantly contributed to changes made in the process of diversionary works planning, by bringing about greater simplicity and control. The success of BECS can be explained, in part, in terms of how the tool has shaped and been shaped by the end users.

Introduction

This article describes the implementation of a software planning tool from initial specification to product support, highlighting essential lessons about how to bring about successful software development.

The starting point is the examination of the product life-cycle diagram (Figure 1) found in many text books¹. This life cycle is applied to the software planning tool BECS (budgetary estimate calculation system).

In reality this simplistic theoretical approach is untenable. There are many feedback loops between the boxes, and the development and implementation proceeds in a more iterative way. Unfortunately this would yield a more complex diagram depicting a network of interactions. The success of BECS can be explained, in part, in terms of how the tool has shaped and been shaped by the end-users. Those unseen loops involving end-users have had most impact on the development. It is in those loops that the integration of process and systems comes about.

Figure 1 is a useful framework for understanding development, although the separation of the boxes and the existence of a 'left to right' ordering actually reveals more about the way products and software are introduced rather than the way they can be most successfully deployed.

The BECS development, culminating in the roll-out of the tool to all repayment project engineering offices,

has significantly contributed to changes made in the process of diversionary works planning, by bringing about greater simplicity and control.

Determining Need

In 1995 a project was started to look for opportunities to automate the diversionary works process within BT. Several months were taken to look at the process, which was a mixture of top-level legislation, and BT corporate and local interpretation.

Diversionary works process

Within the New Roads and Street Works Act 1991, BT and other utilities and service providers have an obligation to move their plant when requested by external agencies, such as local councils, developers and highways authorities. In the majority of cases, BT provides an estimate for like-for-like replacement of its plant which is met by the client requesting the diversion (Figure 2).


BT plant mark-up

The first part of the process is to determine whether any BT plant is affected by the development, and this is achieved with reference to plant maps. A developer will often provide a site map onto which BT marks the approximate position of its overhead and underground plant.

Budgetary estimate

If plant diversion is necessary, BT calculates a budgetary estimate to

Figure 1—Theoretical life cycle for BECS






Figure 2—Diversionary works process

indicate to the client the approximate cost of moving the plant within the development area. BT has a statutory obligation to produce budgetary estimates '...as far as can reasonably be assessed at the draft design stage' as defined by the appropriate code of practice².

As part of the statutory requirement, the cost of the work marking up maps and producing the first budgetary estimate is free to the client. Although, as we shall see later, the costs to BT can be high, the free initial stages encourage developers to contact BT rather than risk damaging plant.

If the client accepts the budgetary estimate cost the next stage is for BT to produce a detailed client estimate (DCE).

Detailed client estimate

The DCE consists of a series of documents detailing the work to be carried out and the way in which it will be conducted. In addition, a refined cost estimate is supplied with any benefits and discounts calculated as explained in the legislation. This is recognised to take more work than the budgetary estimate stage, and so the time can be charged to the client.

Work execution and review

Once the client agrees to the DCE, the work is executed using either BT direct labour, contract labour or a mixture of both. Costs are monitored during the execution stage and reviewed at the end to ensure that the out-turn is not significantly greater than the DCE estimate. During the execution stage, there may be requests to deviate from the specified work due to unforeseen circumstances, such as finding difficult soil conditions when installing a new manhole. These events usually raise additional costs that the client agrees to as they arise. Therefore the out-turn can exceed the DCE estimate cost when the scheme is finally reviewed.

Project management tasks

There is a wide diversity of schemes from a simple request to move a telephone pole, through to plant diversions along the entire Channel Tunnel rail link. This diversity means that every scheme is unique, and the process that BT operates must be flexible enough to cater for all eventualities.

The project engineers manage the diversionary works process within the legislation, code of practice, BT ISIS documents, local knowledge and with reference to their line managers as necessary. There are approximately 60 diversionary works offices with more than 250 project engineers now within BT's *networkbuild* trading unit.

Process flow and control

After visiting several offices it became clear that the diversionary works process was being run with different team organisation, use of documentation and interpretation of the calculation techniques for estimate preparation. It was believed that this had little overall effect as the costs were only 'estimates' and it was recognised that the execution stage could be relied upon to modify the costs to the client. With this in mind, the budgetary estimates tended to be high because generally at the C3 budgetary stage there is not enough information from the client to enable a definitive costing to be identified.

On examination of the work flow, it was evident that on average 70% of the budgetary estimates never progress to the DCE stage. This was because the clients changed their minds about the development, or BT received multiple requests for the same piece of land, or the highways authority scheme was not viable due to the diversionary cost. The consequence was that BT was spending time and effort producing free budgetary estimates that often resulted in no chargeable work at all.

There are often mixed feelings about rejected diversionary schemes. In some cases BT may not want to move its plant but is obliged to do so under legislation. It is not unknown for a developer to change the course of a road to save the expense of moving a BT cabinet. For both parties, the rejection of any work through this negotiation is a success as the client incurs no cost and BT can keep its plant in its original position. In other situations, BT can benefit from plant diversions, as a new like-for-like replacement is required thereby improving the affected route, although the client will not be expected to pay the full cost of the upgrade.

Process automation options

It was concluded that the provision of a budgetary estimate was most suitable for automation. This was because the work had to be done quickly and effectively thus releasing project engineers to spend time managing the chargeable parts of the process. Spending time producing free budgetary estimates had to be minimised, especially as so many did not result in chargeable work.

Many project engineers had realised this and because of the lack of a nationally supported system had devised spreadsheets to calculate budgetary estimates. While this was to be welcomed in terms of innovation and potential efficiency savings, it meant that many solutions were to remain locally optimised. The optimisation usually relied upon interpretation of the calculation ingredients to produce local recipes, as discussed in the next section.

Calculation ingredients and recipes

In order to explain how BECS has evolved to handle budgetary estimate costing it is necessary to provide a short history of BT pricing policy for diversionary works. Change is inevitable and the challenge has been how to keep BECS up to date.

Work execution

The work executed for diversionary works can be divided into two main activities: duct /external plant and cable operations.

Duct work

External contractors carry out duct work, which includes the demolition and provision of manholes, jointboxes and other underground structures. The work is covered by a contract put out to tender across the country. One or several contractors bid for the work, and, once the rates are known, calculating the cost of a new manhole is as easy as looking at the local list. When stores and overheads are applied, the estimate total can be provided to the client.

Cable work

The cable work is calculated in a more elaborate way and is open to a degree of interpretation depending upon scheme circumstances. Back in 1995, BT direct labour staff carried out the majority of diversionary cable work. The way of calculating the total cost relied upon breaking the cable work into activities, such as opening a manhole, jointing a cable, and fitting a closure. Each activity was assigned a time, and by adding these up and multiplying by the man-hour rate could be converted to a cost. A typical cable changeover could contain up to a dozen activity codes, but this was not a great problem as the technique, although complex, was mature and well understood. The allocation of time also included associated activities such as meetings with line managers and gaining access to site.

This meant that analysis of the executed work sometimes revealed little about what really happened because the system was not designed to reveal variances to planned work.

Quality task times (templates)

In 1997, the view was taken that activity-based costing should be improved by taking out the variances to leave the actual time to carry out

the task. The new set of planning ingredients was called *quality task times* (QTTs)³ and allowed activities to be radically rearranged. The concept of activity codes and work categories remained, but more detailed descriptions of the activity scope were included. Upon introduction in 1997, the project engineers noticed that certain activities had been cut in terms of time. What had occurred was that the idealised jointing time was now being used and a separate list of variance codes (such as 'access to site') could be booked. The inclusion of variances meant that it was possible to see whether the activities or the variances dominated the scheme cost.

Group quality task times

Gradually the use of BT direct labour began to diminish and external contractors were being used for cable work as well as duct work. Previously, the contractor cost was based upon an activity category man-hour rate or on a per-QTT cost basis. These calculation schemes were essentially bolted on to the original processes to enable a cost to be determined, and a dedicated pricing scheme was necessary.

Improvements to the contracting out of cable works led to the formation of a list of group quality task times (GQTTs) in late 1997⁴. The idea was to bundle up QTTs into tasks and then obtain a price from contractors. The task would be defined and all eventualities would be the responsibility of the contractor.

This is where the story should end with the introduction of a 'recipe'-based scheme rather than ingredients open to interpretation. However, where execution is still carried out by BT direct labour, the field recording methods are based upon time and not cost. This means that measuring field performance using GQTTs that are based upon cost is very difficult. To reconcile the two methods each GQTT bundle has been assigned a time. Field staff record time per GQTT which can enable project

engineers to monitor scheme progress.

Relevance to BECS

BECS was introduced during the first costing regime and has had to keep up with policy changes. This is very important and relies upon continued communication between the developer and the Diversionary Works organisation (now within BT's *networkbuild*).

Software tools such as BECS ultimately support policy initiatives, and synchronising the policy and tool deployment is vital to system success.

BECS has always employed a rigid approach to calculation in that users enter the real world components (manholes, cables etc.) and BECS applies the appropriate calculation. This means easier data entry and quick estimate creation.

Research and Software Development**How project engineers coped**

The previous section regarding calculation policy neglects to mention that project engineers were presented with new ways of working on a regular basis. They had to understand and apply the changes because the requests to move BT plant still flowed despite changes going on. The project engineers used locally produced spreadsheets that enabled any policy changes to be incorporated quickly. The lack of standardisation was, however, proving to be a problem as audits revealed.

Spreadsheets are a very useful tool, but suffer from problems such as the following:

- Although cells can be protected, it is often easy to modify or erase underlying formulae.
- The majority were not version controlled nor conformed to data security practices.

But none of this can be brought into effect without end-users. Their views are often neglected but they hold much power as agents of change.

- They can be difficult to maintain, as 'owners' tend to use cell references instead of cell naming. This meant that a formula for moving duct could contain more than 20 cell references such as (C3*C4^2)+... rather than (Width * Depth^2). Maintaining this kind of code without documentation can be time-consuming and hence costly.
- Often no one person owned the spreadsheet, which made support and upgrades difficult to obtain.

Solving the spreadsheet problem

It was recognised by the diversionary works managers that using some type of automation for budgetary estimate production was needed. Before BECS was conceived an attempt was made to promote the use of a spreadsheet called *JRET*. The general result was good but soon the program was being customised for local needs. The net effect was that no two versions of *JRET* would give the same result. However, to the users of the spreadsheets, everything appeared to be fine. Clearly, taking time to determine the need was going to be important to achieve anything beyond local success.

What are the needs of the users?

Some people would view technological development as not only necessary to productivity improvement but as desirable in itself. Others, offering an alternative to technological determinism, argue that society shapes technology, improvements being driven by social need other than productivity. BECS has occupied the ground between these two views. In the world of diversionary works, BT has to have standardised, defensible and legislatively consistent budgetary estimation. Technology allows this at the lowest cost. But none of this can be brought into effect without end-users. Their views are often neglected but they hold much power as agents of change.

Alan Cooper, the designer of Microsoft™ Visual Basic, lists the following fundamental users' goals and how most commercial software is designed (Table 1)⁵.

Designing for users' goals

The most obvious step therefore is to consult with the user group and make the design process iterative until agreement is reached. However, a starting point is required and Microsoft Visual Basic™ (VB) was chosen as the development environment. Back in 1995, Visual Basic version 3, a 16-bit development language, was available to produce programs that ran on the 386 and 486 PCs that BT used in planning offices.

VB consists of a toolbox of controls such as buttons, text-boxes, labels and drop-down menus. The controls have properties to adjust their dimensions, font colour, size and actions and are placed upon a blank form. By adding controls on the form, a user interface can be constructed without the need for any code. This enables a quick prototype to be created, although until the code is written the program will just show how the interface will look rather than how it will work.

Prototyping was used to construct the initial user interface for BECS and was demonstrated to a group of project engineers in London. When they realised that the controls could be adjusted, a flood of ideas emerged resulting in the complete redesign of the system over the following two weeks. The resulting dialogue is part of the

route to success. The metaphor most suited to software development is that it should be 'grown rather than be built'⁶ so that at every stage in the process you have a working system. During the two weeks in London, there was an opportunity to see how people used other software systems, and this knowledge was applied when designing for usability.

Designing for usability

The main reason why 'usability' should be considered is to avoid making users look stupid and to minimise wasted time. An additional advantage is that there are financial benefits as well. A study⁷ conducted on a number of software development projects revealed that significant financial benefits could be made through increased user productivity brought about by enhanced usability. Benefits are accrued through productivity and decreased use of help desks and lower maintenance costs. So, by taking time to involve the user in the iterative design loop both user and company goals can be achieved.

Usability is not just about the interface though; it also includes the documentation, installation, training and system performance aspects. From observation it was decided to abandon a printed user manual for BECS as this would only add to the project engineer's growing documentation.

The approach taken was to include an on-line user's manual accessed via 'help' buttons, or context sensitive so that users pressed the 'F1' key when help was

Table 1 User Goals⁵

User Goals	Commercial Software
Not looking stupid	Making the user look stupid Software is often rude
Not making any big mistakes	Causing the user to make big mistakes
Getting an adequate amount of work done	Slowing the user so he/she does not get an adequate amount of work done
Having fun (or at least not being too bored)	Preventing fun and boring the user

Figure 3—BECS on-line help system

wanted on a specific text box or button (Figure 3).

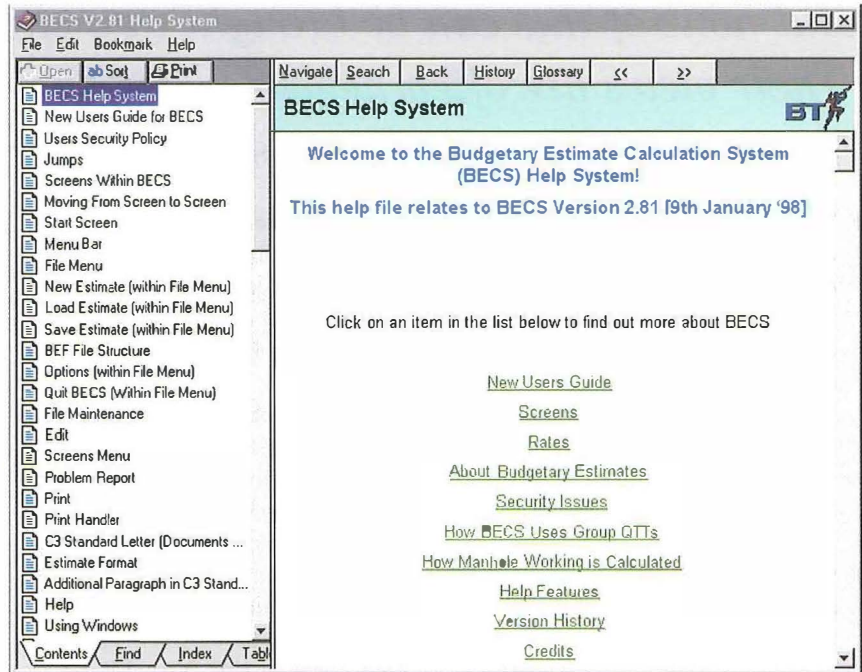
The use of a built-in manual should make getting help an immediate process rather than a lengthy search for the user manual. There is, however, a danger lurking within help systems contributing to the belief that software can often be rude.

Rude software

When BECS was developed the introduction of PCs within planning offices was still underway. Those users who had a PC at home readily accepted the use of software tools at work, while others realised that they had a whole new set of skills to master. Being a developer does not help understand how people use computers and only observation can solve this. For example, with the on-line help system, all users' problems can be solved as long as the operation of a context sensitive system is understood. To a seasoned PC user the idea of pressing the F1 key when in trouble is obvious, and yet this is hardly intuitive to the PC novice. This means that there should be several ways to achieve the same function in the software. While computer experts can cope with the F1 key and short-cuts, the novice needs a button marked 'help'.

Once in the help system the next trap to avoid is the 'obvious yet uninformative statement', such as explaining that the 'Cancel' button 'cancels the current operation.' This is correct but leaves the user pondering as to what the current operation actually means, and what is cancelled.

Error messages are a great source of rudeness in that they often scold the user for trying to do things that are deemed to be wrong. Automation should allow the user to concentrate on the task in hand and should inform rather than instruct. BECS was designed with usability in mind, and yet looking back it would seem that some of the messages intended to help were actually overwhelming the user with information (Figure 4).



While trying to be very helpful by giving a warning, asking a question in the same dialog box can be confusing. It could be imagined that after seeing this message several times it would simply be ignored as users learnt that clicking 'Yes' cancelled the screen as they first intended.

The temptation to include unintentional rudeness is present because the developer is too close to the system, which is why usability testing is important.

Testing usability

As mentioned before, observing users as they operate the program gives vital information about the usability of the design. This can be done at local offices or by arranging workshops, or even by issuing questionnaires. The developers can also carry out usability tests, as this will give a feel for the problems that the users may encounter. The disadvantage is that the developers

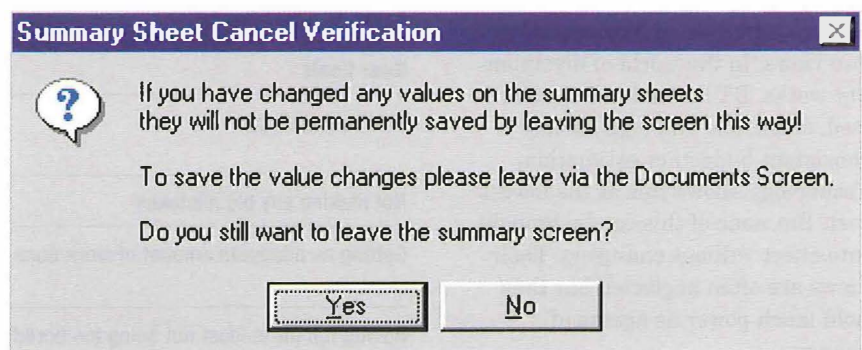
can often miss interface problems, as they know the system too well.

BECS evolved over the course of several months through the 'development/user test' iterative loop and while this was considered to be the best strategy for long-term success several potential problems emerged.

Losing sight of the goal

Development of BECS was largely geared towards satisfying users' needs, but as explained before, there exists a wider remit. At regular intervals it is necessary to check that the operational goals are still being met because a purely user-designed system may turn out to be of diminished use to the company. As BECS is just software that performs calculations it is easy to become absorbed in levels of detail inappropriate to the task at hand simply because it is possible to do so. BECS produces a budgetary cost estimate and yet, during development, users were

Figure 4—BECS being very helpful






Figure 5—Rounding options for estimate total

asking to include components whose costs were insignificant to the overall outcome. For schemes costing several thousands of pounds quibbling over the cost of pole sundries worth £12 was irrelevant. To the user, this level of detail seemed to matter although they requested to be able to round up the final budget cost to the nearest £100, £1000 or £10,000 (Figure 5).

The fact that the software allowed the introduction of an ever-increasing collection of smaller and smaller costed items did not imply that it was desirable. So beware: just because something is possible does not make it essential when considering the main business objective!

Taking longer

Software development does not actually require any user involvement at all and can be carried out by a single person in a dark room with a supply of strong coffee. In fact, talking to users takes longer and may mean that changes are necessary. Any iterative process will take longer than taking the linear path from start to finish, but the outcomes will be different. Without user involvement the business objective will be reached while in the laboratory, only to be scuppered upon implementation. The act of 'throwing the product over the wall' and letting the user community cope has to be abandoned, and in the whole-life analysis of the system, the consultative approach wins every time. This is not to say that users are involved in every aspect of the development because building in business compliance is also necessary. Although it might appear that the iterative process takes longer, it actually saves time because the outcome is the desired one. Having an undesired outcome which then has to be reworked to a desirable one is a guaranteed route to delay and extra cost. This central message of concurrent engineering methods has been

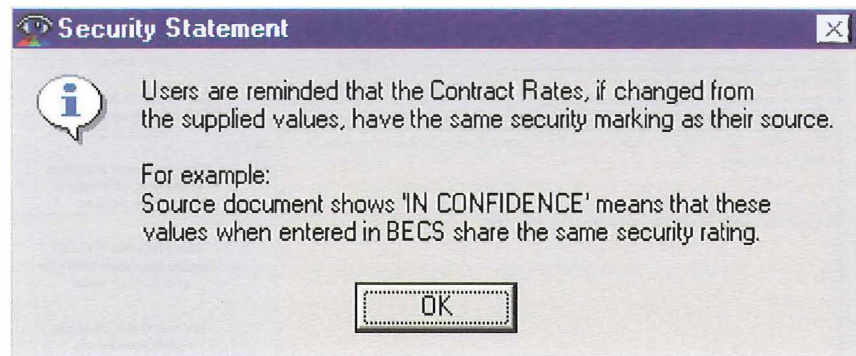


Figure 6—Data security screen in BECS

taken to heart by many companies, and reports of huge savings in cost and time in product development have been reported⁸.

Building in compliance

Compliance issues usually involve aspects that do not directly concern the user and yet without them the system would not be ready for implementation.

Quality

When writing software it is easy for things to get out of hand as the code grows at an alarming rate. Applying basic quality rules during development should be considered beneficial rather than just mandatory. Having development notes before writing

code, making regular back-ups, commenting code, and applying version control will make a big difference to the end result without too much overhead. It is only when the last day's work is accidentally deleted, or you have no idea how a code segment you wrote one week ago works, that you realise the power of quality procedures. The main advice is to scare yourself early in the development so that you use quality to your advantage.

Security

A map displayed on a screen is somehow more 'accurate' than its paper cousin, while data on a hard-drive is unlike a document marked with the words 'in confidence'. All systems holding sensitive data are a potential security risk as data can be

readily available, altered and copied by anyone with suitable access. Users need to be reminded that security is an important part of their work (Figure 6).

Building in security into software protects data and should be considered at the early stages of development, as bolting it on later will be far more difficult. Fortunately, there is a considerable amount of help available within BT to ensure compliance and this imposes the minimum overhead on the main development.

Millennium

When BECS was started in 1995, the millennium issue was not prominent, but as the year 2000 approaches the impact has become more apparent.

The main advice is to scare yourself early in the development so that you use quality to your advantage.

The experience of testing BECS for millennium compliance has demonstrated that fixing a deployed system takes much longer than building in compliance during development (if it had been known at the time!) The point is that even when the year 2000 appears on our calendars there will still be a need to build in compliance for software. Developers may use environments and add-in components that are not compliant and it will take several years for the software and hardware industries to solve the problems.

Testing times

BECS is a modest-sized program of approximately 30 000 lines of code and once complete was tested thoroughly before national implementation. Users were selected from

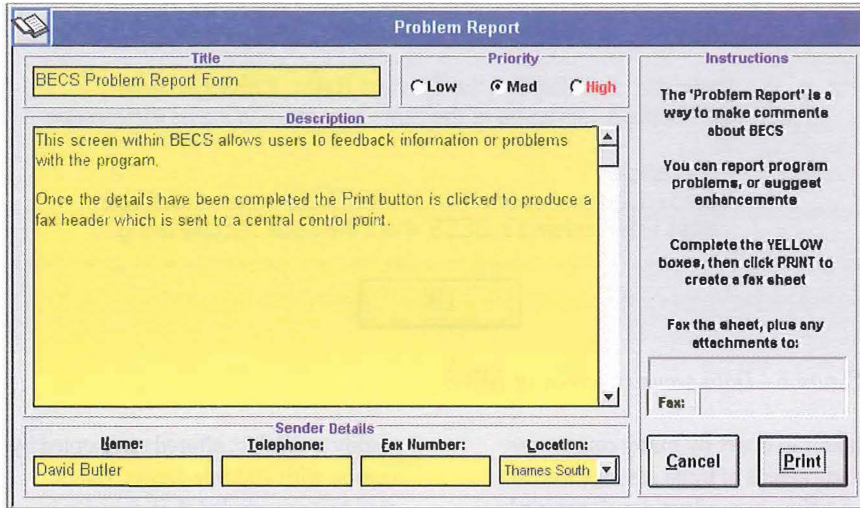


Figure 7—BECS problem report facility

around the country and asked to use the program and feed back any problems for a trial period of two months. To facilitate this process a ‘Problem Reporting’ screen was built into BECS (Figure 7).

This screen enabled the user to enter a problem title, assign a priority level and explain the problem that was encountered. Once user details were completed a fax header sheet was printed and sent back to the development team to be recorded and solved. The inclusion of an issues management system in BECS enabled problems to be captured quickly (rather than have to leave BECS and report back using another method) and recorded so that this would form part of the quality audit process. The fax sheet method was used as many planning offices did not at the time have e-mail or intranet facilities. This method of feedback was found to be so useful that it remained in the program when it was launched thus allowing users to feed back bug reports, suggest new features or comment upon any aspect of the program.

Towards the end of the trial stage, the user groups were given a detailed questionnaire for assessing their views regarding BECS in categories such as workflow, acceptability of calculations, ease of use, and time to create estimates pre- and post-BECS. This information would contribute to the measures set to determine whether BECS had met the business needs.

During testing, all users’ comments were carefully considered and appropriate changes made to the software before a one-month pilot stage commenced. A two-stage review process was adopted to apply ‘quality

gates’ between the different purposes of testing. The trial stage ascertained whether BECS was ‘fit for purpose’ as a budgetary calculation tool and found major problems such as installation difficulties, non-functioning code modules and significant calculation errors. Once addressed, the software progressed through the quality gate to the pilot stage where interface design, program features and calculations were refined. Once through the two-stage testing BECS was ready to be distributed.

Some important lessons were learnt during the test stages, especially about software complexity and associated test methods. Even with a relatively small program there are many aspects to check, such as the installation process, correct operation of the user interface, checking the calculations and testing combinations of inputs. One of the essential properties of software is that it *can* handle complexity and yet this feature makes it difficult to test even with rigorous mathematical methods. It has been suggested that one of the difficulties is that software is ‘*invis-*

ible and unvisualisable..... unlike a floor plan of a building..... where contradictions and omissions become obvious’⁶. Therefore the role of the software tester is very important to the overall success of the project and reinforces the view of collaborative software development and implementation. In conclusion, it would seem that error-free software is virtually unattainable, although this should be seen as a challenge rather than an inevitable obstacle.

Ready to implement

Having considered users’ needs, checked that the business goals have been satisfied and achieved compliance, the software is ready to distribute. This phase is equally as important and it marks a shift from ‘technical’ to ‘logistical’ thinking.

Implementation Strategies

There is a range of choices when deciding how to implement products (Figure 8).

The choice of the strategy can facilitate or frustrate the main objectives of the product and so careful implementation planning is required⁹.

BECS relied upon ‘Zone implementation managers’ (ZIMs) to act as local focal points when installing and supporting the national roll-out process. The ZIMs were local line managers who were invited to a national launch event where they were

Figure 8—Implementation strategies⁹

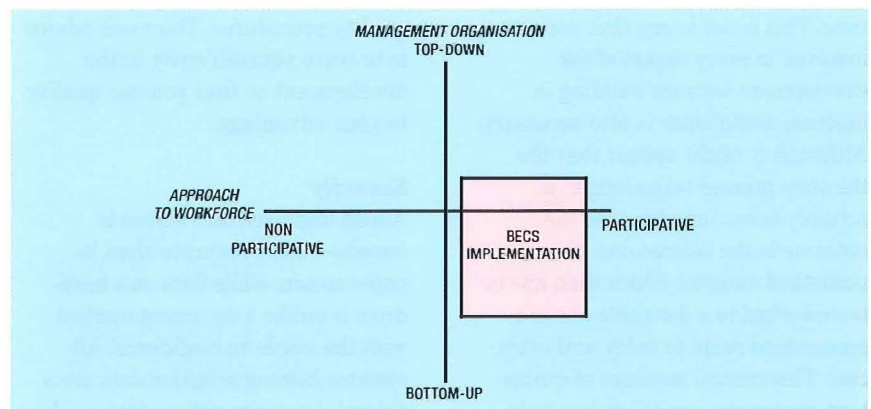




Figure 9—BECS installations

shown the program and briefed as to their role. The main function of the ZIMs was to distribute the floppy disk and installation instructions, act as front line operational support and feed back any issues to the development team. Any technical issues during installation were handled by the development team directly. ZIMs were selected to enable a participative, bottom-up implementation strategy to be used, with some top-down technical support if required.

Users were also encouraged to give their views via team meetings and through the BECS Problem Report screen. A few months after the national launch BECS could be found across the country with more than 200 users (Figure 9).

It can be argued that there is no best way to introduce technology into the work place and that each implementation approach has its benefits and problems. Implementation approaches are often chosen by considering factors disassociated with the technology, such as timescales, commercial survival, business performance and the

influence of product 'champions' to overcome any resistance to change. By taking a predominately bottom-up approach the ZIMs became the product champions with the flexibility to introduce BECS in the most appropriate way in their geographical zone. The highly participative adoption was designed to make BECS the budgetary estimate tool of *choice* because it provided an easier, and quicker calculation method than local spreadsheets. The choice of strategy may have taken slightly longer to achieve the business aims than using a more directive approach, but it did enable the social shaping of BECS to begin, which has led to benefits.

Coping with change

One of the best features of software is that it can be more easily changed and at lower cost than other products such as cars or buildings. Frederick Brooks⁶ suggests that '*all successful software gets changed*', and cites two main reasons. The first is that end-users become comfortable with the software and once found to be useful

will try it in new cases at the edge or beyond the original intended specification. Users will optimise the way in which they use the product and will start to demand new features.

The second change process is that computer hardware evolves rapidly and constraints such as memory costs, or processor speeds are overcome. New operating systems such as Windows 95 introduce features that aid productivity (for example, use of long filenames), and programming languages become more powerful (for example, Visual Basic 5).

There is a third reason why software gets changed and this is due to business process refinements. When BECS was introduced the costing mechanism was based around activity codes, but as previously noted the introduction of QTTs and then GQTTs meant that BECS had to be rewritten. The key message is to realise that change is good and does not imply that the initial requirements capture was flawed. Understanding and keeping pace with process change is the vital element to product support and is assisted when good communication exists between end-users and developer.

Effects—brief discussion of social shaping

Having progressed through the many iterations of the product life-cycle diagram as shown back in Figure 1, it is useful to compare the prototype of BECS with the released version. Over the course of many months the program evolved due to end-user involvement. This social shaping determined the way in which the program accepts inputs, how it applies calculations and then presents the answers while preserving the business need. It is postulated that a major part of technology success is that end-users *want* to use products because they achieve both user and business goals (see Table 1) and an understanding of social shaping is relevant. If the product life cycle is to be retained as a useful model for technology development

and implementation then 'all the stages and all of the feedback loops' must be considered¹. It has been suggested that projects that succeeded were likely to contain engineers who were good practical social scientists, even if they didn't realise so at the time¹⁰.

Conclusions

BECS has been successful, in part, due to the collaborative and participative approaches taken during development and implementation. The program has helped to shape working practices and has been shaped by the interaction of the end-users who have played a vital role in adopting BECS as their budgetary estimate calculation tool. The experiences learnt are currently being used in the development of a suite of programs that will ultimately supersede BECS to provide end-to-end automation for the whole diversionary works process.

Acknowledgements

The author wishes to thank all his BT colleagues within Change Programmes (particularly Mick Saunders and Ray Garlick) and *networkbuild* (Denis Taylor and his team of sector managers, in particular Roger Fagg, and all project engineers) for their cooperation, support and help during the development and deployment of BECS. Local management at BT Laboratories, specifically John Mellis and Ray Hooper, provided support and assistance.

Glossary

BECS Budgetary estimate calculation system
DCE Detailed client estimate
GQTT Group quality task times
JRET Jerrard rapid estimation tool
MEBF Man-hours per estimated basic-hour factor
QTT Quality task times
VB Microsoft Visual Basic™
ZIM Zone implementation manager

References

- 1 EDGE, D. The Social Shaping of Society. *Information Technology and Society*, 1995.
- 2 HMSO. Measures Necessary Where Apparatus is Affected by Major Works (Diversionary Works). A Code of Practice, June 1992.
- 3 BT. Access Network Synthetics (Quality Task Times), BT ISIS document PRO|SOA|A410, Issue 3.
- 4 BT. Group QTT Portfolio—Contract Installation., BT ISIS document PRO|SOA|A904, Issue 1.5.
- 5 COOPER, A. About Face—The Essentials of User Interface Design. ISBN 1-56884-322-4.
- 6 BROOKS, F. No Silver Bullet. *Information Technology and Society*, 1995.
- 7 BT. Corporate Standard Visual Design of User Interfaces. Corporate Relations Department and Human Factors, 1993.
- 8 HOOPER, RAY; and JONES, CAROLE. Lean Operations—The Route to Competitiveness, *Br. Telecommun. Eng.* April 1998, **17**, p. 2.
- 9 McLOUGHLIN, I.; and CLARK, J. Technological Change at Work, *Information Technology and Society*, 1995.
- 10 MacKENZIE, D. Missile Accuracy: a case study in the social processes of technological change. In Bijker, Hughes and Pinch, 1987, pp. 2–6.

Biography



David Butler
Networks and
Information Services,
BT UK

David Butler has worked at BT Laboratories since 1986 on a variety of research projects. He has investigated cable installation techniques such as moleploughing and trenchless technology, and designed test equipment to measure optical-fibre characteristics and has a patent for measuring blown-fibre tube friction properties. He has also researched and presented papers at conferences regarding 3D mapping techniques such as ground probing radar. Productivity improvement projects such as the investigation of in-vehicle navigation systems using GPS technology led to software development for BECS and its support systems. During his time at BT he gained an honours degree studying with the Open University.

Graeme Maxwell

Micromachining Future Networks

Micromachining is an essential area of technology for BT's current and future network. It is used within every optical network in the world, but is largely unseen and unknown, yet without it there would not be the level of integration of the hardware components essential for the network.

Introduction

It is becoming clear to even the most conservative of observers that there is a fantastic growth in both the Internet and data traffic in general. This is having, and will increasingly have, a very significant impact on BT's networks. The demand for bandwidth is forcing network designers to seek techniques and technologies which will provide solutions to their current capacity problems and be scalable for the expected demand in the future. This in turn is forcing engineers to develop ever more effective and cheaper solutions to the technologies which go into BT's core network.

The purpose of this article is to outline the technological developments which will impact on these current and future networks, within the realms of micromachining techniques and technologies which are used to put together these systems. The article outlines some of the techniques used in micromachining and why they are important, using appropriate case studies.

Micromachining is a term which can cover a multitude of techniques. Photolithography and plasma etching are used in every silicon chip and in every laser device, and both are examples of micromachining. Using wet chemical etching techniques which are sensitive to crystallographic planes in semiconductors is micromachining. Depositing materials with precise levels of stress and subsequently releasing them from the underlying substrate is micromachining. Locally doping a material to change its properties when subsequently etched or deposited upon is also micromachining.

The important thing is how can the various techniques which exist and are being developed be combined to produce components or subsystems which will have improved functionality for future networks.

Many people are familiar with the development of computer chips. Moore's law¹ has described the increase in speed and power of the computer chip which has been brought about by the reduction in size of individual transistors and increase in functionality of circuits. However, impressive as this may be, this technological development is of secondary importance to BT's core network. The core network carries its information in the form of photons rather than electrons, down optical fibre made of silica rather than wires made of copper, using transmitters and receivers based on indium gallium arsenide, rather than transistors made of silicon. The technological challenges for the optical network have not been 'how can we make this smaller and faster?', more 'how can we put together these components which are derived from totally different material systems and when individually optimised are bordering on being incompatible?' This is the area from which the micromachining techniques being investigated have arisen. That is to say it is primarily from the needs of packaging components that the micromachining techniques have been investigated at BT. However this is not the sole use of micro-machining by any means. Many devices—gyroscopes, accelerometers, magnetometers and switches are being developed using micromachining². Furthermore, as optical processing techniques

Micromachining will play a role in producing functional elements in its own right, rather than as a packaging aid as it is currently used.

currently being demonstrated mature, then optics will move further into switches and routers, and the importance of micromachining will become even greater.

Micromachining will then play a role in producing functional elements in its own right, rather than as a packaging aid as it is currently used.

This article consider as an initial case study the laser-to-optical-fibre interface which is present in every transmitter in every optical network. This demonstrates the packaging benefits of micromachining as well as outlining some of the techniques used, which are applicable elsewhere.

Laser-to-Fibre Interface

The lasers currently used in the core network are made using semiconductor materials such as indium gallium arsenide (InGaAs). The refractive index of this material is around 3.4, with a refractive index difference between core and cladding of around 0.24. The core size of the device is around 1.5 microns wide and 0.2 microns thick. These numbers determine the size of the optical beam (or mode) which exits the laser facet and how rapidly the beam diverges. This directional characteristic of the emitted mode is referred to as the *far-field distribution*. The core is rectangular, which leads to an elliptical beam exiting the laser. If we now compare this with the equivalent values in an optical fibre, we see a core index of around 1.45, with a cladding index of 1.444. The core size

is around 8 microns in diameter, and the mode is circular.

The object of the exercise is to couple as much of the laser output as possible into the fibre. As it stands in Figure 1, the amount of light coupled from laser to fibre is very low—around 10–12%

In order to improve the coupling (to increase system power budgets), the options are to redesign the laser to make its far-field look more like that of a fibre; put a lens on the end of the fibre; or put some optics between them to improve coupling. The latter two options are what are used in manufacturing. The problem with this is that it is expensive to fabricate the fibre lenses, or add other components, and most importantly, the packaging tolerances are very low—requiring positional accuracy of less than 1.5 microns transverse to the facet to obtain around 60% coupling efficiency. This means that these components have to be actively aligned, requiring someone to position the components together while monitoring the coupled power, and then glue them in place. This is an expensive and time consuming task. Figure 2 shows two

Figure 2—Two lensed-end fibres aligned to semiconductor devices




Figure 1—Comparison of laser and fibre farfields





Figure 3—Schematic showing tapered laser structure and cross-section detailing precision cleavage planes fabricated using selective crystallographic etches

such lensed-end fibres aligned to semiconductor devices.

The more attractive approach adopted here has been to redesign the laser, as in Figure 3, and combine the new laser with a micromachined motherboard to passively align all of the components. This gives the same levels of coupling efficiency without the need for active alignment.

Here the micromachining has been used in both the laser fabrication and the motherboard. The result of redesigning the laser is to produce a device with a large spot. This has the result of making the laser far-field look more like the fibre, and making the alignment of the two components much more tolerant of positional error. This then brings the whole process into the realms where micromachined motherboards and passive alignment can play a role because the micromachined structures can meet the tolerances required, when previously they could not.

The motherboard fabrication has involved using selective etching of different crystal planes within a silicon substrate to produce 'V'-grooves for the fibre to sit in, as well as patterning an end stop against

Figure 4—Schematic of micromachined silicon motherboard

which the precision cleaved laser chip is butted. (See Figure 4.)

This still requires the fibre to be held in place while being glued. A further refinement of the technique involves some additional steps of micromachining to integrate silicon nitride (SiN) clips³ to allow one to push the fibre into position (see Figure 5). The requirement here is to be able to deposit a SiN layer with very low stress, and to be able to selectively etch away the silicon below the SiN material.

This now provides a simple, lower-cost approach to the fabrication of transmitters. These techniques are now being pursued by manufacturers with a view to simplifying their packaging process and reducing costs.

LIGA—TOAD

A second case study looks at a more complex system. This uses a different technique for the fabrication of a motherboard into which many more components are placed. The interest in this example arises from the use of a process known as LIGA⁴, which is designed to produce a very-high-precision metal master from which replicas can be made using either injection moulding or hot embossing in much the same way as CDs are stamped out from an original master. The difference in this example also stems from assembling a subsystem rather than dealing with a single interface, in this case a terahertz optical asymmetric demultiplexer (TOAD)⁵.

The TOAD forms a functional building block for the fabrication of more complex circuits such as all optical regenerative memory with full read/write capacity and all optical binary half-adders. The basic device is a loop interferometer, with a non-linear semiconductor optical amplifier (SOA) placed asymmetrically within the loop. There are additionally directional couplers to send control pulses used to switch the data. An optical fibre version is shown in Figure 6. These are the





Figure 5—Cross-section of SiN³ clip holding a fibre, with SEM photo showing the fibre/laser interface with SiN fingers holding the fibre

types of component which are being investigated with the aim of developing optical signal processing elements which will feature more strongly in future generations of high-speed networks⁶. It is predicted that these devices will be capable of operation at 100 Gbit/s.

It is beyond the scope of this paper to detail the operation of the TOAD, and the reader is directed to the references for more information. However, the device makes use of nonlinear effects inside a semiconductor optical amplifier (SOA) contained inside an interferometer to enable high-speed switching. For our purposes we are considering a micromachined equivalent, with the potential for mass production.

The LIGA technique is a method of fabricating high-aspect-ratio features with a high spatial accuracy. A collimated beam of X-rays exposes a thick layer of poly-methyl-methacrylate (PMMA) which has been spun on to a metal substrate. This is developed, giving deep vertical features. The sample is then electroplated, allowing

the nickel to grow up from the metal substrate, through the channels which were removed at the development stage. Once the plating is complete, the remaining PMMA is removed. One then has the inverse of the pattern originally exposed—this is the metal master from which copies can be made.

In this example, the motherboard was 500 microns thick with spatial accuracy of the features of 1 micron. This then allowed commercially

Figure 6—Optical-fibre based TOAD

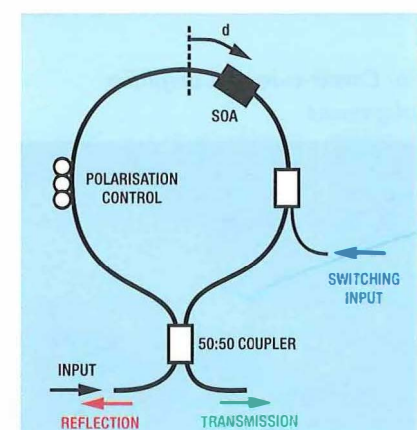


Figure 7—LIGA motherboard detailing SOA, ball lenses, fibres and prisms

available components to be dropped into place within the LIGA motherboard. Figure 7 shows schematically the motherboard design, with Figure 8 detailing a picture of the motherboard itself, and Figure 9 showing a partially assembled motherboard. Additional components required to assemble the optical components in place were fabricated at the same time, so any systematic process variations affecting the motherboard will also affect the piece parts and so the overall structure can still be assembled.

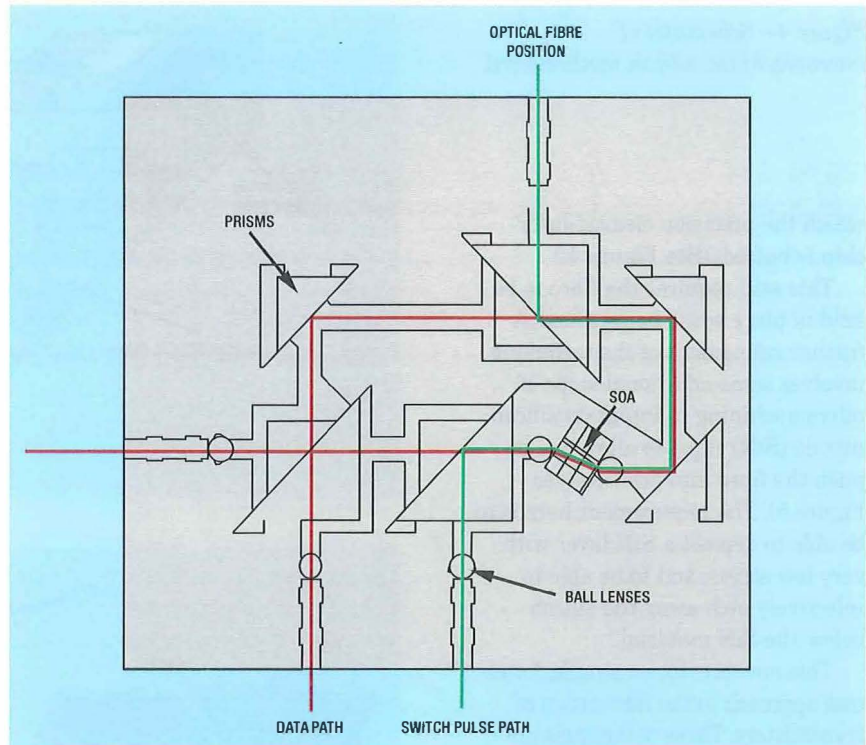
While this technique is very attractive, it requires a cyclotron to produce the collimated X-ray beam which is used to expose the thick polymer layer. This is not available in most laboratories, and BT currently has access to a facility at IMM-Mainz in Germany. However, similar high-aspect-ratio structures can be achieved with alternative machines producing deep etched features in silicon², and these show a great deal of promise for this type of work.

Figure 8—Details of the LIGA motherboard

(a) Fibre and ball lens alignment



(b) Corner-cube and amplifier alignment



Concluding Comments

It is hoped that the reader has gained some idea of one or two of the technologies involved in micromachining, and why the industry is interested in it. These technologies are continually being developed in response to novel materials and components which bring with them their own specific requirements for integration⁷.

Acknowledgements

The author acknowledges the assistance of various people for the pictures used in this article; namely, Colin Ford, Ian Lealman, Arbinder Pabla and Alistair Poustie.

References

- 1 <http://www.ccta.gov.uk/25years/moore.htm>
- 2 <http://www.stsystems.com/mems.html#mems>
- 3 <http://www.camcon.co.uk>
- 4 <http://www.imm-mainz.de/contents.html>
- 5 http://www.hhi.de/englisch/PRESSE/Presse96/Jahn_en.htm
- 6 <http://www.labs.bt.com/people/cotterd/index6.htm>




Figure 9—LIGA-TOAD with fibre and ball lens assembled

7 BURR, M. Quantum Dots—Designer Atoms. *Br Telecommun. Eng.*, Oct. 1996, **15**, p. 230.

Biography

Graeme Maxwell
Networks and Information Services, BT UK

Graeme Maxwell received his B.Sc. and Ph.D. degrees at Strathclyde and Glasgow Universities respectively. He joined BT in 1989 to work on planar silica waveguide devices in the Oxide Glasses and Fibres group. He is currently responsible for the Technology Research group looking at semiconductor devices, silica waveguides and fibre, packaging and micromachining.

Chris Gibbings, Mark Barrett, Margarida Correia, Uma Kulkarni and Kevin Smith

Broadband Multicast on BT's Futures Testbed

Internet video and audio will be revolutionised by IP multicast, which allows a personal computer to send data to tens of thousands of destinations at the same time. Replication by the network means that only one copy is needed on any link, so that congestion is much less likely. This article describes work on broadband video multicast over the BT Futures Testbed advanced network.

Introduction

The Internet can give global reach to a child's Web page, a professional news site, a 'make money fast' scam, an airline timetable or an on-line shop. Finding resources may not be trivial (although Web searching is becoming more sophisticated all the

time), and believing the information on some pages may not be wise, but it has never been easier to upload text and images for worldwide distribution.

Video and audio material is also coming on-line. However, link sizes on the Internet are a major obstacle to widespread deployment. Only a big server is capable of serving data

The Futures Testbed

Over 700 researchers are connected to the Futures Testbed network at BT Laboratories^{2,3}. The network was set up to provide a high-speed backbone connection between eight buildings on the Martlesham site with high-quality switched desktop connections. Currently over 1200 end systems (desktop computers, printers, servers, etc.) are supported. The testbed has been gradually upgraded to have a mixture of 155 Mbit/s asynchronous transfer mode (ATM) and Gigabit Ethernet in the backbone. The majority of desktop connections are switched 10 Mbit/s Ethernet, with an increasing number of 100 Mbit/s connections to servers and high-performance PCs. The network is fully enabled for Internet protocol (IP) multicast, and a wide range of multicast applications have been used over the campus network.

The testbed has also been extended over the wide area, with connections to London, Cambridge, Norwich and Colchester via LEANet

(London and East Anglia Net). Wave-length division multiplexing

(WDM) is used to make full use of the fibre pairs, and an STM-16 (2.5 Gbit/s) synchronous digital hierarchy (SDH) ring allows 155 Mbit/s links between individual sites to be set up using add-drop multiplexors on the network. The LEANet infrastructure has been used for IP directly over SDH and IP over ATM running on the SDH links (both 155 Mbit/s) and IP over gigabit Ethernet. Recent work has included ground-breaking advances in sending gigabit Ethernet over long distances (>180 km)⁴. Various experiments with IP multicast have been carried out over the wide area network, including broadband IP video transmission over SDH, ATM and gigabit Ethernet.

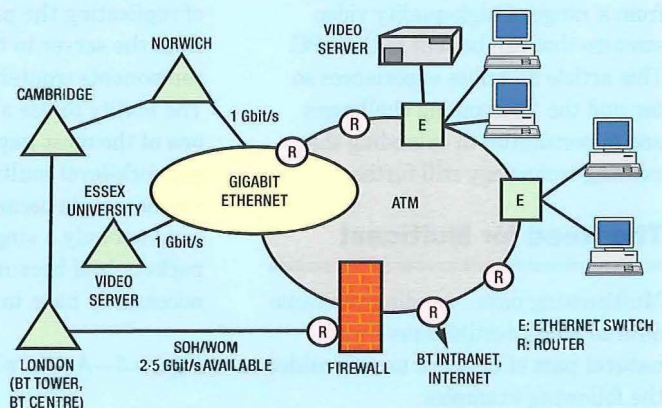


Figure 1—Futures Testbed

Figure 2—(a) Application-level, and (b) network-level multicast

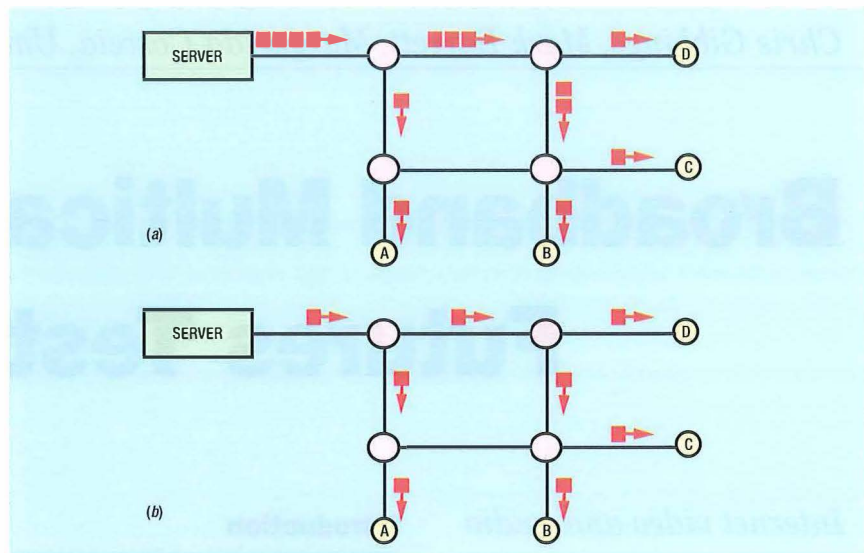
streams to large numbers of viewers. The link to the Internet has to be big enough for all these streams. They then have to fight through (and contribute to) the congestion on the Internet.

An emerging technology, IP multicast, is revolutionising this process, by distributing the copying task across the network rather than centralising it on the server. The attraction of multicast for a wide area, relatively congested network is that no more than one copy of the data is sent over any individual link, however many copies are made downstream. The overloaded public Internet would certainly benefit from such efficiency. In advanced networks with high-capacity backbones the selling point for multicast is that it allows simple cost-effective servers to be used. One such advanced network is the Futures Testbed, which supports over 700 researchers at BT Laboratories. They have been enthusiastic early adopters of multicast technology, and the network is now fully enabled for broadband multicast. Users can select from a range of high-quality video streams that can be sent to their PC. This article describes experiences so far, and the forthcoming challenges and opportunities in extending this exciting technology still further.

The Need for Multicast

Multicasting data—sending the same data to many destinations—is a natural part of network use. Consider the following examples:

- video broadcast to employees in a company,
- software updates downloaded over the network,
- share price information sent to dealing screens,
- radio across the Internet, and
- many-to-many videoconferencing.



In each case the same data is sent to a number of parties. Many of these applications involve 'push' technology, which has recently become a very important aspect of communication over the Internet and company intranets¹. There are several ways to send out the data.

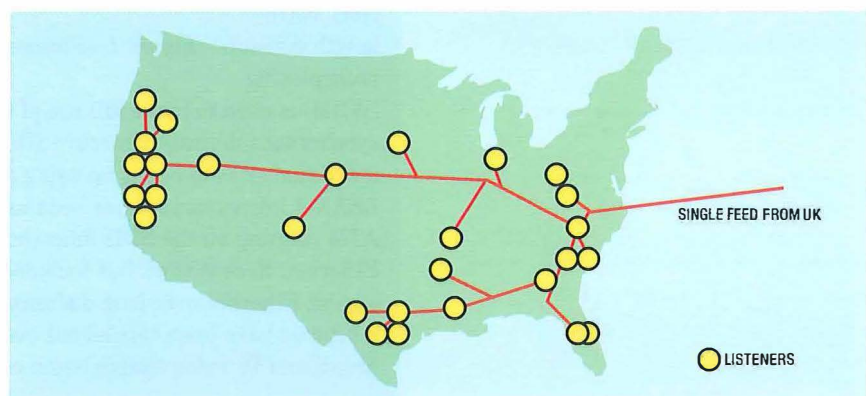
Application-level multicast sends out packets to each of the required destinations, as shown in Figure 2(a). Note that some links will be carrying many copies of the same data, which is a big disadvantage when link capacities are small.

Multicasting at the network level (Figure 2(b)) ensures that only one copy is sent down each link. The task of replicating the packets is moved from the server to the network components (routers, switches, etc.). The ability to use a simple server is one of the most important aspects of network-level multicast. The server can be simple because it needs to send out only a single stream of packets, and because it does not necessarily have to concern itself

with users choosing to subscribe to the stream or to stop subscribing. As the numbers get bigger and bigger these factors are ever more important.

The scalability issue is highlighted in Figure 3. In a few years time when large sections of the Internet are multicast-enabled, it will be quite possible for a European radio station to feed a compressed stream (perhaps 10 to 20 kbit/s) to their Internet service provider (ISP). Despite the small investment needed to keep this going, the radio station could have global coverage using multicast at the network level. Even with large numbers of listeners in the United States the loading on the transatlantic link is very small. Application-level multicast can become more network-friendly by putting in proxy servers in the USA to relay the audio stream, but this requires extra investment and ongoing management. Flexibility is another area in which network multicasting wins. Demand can be dealt with immediately, whereas there will always be a time delay in commis-

Figure 3—A future Internet radio application in action



With so many advantages it is clear that copying data in the network is of great interest.

sioning or removing proxy servers as listener numbers go up or down.

With so many advantages it is clear that copying data in the network is of great interest. The Internet protocol (IP) is the global standard for data networking. Even traditionally separate areas such as voice communications and video conferencing are beginning to migrate to IP. A huge amount of development work has been going on in IP multicast in recent years, leading to the applications and network functionality described in this article. Other data protocols such as Appletalk and IPX are likely to linger on for many years but will not be considered further.

IP Multicast

Figure 4 shows the bare essentials of how IP operates. Data is sent in packets with a source address and destination address (and a few other bytes of header data with various options). Conventional unicast IP, upon which the World Wide Web, e-mail, file transfer and all other Internet applications are built, travels across the network on a hop-by-hop basis. In a single building there might be one or two hops, while a journey through 20 routers or so might be typical over the Internet in general. Routers decide where to send the packets on the basis of the destination IP address. They have quite specific information about the IP numbers near them and more general routes for other ranges of numbers.

Printers, computers and other network-connected machines all have a 32-bit IP address, allocated out of

Figure 4—Unicast data transmission




Table 1 IP Number Ranges

First address	Last address	Available	Use
0.0.0.0	223.255.255.255	~4000 million	Unicast addresses (Class A, B and C)
224.0.0.0	239.255.255.255	~250 million	Multicast addresses
240.0.0.0	255.255.255.254	~250 million	Reserved for future use

the unicast range shown in Table 1. The distinction between classes A, B and C should really be of historical interest only, but some network software still lives in the past.

Changing a few bytes in an IP packet so that it goes to a multicast rather than a unicast address makes an enormous difference. Multicast addresses are also known as *groups*, and a network-connected host can subscribe to one or many groups. In Figure 5 several hosts have subscribed to multicast group 239.255.0.2. When a PC sends a packet with destination IP address 239.255.0.2 the packets are copied by the network to all the subscribed hosts. This simple change in behaviour has profound implications for both the efficiency and scaling of many Internet services.

Sending a multicast stream is very easy—the source simply puts the appropriate multicast destination address in the IP packets it sends out. Potential recipients request this multicast stream using the Internet group management protocol (IGMP). The computer sends an IGMP ‘join’ packet to its local router to request a stream. If the router is already receiving the group it simply starts to copy packets to the computer. If it is not currently routing that group

then it has to request it from another router. This is where the *multicast routing protocol* takes over, as described later in this article.

The router sends out periodic IGMP messages to check whether the stream is still wanted. The computer must actively confirm that it needs the stream. If three periodic requests are unanswered, the router stops sending out the stream. IGMP has been designed to work well when several computers are on the same Ethernet. Only one reply is needed to ensure that the stream continues to be forwarded. Once the other computers have seen this reply they know that they do not have to answer. A random timer in the software ensures that they do not all reply at once.

Sending multicast data is remarkably easy, and receiving a particular group is straightforward. The only other aspect to consider is how receivers know which multicast groups are available and what kind of content they provide. A central server which can send out a list of groups is one answer, while a more distributed solution is provided by a session directory which is multicast on a well-known address.

Multicast Applications

This section describes some of the applications that can use IP multicast once the basic building blocks are in place and a well-engineered routing protocol is running between the routers. The pioneering efforts in IP multicast took place in the *Mbone* (or *multicast backbone*). This is a wide area IP multicast network which runs over the public Internet, which has the advantage of connecting together a

Figure 5—Multicast data transmission

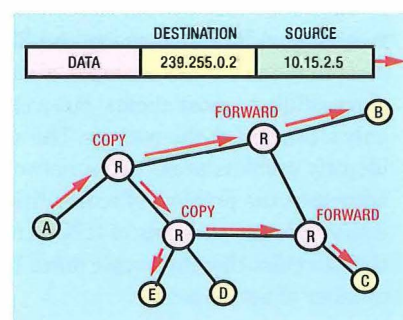


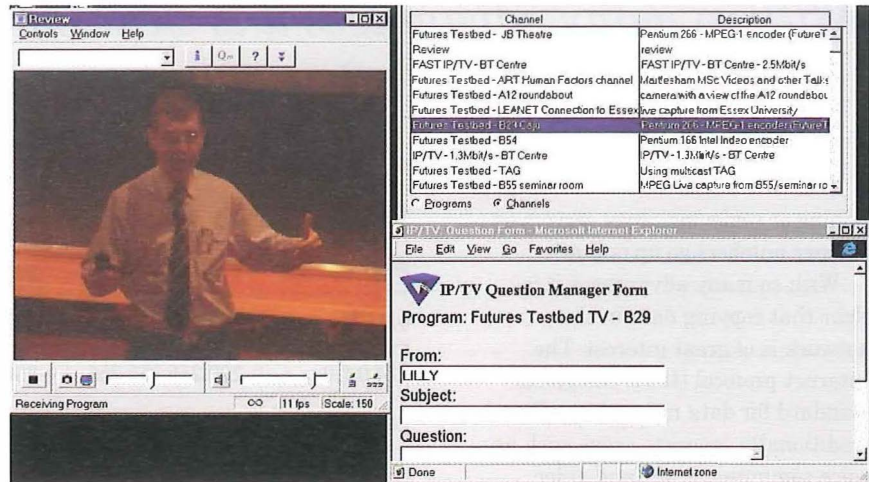
Figure 6—IP/TV viewer

global community of university and industrial researchers but also has the disadvantages of low bit-rates, high packet losses and restricted content due to copyright considerations. The emerging generation of multimedia applications builds on the experience of the Mbone workers and takes advantage of the high-capacity local-area and campus networks that have recently become possible. A good example of this kind of network is the Futures Testbed, a campus network supporting 700 researchers at BT Laboratories, which is fully multicast enabled. The technology in use is described in more detail later in this article.

Three IP multicast applications have been used extensively on the Futures Testbed network. IP/TV from Cisco⁵ and NetShow from Microsoft⁶ use multicast for the distribution of video and audio channels, and Cu-SeeMe from White Pine⁷ has the capability of establishing multipoint videoconference sessions using multicast.

The distribution of video and audio to corporate desktops is usually called *intranet-TV*. Both IP/TV and Netshow are PC-based solutions that can deliver live and pre-recorded content to the desktop. These applications consist of content servers, content managers and viewers. IP/TV enables the use of different video and audio codecs like VxTreme, H.261 and MPEG-1, which make available different bit rates and consequently offer differing image and sound quality. Among these, MPEG-1 is the codec that gives the best quality and is the one used in the BT Laboratories' trial.

The IP/TV software consists of three modules: viewer, server and programme guide. At the moment there are 12 IP/TV servers on the Futures Testbed network providing live and pre-recorded content to about 80 desktops with IP/TV viewer software. The servers for live content are in lecture theatres and seminar rooms, with a LEANet fibre link to Essex University running at 155 Mbit/s (IP over SDH) or 1 Gbit/s (IP over Gigabit



Ethernet). The content varies from talks and lectures to general information about the company and the world. The server transmits its content based on parameters specified in the programme guide. The programme guide is a Web-based tool that is part of a Web server (NT or UNIX) for administrating and managing the IP/TV servers. The viewer can then use the programme guide information to request streams from the local router using IGMP.

The viewer (Figure 6) displays the list of programmes available, including Mbone sessions. The user can then select a channel. When viewing MPEG-1 encoded channels, the IP/TV viewer can give full-screen images with quite good quality and a high frame rate. Slidecast is an additional feature used when transmitting lectures or presentations to show any slides being presented to viewers in a

separate window. Two-way interaction is often useful during a live transmission (for example, a lecture or section briefing), and IP/TV allows the user to type in questions and send them to the speaker. IP/TV also provides a quality-of-service indicator, which shows whether any reception problems are local to the machine itself or are more widespread. This information is gathered from RTCP packets (see panel) multicast by other viewers.

Netshow is the Microsoft solution for distributing video on intranets and on the Internet. It can stream video over a range of bit rates. BT's Intranet-TV service uses NetShow extensively. The servers are based on the Futures Testbed network, which is fully multicast enabled, allowing very efficient distribution of live content to testbed-connected clients using IP multicast. The application can automatically switch to unicast

New Protocols for Real-Time Data

IP/TV and Netshow use real-time protocol (RTP)⁹ to provide information for sequencing and synchronising data streams. The standard RTP header includes a sequence number and a timestamp. The receiver can use this information to reconstruct a real-time stream and play it out. Unlike transmission control protocol (TCP), there is no mechanism for retransmitting packets that are lost, as this is usually pointless for a video or audio application. However, it is important to find out when packets have been lost or delayed, and the application can get this information by looking at the sequence numbers and timestamps.

RFC 1889 also defines the real-time control protocol (RTCP). Clients receiving an RTP stream can use RTCP packets to feed back information on reception quality. For example, a multicast stream might be transmitted successfully to most clients, but a client on a congested network might see only a fraction of the packets. The server can monitor RTCP feedback to identify problem areas of the network. The specification for RTP/RTCP also addresses the problem of scalability. There might be thousands of clients, and if they all sent back lots of RTCP traffic it could lead to network congestion. Clients must therefore scale down their transmission rate in line with the number of participants.

Figure 7—CU-SeeMe application during a multicast multiparty session

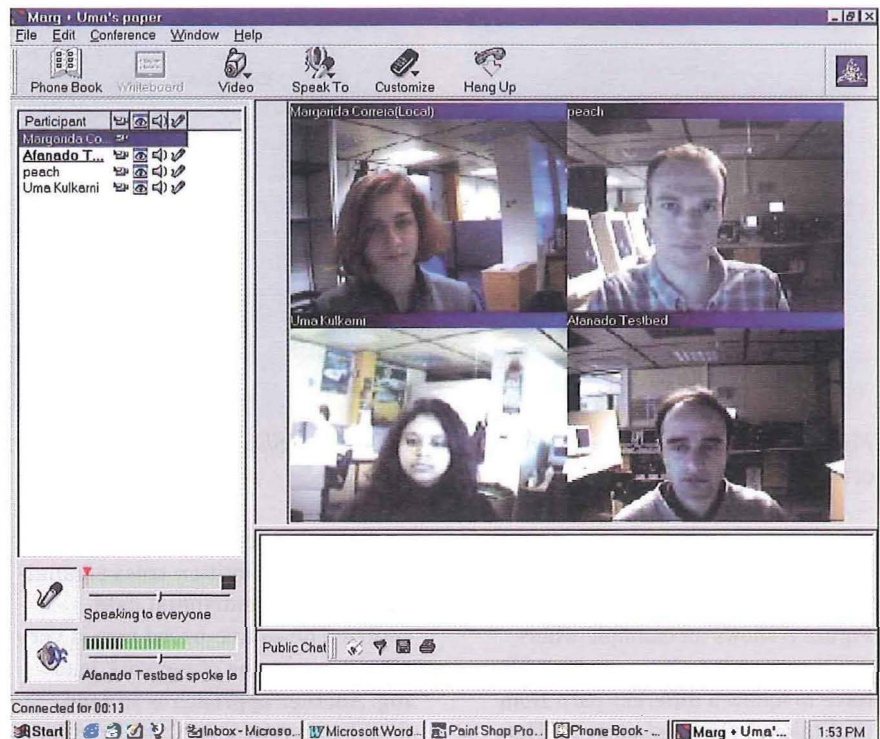
transmission on networks where multicast is not available. However, whereas multicast is highly scalable the maximum number of unicast clients is usually fixed by network constraints. A good description of Intranet-TV is given in a recent article in *British Telecommunications Engineering* journal⁸.

Videoconferencing over IP networks is already available from different software and hardware vendors such as Microsoft, White Pine, Intel, and PictureTel. All these products allow point-to-point video, audio, written chat, and white board. The Microsoft software enables sharing and collaboration of Windows applications and is used on other brand applications for this purpose. Previously a conferencing session with more than two participants required a server to switch the audio and video channels between the users. White Pine's CU-SeeMe application (Figure 7) now uses IP multicast and allows the setting up of video-conference sessions without a server. This application allows any user to create a multicast session that is then advertised to other users⁷. When another user joins the same session it uses the same multicast addresses, one for the video and another for the audio. It is possible to have up to 12 participants in the same conference receiving everyone else's video and audio as well as sharing the whiteboard. CU-SeeMe complies with H.323, the new standard for IP videoconferencing, but also has its own proprietary M-JPEG software codec.

These three applications demonstrate the potential for high-quality video to the desktop over a campus network. They have also given useful experience in running multicast routing and in distributing broadband data across local area networks. The next two sections describe some of the experiences of BT Laboratories in these areas.

Multicast Routing

Many years of work have been dedicated to the processes for routing



unicast IP packets from A to B without problems. For example, one router might send a packet down a particular link that it thinks is the best route to B. However, the router at the other end is configured to route B-destined packets in the opposite direction. The packet bounces back and forth, generating unnecessary congestion. This kind of behaviour is equally possible with multicast, with the added danger of packet replication at every stage (Figure 8). As in the simple point-to-point case, the answer is to use a well-designed routing protocol, with which routers can swap information about routes and so ensure that routing loops are avoided.

Rather than re-inventing the wheel and generating a completely parallel set of multicast routes, it is often best to work with the well-established point-to-point routing information. This is the approach taken by protocol independent multicast (PIM), a popular multicast routing protocol. Figure 8 shows router A receiving packet 'X' and then multicasting it out again, with disastrous consequences. A router running PIM will only accept packets coming in on the 'right' interface, and will discard others. The 'right' interface is the one it would use to send packets to the source, using its own unicast routing table. In Figure 8 the router will decide that packet 'X' has come in on the 'wrong' interface and

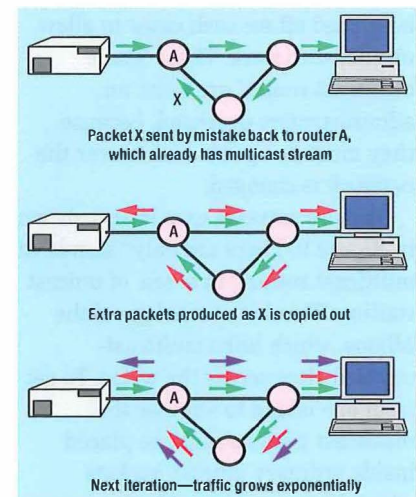


Figure 8—Multicast routing goes wrong

discard it, because the router clearly has a much more direct route back to the source. The packet is said to have failed *reverse path forwarding* (RPF). Incoming packets on the interface connected to the source will pass the RPF check and be forwarded on. Note that the router only has information on the best way from itself to the source, rather than the best route from the source to the router, but normally routing paths are symmetric and this distinction can be ignored. Asymmetric routing can bring in some interesting difficulties, as described later in this article.

Basing multicast routing on the unicast routing protocol works very well unless there are significant




Figure 9—Multicast routing with areas of unicast-only routing

parts of the network where, for technical or political reasons, multicast packets cannot be routed. Figure 9 shows an example where multicast packets between A and B have to follow a different path from that taken by unicast packets. The reverse path forwarding checks must be turned off for such cases to allow multicast to work. These 'static multicast routes' are quite an administrative overhead, because they must be updated whenever the network is changed.

A more extreme case is also shown in Figure 9. There are only 'islands' of multicast routing in a 'sea' of unicast routing. This is the topology of the Mbone, which links multicast-enabled sites across the world. To get from one island to another the multicast packets must be placed inside ordinary unicast packets which have a source address on one 'island' and a destination address on another. Once they reach the destination island the unicast wrapper can be discarded and they can be routed as multicast packets. These point-to-point links between islands are known as *tunnels*. They are very important in setting up a multicast overlay network in a predominantly unicast environment.

Multicast routing protocols that maintain their own routing tables are favoured when multicast and unicast routing differ significantly. The Mbone is an excellent candidate for such treatment and indeed uses distance vector multicast routing protocol (DVMRP), a multicast routing protocol that is independent of unicast routing.

Routing protocols also differ in their strategies for replicating

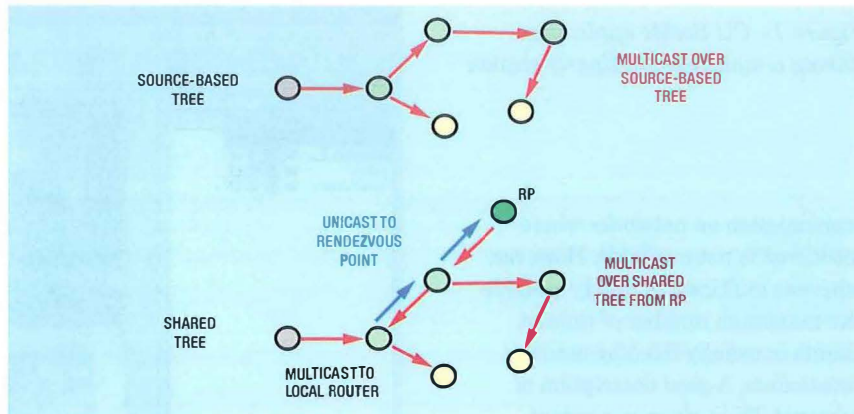
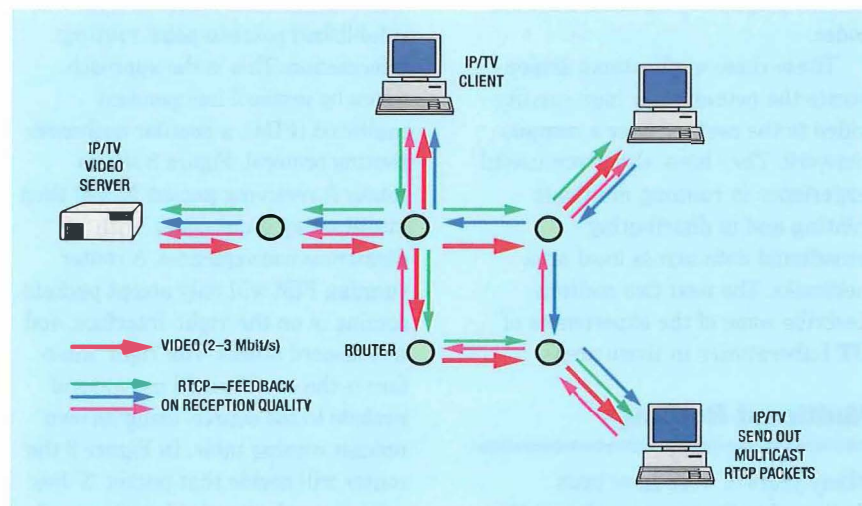


Figure 10—Shared tree and source-based tree

packets. *Source-based trees* are multicast distribution trees emanating from each individual data source. The source-based tree is very efficient in terms of network loading. Another approach is to send data by unicast to a central rendezvous point, where it can then be distributed over a *shared tree*. The shared tree and source-based tree approaches are contrasted in Figure 10.

The network efficiency gains from the source-based tree are most pronounced when there are a few large streams, as in the case of broadband multicast video. Figure 11 shows IP/TV being multicast over source-based trees. Routing becomes quite complicated due to the small flows coming from the IP/TV viewers. These data streams feed back information on quality of service, and are typically less than 1 kbit/s. With a large number of viewers, routers may have to set up very complicated routing tables.

Figure 11—IP/TV over source-based trees



The shared-tree approach trades off efficiency against complexity, and is especially useful with large numbers of small streams. A simple multicast distribution tree can be set up. Sources send data to their local router as normal, but the router then forwards the traffic in unicast packets to the rendezvous point (RP) for distribution. The *core-based tree* (CBT) routing protocol uses source-based trees. The *sparse-mode PIM* routing protocol is slightly more complicated. It can switch from shared to source-based tree once a configurable traffic threshold is exceeded, and so in theory should get the best of both worlds. Figure 12 shows that the shared-tree approach can be quite efficient even for broadband sources if the RP is chosen to be near to the source.

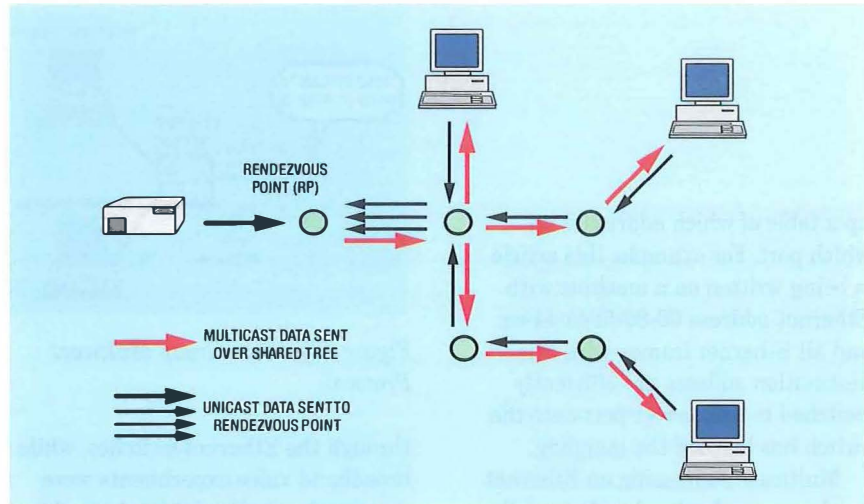
Another important mechanism in IP multicast is 'flood and prune', which is used by both DVMRP and *dense-mode PIM*. A new multicast

Figure 12—IP/TV over a shared tree

group is copied to all parts of the network. Routers that do not need that multicast group send a prune message to their upstream neighbour, which then stops copying it to them. However, when a subsequent request comes along, the router sends a join message upstream to restart the flow. In a protocol such as sparse-mode PIM an RP has to be configured into all of the routers so that they know where to go for information about particular groups. There is more flexibility when using *auto-RP*, which distributes RP information for different group ranges over a well-known multicast address, but routers still need to be told where to obtain the auto-RP multicast group itself! 'Flood and prune' avoids this bootstrapping problem completely.

The routing protocols described so far—DVMRP, dense-mode PIM, sparse-mode PIM and CBT—are equivalent to the unicast routing protocols run within a company. In the unicast case another protocol—*border gateway protocol* (BGP)—is used when exchanging data between different Internet core networks. BGP allows the high level of control and policy setting needed in a commercial environment. An extension to this protocol—*multicast BGP* (MBGP)—is now coming into service as the first steps in the commercial exchange of IP multicast are taken. The Futures Testbed is currently involved in trials of MBGP with major network operators and vendors.

In summary, the various IP multicast routing protocols are distinguished by the use they make of unicast routing information, by the way they set up trees for multicast distribution and by the method by which routers join a particular group. The various choices make them more or less suitable for wide area or local area networks, for large numbers of groups or for setting policies. Equipment vendors also tend to support some protocols more than others. It is likely that several protocols will co-



exist in future, just as in the unicast area. For example, the Futures Testbed uses sparse-mode PIM for IP multicast across the high-capacity backbone, but interworks with DVMRP for Mbone data and with MBGP and dense-mode PIM for experimental work on commercial interconnection.

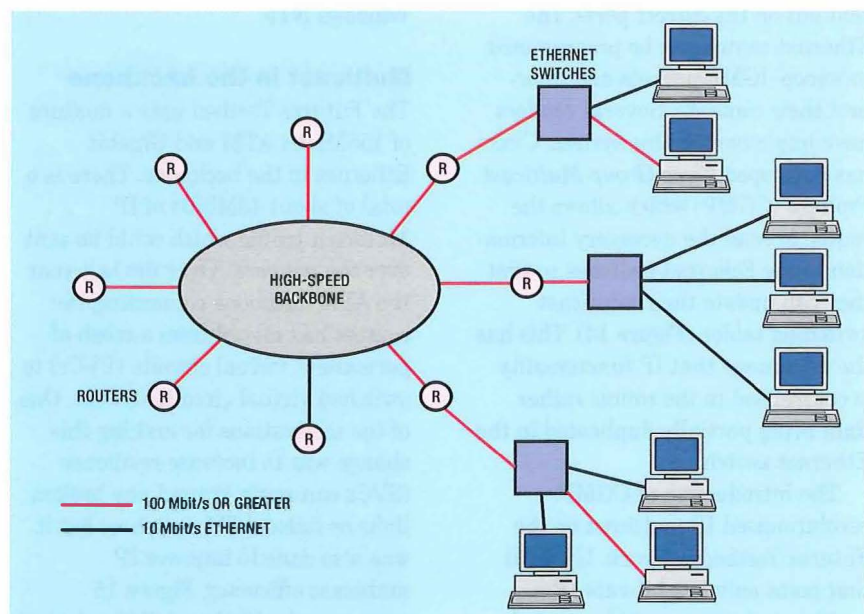
Multicast Replication

This section considers the use of IP multicast in an advanced campus network. As mentioned above, the Futures Testbed is a good example of a broadband IP network. Figure 13 shows a schematic of the network at BT Laboratories. Users have a 10 Mbit/s or 100 Mbit/s switched Ethernet connection to a router running IP and Appletalk. There are eight main routers interconnected by a high-speed backbone.

Multicast to the desktop

Ethernet has been receiving a lot of press attention recently, as it has a large installed base of machines (and perhaps more importantly of trained people), but also now provides an upgrade path to 1 Gbit/s or higher. Ethernet switching technology, based on fast custom-designed silicon integrated circuits, can cope with very high packet throughputs. Ethernet switching is an enormous improvement over the original concept of a shared network. Frames are only sent out on the port to which the destination machine is attached, so that users do not have to compete with a lot of other people's traffic—they effectively have their own Ethernet segment. The means by which this is achieved is quite simple. The Ethernet switch listens in to the 48-bit Ethernet source addresses on frames passing through it and builds

Figure 13—The futures testbed network at BT Laboratories



up a table of which address is on which port. For example, this article is being written on a machine with Ethernet address 00-80-5f-ca-44-ee, and all Ethernet frames with this destination address are efficiently switched to the correct port once the switch has learned the mapping.

Multicast addressing on Ethernet has been standardised to fit in well with the IP multicast addressing scheme. The last 23 bits of the IP multicast address are appended to a standard prefix, so that 224.0.0.2 maps to 01-00-5e-00-00-02. In theory the Ethernet switch can set up a table of ports for multicast in an analogous way to its unicast mapping table. For example, packets with destination address 01-00-5e-7f-20-06 might be sent to ports 4 and 9, where users have requested that stream, while 01-00-5e-02-df-c1 might be sent to ports 1, 3 and 7. There is a serious difficulty here, which is that all the negotiations to get a multicast group are carried out in IGMP packets. As the switch does not know which destinations want which group, it defaults to sending out multicast groups on all ports. This means that broadband (2–3 Mbit/s) video streams are not a good idea in such Ethernet switches!

There are two possible solutions to get the multicast Ethernet frames sent out on the correct ports. The Ethernet switch can be programmed to 'snoop' IGMP packets and interpret their contents. Several vendors have implemented this system. Cisco has developed *Cisco Group Multicast Protocol* (CGMP) which allows the router to send the necessary information to the Ethernet switches so that they can update their multicast switching tables (Figure 14). This has the advantage that IP functionality is centralised in the router rather than being partially duplicated in the Ethernet switch.

The introduction of CGMP revolutionised IP multicast on the Futures Testbed network. Up until that point only low-bit-rate IP multicast streams could be passed

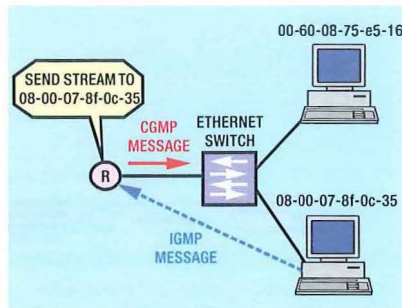


Figure 14—Cisco Group Multicast Protocol

through the Ethernet switches, while broadband video experiments were restricted to isolated 'islands' in the network. Once Ethernet switches were able to deal with broadband streams, applications such as Cisco IP/TV could potentially be deployed on any of the Ethernet connections to the desktop.

The next improvement in IP multicast performance in Ethernet switches will be the widespread adoption of IGMP version 2. The introduction of an explicit IGMP 'leave' message greatly improves IP multicast for 'channel-hoppers' who move from group to group. With IGMPv1 the router continues to send the group until three periodic updates are missed. Hence the minimum time for a stream to disappear is about 30 seconds or so. However, if someone else is watching the same channel his/her IGMP responses will keep the stream up and the unfortunate channel-hopper will continue to receive the stream even though they do not want it. IGMPv2 is part of the network software for Windows 98 and Windows NT5.

Multicast in the backbone

The Futures Testbed uses a mixture of 155Mbit/s ATM and Gigabit Ethernet in the backbone. There is a total of about 40Mbit/s of IP Multicast traffic which could be sent over the network. Over the last year the ATM backbone connecting the routers has moved from a mesh of permanent virtual circuits (PVCs) to switched virtual circuits (SVCs). One of the motivations for making this change was to increase resilience (SVCs can route around any broken links or failed ATM switches) but it was also done to improve IP multicast efficiency. Figure 15 compares the PVC and SVC solutions

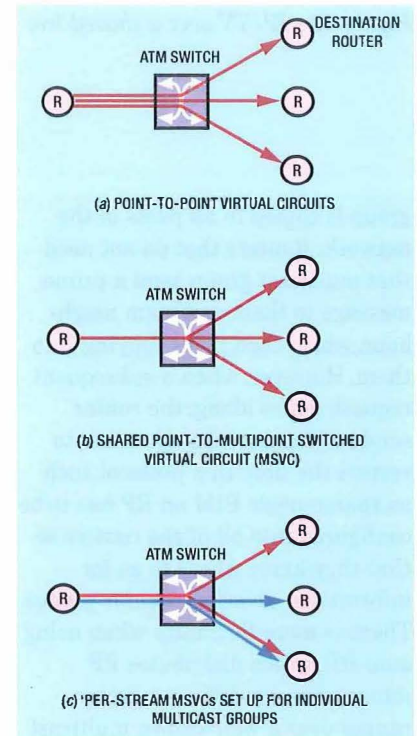


Figure 15—IP multicast over a PVC mesh, a point-to-multipoint SVC mesh, and a dynamic group-specific SVC

from the point of view of one of the eight core routers on the network.

In principle, sending out multiple copies of the same data over the same network link is not desirable, and this counts against the PVC solution. A point-to-multipoint SVC (MSVC) carrying all multicast traffic avoids this problem, as replication is carried out on the ATM switches, ensuring that only one copy is sent over any particular physical link. The disadvantage of using a common MSVC for all traffic is that not all routers want the same groups, so in many cases the data is being copied to destinations where it is not needed. However, it is wrong to focus on network capacity (which, after all, is relatively cheap to upgrade in a campus network) to the exclusion of other factors. An important factor here is the relative efficiency of various network components in replicating IP multicast streams.

Figure 16 shows the processor load on a mid-range router and a high-end router when replicating 1.7 Mbit/s IP multicast streams, made up of MPEG-1 compressed video data. For, say, four video servers and five destinations the mid-range router is beginning to look

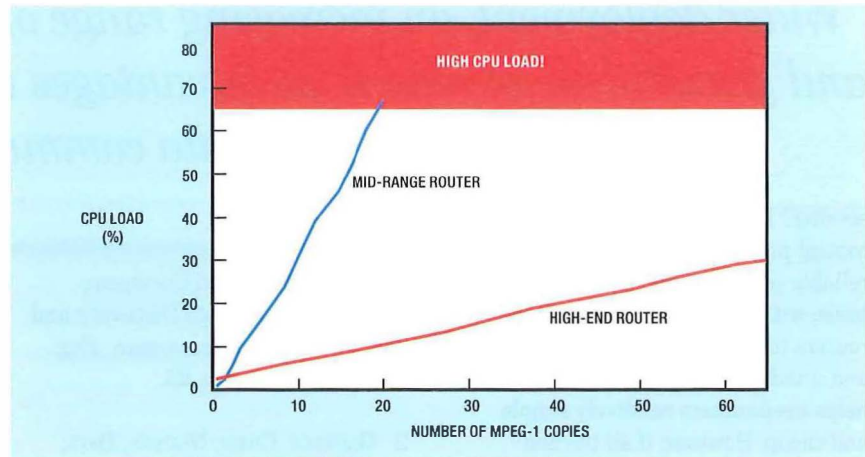
Figure 16—Router processor load during replication

an uncomfortable choice, and an upgrade to a much more expensive router is needed. However, using the Layer 2 infrastructure between the routers to do the copying instead means that the router only ever needs to send one stream. The only question then is the ability of the backbone to copy the data to all the other routers. Ethernet switches copy multicast data very efficiently, as can be seen from their performance at the edge of the network. Replication in ATM switches using MSVCs is also very efficient. In order to test this, a 120 Mbit/s stream has been copied to seven ports on an ATM switch without problems, which shows the impressive capabilities available at Layer 2.

The Futures Testbed ATM backbone is currently configured for each router to send IP multicast over a single MSVC, ensuring that the processor loads on the current mid-range routers are acceptable. A more targeted approach using separate MSVCs for specific groups has been successfully tested, and is likely to be installed in the future. Multicast over Gigabit Ethernet is at a relatively early stage of development, but, as in the ATM case, network capacities are such that copying all the streams to all the routers on a common IP subnet is likely to be an acceptable solution.

The Multicast Future

There are many reasons why IP multicast will flourish in the next few years. Multimedia applications such as video to the desktop and videoconferencing will be developed further. Multicast-compatible Ethernet will allow broadband multicast streams on the local area network. Cost-effective IP multicast video will complement the more expensive video-on-demand (VoD) servers. More research is needed in several areas. Firewalls and network address translation (NAT) are used by many companies to get relative isolation from the Internet, keeping



hackers out and private address ranges in. At the moment firewalls do not support multicast, and setting up a tunnel to allow multicast through the firewall can introduce a security hole. Developments in firewall technology should allow IP multicast to pass natively and securely.

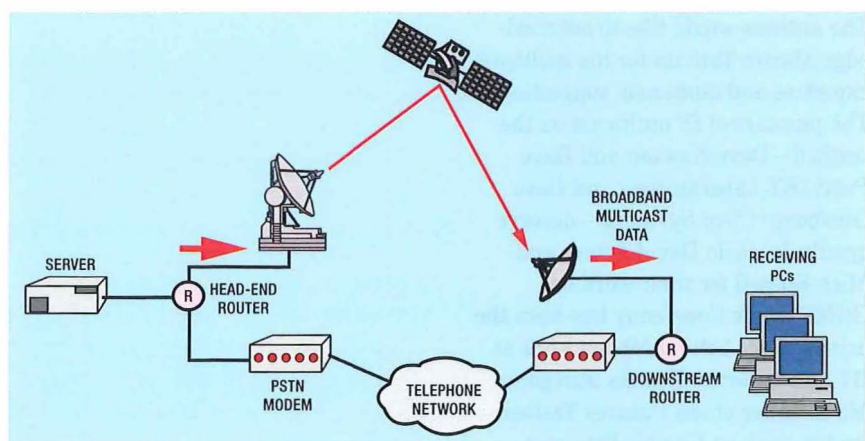
Quality of service is another area in which progress is needed. RTCP allows feedback on packet loss and jitter as seen by multicast recipients. However, it is not easy to ensure that time-sensitive streams take precedence over other traffic. The resource reservation protocol (RSVP) was developed to address this problem, but may not be scalable for large numbers of multimedia streams. The 'class of service' approach which sorts IP packets into various classes (high, medium, and low priority) is likely to be easier to implement in a large network.

Another interesting area of multicast development is that of asymmetric networks. For example, data may be broadcast by satellite, with a PSTN modem connection used as a return path from the remote

user to the central site (Figure 17). Satellite transmission is another good way of getting the same data to many people, of course. When IP multicast data is sent by satellite (perhaps for redistribution to many users at the remote site) there is a problem that RPF (reverse path forwarding) fails. Data comes in over the satellite connection even though the receiver routes traffic back to the source via the modem. Multicast tunnels are one way of solving this problem, but it would be desirable to broadcast multicast packets natively.

None of the multicast protocols described here have any means of asking for retransmission of packets when they are lost. This reflects the nature of real-time data, where there is no point in receiving part of a video frame well after the rest of the frame has been displayed to the user. *Reliable multicast* is of great interest, and deserves a separate article. Designers of such a protocol have to make several interesting design choices. For example, to what extent should devices in the multicast tree keep track of the end systems they are

Figure 17—Asymmetric data transmission over satellite



Wider deployment, an increasing range of application software and growing awareness of its advantages will make IP multicast an important part of the data communications world.

...serving? The unicast transmission control protocol (TCP) achieves reliable transmission on an end-to-end basis, without asking intermediate routers to keep track of acknowledged and unacknowledged packets. This helps keep routers relatively simple and cheap. However, if all the end systems were to communicate directly with a multicast source it might be deluged in negative acknowledgement packets. A lot of effort is being put into finding a good solution to this and other problems in reliable multicast.

In this article we have described some of the uses to which IP multicast is already being put. Wider deployment, an increasing range of application software and growing awareness of its advantages will make IP multicast an important part of the data communications world. We have shown that, once Ethernet switches are made multicast-aware, campus networks are very well suited to IP multicast of several 2–3 Mbit/s video streams. The backbone infrastructure (Ethernet or ATM) can replicate broadband multicast data streams very efficiently, and can remove the need to upgrade routers. As IP multicast becomes more widespread—as broadband multicast on campus networks and narrowband multicast over wide areas—most people will soon be part of a multicast tree, whether at home, on the road or in the office.

Acknowledgements

The authors would like to acknowledge Martin Tatham for his multicast expertise and numerous suggestions. The pioneers of IP multicast on the testbed—Dave Newson and Dave Pratt (BT Laboratories) and Dave Ginsburg (Cisco Systems)—deserve gratitude, as do David Armes and Mick Russell for their work on CGMP. Mark Courtenay has been the driving force behind Mbone work at BT Laboratories. Thanks also go to Mark Salter of the Futures Testbed for his work on Gigabit Ethernet.

References

- 1 DAVIES, JOHN; and COCHRANE, RALPH. Knowledge Discovery and Delivery. *Br. Telecommun. Eng.*, April 1998, **17**, p. 25.
- 2 GIBBINGS, CHRIS; NEWSON, DAVE; HENNING, IAN; BARNES, JEREMY; PRATT, DAVE; and COCHRANE, PETER. Superhighway... Supergridlock? *Br. Telecommun. Eng.*, Jan. 1997, **15**, p. 315.
- 3 HENNING, I.; SIM, S.; GIBBINGS, C.; RUSSELL, M.; and COCHRANE, P. A Testbed for the Twenty-first Century *Proceedings of the IEEE*, Oct. 1997, **85**(10), p. 1572.
- 4 SALTER, MARK. Quote in *Network Week*, **3**, no. 43, 5 Aug. 1998, p. 7.
- 5 <http://www.cisco.com>
- 6 <http://www.microsoft.com>
- 7 <http://www.wpine.com>
- 8 GRACE, ANDREW; JACOBS, RICHARD; COX, JON; and BARROW, ANGELA. Intranet-TV—Video Streaming for the World Wide Web, *Br. Telecommun. Eng.*, April 1998, **17**, p. 49.
- 9 SCHULZRINNE, H.; CASNER, S.; FREDERICK, R.; and JACOBSON, V. RTP: A Transport Protocol for Real-Time Applications. *IETF RFC* 1889, Jan 1996.

Biographies



Chris Gibbings
Networks and
Information Services,
BT UK

Chris Gibbings gained a Ph.D in Semiconductor Physics from the University of Cambridge in 1986. He joined BT in 1985 to study the growth of silicon-germanium strained layers using molecular beam epitaxy. In 1992, he moved to the security research area and in 1996 joined the Futures Testbed team. He has worked on IP multicast over ATM in the City Media trial run by BT Broadcast Services, and developed the IP/ATM configuration of their Medianet8 satellite data service. Recently he has been working on network design for Vio, the BT/Scitex joint venture for the printing industry.



Mark Barrett
Networks and
Information Services,
BT UK

Mark Barrett joined BT in 1986 and gained a HNC in Electronic Engineering in 1992. He has worked within the LAN/WAN data networking field since 1989. In April 1997 he moved from a network support role to the Futures Testbed and has specialised in IP multicast routing, protocols and management. He has provided consultancy to various projects on IP multicast and is currently working on inter-domain issues.



Margarida Correia
Networks and
Information Services
BT UK

Margarida Correia graduated in Electronics and Computer Engineering from Instituto Superior Tecnico, Lisbon, Portugal in 1995. In 1996 she gained an M.Sc. in Telecommunications and Information Systems from the University of Essex. In the same year she started an industrial placement in BT Laboratories with the Futures Testbed team, subsequently joining BT on a permanent basis in April 1997. She has been working on multimedia distribution over IP and ATM networks. She is part of the team responsible for multicast delivery over the Futures Testbed network, and she has had a major role in the IP/TV trial. As well as regular presentations to BT customers and BT senior management she has presented a paper on IP Multicast at INET'98 in Geneva.



Uma Kulkarni
Networks and
Information Services,
BT UK

Uma Kulkarni graduated in Aeronautical Engineering from Imperial College, London in 1996 and, in 1997, gained an M.Sc. in Spacecraft Technology and Satellite Communications from University College London. She joined BT Laboratories in 1997 and started work with the Futures Testbed team where she has been working on multimedia distribution over IP and ATM networks. Since joining BT she has given regular presentations to BT customers and participated in several high-profile demonstrations to senior managers.



Kevin Smith
Networks and
Information Services,
BT UK

Kevin Smith received his Ph.D. from Imperial College in 1984. After postdoctoral fellowship and assistant lectureship positions at St. Andrews University, he moved to AT&T Bell Labs., Holmdel, NJ, USA. There he designed and built the first experiments demonstrating the transmission of ultrafast (Tbit/s) optical pulses over transcontinental distances using optical amplifiers. In 1989, Kevin joined BT's research laboratories, where he continued researching the potential of ultrafast optical systems including laser sources, transmission, optical switching and processing elements. Together with Julian Lucek new techniques for all-optical signal processing were developed, including the world's first all-optical signal regenerator. Kevin has published more than 100 papers in international journals and conferences, and holds 10 patents. Since December 1996, Kevin has been the Group Leader of the Futures Testbed Group within the Advanced Research and Technology Department.

Engendering Trust in Electronic Commerce

The Internet is widely predicted to revolutionise commerce over the next few years. However, the full potential of electronic commerce will only be realised if buyers and sellers have the confidence to trade in electronic environments. This article outlines some of the major trust-related issues associated with electronic commerce and identifies key roles for trusted third parties. BT's new Trustwise service is presented, together with a discussion of possible future directions for trust services.

Introduction

Within the last two years, electronic commerce has emerged as one of the most talked-about phenomena since the Internet itself captured popular attention. While not an entirely new concept—large businesses have been carrying out electronic data interchange (EDI) and other business-to-business transactions over proprietary networks for several years¹—the Internet has brought electronic trading within reach of smaller businesses and individual consumers for the first time and looks set to revolutionise all aspects of commerce over the next few years.

To facilitate this revolution, however, a number of fundamental barriers must be addressed. These barriers impede commercial cooperation, and within the electronic domain there are no established mechanisms for overcoming this impedance. One of the key mechanisms used in the physical domain is trust, but trust is a complex phenomenon which is based upon a multitude of factors and does not translate directly to the electronic domain. This article outlines some of the major trust-related issues associated with electronic commerce, overviews work currently being carried out in BT to address these issues, and identifies a number of potential new roles for BT as a trusted third party (TTP) in acting to create the confidence required for buyers and sellers to trade in electronic marketplaces.

Throughout the article, the term *trustor* is used to refer to a party that needs to exercise trust, while the term *trustee* refers to a party that needs to be trusted.

Analysis of Cooperation

The physical world

The act of physically entering a store and purchasing a product is so commonplace that it is rare for most of us (except marketing experts!) to give conscious consideration to the process involved. However, whether consciously or subconsciously, the decision to make a purchase is influenced by a wide range of factors including:

- the brand strength and reputation of the product and/or merchant, as determined by advertising, previous experience, personal recommendations, 'expert' opinion, etc;
- the customer's experience while in the store, including ease of finding required products, nature of interactions with sales staff, general impression of the store, etc; and
- awareness (albeit often vague) that redress is possible if the product turns out to be faulty or otherwise fails to meet expectations, including statutory consumer rights, warranties, perceived permanence of the store, etc.

These factors increase in importance when considering high-value or high-risk purchases. When purchasing a used car, for example, a buyer is likely to make a careful, conscious evaluation of the risks involved and the trustworthiness of the seller before entering into the transaction.

Purchasing from an established dealer presents a relatively low-risk option to the buyer: knowledge that the

Trust is an essential part of everyday life and plays a role, whether conscious or sub-conscious, in almost every transaction.

dealer has a brand and reputation to protect provides the buyer (rightly or wrongly) with a degree of assurance in the quality of the used car. This assurance may be further backed by a warranty and/or enhanced by third party accreditation of the dealership. However, the added assurance comes at the expense of a higher price tag.

Alternatively, the buyer may choose to purchase from an individual seller or back-street dealer. This is a higher-risk strategy, since the seller has no known brand or reputation to maintain and there is little likelihood of redress if things go wrong—the back-street dealer may go bust by tomorrow! The buyer must therefore work harder at ascertaining the trustworthiness of the seller as part of the dialogue accompanying the buying process. Among the things considered by the potential buyer are:

- seller's motivation (Why does the seller want to do this deal? Is there anything suspect?),
- strength of commitment (Is the seller committed to completing the deal?),
- openness (Is the seller being open and honest? Or is he hiding something?), and
- integrity (Is the seller perceived to be of good character?)

In addition to relying on his/her ability to assess the trustworthiness of the seller, the buyer may choose to improve his/her level of confidence by getting a knowledgeable colleague or a professional (for example, from the AA) to check the car over before completing the deal. However, again the added assurance comes at a price.

In each case the buyer will only demonstrate cooperative behaviour (for example, enter into the deal) when he/she is satisfied that the risk is acceptable, for the specified price, and has sufficient confidence (for example, trust) that the deal will be successful. There are barriers to cooperation in the

physical world, but we have a number of well-established mechanisms which allow us to get over them.

The electronic world

As illustrated in the scenario above, there are many complex considerations that are taken into account when participating in commercial activity in the physical world, but these complexities are amplified when commerce moves to the Internet². The geographical separation of buyers and sellers, often coupled with a lack of real-time visual or oral interaction, radically changes the customer's buying experience and creates new challenges for communicating trustworthiness.

To create an environment where users feel confident to trade, it is necessary to provide facilities that:

- allow the creation, maintenance and evolution of persistent identities (for example, digital identities (IDs));
- allow users to communicate their trustworthiness and to evaluate the trustworthiness of others against personal requirements (for example, trust profiles);
- alleviate privacy concerns—users should be assured that personal data collected as part of a commercial transaction will not, or cannot, be used for inappropriate purposes (for example, data protection policies, anonymity services);
- alleviate security concerns—mechanisms should be provided to guarantee the integrity and confidentiality of data passing between users (for example, digital signatures, encryption);
- support non-repudiation—evidence should be generated during transactions that can be used in the event of a dispute (for example, audit trails, digital receipts);

- offer redress—users should have rights to redress if purchased products turn out to be faulty or otherwise fail to meet expectations (for example, international warranties); and
- create a good customer experience—user interfaces should be easy to use, with ready access to all the functionality and information required to support the purchasing process.

The barriers to cooperation in the electronic world are effectively higher than in the physical world. If electronic commerce is to be successful, it will be necessary for new intermediaries (TTPs) to develop and offer mechanisms which address specific barriers and effectively engender trust in electronic commerce. The remainder of this article presents a model of trust developed within BT's Electronic Commerce Research team and looks at roles for TTPs.

A Framework for Trust

Trust is an essential part of everyday life and plays a role, whether conscious or sub-conscious, in almost every transaction. Despite its ubiquity, however, there are no generally accepted definitions of what 'trust' actually is or how it works. A simple dictionary definition³ describes trust as:

'...complete assurance and certitude regarding the character, ability, strength, or truth of someone or something.'

This is a very broad definition. The real meaning and importance of trust is specific to particular domains and transactions.

In the context of a society, trust can be viewed as a mechanism for reducing complexity and a means of coping with the freedom of others: it is an aspect of all social relationships and implies some form of expectation

it is perhaps not surprising that most business-to-business commerce is currently conducted within closed trading communities or extranets, rather than over the open Internet

about the future. Specifically, when applied to electronic commerce, trust is an essential ingredient in generating the cooperation required for transactions to be conducted between buyers and sellers.

In order to gain a deeper understanding of how trust can be engendered in electronic commerce, we will consider two distinct aspects of trust:

- *environmental trust*—general trust that applies to all situations; and
- *transactional trust*—trust that is specific to particular situations.

Environmental trust

Environmental trust refers to the trust that exists in a particular domain made up of participants with some sense of commonality or shared values; it provides a foundation for interaction between the citizens of a domain. This is a basic, general component of trust that forms the basis of everything we do⁴. It is essentially trust based on induction: because we have observed the world around us in action, we have come to expect (that is, trust) that certain laws will be followed. This aspect of trust can be viewed from the perspective of society in general, or, as will be shown later, in the context of a well-defined environment.

To illustrate the importance of environmental trust, it is instructive to consider some of the issues currently limiting the widespread adoption of international electronic commerce. The issues are not primarily technical—the Internet now offers the potential for buyers to purchase goods from merchants located anywhere in the world. Rather, the issues are associated with the differences in culture and legal systems that exist around the globe. A typical customer contemplating an international transaction on the Internet today is likely to have little confidence in his position with regard to key issues such as privacy, security, non-repudiation and redress. International commerce bodies, such as the ICC and

Organisation for Economic Co-operation and Development (OECD), have important roles to play in creating frameworks that will increase the sense of commonality and shared values around the world, and thus enhance environmental trust at an international level.

The importance of environmental trust is also apparent in more restricted domains, such as closed trading communities or extranets. In these environments, a trustee may be motivated to meet the expectations of trustors through the existence of contractual agreements or informal community standards, typically backed up by the threat of penalties or ostracism from the community. Another important factor within restricted communities is the potential for reciprocation: studies have shown that mutually beneficial behaviour is more likely to occur where there is a high probability of further interactions between the trustor and trustee in the future⁵. With these considerations in mind, it is perhaps not surprising that most business-to-business commerce is currently conducted within closed trading communities or extranets, rather than over the open Internet.

Transactional trust

Transactional trust is situational and generally involves action and behaviour; it is specific to the participants in a particular transaction. Two key components of transactional trust have been identified⁴ as:

- expectation of competent role performance on the part of the trustee, and
- expectation that fiduciary obligations will be fulfilled.

The expectation of competent role performance is a specialised aspect of trust based on the assumption that a trustee possesses the required expert knowledge or technical skills to perform a particular function satisfactorily. In this sense, the trustee's

trustworthiness is largely based upon his/her ability or competence. Trustees often seek to enhance this aspect of trust by presenting evidence of accreditation by an appropriate professional or industry body (for example, a travel agent may demonstrate membership of the Association of British Travel Agents (ABTA); an independent financial advisor may demonstrate that he/she is regulated by the Personal Investment Authority (PIA)). It is perhaps worth noting that this specialised aspect of trust cannot be generalised: the trust placed in a trustee is based upon perceived competence in a particular situation (among other things) and that competence may not be applicable to any other situation. Thus, a taxi driver may be trusted to drive to the airport, but not to fly the plane!

The expectation that fiduciary obligations will be fulfilled implies that the trustee has some moral obligation to place the trustor's interests above his/her own. This aspect of trust goes beyond competence and ability to consider morality and ethics. Politicians, for example, have certain fiduciary obligations arising from the trust placed in them by the electorate. Perhaps a more credible example is provided by a Web-based financial information service that prohibits the holding of shares by any of its employees⁶. The company thus assures customers that its fiduciary obligations are met by removing the threat of employees obtaining personal financial benefit from the advice they provide.

When considering the dynamics of transactional trust, the level of trust required for a transaction to take place will be determined by certain characteristics of the trustor and certain characteristics of the trustee. In other words, the level of trust for a particular type of transaction will be specific to a particular combination of trustor/trustee. Perhaps the most important class of considerations is associated with the trustor and his/her propensity to trust: to what degree is a particular individual

inclined towards displaying trusting behaviour? For example, in the used car purchasing scenario discussed earlier, the likelihood of the sale being completed is likely to be more dependent on the gullibility of the buyer than on the trustworthiness of the seller!

The second class of considerations is associated with the trustor's perceptions of the trustee. These are the factors that are examined in the process of assessing the trustworthiness of a potential trustee. In the used car example these included: reputation, motivation, strength of commitment, openness and integrity. These perceptions are highly subjective and difficult to specify. In the electronic world, the deciding factors are likely to be more specific and formal.

Relationships between environmental and transactional trust

Environmental trust and transactional trust are closely related: the degree of transactional trust required in a particular exchange will depend upon the amount of existing environmental trust. In some cases, the environment may be well defined and narrow in scope so that the environmental trust is sufficient to allow exchanges to occur without careful analysis of each transaction, while in other cases the environment may be broad and inclusive so that the environmental trust provides only a foundation. In the case of electronic commerce, an EDI system is an example of a well-defined and narrow environment: the participants are well known and transactions are underpinned by bilateral agreements. In an open trading environment, on the other hand, there may be a set of trading standards that create a basic level of environmental trust, but the participants must evaluate the trustworthiness of potential trading partners on a per-transaction basis.

Environmental trust will be reinforced by successful transactions based upon transactional trust, and will be undermined by unsuccessful

transactions. As users enter into transactions and find that their expectations are met, their trust in the system in general will tend to increase. When expectations based on a particular transaction are not met, however, environmental trust may be undermined to some extent. Unfortunately it is the case that trust is very easily undermined and very difficult to rebuild, so it is particularly important to understand how to prevent malfeasance.

Alternatives and complements to trust

Trust is an instrument that has both functional complements and functional alternatives⁴. That is to say, in a particular situation there may be other functions that can aid trust (that is, functional complements) or replace trust (that is, functional alternatives). For example, in our used car purchasing scenario the issuance of a warranty acts as both a functional complement and a functional alternative: it aids trust in the dealer; since the dealer is unlikely to offer a warranty if the car is suspect, but also reduces the degree of trust required—the car will be fixed for free anyway if there is something wrong with it. Beyond understanding what trust is, it is important to understand what the functional complements and alternatives to trust are: these are the tools of the TTP that will be examined in more detail below.

Roles for a Trusted Third Party

It is clear that the issue of trust is a difficult one to address through simply creating technical infrastructure. In the electronic world, TTPs will have vital roles to play in facilitating the trust required for commercial transactions⁷. TTPs provide a mechanism whereby users can transitively trust one another: if a buyer trusts a TTP who, in turn, trusts a merchant, the buyer can have a degree of trust in the merchant. However, it should be noted that trust is not a purely transitive

instrument: the degree of trust between end-users will become progressively diluted as the length of the trust chain increases.

In general, TTPs provide functional complements to trust: they provide mechanisms that help to enhance trust and thus increase the likelihood of cooperative behaviour between buyers and sellers. However, in some cases, particularly where the TTP is prepared to accept a degree of liability for the transaction, TTPs may also act to reduce the risk of entering into a transaction, thus providing a functional alternative to trust. Potential roles for TTPs, both in engendering trust and in reducing risk, are discussed further below.

Engendering environmental trust

The initial step in engendering environmental trust is in establishing and defining the environment and the requirements for community membership. A carefully defined and controlled environment with a controlled user community will have a greater sense of shared values and hence a higher level of environmental trust than an open community with more flexible membership. The TTP can vet users prior to membership to ensure that they meet certain minimum criteria. This will allow users of the system to make basic valid assumptions about the trustworthiness of all other users.

In order to maintain the level of environmental trust created by defining the environment and vetting the members, it is necessary for the TTP to exercise some form of governance to enforce trading rules and community standards. Establishing community standards and then policing them will also create an incentive for cooperation and will act as a clear perceived motivation for trustees to fulfil expectations.

Engendering transactional trust

The TTP's role in engendering transactional trust will be based

upon providing mechanisms to enable trustors to specify their trust requirements and to enable trustees to communicate their trustworthiness. This will require a trust profile to be established for each user.

A user's trust profile will be made up of a number of components. On entry into the system, and as part of the vetting procedure, the TTP can create a basic trust profile based upon information from third-party databases and from published accounts. Credit ratings of individuals (or of company directors) can be accessed and used to gain some idea of the credit worthiness of the user, and information sources such as Dun & Bradstreet⁸ can be consulted to obtain a supplier evaluation report that contains information about the risk associated with a particular supplier, as well as public filings and judgements. The information from third-party databases can be used to create a trust profile that can be augmented and amended according to the success or failure of transactions conducted within the trading environment.

Reducing risk

It is clear that the degree of trust required in any given transaction is dependent on the risk involved. In some cases, the risk associated with a transaction may be so high (or the trustworthiness of the other party so low!) that mechanisms available for engendering trust will fail to produce sufficient confidence for the buyer or seller to enter into the transaction. In these cases, the TTP may have a role to play in reducing risk (providing a functional alternative to trust) rather than somehow trying to create more trust. Returning to our example of the used car purchase, the use of a TTP (for example, the AA) to provide independent assessment of a car reduces the risk for the buyer and, consequently, means that the buyer has less need to trust the back-street dealer.

If a TTP is prepared to underwrite a transaction, the need for trust

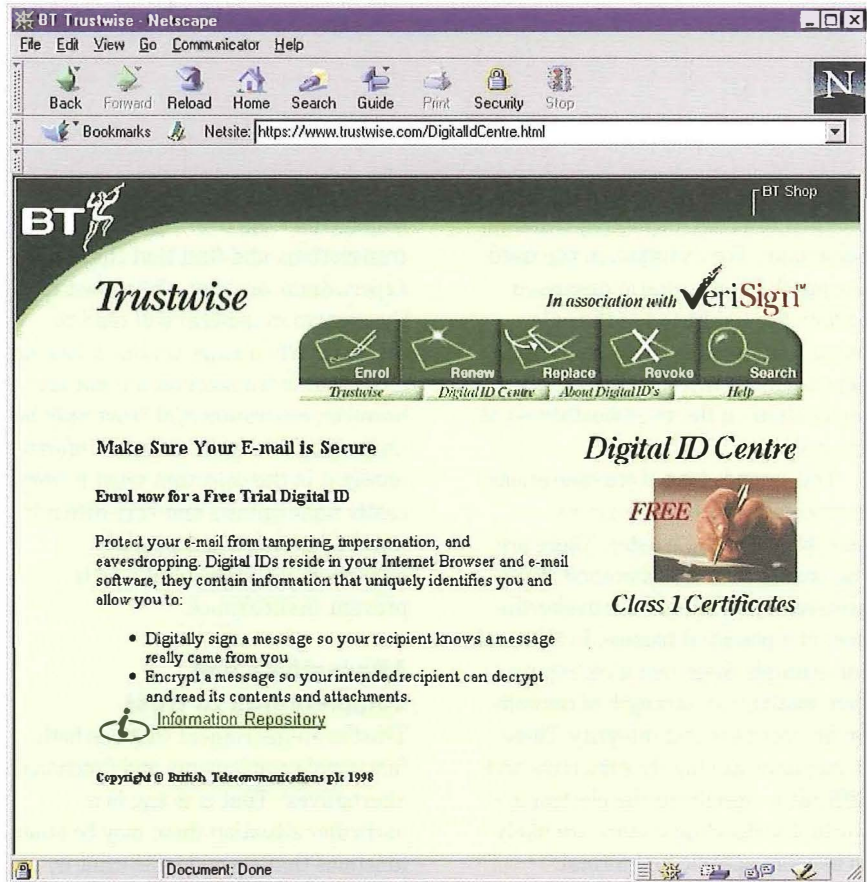
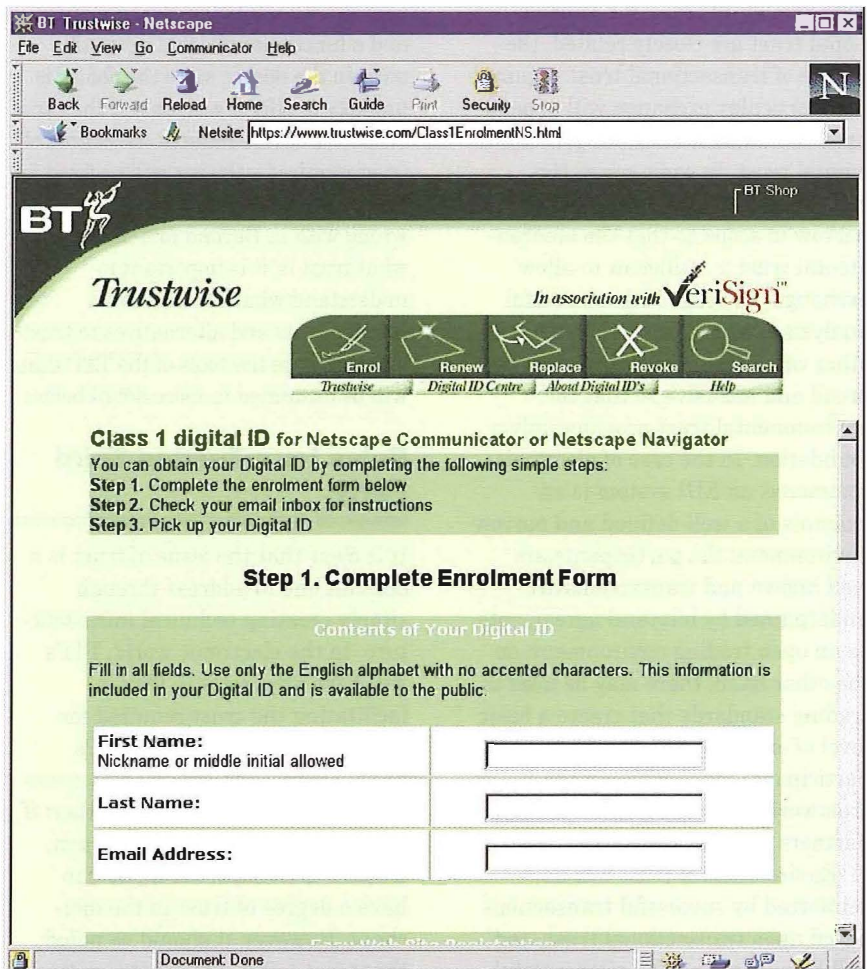


Figure 1 – BT Trustwise screenshot

Figure 2 – BT Trustwise screenshot



Proof of identity is an important issue for electronic commerce, since faking an identity is much more easily achieved online than in the physical world

between the transacting parties is effectively eliminated—trust placed in the TTP replaces the need for trust between the transacting parties. Bank payment systems provide a good example of this: a merchant can readily accept a cheque from an unknown (untrusted) purchaser, safe in the knowledge that the payment will be honoured up to the value specified on the purchaser's cheque guarantee card. Credit card systems also guarantee payments to the merchant (upon presentation of a chit signed by the cardholder) and may, in addition, provide guarantees to the purchaser; for example, some Platinum VISA cards provide refunds on faulty purchases and even extend manufacturers' warranties for a year! In each case, a TTP is guaranteeing aspects of the transaction, reducing the need for trust between buyer and seller.

Trustwise

BT has recently entered the electronic commerce TTP market with the launch of a service called *Trustwise*⁹. BT is in an excellent position to play TTP roles owing to its highly visible brand associated with integrity, reliability and global reach. Trustwise, illustrated in Figures 1 and 2, offers facilities that engender environmental trust on the public Internet by alleviating some of the concerns associated with user/server authentication and security of communications. These are concerns that have been widely publicised as inhibitors to the uptake of Internet commerce¹⁰.

Trustwise issues digital IDs (also known as *digital certificates* or *X.509 certificates*) that provide a means of proving identity in electronic transactions⁷. Proof of identity is an important issue for electronic commerce, since faking an identity is much more easily achieved online than in the physical world—'On the Internet no one knows you're a dog!' Digital IDs support a range of security functions based on public

key cryptography¹¹, including user/server authentication, digital signatures and data encryption. A brief description of the technology

underpinning digital IDs is provided in Panel 1.

The initial Trustwise service offers two types of digital ID:

Panel 1—How do Digital IDs work?

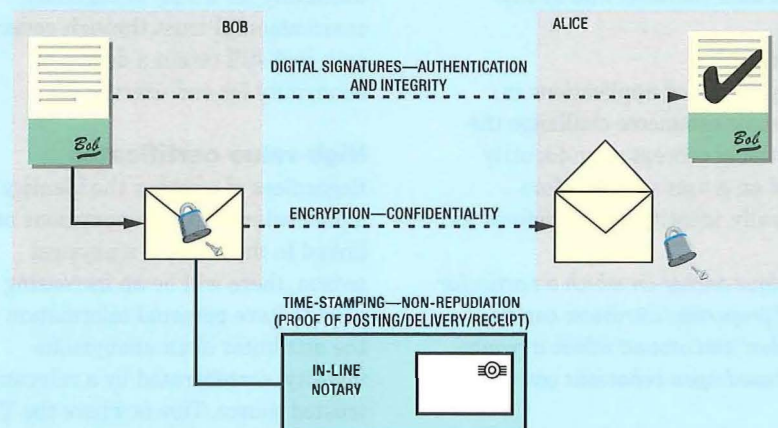
Digital IDs support a range of security functions based on public key cryptography. In a public key cryptography scheme, each individual is represented by a *public key* which corresponds to a specific *private key* that is kept secret by the individual. The public and private keys are used as a pair: typically, a message encrypted with either of the keys may be decrypted using its counterpart. The key pair may be used to provide:

- **Confidentiality** The public key encrypts a message such that it can only be read by the owner of the corresponding private key.
- **Authentication and Integrity:** The public key validates that a digitally-signed message was signed by the owner of the corresponding private key and that the message has not been changed since it was signed.

For a public key scheme to operate effectively, the public keys must be made available to all users of the scheme. They may, for example, be stored in a database or published on a Web site. However, if the public keys are published as raw data objects, users cannot readily determine whether they are genuine or whether they have been tampered with. Digital IDs overcome this problem.

A Digital ID 'binds' certain information about an individual (for example, name, e-mail address) to the public key that represents the individual in the on-line world. The digital ID must be issued and digitally-signed by a third party, generally called a *certification authority* (CA), who is trusted (directly or indirectly) by both transacting parties. In signing the digital ID, the CA guarantees the *integrity* of the information contained in the digital ID (so digital IDs can be freely distributed over insecure networks, safe in the knowledge that any tampering will be detected), and vouches for the *accuracy* of the information contained in the digital ID, to within the constraints laid down in the CA's certification policy. The use of trust services for messaging is illustrated in Figure 3.

Figure 3—Trust services for messaging



- *secure server IDs*—used by merchants to assure customers that their Web sites are genuine and to enable confidential communication of personal information (for example, credit card details) between Web browsers and servers using the secure sockets layer (SSL) protocol¹²; and
- *personal IDs*—used by individuals to authenticate themselves to Web services (for example, to provide access to personal bank account details) and to secure e-mail messages using digital signatures and encryption.

It is envisaged that the Trustwise product portfolio will be extended to include a wider range of digital IDs; for example, digital IDs for use in closed trading environments, plus time-stamping and notarisation services to support non-repudiation⁷.

Future Directions

The basic trust services that are currently being rolled out under the name of Trustwise provide functions that enhance the security of the communications environment and form the basis for establishing an on-line identity. To begin to create a richer environment, research work is also being carried out into enhanced services that will allow businesses and individuals to go beyond saying who they are and start saying something about their character and ability.

Anonymity

Many advanced applications in electronic commerce challenge the traditional concept of an identity based on a user's name. More generally, identity can be defined as:

'A unique anchor on which a particular set of properties / attributes can be attached and around which a reputation based upon behaviour can be built.'

One of the attributes associated with the identity can be a person's

name, but in many cases this will be undesirable. The increasing popularity of the Internet as a medium for human-human interaction and for commerce has raised many worries about privacy. With the power and efficiency of modern database tools, it is possible to build a sophisticated profile of particular users based upon what sort of things they purchase, what sort of information they search for on the Internet, the sort of people that they interact with on the Internet and the amount of time they spend on-line. The sophistication of on-line information gathering will increase the demand for anonymity, but users will still need some way to maintain a persistent state on the Internet. Anonymous identities will allow users to build up profiles and participate in activities without revealing their true physical identities.

In many cases it is not necessary to know who a person actually is as long as it is possible to verify that he/she has particular properties (for example, is authorised to enter into commercial transactions up to the value of £100) or has a particular reputation (for example, this is the same anonymous person I dealt with last week). These types of services can be provided by a TTP that has a strong relationship with its subscribers and a sufficiently strong reputation to encourage the trust of relying parties. The provision of these types of services will help to create an environment where there is sufficient continuity for a TTP to engender environmental trust through governance and still retain a degree of anonymity for end users.

High-value certificates

Regardless of whether the identity of a particular entity is anonymous or linked to the name of a physical person, there will be an increasing need to have personal information (or the attributes of an anonymous identity) corroborated by a relevant trusted source. This is where the TTP acts to engender transactional trust through enabling the transacting

parties to prove various attributes to one another.

On a simple level, these services could be viewed as static and one-way. A TTP, in conjunction with the relevant authority, might issue a certificate that testifies to the creditworthiness of a particular individual up to a certain level. This would function rather like a cheque guarantee card, providing assurance for transactions made by the individual up to a fixed value. A more sophisticated service might authorise a user to participate in commercial activity up to a variable amount, specified dynamically. This would be analogous to having a cheque guarantee card for which the guaranteed value was determined on the fly for each transaction.

The whole area of corroboration is based upon the need of one party to prove something about themselves to another party. This leads to a range of applications that can be loosely grouped together under the banner of *high-value certification*. High-value certificates will be dynamically issued, may be for limited use (for example, one-time only certificate which states that the user can pay £1000) and will be specific rather than generic.

The use of high-value certificates for the corroboration of data provides a basis for a range of commercial activities. Some of the more advanced electronic commerce applications of the future, for example, will allow users to source/supply complex goods and services where trust may be the most important characteristic of a potential supplier. A contract to supply a certain material that is essential to a manufacturing process will only be agreed with a particular supplier if the buyer is certain that deadlines will be met and that the supplier is likely to remain in business for the foreseeable future. This leads to the concept of *trust-based matching* in which the matching of suppliers to buyers is based on a sound understanding of the buyer's idea of a trustworthy supplier and an evaluation of a supplier's verifiable reputation within the trading environment.

Certification of intelligent agents

In the trading environments of the future, an increasing amount of activity will be undertaken by intelligent software agents that act autonomously on behalf of users. These agents will have the ability to enter into transactions on a user's behalf. However, if users are to have the confidence to employ such agents, they will require services that allow them to control the agent's activities and limit their own liability.

For example, within an agent-based trading environment, each agent might be certified by its owner, who in turn would be a certified user of the system. The agent's identity certificate would be made to reflect its ownership, either explicitly or anonymously. This would form a basis for agent identity and, hence, a foundation for trust in user-to-agent and agent-to-agent interactions.

Having proved that it belongs to a valid owner, an agent would need to prove that it is authorised to act on its owner's behalf in each particular activity. A TTP (or trading environment owner) could provide facilities for users to certify their agents based upon personal properties and requirements. For example, a particular user who is certified to enter into commercial transactions up to the value of £1000 may decide to sub-certify his/her agents to enter into transactions up to the value of £100. The TTP would still provide the anchor point for trust in the system, but within the guarantees offered by the TTP a user could certify the agents that work on his/her behalf.

Conclusions

The Internet is widely predicted to revolutionise commerce over the next few years. However, the full potential of electronic commerce will only be realised if buyers and sellers have the confidence to trade in electronic environments. TTPs have vital roles to play in engendering trust between trading parties and reducing the risks associated with carrying out

electronic transactions. BT is well positioned to take on these TTP roles and has recently entered the TTP market with the launch of Trustwise. Opportunities exist for a wide range of advanced trust services to be developed over the next few years.

References

- 1 EDELHEIT, J. A.; and MILLER, M. R. Electronic Commerce comes to the Net. *BT Technol. J.*, April 1997, **15**(2), p. 24.
- 2 STEWARD, S.; and VIDELO, I. D. E. Intelligent On-Line Purchasing. *Br. Telecommun. Eng.* April 1998, **17**(1), p. 36.
- 3 Miriam Webster Dictionary. <http://www.m-w.com/>
- 4 BARBER, B. Logic and Limits of Trust. New Brunswick Rutgers U.P., 1983.
- 5 AXELROD, R. The Evolution of Co-operation. Basic Books, 1984.
- 6 The Street. <http://www.TheStreet.com/>
- 7 SKEVINGTON, P. J.; and HART, T. P. Trusted Third Parties in Electronic Commerce. *BT Technol. J.* April 1997, **15**(2), p. 39.
- 8 DUN & BRADSTREET. <http://www.dnb.com/>
- 9 Trustwise. <http://www.trustwise.com/>
- 10 PHOENIX, S. J. D.; and REGNAULT, J. Fighting the Information Terrorists. *Br. Telecommun. Eng.*, Jan. 1998, **16**(4), p. 301.
- 11 PHOENIX, S. J. D. Cryptography, Trusted Third Parties and Escrow. *BT Technol. J.*, April 1997, **15**(2), p. 45.
- 12 Secure Sockets Layer protocol. <http://www.netscape.com/newsref/std/SSL.html>

Biographies



Tim Rea
BT Networks and
Systems

Tim Rea works in the Electronic Commerce Research group in Applied Research and Technologies and is involved in a number of activities including trust and security as applied to electronic commerce as well as international issues in electronic commerce. Tim was awarded an Engineering degree from Sheffield University in 1993, he completed the BT M.Sc. in 1997 and is currently working on an Executive MBA at London Business School.



Peter Skevington
BT Networks and
Systems

Peter Skevington manages the Electronic Commerce Research group in Applied Research and Technologies and leads BT's research work on trust services. He is also involved in a wide range of other activities associated with the future of electronic commerce, including work on trading environments, intelligent agents, internationalisation and novel applications. Peter joined BT Laboratories in 1985 after graduating in Physics from the University of Cambridge. He initially worked on the growth and analysis of semiconductor devices for optical transmission systems and was awarded a Ph.D. for this work from Cardiff University in 1993. He subsequently spent three years working on systems for managing change in complex enterprises before taking up his current interests in electronic commerce in 1996.

Whither Video?— Television and Telepresence

The telephone is one of the most successful and popular inventions of the modern world. Television is also enormously popular and successful. Yet videotelephony remains a relatively unpopular technology, despite extensive study and major advances. This is the first in a series of five articles about the future role of video pictures in communications, exploring commercial, engineering, human, artistic and cultural factors.

The Magic Box

An eight-year old boy once stood in awe of his grandparents' Bakelite radio. Inside there were glass bottles, with curious little fires that burned without burning away. That radio could do magic. With a long piece of wire and something called a *short-wave band*, it could collect voices from faraway lands. The boy was intrigued to hear how an angry torrent of chilling words could be followed at once by a warm and happy song from the same place. What kind of contradictory people lived in this place called Moscow? His body was in the upper Neath valley, but his imagination was somewhere else.

His adult imagination is today transported less effectively by the more advanced technology of television. How can this be? Everyone knows that television is more powerful than radio ever was. Nothing escapes the video camera, which can instantly show the smallest detail of distant events. Television is now so powerful and pervasive that it shapes our attitudes to life and the world. The visual conventions of television permeate our unconscious minds so deeply that recognising what those conventions are can be difficult. Television practically defines our expectations of what electronic pictures can do.

Picture Paradox

The American comedian Ernie Kovacs said that 'television is called a medium because it is neither rare nor particularly well done'. The electronic mirror that reflects the world so sharply has nonetheless

acquired a mixed reputation. Television accurately portrays events yet can put things out of context. It readily creates intrigue but sometimes struggles to create involvement. We can watch a sporting event on television and see exactly what happens without sharing in any of the atmosphere of the place. Television is able to take us round the world while keeping our minds firmly fixed at home. Is this paradox an inevitable consequence of video technology?

Only Speech

Speech is always someone's opinion or belief. Pictures have an independent claim on reality and truth. Yet television is full of fantasy and radio is where facts weigh more heavily. In news programmes and in soap operas, television is about conflict yet telephony has a reputation for bringing people together. So could speech, by telephone or radio, be more effective at transporting the imagination than speech combined with pictures?

Role of Pictures

This question is not esoteric but commercially important in telecommunications. Electronic pictures figure prominently in the future convergence of telecommunications, computing and entertainment. Recent developments promise an interactive mixture of telephony with pictures and data, for work and for play, that might revolutionise the business of telecommunications. What lessons can the development of television teach about the future role of pictures in communications?

Both the technical and psychological needs of telepresence are very different from those of entertainment, because of the pivotal role played by interaction between the participants

Commercially-feasible entertainment, broadcast one-to-many, gives viewers no influence over what they see. But face-to-face communication is always interactive and often one-to-one, where viewers can strongly influence what they see. Some sense of 'being there' can be valuable in entertainment but that sense is vital for telepresence communications. Both the technical and psychological needs of telepresence are very different from those of entertainment, because of the pivotal role played by interaction between the participants (Figure 1).

Social Serendipity

The success of the telephone has been a lucky accident in more than one sense. The electrical signal, from a microphone held close to the mouth, happens to be an excellent representation of the content and character of human conversation. Technical signals and cognitive sentiments are closely related in telephony. But remote communication by telephone is successful for more complex and subtle reasons, too.

It is not socially acceptable to walk up to a total stranger and whisper in their ear. When we meet a stranger face-to-face, we do not speak in a casual and direct way like old

friends. Yet these kinds of behaviour are entirely acceptable when using a telephone. Natural human behaviour in face-to-face conversation has proved very adaptable to the different circumstances of remote speech. When we make a telephone call there is an assumed fraternity that makes informality more acceptable, as well as a shielded intimacy that makes direct contact unthreatening. These social and psychological adaptations have become so well established and widely understood that they are now taken for granted in telephony.

Fragility in Pictures

The social norms and adaptations for remote communication with pictures are not well developed or understood. Any sense of fraternity or of shielded intimacy develops less easily, or not at all, with a picture. Furthermore, picture signals can be a misleading representation of the intended sentiments. In all sorts of ways, pictures are not like spoken words. They seem to be more fragile than speech, in technical and psychological terms, and are not as robust to dislocation by distance, despite their greater expressive potential. Extending and expanding the experience of telephony through the addition of

pictures means addressing a range of interdependent social, psychological and technical design issues.

Outlook

This series of articles examines the future use of pictures in communications, putting task above technology to link disconnected areas of knowledge and challenge preconceptions. This article reviews television technology, and provides a glimpse of some of the dreams of the pioneers. Subsequent articles will discuss videotelephony, the nature of human communications, perceptual issues in telepresence and artistic aspects of pictorial culture. Finally, some choices in the capture, display and use of pictures in communications will be illustrated.

Video History

At the end of the nineteenth century, people all over the world were experimenting with the *radioscope*, the *telephone eye* or *visual listening*. Paul Nipkow and Jean Weiller described mechanical devices to scan pictures in the 1880s. In 1907, the Russian Boris Rosing proposed the use of a cathode ray tube for displaying pictures. In 1908, Alan Campbell Swinton was the first to suggest a wholly electronic picture system¹, with magnetically-scanned cathode ray tubes for both camera and display.

It was John Logie Baird (Figure 2) who made the first public demonstra-

Figure 1—Human interaction, grouped by number of originators and recipients

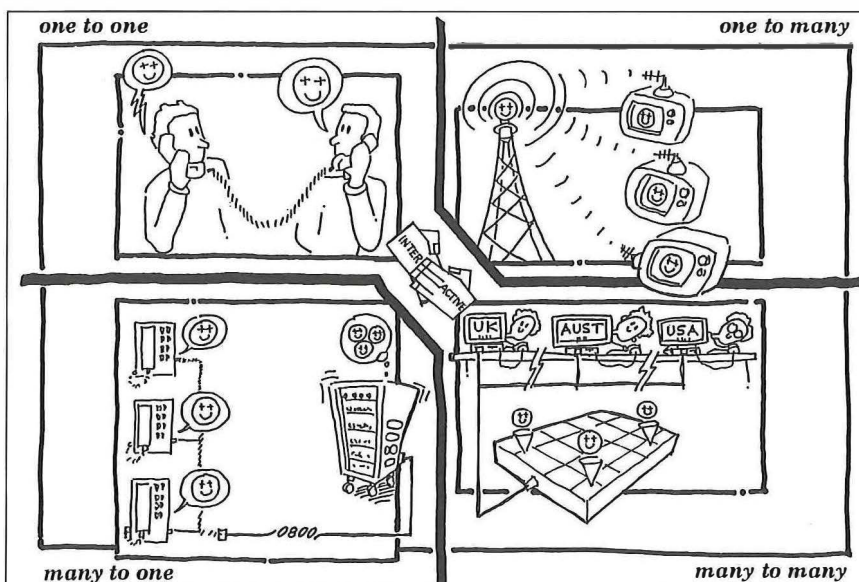


Figure 2—John Logie Baird, adjusting the receiver of his early wireless vision apparatus, 1926 (Credit: DHA/NMPFT/Science and Society Picture Library)




Figure 3 – Philo Taylor Farnsworth



tion, in April 1925 at Selfridges's Department store in London, of television pictures of naturally-lit objects². His system used 16-line mechanical scanning, producing red-and-black pictures without intermediate shades. Baird demonstrated continuous-tone 30-line television pictures in January 1926, at a meeting of the Royal Society. On 9 February 1928, Baird transmitted pictures across the Atlantic ocean, 34 years before the Telstar satellite. Later that year, the irrepressible Baird demonstrated colour television with sequential red, green and blue line scans. By 1930, he had experimented with infra-red *Noctovision* and stereoscopic television.

Through Amateur Endeavour

With far smaller resources than the silicon revolution deploys today, television developed much faster. In a staggering burst of inventiveness, television went from crude curiosity to high-definition broadcasting in little over a decade. Amateur enthusiasts, in the UK and America, played a major part in these developments³.

In America, Ernst Alexanderson experimented with 48-line mechanically-scanned television in 1927. Multi-standard receivers, for 24, 36 and 48-line pictures, were commercially available in 1928, when Charles Jenkins started experimental transmissions on shortwave radio. In the same year, Ulysses A. Sanabria invented interlaced scanning to reduce picture flicker. But the pioneer of electronic television was Philo Taylor Farnsworth Figure 3, an American inventor and polymath who described his ideas for capturing images with electrons to his school-teacher at the age of fifteen^{4,5}. He first demonstrated his *image dissector*, a magnetically-scanned electronic camera tube, seven years later in 1927.

Law and Commerce

Vladimir Zworykin had applied for a patent of a similar idea in 1923 but

had been unable to make it work. A legal battle ensued and in 1935 the court, in a 47 page judgement, awarded priority of invention to Farnsworth. Researchers at the Radio Corporation of America tried long and hard to avoid infringing Farnsworth's portfolio of over 100 patents on fundamental aspects of electronic television. In 1939, RCA conceded defeat and negotiated their first ever agreement to pay third-party royalties. The giant of corporate America bowed to the farm-boy from Idaho. Despite this clear outcome, record books as eminent as the *Encyclopaedia Britannica* still attribute Farnsworth's numerous discoveries to RCA's more heavily-publicised developer Zworykin.

Testing Testing

After much wrangling, Baird was first granted the use of BBC transmitters in 1929, to radiate his audio-band-width video signal after the day's radio broadcasting had finished. By 1935, several thousand of Baird's 30-line *Televisors* had been sold to receive a tiny, flickering picture of the half-hour experimental programme, transmitted around midnight.

The first television transmission by a radio amateur was in January 1933. Reception reports of Baird's high-frequency test transmissions, at 43 MHz from Alexandra Palace, were received from Europe, Iceland and Morocco in the early 1930s. The first outside broadcast came from the finishing line of the Derby racecourse in 1931. The first television transmission from an aircraft in flight was made over Hendon aerodrome in 1939. Workers at EMI used a development of the image orthicon camera tube to make an all-electronic 405-line system, which was tested alongside a 240-line refinement of Baird's mechanically-scanned approach.

Broadcasting

This profoundly creative period of experimentation was replaced by

professional development when the BBC, from a single transmitter at Alexandra Palace in London, launched the world's first regular television broadcasting service. On 2 November 1936, after the opening speeches, Adele Dixon sang a eulogy in round vowels:

*A mighty maze,
of mystic, magic rays,
is all about us in the blue,
and in sight and sound they trace,
living pictures out of space,
to bring our new wonder to you.*

In America, broadcast television was launched by RCA at the New York World's Fair on 30 April 1939, without full permission from the Federal Communications Commission.

Television Develops

Some key events in the history of television in the UK are:

September 1939 BBC television stops abruptly, in the middle of a Mickey Mouse cartoon, shortly before the outbreak of war. Mickey's final words were 'I tink I go home now'.

December 1940 Baird demonstrates the *Telechrome*, a 600-line electronically-scanned colour display (Figure 4), but war prevents development.

June 1946 BBC television re-opens and the ending of the Mickey Mouse cartoon is shown.

October 1954 The first test of 405-line colour television, using the

modified NTSC system, is transmitted from Alexandra Palace to a single colour receiver.

July 1962 High-definition television pictures are sent across the Atlantic ocean, via the low-orbit Telstar satellite.

1964 625-line television transmissions start, in the London and Midland areas.

July 1967 Colour television is launched, using the PAL system.

A Glimpse of the Past

Baird had attempted to record his pictures on gramophone discs, applying for a patent for his Phonovision process in 1926, but he never achieved satisfactory playback. Some enthusiasts attempted to record Baird's experimental transmissions with the *Silvertone* audio disc recording system, commercially available in 1930.

Recently, Don McLean⁶ has applied complex digital signal processing algorithms to replay a number of these recordings, with excellent results. It is now possible to get a glimpse of what was on television in the 1930s. One privately-made *Silvertone* recording is particularly revealing⁷. It shows almost four minutes of the first television revue,

Figure 4—Dufaycolour photograph of Miss Paddy Naismith, the aviator, taken at Sydenham on 14 January 1941 from the screen of Baird's 600-line, two-phosphor Telechrome television display tube (copyright Miller-Freeman Publications, used by courtesy of Electronic Engineering magazine)



starring the Paramount Astoria Girls, broadcast on 21 April 1933.

Fitting Art to Science

Despite the technical limitations of the 30-line system and the difficulty of performing in front a camera that scanned a dazzling spot of light around a totally dark studio, the show is slick and vibrant. The production is unusual, because solo artists enter and leave by standing up and ducking down through the bottom edge of the picture. This seemingly strange behaviour minimises the visibility of the coarse vertical scanning.

This shows that the pioneers of television, immersed in engineering problems, were aware how appropriate artistic behaviour was a fundamental part of fitting technology to a purpose. This is still true today, for both television and telepresence. Art and science complement each other at the edge of the unknown.

Ways of Thinking

Much of the development in the early days of television depended on synthesis not analysis and was motivated by dreams not logic. This does not fit with the conventional view of engineering as a deductive, applied science. Words like judgement, intuition, creativity and metaphor do not belong in that view, yet those qualities are essential to the engineering of unprecedented things. It would be wrong to advocate fervent dreaming instead of academic proficiency, yet that is just what gave birth to television and to the telephone.

Long ago, an amateur was someone rich enough to do a thing for love not money. Calling somebody an amateur was a compliment—professionals had to earn a living and had lower standards. Today, a resurgence of the inspired enthusiasm that gave birth to television is directing new applications of electronic pictures. Much of this resurgence is amateur, coming from outside the established hierarchies.

New Roles

Recent initiatives have explored a new role for television, linked to new forms of cultural experience in the arts and entertainment. Inhabited TV is a multi-user virtual environment, mixing the ideas of social interaction, audience participation and content production to create a new kind of television programme. The *Mirror*⁸ was a groundbreaking experiment, created by BT, Sony, Illuminations and the BBC, shown on BBC2 in February 1997. More than 2000 viewers shared six imaginary worlds, becoming a temporary community through both scripted and unscripted encounters.

Broad What?

The television industry is now in technical turmoil over what comes next. International agreement on a single standard for high-definition television (HDTV) seems little closer after 10 years of development and debate. There is no consensus about whether cable, satellite or terrestrial transmission is the best method of programme delivery.

Critics say that television is evolving away from the high, if condescending, public-service traditions of *broadcasting* towards the vitality and lower standards of the free market, where thousands of channels will amount to *broadcluttering*. If the role of performer and viewer become interwoven, as they were in the *Mirror*, perhaps the diversity of the Internet might combine with the ubiquity of television, in an era of highly fragmented but interest-specific *broadcatering*.

Video for All

The price of televisions, video recorders and camcorders has remained numerically stable, but the price of computer video technology has fallen rapidly—by a factor of almost 10 in six years. It is now possible to buy a simple camera for less than \$100 and do domestic-quality video editing on a high-specification personal computer.

The quality of Internet videotelephony connections is severely limited by bandwidth and latency, but several vendors offer hardware and software⁹.

Real-time MPEG encoders and accelerator devices for 3-D rendering are already available and digital video disks (DVD) have seen commercial release. New types of video image sensors promise both lower power consumption and lower costs¹⁰. CMOS digital micro-mirror arrays and organic light-emitting polymers^{11,12} are just two of the competing technologies for flat displays. Electronic pictures are becoming increasingly affordable and ubiquitous.

Looking Ahead

Yet videotelephony remains paradoxical. The authors believe that thinking of a videotelephone as either a telephone with pictures or as a two-way television has been traditional but is misleading. Too little is understood about the psychological and social differences between interactive communication with words and with pictures. The next article in this series will examine how well videotelephony supports natural behaviour in human communications.

References

- BRIDGEWATER, T. H. 'A. A. Campbell Swinton'. Royal Television Society monograph, 1982.
- MOSELEY, S. A. John Baird: The Romance and Tragedy of the Pioneer of Television. Odhams, London, 1952.
- HERBERT, R. Radio Amateurs and Early Television. *Radio Communication*, Oct. 1996, pp. 17–19.
- FARNSWORTH, E. G. Distant Vision: Romance and Discovery on the Invisible Frontier. **Pemberly**Kent, Salt Lake City, 1990.
- <http://songs.com/noma/phil/index.html>
- MCLEAN, D. F. Computer-Based Analysis and Restoration of Baird 30-line Television Recordings. *Journal of the Royal Television Society*, April 1985, **22**(2), pp. 87–94.
- <http://www.dfm.dircon.co.uk/>
- WALKER, GRAHAM. The Mirror—Reflections on Inhabited TV. *Br. Telecommun. Eng.*, April 1997, **16**(1), pp. 29.
- SEACHRIST, D. See and be Seen Over IP. *BYTE*, Sept. 1997, **22**(9), pp. 104–108.
- KAO, S. Digital Images. *BYTE*, Sept. 1997, **22**(9), pp. 21–24.
- YOUNSE, J. M. Mirrors on a Chip. *IEEE Spectrum*, Nov. 1993, **30**(1), pp. 27–31.
- DETTMER, R. Conjugal Lights. *IEE Review*, July 1996, **42**(4), pp. 135–137.

Biographies



Alwyn Lewis
Networks and
Information Services,
BT UK

Alwyn Lewis is an advisor in Advanced Applications and Technologies, in Networks and Information Services, BTUK. He gained a B.A. from Cambridge University in 1971, and an M.Sc. from Essex University in 1972. After a spell with Plessey working on format-tracker vocoders, he joined BT Laboratories to work on acoustics, telephotonometry and telephone design. In 1984, he became head of a team developing a duplex hands-free telephone with a custom VLSI chip-set and a DSP microprocessor. He then led the speech coding team, helping develop and test DSP software for CallMinder and for Skyphone facsimile and data services. His interests include beam steering

microphones, speech coding and enhancement, adaptive audio signal processing and the social impact of electro-technology. He is a Chartered Engineer and a Member of the IEE and IEEE.



Graham Cosier
Networks and
Information Services,
BT UK

Prior to 1984 Graham Cosier worked in the area of electro-acoustics and published several papers on auditory human perception which contributed to the British Standard for hands-free telephony. In 1986, he became involved with the emerging digital mobile radio system in Europe and was elected chairman of one of the technical groups that standardised GSM. In 1988, he became head of the group concerned with the assessment of coding methods, channel optimising techniques and echo control systems for increased utilisation of services using telephone networks. In 1994, he was appointed UK coordinator for ITU SG12 where he led the formation of the international standard for digital telephony. In 1995, he was appointed as the BT technical manager for the MIT Media Lab and was involved in establishing the UK Creative Art and Technology Centre. He is currently head of one of the units within the Applied Research and Technology Department at BT Laboratories where he leads several research teams looking into future tele-presence systems, perceptual modelling, media environments, audio and vision based perception, gesture and affective communications. He is also the business support manager for BT Global marketing, championing a vision of 'through the screen telepresence'—the ultimate collaboration space, that will provide a 'better than real life' experience between remote locations with the vision that 'you don't have to be there, to be there'.

The Future World of Telecommunications, 2005?

This article presents one perspective of what the world will be like in the year 2005. It considers political, economic, social and technological trends in order to paint a 'vision' of the future, and examines how these trends will affect the different roles within the telecommunications industry. The question this work is leading to is 'how will operational support be different?'

Introduction

In the year 2005, the telecommunications market is likely to be quite different from today. For a company like BT, which invests in systems that offer high-quality services and service surround to its customers, it is important to understand what such a world may be like, to help plan for the long-term system investments. This is not just for technology aspects, although changes here are very important to service companies based on infrastructure. The political, economic and societal aspects of the world in 2005 are also important as they structure the marketplace in which the players do business. It is also important to consider the types of role that will form the industry and how they will interact. Players may take on many roles and need to understand who are the suppliers and customers of each of these roles.

This exploratory article describes one possible perspective on that world of 2005, covering the political, economic, societal and technological views and their impact on roles involved. This article tries to represent a significantly futuristic view—2005 is a label for a time in the future

rather than an exact date to which plans should be constructed. This view of 2005 is part of a high-level methodology under development to enable analysis of the effects of market and technology changes on the role of operational support systems (OSS).

Initially the article discusses a model based on four roles within the future telecommunications industry and then prepares an outline of the future economic and political world. The view is a very western view, one based on strong economic-political linkage rather than one built from a connection between society and politics (for example, religious) or one dominated by political influence (for example, communism). After considering society and the nature of technology in 2005, the discussion moves on to conclude on five derived key points that describe the customer marketplace for telecommunications. Finally the article presents some possible implications that affect the systems within the roles discussed.

The Industry Model

The telecommunications industry can be broken into four top-level roles. Table 1 considers these four types of

Table 1

Customer	To use services to contract terms
Service Provider (SP)	To provide addressed services to customers by using network and access connectors; for example, intelligent network (IN) services, information storage and processing, conferencing, etc.
Network Connector (NC)	To deliver addressed 'streams' to correct address (network connection point)
Access Connector (AC)	To provide physical connectivity between customer premises and network connection points. Logically represent the customer at the network connection point and the network at the customer's location.

Figure 1—(a) A model for the role interactions in the telecommunications industry, and (b) three configurations that may exist

role, and details their descriptive responsibilities. These are roles not players; a player may take on one, a few or all of the roles. This is not a business model.

The service-provider role covers a large range of possible services, from 'IN-like freephone' or 'play when dialled' services to multimedia conferencing. This SP role has sometimes been split into two: a network service provider (NSP) for the IN-like services and the higher 'value' service provider (HSP) for the conferencing services as shown in Figure 1.

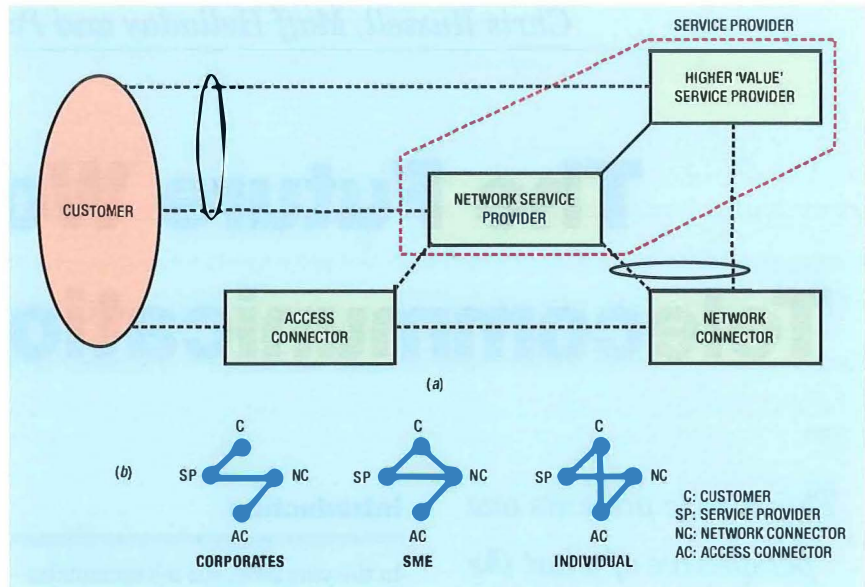
The service-providing industry is potentially very complex with service-provider players buying in and using other services from other suppliers. Modelling such a 'hierarchy' is beyond the scope of this article¹. The model presented here is based on the assumption that all such SP models will be enabled by communications infrastructure services, thus they will require users to have access to telecommunications networks. The access-connector and network-connector roles (and possibly the network service-provider role) will therefore be present for all situations.

Customer types and interaction diagrams

There are many ways in which these roles may interact, and players may adopt one, or all, of these roles. The interaction of the three providing roles (SP, NC, AC) with the end-customer has seven possible options. The marketplace may result in all options occurring but there are perhaps three strong configurations, which are shown in Figure 1.

The first is where the end-customer interacts only with the service-provider role, which arranges the necessary network and access facilities on the customer's behalf. This configuration may be particularly suited to the *global corporates*.

The second is where the end-customer interacts with both the service-provider and the network-connector role. The network-connec-



tor in turn arranges the necessary access-connector facilities. This configuration may be suited to the *small and medium enterprise (SME)* customers who may have a wide geographical spread of locations and rely on networks for communications between these locations.

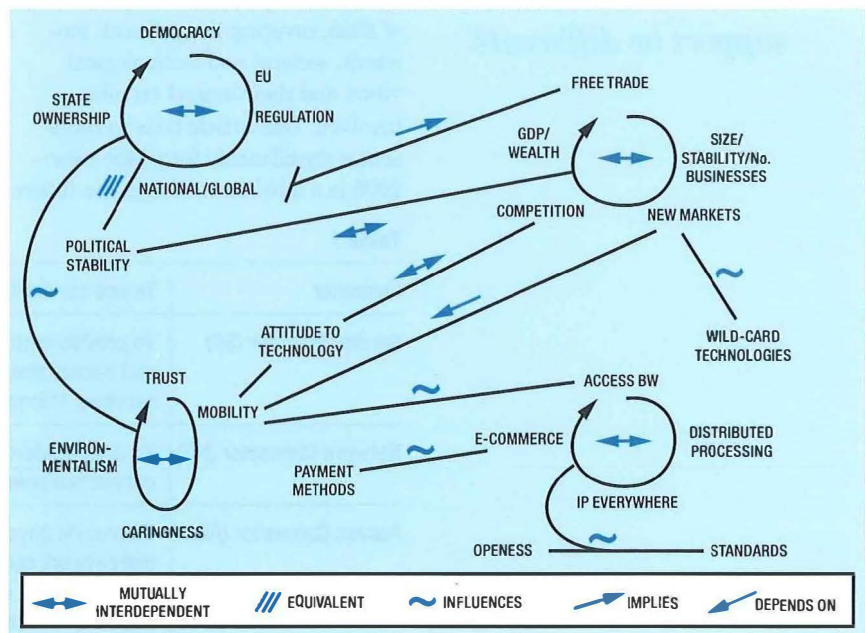
The third configuration is where the end-customer interacts with the service-provider role and the access-connector role. The network-connector facilities could be arranged by either the access connector or the service provider in this case. Such a configuration may be very suited to the *individual* customer who generally has a single location and is best

placed to understand the cumulative requirements from the many services that they may be using.

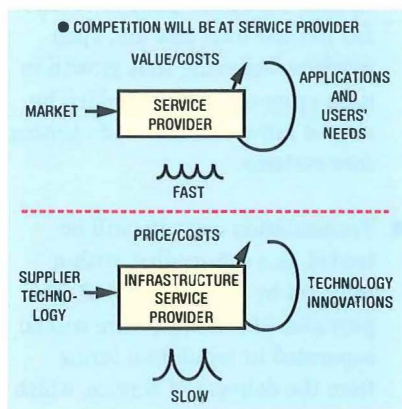
In the Year 2005

So what factors will impact on these roles? What should a player be concerned about in 2005 if they are planning to adopt any one of them? The next four sections present opinions and information on present trends that describe views of the future world and the impact that they may have on the roles described. In all, 24 different factors, six in each view, were considered and the linkages between these are shown in Figure 2.

Figure 2—Factors considered in developing the four views



The economic world



The migration of power from state to companies will have continued and the power of the global economy will be significantly stronger than it is today.

- Already 51 out of the top 100 economic entities are companies not states².
- If the world's wealth is measured with the wealth of Italy today as a unit, the USA is eight Italys, Western Europe another eight, and Japan a further five. The rest of Asia, South America and Russia constitute the remaining five.
- Five Italys are 'spent' by companies each year purchasing other foreign companies, buying into those countries' markets.
- Today is the age of skills, and the service economy (the GDP of a western developed nation is divided approximately 5% agriculture, 20–30% manufacturing and the remaining 60–70% the service sector).
- The USA has seen a 23% increase in the number of jobs since 1980 which have all been in the private sector. The growth in SMEs (less than 100 workers) has been twice the average².
- The developed countries' collective productivity (the ability to utilise technology to integrate the value and supply chains together) has

been and will be the key that enables them to retain wealth and continue to grow their economies as other nations develop (for example, China).

Competition will be most vigorous in the provision of tailored services to end-consumers. Value to end-customer and cost of production are the key considerations. Potential markets for service providers will come and go as the needs and aspirations of the many individualised consumer types change.

- The service economy will be segmented across all sectors into small and medium enterprise (SME) niche providers able to respond rapidly to the frequent shifts and developments in their markets.
- These SMEs are either virtual (structured from networked 'self-employed' skilled individuals) or very localised geographically but will be able, through the development of open-networked technology, to offer their services to the global market.
- They will however be very sensitive to market change and less stable as a result—the services that the SMEs offer change as rapidly—and employment will be dependent on people's skills and experience.

The world of infrastructure provision will be less dynamic and competitive. Price and cost are the key concerns. Bandwidth exchanges³ will have become quite well developed and capacity will be a commodity trading at a published, and fluctuating, market price.

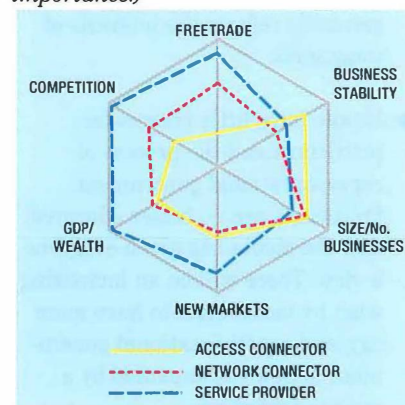
- Competition in the provision of network or access infrastructure occurs when technology opportunities arise.
- The market has reduced drastically the profitability of the

connectivity supply business and only a few, well established and stable players exist as a result. Large-scale geographic networks are operated by large stable players or alliances.

- The access environment may be a little more complex. For example, housing associations and developers may opt to install communications infrastructure on estates that they then own and maintain. As access connectors they then sell connection to this 'valuable' access infrastructure to the players offering networks.


Figure 3 describes some impressions of how these economic factors may affect the different roles in the industry model (Figure 1). This is qualitative information, based on perceptions and thoughts gathered using a questionnaire from several people. Strong relevance is represented by intersection of the line near the edge of the matrix. The access connector and network connector are similar; but the access connector is more affected by the stability of business while the network connector is more affected by competition in the service-provider market and the global trade

Figure 3—The impact of economic factors on the telecommunications industry supply roles. (Note: Strong relevance is represented by intersection of a line near the edge of the matrix; the centre indicates no importance.)



situation. The service provider is influenced by the size and stability of businesses as a secondary effect to the wealth and competition in new markets.

The political world



The increasing role of business and the dominance of the global economy have diminished the role of government, reducing its influence. Government will become ever more involved with regulation and control rather than leadership.

- Increased globalisation will focus government at a world level (supranational governments) and support regulation on a global scale (World Trade Organisation, World Intellectual Property Organisation, International Monetary Fund are existing examples). Countries are now just members of these global government bodies, with ministers acting as representatives, and have to agree to consensus views.
- The conflict between companies' and peoples' interests continues but a nation's representation generally reflects the interests of companies.
- People have little respect for politicians and the process of representational government. Politicians are no better educated, and the media has given everyone a view. There will be an increasing wish by individuals to have more say, and representational government is now accompanied by a series of referenda.

Governance will therefore move more 'local', focussed on supporting the natural needs of individuals. To support this community networks become the norm, government (local and national) being a key player, and this will represent a new market opportunity for players in the telecommunications industry.

- The availability of affordable computing tools in the home will help local government, via these community networks, to communicate better with their constituents.
- The move to local government will be hastened by society's recognition that people elsewhere have an effect on their local circumstances (wealth, environment).

The enhancement of the regulatory role of government and its localisation will result in increased political stability within nations. Globally, the growing economy brings political stability; there is less scope for different government policy between left and right. The destructive forces of nationalism will not have developed as people recognise that stability requires cooperation and sharing (for example, the Euro stabilising the European economy).

Privatisation of telecommunications will have occurred in all major markets around the world. The way the UK and European market is regulated has changed significantly from the picture of the late 1990s. The regulation will have forced significant openness.

- Access by competitors to facilities within infrastructure such as switch interfaces, main distribution frames, and cross-connection equipment will be normal.
- The European Union (EU) regulation for open network provision and convergence will be in place.

- Developing nations will adopt similar (identical) policies to the EU in their own new and open markets, believing that growth in their economy is best enabled by shared infrastructure and opening core systems.
- Transmission capacity will be traded as a commodity³ with a price set by market demand. The provision of infrastructure will be separated in regulation terms from the delivery of service, which is regulated only by market practices (fair trading etc.). The role of the telecommunications regulator will be radically different from that of today.

Privatisation works!⁴ The role of the state in organising, investing and managing the utility sector and other 'service' activities is all but gone. Government regulates and manages those parts of the economy that cannot be privatised. Privatisation programmes only last for a short period as society loses interest, and so by 2005 the western economies will be fully privatised.

- There are presently privatisation programmes in 100 countries throughout the globe expected to generate \$200 billion⁴.
- In the UK the subsidy paid out by government has been turned into income for Treasury. A £1.5 billion subsidy and loan finance to 33 companies in 1979 turned to £8 billion contribution to the Treasury in 1987.

Figure 4 describes some impressions of how these political factors affect the different roles in the industry model (Figure 1). All roles are impacted similarly, the service provider more by the ability for people to choose (democracy) and the absence/presence of any political hurdles to entering markets (nationalism). The access and network connector roles are affected most by regulation and, especially for the

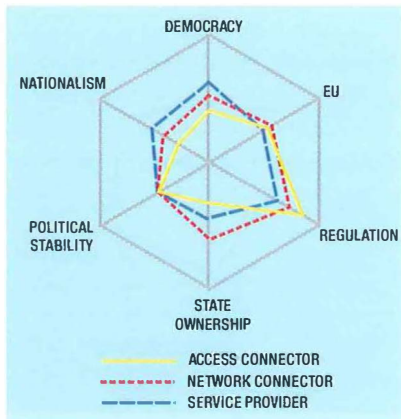
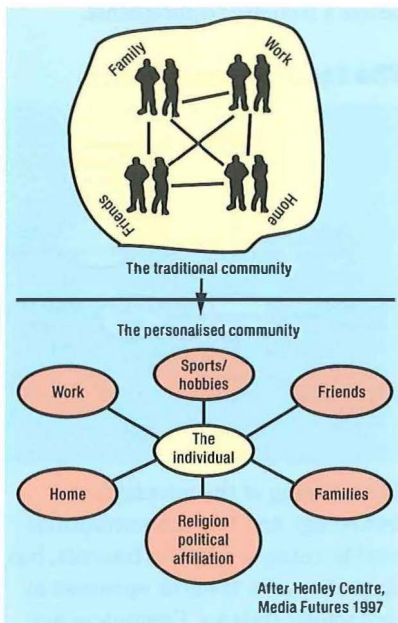


Figure 4—The impact of political factors on the telecommunications industry supply roles. (Note: Strong relevance is represented by intersection of a line near the edge of the matrix; the centre indicates no importance.)

network connector, by whether the market is open to fair competition (state ownership).

The societal world



Today individuals are taking on more and more responsibility for their futures (health, pensions, insurance etc.) and the service sector marketing is starting to become 'personalisation' dominated. In 2005 the individual will be expected to be even more responsible and society will be even more personalised and multifaceted.

- Traditional community-based links will have been supplemented (or supplanted) by links chosen by the individual. People's work, hobbies, affiliations etc, which may not be so local, will drive these links. The 'family' is likely to have a very wide meaning covering many forms, with people's interest networks (friends, hobbies etc) becoming important.
- Local interest groups will also be encouraged by the existence of community networks supported and encouraged by government. Home and work will have become less distinct, teleworking will have become more prevalent as technology and the SME skills-based market develops.

People's concept of telecommunications services will have changed significantly. Little distinction in value will be made between voice, video and data. Users will have begun to accept that quality is balanced with price and that premiums will apply to personalisation of factors (such as reliability or additional requirements) that are beyond the basic service package. The present trends towards regular, informal, and instant communication interaction will have, by 2005, become the expected norm.

- To support people's society networks communication will be more network-based, increasingly in 'electronic' form both for written and spoken interaction.
- Instancy will be vital, people will not want to wait long; cellphones are instantly enabled for mobile telephony services whereas more traditional voice 'services' have long lead times at present.
- Through experience with the Internet, people will have got used to only waiting a few seconds/minutes for the information they have requested. They will expect

the power of the technology they have invested in (computers etc.) to be enough for any task that they undertake.

Societies attitude to computing technology will be very different from that adopted with the telephone. Unlike the telephone, which has been around for many years and has just reached approximately 95% consumer penetration, the computer gets embedded into the lives of people from school age.

- Recent studies^{5, 6} indicate that only about 60% of the population are comfortable with the telephone but some 90% of UK children (11–18) in schools are confident in using computers/PCs.
- In addition, some 70% of parents, regardless of social class, felt it was critical that their children learn to use a PC. Assuming that those aged 18–25 today had a similar experience, by 2005 one third of the UK working population will be computer aware and confident in their use.
- People across all age and social groups are already of the opinion that computers are 'essential or important in my working life'⁷—in 1997, 53% of all people felt computers essential, with 60–70% in the higher earning groups.
- The PC is penetrating all sections of society^{8, 9, 10}—about 30% of people aged between 16 and 59 had a PC at home in 1997, a similar figure to that in 1992 but the machines are more powerful. For the higher earning groups between 55–65% own home PCs. 33% of small firms have portable laptop PCs. Ownership seems to be shifting emphasis away from games and entertainment towards education, skills and training, echoing the socio-political emphasis on individual responsibility (the risk society).

- As with the culture of the computing peripherals, the young workforce consider the mobile telephone to be the norm. They have grown up with the telephone and have no fear of mobile technology. Already the penetration into AB social classes is high with 43% owned in 1996 and 68% of small firms having a mobile telephone¹¹.

Society's approach to payment methods, something linked heavily to people's trust in both technology and companies, is also starting to change. In 2005, people will have some trust in, and will start to use, electronic methods of payment as the risks to them (fraud, error etc) have been minimised and a track record established.

- There is an increasing percentage of people who are happy to use the telephone as a purchasing channel. Studies¹² suggest a recent increase of 10–20% for a wide range of goods/services (insurance to shopping to holidays).
- Many people are happy to use credit cards on the Internet. American Express estimate that \$4–6 billion was spent using credit cards in 1997¹³ compared to \$68 billion private consumption in the US.
- New payment technologies like smart cards need to identify the value and benefit they give to the user. The MobilTM Speedpass is an example showing that users will take up the technology if the incentive is right—it makes life simpler and does not cost a lot¹⁴. The lack of widespread ability to use electronic cash is the present hurdle to adoption by society. Initiatives, such as the electronic payment system for London Transport that was recently in the British news will stimulate market adoption by generating a large captive high-density user base.

- Business-to-business commerce, e-business, can be efficient, as it can remove a number of intermediaries from the supply chain. By 2002 it is predicted that the number of businesses participating in business-to-business interactions on the Internet will rise from around 400–900 in 1997 to 245 000 in the EU (\$16 billion revenue from providing e-commerce systems) and 640 000 in the USA (\$37 billion revenue). Spend on implementing systems is expected to have peaked in 2001¹⁵.

There is strong evidence supporting the increased personal mobility of people. By 2005 one-third of the workforce will expect to use mobile technology at work and will increasingly own and use such equipment in their social time.

- Figures from the Department of Transport suggest that there will be a 20–30% growth in air travel compared to 1997. The growth is expected for both business travel and tourism in general¹⁶: UK air business departures will reach 9.7 million short haul and 4.6 million long haul and arrivals will reach 10.7 million and 5.2 million by 2001.
- Mobile telephone connections are also expected to rise rapidly, and by 2003 there will be 23.6 million mobile connections in the UK¹⁷. The same report predicts that the rest of Europe will have 131 million mobile connections. The situation is similar in North America and Asia-Pacific.

Figure 5 describes some impressions of how these societal factors affect the different roles in the industry model (Figure 1). The service provider and access connector are significantly affected by society's accommodation of new technology and payment techniques, and therefore on the level of trust, and the level of inclusion and universal service (Care).

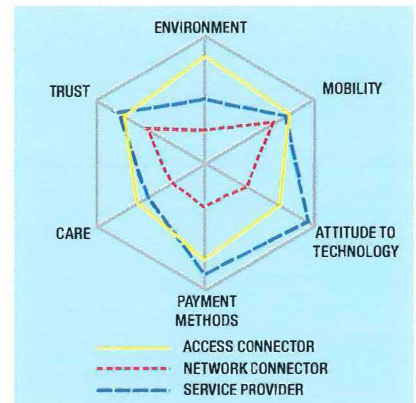
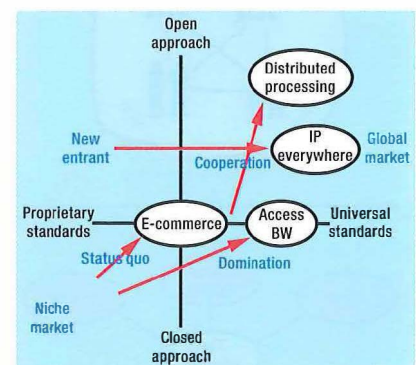


Figure 5—The impact of societal factors on the telecommunications industry supply roles. (Note: Strong relevance is represented by intersection of a line near the edge of the matrix; the centre indicates no importance.)

In addition, the access connector needs to consider society's view on the environment. The network connector role is affected least by societal factors, the primary influence is people's mobility requirements.

The technological world



The merging of the information technology and telecommunications worlds, catalysed by the Internet, has started a move towards openness in technology systems. Computers are already becoming a personal commodity item. A greater penetration of computing devices into business (and the home) will be assured.

- Adherence to the World Trade Organisation Information Technology Agreement that removes tariffs on IT products by 2000 (which already has 40 countries

representing 90% of world supply signed up¹⁸) will result in continued reduction in the price of computers.

Distributed processing has always striven for openness through the Object Management Group and Open Systems Foundation. Present standards work is aimed at increasing the openness and portability of object-request broker-based solutions.

- Distributed processing standards will be well developed and open, hardware-independent software (for example, Java) will be the norm. Design and development of distributed systems will be a normal skill set of development engineers in the industry.
- A standard browser protocol capability will be agreed, forced by the market so as to allow user mobility and services of higher perceived customer value (like interactivity or personalised accounts accessible from anywhere).

The term 'IP everywhere' includes advanced functions enabled by Internetworking protocol (IP) based features under development such as in the areas of mobile-IP, quality of service, session handling (SIP/SAP), IP version 6 and multicast. At present the nature of the capabilities and features is not universal. In 2005 standard solutions for IP technologies covering all of the facilities will be deployed globally.

E-commerce for individual users has many different approaches undergoing trial at present with different players adopting competing approaches.

- In 2005 progress towards the standardisation of customer premise equipment e-commerce technology will have been made. Standards for encryption, watermarking, authentication and signatures will be stable and deployed.

- Smart cards will not have significantly penetrated the market—Reference 19 predicts only a 2.7% share (4.7 million) of the North American automatic teller machine card market in 2002—even though the major credit card companies are active²⁰.

- There are presently a variety of 'card' technologies in use and two main standards for the operating system exist for them (MasterCard™—*Multos* and Visa™—*JavaCard*).

- The cost to merchants of supporting smart card technologies (\$120–250 per terminal) mean that high usage of the terminals is required to justify the investment²¹.

- Regulation (copyright and security) will have allowed free use of e-commerce technology tools internationally.

Access bandwidth technologies, and the associated core network standards that are needed (and implied), will have matured, and regulation and standards will have enabled both global interconnect as well as some degree of open interfaces onto network capabilities.

- Access¹⁷ (and network) technologies are still 'low' in bandwidth. ISDN is prevalent within Europe (28 million lines) with xDSL only starting to pick up.

- Access technology solutions in North America will be dominated by xDSL and cable modems. The Asia-Pacific market will be dominated (50%) by radio access lines.

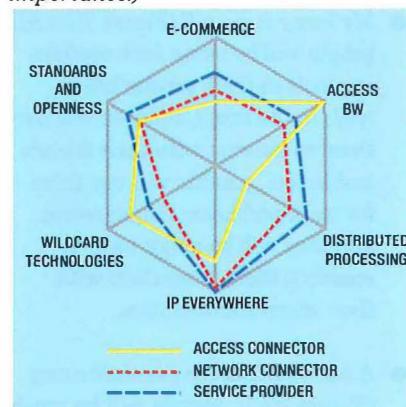
- Within the European core network, asynchronous transfer mode (ATM) technology will be dominant but with a significant and growing use of IP as the networking protocol. In North America, IP and related technologies will be dominant.

- Universal mobile telecommunications services (UMTS) technology will be maturing and deployed on a small scale. With economic factors neutral and societal forces strongly for the paradigm, the slow development will have been due to political regulation (such as spectrum allocation problems).

- Global systems for mobile (GSM) and code division multiple access (CDMA) systems proliferate along with mobile-IP infrastructure. Predictions¹⁷ suggest that mobile infrastructure in the west will be dominated by digital GSM (30%) and CDMA (10%).

Figure 6 describes some impressions of how these technology factors affect the different roles in the industry model (Figure 1). All roles are strongly affected by technology but in different ways. The network connector and service provider are both strongly influenced by the types of protocol that will be used by end-customers (IP everywhere) and by technology that provides an independence from location (distributed processing). The access connector and service provider need to consider access bandwidth developments and the opportunities and threats from a

Figure 6—The impact of technology factors on the telecommunications industry supply roles. (Note: Strong relevance is represented by intersection of a line near the edge of the matrix; the centre indicates no importance.)






Figure 7

variety of 'wildcard' technologies from other areas.

Key Points

So what will the world be like in the early part of the next century? The following five key points are conclusions inferred from the information presented. The starting point was the societal view.

- *All companies rely on computer peripherals (Figure 7(a))* The year 2005 will see the majority of the working population using computing peripherals (CPs) as the tool of work. All companies will use them extensively, the role of the telephone within businesses has reduced significantly although voice communications is still strong.
- *My home is my CP (Figure 7(b))* All people will be using such computing tools in their private life as well. In the home, children will use them to interact with their friends and school, parents will use them for work and home management, and they will be major tools in enabling their interaction with their many communities.
- *A mobility tsunami is on the way (Figure 7(c))* People will be much

more mobile and use voice communications during this mobility. Mobile data services, beyond the short message services currently in GSM, will not be widespread. Deployment of UMTS infrastructure will not be widespread and these services will primarily be taken up by global companies for their highly mobile workforce. However, the mass-market take up of such services is about to start, the value is clear, the technology developing and the market for these services will flourish quickly.

- *An electronic payment tsunami is about to hit (Figure 7(d))* Electronic commerce between businesses will be ubiquitous to support a dynamic and changing service economy. The lack of universal standards has not yet enabled cost-effective solutions for the mass market but things are aligned ready for a mass take up and use in the near future.
- *SMEs are the bees knees Figure 7(e))* Small and medium enterprises are the major driving force to the service economy. They are specialised, highly dynamic companies responding to the demand of end-customers and the global

corporates they supply. The days of the 'one-man-band' are numbered.

So What Might Happen to Operational Support?

The purpose of the vision is to understand possible impacts on the operational support systems within BT. The following are some possible implications that result from the world described above.

- Only the service provider maintains end-customer data.
- Customers themselves are in control, owning their own data and deciding who they allow to use it.
- Players adopting the access connector and network connector roles have to open their business processes so that their customers (players taking on the service-provider role) can meet timescales expected by the end-customer market.
- Systems are open to third parties, testing systems as well as billing and charging systems. Any hesitancy because of the lack of trust is countered by the risk of losing business and the protection provided by electronic commerce.

- Separate (niche) companies set-up in business to offer generic support services to the industry.
- Customers take on more and more responsibility for service support as the required processing power becomes cheaper and cheaper.

Further work is needed to consider these initial ideas and other scenarios, and assess what impact they may have.

Conclusions

A world of the future, very different from that of today, has been built from realistic options and/or developing trends. Its relevance and accuracy will be dictated by time. But the implications on service-providing companies like BT need to be considered now, as they are reliant on understanding and meeting the needs of such a future marketplace. Companies that have to make long term, large investments to build and manage service systems need to think ahead. The ideas presented here suggest a large shift from today's status quo. A significant redistribution of responsibility is possible, implying key changes to where value (and hence business opportunity) may reside. The thinking has started: the results will follow.

References

- 1 European Information Infrastructure. ETSI Sixth Strategic Review Committee, May 1995.
- 2 Colchester, Nico. The World We're in. *Prospect*, Nov. 1996. <http://www.prospect-magazine.co.uk/>.
- 3 Commoditising Telecoms. Pork bellies calling. *Economist*, 30 Aug. 1997 or at <http://www.band-x.com>
- 4 The end of privatisation. *Economist*, 13-18 June 1998.
- 5 Henley Centre/BT Telecare Futures 1996.
- 6 The Henley Centre, Media Futures 1997, pp. 226-236.
- 7 Henley Centre Planning for Social Change 1998, chapter 19 section 5.

- 8 Henley Centre Planning for Social Change 1998, chapter 19 section 4.
- 9 Henley Centre, Media Futures 97, Internet Survey UK and US p. 188/190.
- 10 Henley Centre Media Futures 97, pp. 226-236.
- 11 Henley Teleculture Futures 1997, Small Firms Futures 1996.
- 12 Teleculture Futures 1997.
- 13 Paperless? Forget it! *Economic Times*, 15 June 1998.
- 14 FORRESTER. *Money & Technology (M&T) Strategy*, 3(12).
- 15 DataMonitor. Business—business e-commerce. Exec. summary and Ch. 5, Aug. 1997.
- 16 KeyNote. *Business Travel*, Aug. 1997 or ISBN 1-85765-7187.
- 17 OVUM. Key Telecoms Markets section A6 1998
- 18 *Information Week*, 17 Nov. 1997, p. 304.
- 19 FORRESTER. *M & T Strategy*, 3(4).
- 20 FORRESTER. *M & T Strategy*, 2(18 & 10).
- 21 FORRESTER. *M & T Strategy*, 3(7).

Biographies



Chris Russell
Networks and
Information Services,
BT UK

Chris Russell is a senior member of the OSS research team within Applied Research and Technology (ART), BT. His current research is concerned with enterprise modelling and how enterprise analysis can be used to inform the architecture and design of IT systems. He gained a B.Sc.(Hons.) in Physics and Computing from the University of Hull, and joined BT in 1987. He joined ART in 1997 as part of the team that produced the experimental service platform (ESP). Before joining ART he worked as a systems engineer within Systems Engineering. He gained a Diploma in Management Studies (DMS) in 1993 and is a member of the Institute of Management.



Maff Holladay
Networks and
Information Services,
BT UK

Maff Holladay joined the Post Office in 1979 and has held a variety of posts in the PO and BT. These have covered migration from paper to computer records and from analogue to digital exchanges. Prior to joining Applied Research and Technology (ART) Maff was managing the development of interfaces to the BT charging data. His current work within ART is researching the changing requirements of BT's Operational Support Systems (OSS) as we enter the next millennium.



Pete Barnsley
Networks and
Information Services,
BT UK

Pete Barnsley joined BT in 1987 after completing a M.Sc. in Optoelectronics and Lasers at St Andrews University. He spent the next eight years investigating developments in high-capacity optical transmission systems. During this time he published more than 40 research papers and generated six granted patents and completed his Ph.D. Since 1996 he has been involved with strategic study projects looking at the simplification of network structures and the development of the global information infrastructure. In 1997, he joined a team researching flexible service platforms for fast service creation, based on ideas from TINA-C (Telecommunications Information Network Architecture Consortium). In early 1998, the team conducted the world's first operational TINA-like public Internet-based platform (ESP). His present research interests are the study of systems that support services in a dynamic competitive and regulated market, and the tools and methods to analyse problems associated with building them.

Future Switched Network Strategy—Profiting from an Uncertain Future

Traditional telecommunication network switching technologies and architectures are increasingly being challenged by edge-of-network solutions based on technology drawn from the IT industry. The increase in data traffic and the rise in importance of 'voice over IP' technology are significant drivers of change. This article reviews the business drivers and deduces the key characteristics that any future nodal switching function must support. Particular initiatives include collaboration undertaken to accelerate the introduction of appropriate standards and Next Generation Switch, an initiative dedicated to introducing new technology in BT's core PSTN network.

Introduction

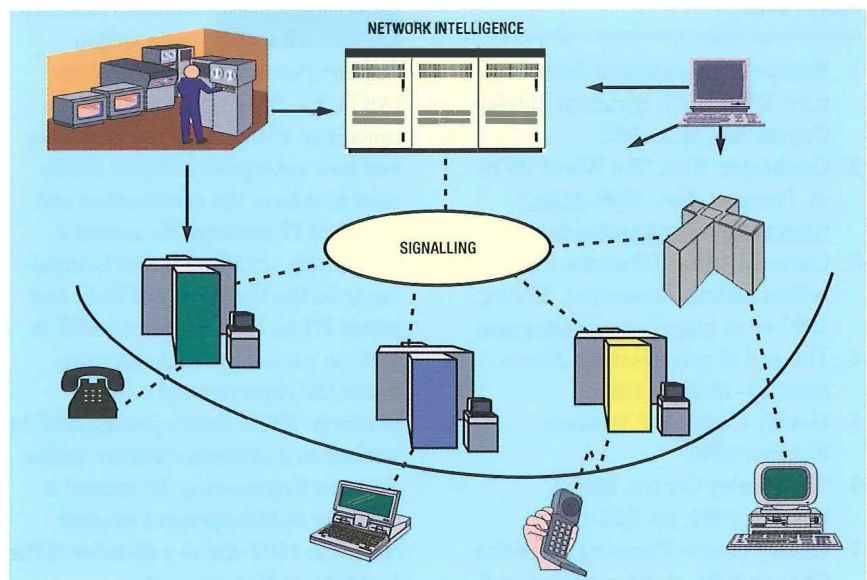
BT's future switched network (FSN) strategy is driven by a complex mix of market, financial, operational and technology drivers. With increased competition and the growth of substitute products, it has long been apparent that BT could not continue indefinitely to upgrade an increasingly complex mix of network platforms by reactively purchasing proprietary vertical 'stove pipe' functionality to meet new and future requirements as they arise. As well as increasing complexity to an unmanageable extent, such an approach would not be cost-effective as it would inevitably lead to repeated development of the same functionality. Like many other operators therefore, BT has embraced the concept of the

intelligent network (IN) as a way of reusing expensively developed service functionality and enabling it to be accessed by a variety of underlying transport platforms.

IN is a key aspect of BT's vision for the cost-effective delivery of the high functionality multi-service networks of the future, as can be seen depicted in Figure 1. However, such a vision is predicated on a network-centric view of the world in which highly customisable service functionality is delivered from the core of the network via service logic deployed on high-performance computer platforms which support a variety of underlying switch technologies.

While the introduction of IN capability has been an important step towards improved service

Figure 1—The intelligent network vision



interoperability and the integration of switch platforms, it is now apparent that such an architecture is not the complete answer. Indeed, with the growth of data and the increasing importance of Internet protocol (IP) technology, the very bedrock of the traditional public telecommunication operator (PTO) network vision is now under challenge. However, as well as challenging the accepted wisdom, such technology will also support a much greater range of network solutions than hitherto envisaged.

To compete effectively in future markets, PTOs must pragmatically embrace the full range of opportunities that new technology now offers. For example, voice over IP offers a particularly seductive solution for voice/data integration that is currently the focus of much attention around the world. However, it is far from clear if and when such a solution might become dominant. If there is likely to be an enduring mix of packet- and circuit-switched technologies for long enough, should PTOs press for upgraded nodal switch functionality as an important contribution to sustaining competitive advantage? One specific development to be considered would be the deployment of larger faster more-flexible switches. Another would be the introduction of open application programming interfaces (APIs) and the use of object-oriented techniques drawn from the information technology industry to both provide access to embedded service logic functionality and effectively 'shrink wrap' legacy complexity. Would such an approach be an appropriate way of cost-effectively meeting customer needs in the highly volatile markets of an uncertain future?

In addressing these issues, this article describes key aspects of FSN technology, strategy and architecture, as it will impact on the nodal switching function. To enable better understanding of the context for current initiatives in this area, there is first a review of business drivers. Such a discussion leads naturally to

a deduction of the essential characteristics of the nodal switching function of the future. One possible multi-service, multi-fabric architecture is described together with a discussion of some of the challenges that must be faced by the industry if such a vision is to be realised. To conclude, there is a brief review of some of the practical steps that are being taken to progress this work. These include the key role of standards and Next Generation Switch (NGS), an initiative that will hasten implementation of new technology in BT's core PSTN trunk network

Business Drivers

Market

New services/markets

While the industry now has much experience in the support of advanced voice and data services¹, there will always be future services and applications which are as yet unknown. This will be particularly true as the external service provider market grows and PTOs extend their operations into hitherto uncharted territories. Although the introduction of IN has provided a certain level of flexibility, it is important that the FSN extends such capability to support application-independent flexible functional blocks of service logic which are able to meet all future service requirements.

Narrow/broadband (data, voice and mobile)

The services of tomorrow will be based on multi-service flexible platforms which can cost-effectively meet the requirements of voice, mobility, data and variable bandwidth/broadband services. With inherently volatile forecasts for new services, investment in FSN technology will have greater longevity if it results in a network comprising highly-flexible modularised platforms capable of cost-effective reuse in different configurations to meet market demand as it

arises. Increasingly, IP will set the benchmark against which multi-fabric network solutions are judged.

Time to market

Edge-of-network service solutions bring to bear the expectations and disciplines of fast-moving consumer markets on core network design and development. Typically in such markets, the number of players and the threat of substitute products are high, while the barriers to entry are low. First-mover advantage can be an important aspect of revenue generation and differentiation. If network intelligence is effectively to compete in this new paradigm, services must be capable of being developed, deployed and modified much more rapidly than at present. Inherently, there will be significant impact on the support- system surround as well as network functionality.

Seamless service

Services introduced on disparate platforms pose problems of interaction and may present a different, implementation-dependent appearance to the customer. The FSN must resolve these issues and provide seamless service to customers irrespective of the complexity of the underlying functionality.

Customer mobility

In the future, the requirements for increased personal mobility and terminal mobility are likely to be significant and widespread. The support of advanced customisable profile-driven find-me/follow-me type services and the introduction of the Universal Mobile Telecommunications System (UMTS) functionality will enable customers to be contacted via any type of appropriate terminal, at any location, at any time, according to customer preference².

Operational

Price performance

PSTN switches currently in the networks of all major PTOs are

Figure 2—Advanced voice platforms in the UK

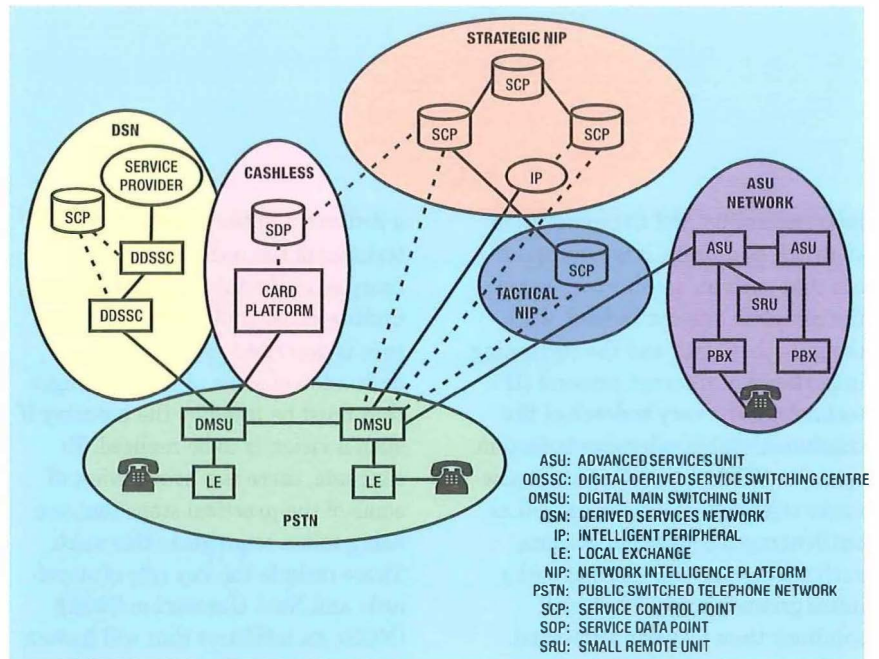
based inherently on the technology paradigm that emerged in the early 1980s with the introduction of digital switching, stored-program control and common-channel signalling. Invariably, the nodal switching function is based on closed proprietary architectures which are inflexible and expensive to modify. Subsequent significant advances in technology are now such that, worldwide, suppliers are beginning to address next-generation solutions. For example, the advances in open processor technology that are taken for granted in the information technology industry act as a salutary benchmark for the sort price/performance improvements that we should be seeing in telecommunication solutions if today's network demands are to be met cost-effectively.

Unit costs

With the twin pressures of competition and substitute products, it is imperative that unit costs continue to reduce to reflect global trends. Such cost reduction will include the reduction in the number of 'board' types (with a consequent reduction in network spares holdings) and the introduction of devices that consume less power and have reduced footprint. Reduced maintenance, with the achievement of zero downtime through increased resilience and self-healing, is also an important consideration.

Build variety

Figure 2 illustrates the variety of switch platforms in BT's PSTN. Such a variety of switches and builds is obviously a barrier to achieving economies of scope and scale. While it would be easy with the benefit of hindsight to suggest that the proliferation of switch types is undesirable, the fact is that in all major PTOs, such investment decisions are made for the right reasons at a particular time. In all cases, PTOs strive to meet market need with the most-appropriate cost-effective technology. However, given the



present growth of network complexity and variety, there is evidence to suggest that the introduction of appropriate open interfaces to allow reuse of existing assets might be a better and more enduring approach than the pursuit of strategies simply based on *post hoc* rationalisation. Indeed, with the continued flexing of organisational boundaries with mergers and acquisitions, it might be argued that such an approach is essential to utilise legacy investment effectively.

Change in call mix

In the highly complex and interdependent networks of today, planners have to face up to highly volatile and inconsistent patterns of traffic behaviour. For example, Internet, calls with long holding times contrast sharply with EFTPOS (electronic funds transfer at point of sale) calls involving a larger number of transactions with extremely short holding times. Increased demands by advanced calls for increased nodal processing and the volatile demands of interconnect traffic complete a picture that demonstrates the challenge facing existing technology.

Capacity constraints

With the growth of interconnect and advanced services traffic, there is an imminent need for additional capacity in the BT core narrowband trunk network. However, significant issues with respect to network

topology limit the continued deployment of additional switches. In contrast, significant benefits can be shown to accrue by flattening the existing network structure and reducing its complexity through the use of fewer, larger switches. The network also needs cost-effectively to handle high volumes of data traffic.

Upgrade cycles

Inflexibility in the switch upgrade cycles is a significant barrier to delivering improvements in time to market. Historically, PTOs have been restricted by the supplier's development cycle for major generic upgrades, increasing the time to market and reducing/delaying revenue generation. This situation is worsened in the face of significant build variety when ubiquitous deployment of a new service requires synchronisation of developments on different proprietary platforms. Though IN has done much to improve new service delivery in such circumstances, there are still too many occasions when the delivery of a particular feature requires an update to switch generic software. There is thus a need to be able to modify existing switch builds rapidly to deliver 'fast features' across platforms coherently. The ready availability of such a capability could also be a driver in increasing the serviceability of each generic switch build.

If the nodal switching function is to endure it must be sufficiently flexible to cope with traffic uncertainty while also achieving significant reductions in whole life costs.

Technology and Architecture—The Key Characteristics

If the nodal switching function is to endure, it must satisfy the business drivers described above. In particular, it must be sufficiently flexible to cope with traffic uncertainty while also achieving significant reductions in whole life costs.

Logical architecture

Historically, developers have attempted to modify voice platforms to support data. Now the emphasis is increasingly on the packetisation of voice in order to carry it at marginal cost on high-capacity data networks. Given this background, is it appropriate to postulate a target vision that encompasses narrowband and broadband switch functionality within a single architecture? One such architecture is illustrated in Figure 3. Could such an architecture provide the opportunity for unified call and service control (albeit that dissimilar call types would be catered for by different instances of call control)? Could the introduction of interfaces between major functional blocks provide for the concept of flexibility

and reuse? Could the incorporation of dedicated standard open operations, administration and maintenance (OA&M) interfaces from the outset avoid the unnecessary development of mediation engines to interface proprietary support systems?

The adoption of a layered approach to nodal architecture draws much on successes within the computer industry. While a detailed discussion of specific detail is beyond the scope of this article, it is appropriate to consider its essential characteristics. For convenience, such characteristics have been related back to the business drivers described earlier, with grouping according to the competitive advantage delivered; namely, the delivery of economies of scope and scale; the delivery of flexibility and openness.

Economies of scope and scale

Large switch

In BT's UK PSTN network, local exchanges (LEs) are multi-parented on a fully meshed core network comprising digital main network switching units (DMSUs). While such a network architecture has great

advantages in terms of network resilience and security, there are limits to the cost-effectiveness and practicality of continuing to add to the total number of DMSUs. However, looking to leading global suppliers, it is apparent that technology can now support much larger switches than those at present deployed. In the medium term, it is foreseeable that such technology will support switches which are up to four times as large as existing DMSUs, with support of 8000 2 Mbit/s ports (256 000 × 64 kbit/s channels) being attainable in the foreseeable future. By deployment of this technology, it will be possible not only to support growth, but also to achieve significant node consolidation. This will provide the opportunity to reduce the number of interconnecting routes between DMSUs which, in turn, will significantly reduce the cost of transmission equipment and other associated costs.

High-speed ports

With advanced transmission systems supporting higher-order synchronous digital hierarchy (SDH) rates (for example, 155 Mbit/s: STM-1 and faster) with optical interconnect, technology is now providing the opportunity to reduce transmission termination costs significantly. By integrating such technology directly with the switch, there will be a significant opportunity to reduce costs, with high-speed ports providing direct connection to optical fibre systems for maximum efficiency and reduction in exchange size without intermediate transmission equipment.

Multi-service and multi-fabric

With the introduction of IP, asynchronous transfer mode (ATM) and cell-switched technology, networks will increasingly be required to support a variety of narrowband and broadband services (for example, voice/ISDN, mobile and variable bandwidth data). As well as delivering such services by edge-of-network grooming and dedicated data networks, there is an opportunity to

Figure 3—Future switched network architecture

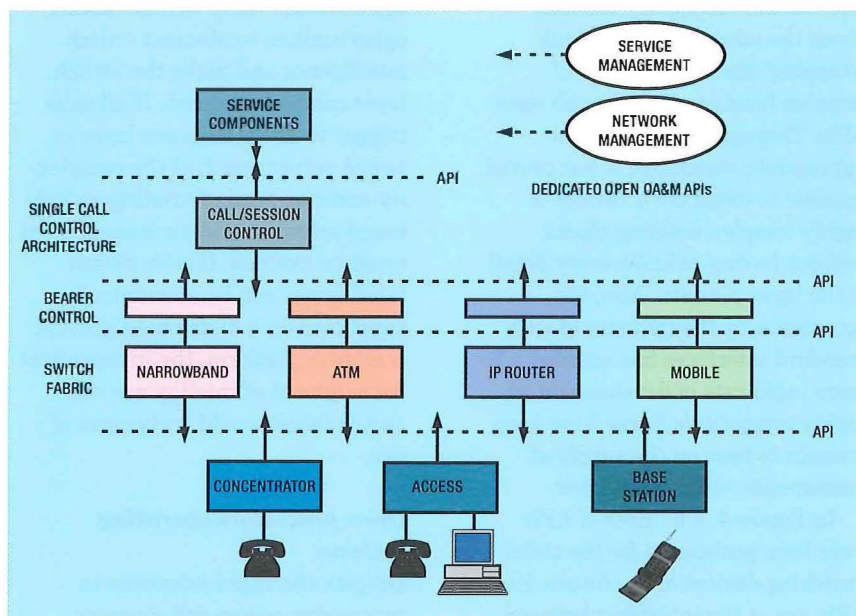


Figure 4—Total service creation

introduce in the core network an appropriate architectural framework which is consistent with a variety of switch fabrics. Such an approach could potentially provide the vehicle which will unify a variety of network environments, while meeting the increasing need for new services and increased bandwidth.

Dial-up broadband data switch

Not all services will be cost effectively handled by a monolithic approach to integrated transmission and switching. Any practical future network will inevitably have to also be sufficiently flexible to support non-integrated components, both legacy and interconnect. For example, there will be an enduring requirement for grooming at an appropriate point in the network (for example, at the remote concentrator unit (RCU)) to provide for the efficient support of non-integrated data traffic. As well as supporting a dial-up data service, there will also be a significant opportunity for 'always on' background IP services, with efficient onward routing to an IP router or to a broadband (for example, ATM) switch.

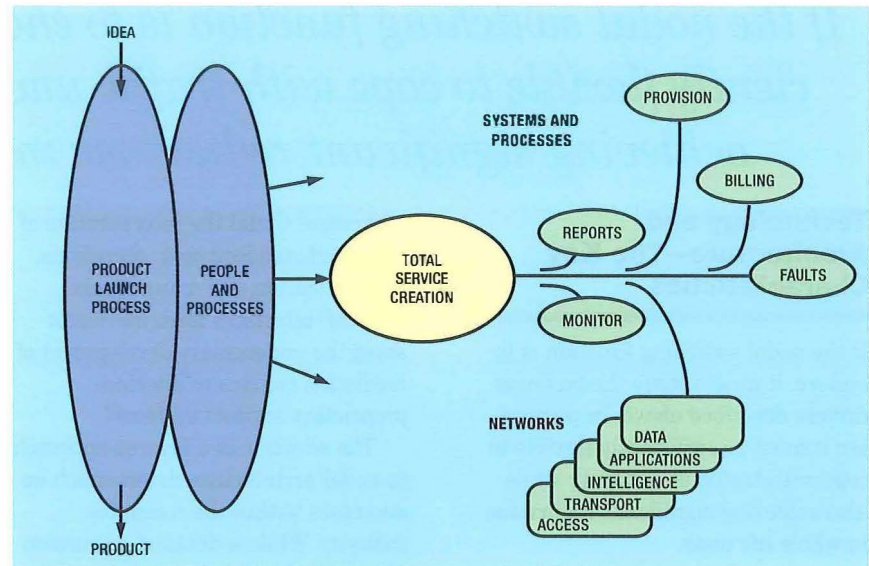
Flexibility and openness

Software reusability

The logical architecture illustrated in Figure 3 could provide the inherent flexibility to allow progression from existing monolithic vertical 'stove pipes' of switch-based functionality to a more cost-effective open architecture based on reusable modular components. Through extensive use of modular software building blocks, it should be possible to minimise any duplication of development effort within the control software stack. Ultimately, this should lead to similar benefits to those encountered in the more general software 'plug and play' environment that has now become common in the office and home.

Fast time to market

Despite the widespread introduction of IN, there are still further opportu-



nities to improve time to market. All too often, a particularly attractive new service is dependent on an IN trigger or event detection point that requires development of switch-based software. Currently, the time required to develop switch-based services can be in excess of 12 months. Such bottlenecks demand fast-track service creation environments for residual switch based logic and the OA&M surround. As is shown in Figure 4, being able to support a feature on a network platform is simply not enough. The adoption of appropriate standard OA&M interfaces from the outset together with the concepts of managed objects can do much here to give the performance improvements required.

Access to internal APIs

Much has been learnt from the information technology industry about the advantages of 'shrink wrapping' successive layers of complex functionality through open APIs. Through the adoption of appropriate standards, it has proved possible to integrate a variety of highly complex building blocks without having to know every detail of the layer beneath. Also, the acceptance by the industry of such standard interfaces has enabled a far more rapid rate of development as highly competitive forces have been brought to bear on the supply of components within each layer.

In Figure 4, a number of APIs have been postulated for the nodal switching devices of the future. Key APIs are at the interfaces between

service control, call control, bearer control and resource control. While the service control/call control layer is a familiar and accepted concept as typified by the INAP (IN application protocol), the lower layers are perhaps less familiar. As a longer-term vision, comparable advantages to those experienced in the computer industry could be expected if the industry can develop appropriate open and non-proprietary standards. Ultimately, it might be envisaged that 'best of breed' commoditised switch components can be purchased independently of one another from suppliers who conform to such APIs.

Separate switch intelligence from switch fabric

While IN has done much to separate service control from call control, it could be argued that with technology advance, there will be further opportunities to abstract switch intelligence and make the switch layer completely dumb. If all calls trigger to an intelligence layer or *switch server*, much of the complexity and overhead of existing switch-based software and its management could be avoided. If such switch intelligence can be economically supported on a distributed general computer platform, the independent development of intelligence and switch fabric would truly come of age.

Open processors/operating systems

Despite the rapid advances in processing power and memory

In future, there must be no discontinuities in the availability of additional processing power, and processing power must more rapidly follow industry trends.

storage capacity seen in the computer industry, telecommunication switching systems tend to be reliant on proprietary dedicated processors which are expensive to upgrade and which lag considerably in terms of performance. Through the adoption of IN standards, the use of adjunct processors has increasingly become accepted for supporting service logic programmes. However, to compete fully with edge-of-network solutions it will be necessary to take advantage of the low costs and rapid advances in general computing technologies and software developments. Although reliability and resilience will continue to be important aspects of the telecommunication processing platform, the price/performance of edge-of-network solutions will set a formidable benchmark. In future, there must be no discontinuities in the availability of additional processing power, and processing power must more rapidly follow industry trends.

'Plug and play' hardware

Increasingly, technology makes it possible for self-configuration of equipment shelves. This is important in reducing the planning overhead and makes for simplified maintenance. Ideally, hardware 'plug and play' would support the capability to put any card in any slot and in any supplier's switch. Initially, however, the aim will be to reduce the number of board types (and hence spares) per supplier and to press to implement open interfaces to make it possible to use other supplier's cards where appropriate; for example, echo-cancellors. Such 'plug and play' concepts can operate at different levels. In the United States, standards are in existence (for example, BellCore GR/TR303) which even enable concentrators from different suppliers to be mixed and matched. In Japan, the specification of a common line-shelf interface has enabled competitive purchase at the line-card level.

The Challenges Facing the Industry

Development costs

In today's competitive environment, any cost premium to achieve an acceptable paradigm shift must be ultimately shared by the whole industry. While certain key technology and architectural characteristics of the FSN will be available at an early date, it is inevitable that others will be developed only with persistence and time. While in the long term there will be significant cost savings, at the outset there will inevitably be a requirement for some strategic investment.

IP telephony bypass

Much is made of this threat but quality-of-service issues and the investment in legacy platforms suggest long-term coexistence of circuit- and packet-switched technologies. In a future where there is a belief that circuit-switched voice will have to compete with voice that is carried over IP, the support of a multi-service multi-fabric nodal switch is important in facing up to the inherent pluralism and uncertainty of the future.

ATM versus legacy investment

For narrowband switches, much investment has already been made in the existing switch-based call-processing software and it will be increasingly difficult to justify the large development costs involved in simply re-engineering existing solutions. However, with markets increasingly turning to broadband multi-service platforms, there does appear to be an opportunity in focusing effort on the development of 'fabric independent' call-processing software for these new areas. ATM switches lend themselves more readily to such an approach as they currently have little, if any, embedded call-control switch-based software. Despite the difficulties of

legacy, it could also be argued that the lower APIs are necessary to 'encapsulate' existing complexity as well as providing the ability to take advantage of the new switch fabrics as they become available.

Open API/integration and test

Service integration and testing of a multi-sourced FSN based on component building blocks must be cost-effective and realistically achievable. Clearly, this is wholly dependent on the APIs being open and unambiguously defined so that, with the necessary tools, it will be possible to carry out appropriate integration and test activities across a range of interfaces.

There are concerns that, however well defined the APIs may be, should a change be required within underlying software, there will still be some aspects of time to market that are implementation dependent.

Performance of distributed real-time operating systems

Ultimately, it should be possible to dissociate call control fully from the switch matrix and place it on a general-purpose computer platform implemented within a distributed computing environment. However, appropriate industry standards (for example, CORBA) would still seem to be ahead of the necessary technology developments. Indeed, according to some industry analysts, scaleable solutions with the required real time performance levels will not be available for another four-five years.

Next Steps

Standards

The development of Next Generation Switching products is not something that can be achieved without collaboration within the industry to define and support appropriate standards. While an individual PTO or supplier might seek to gain advantage by bespoke developments for a particular

narrow market opportunity, market economics dictate that it will be necessary to achieve economies of scope and scale to recover the necessary development expenditure entailed in a true Next Generation product.

IN CS4 contribution

Already BT is supporting collaborative efforts to introduce some of the concepts discussed in this article into appropriate international standards fora. For example, at the service control layer, BT has already submitted a contribution exploring the use of APIs³ to the debate on development of the intelligent network application protocol (INAP-CS4).

At the core of this submission is the assertion that in the evolving world of telecommunications, there is a strong need to keep up with new technologies and network concepts, particularly in the information technology area. A common need across all such architectures is the reuse of service scripts and the ability to send and receive information packaged in a recognisable form. Standardised and proprietary protocols work alongside each other and often a need arises to map between one protocol and another at gateways or interworking mediums.

In particular, the need to describe services using a common format is explored. In conclusion, given that services in the IT domain are utilising techniques associated with APIs, it is recommended that international standards, particularly those associated with the intelligent network, would benefit from this work.

The Next Generation Switch initiative

Early in 1998, BT invited tender for a Next Generation Switch (NGS) product suitable for deployment in its core trunk network. Specifically, suppliers were asked to meet the following objectives:

- reduce total cost of ownership,

- radically improve time to market,
- minimise dependency on individual suppliers,
- support a variety of traffic types and mixes,
- cost-effectively deal with volatile traffic volumes,
- maximise exploitation of existing investment, and
- maximise future flexibility.

Inevitably, opportunities to implement major steps towards any strategic vision arise at times of significant environmental change as well as in reaction to immediate commercial and operational need. Although meeting an immediate need, BT's NGS programme is also an initiative which aims to take a significant step towards delivering the longer-term vision taking account of technology dislocations currently being experienced by the whole industry. Discussions held with the global industry to implement this initial step towards the achievement of FSN strategy have been particularly productive.

Other collaborative opportunities

In addressing some of the key challenges at present facing the industry, there will be many opportunities to build on the ideas outlined in this article. In particular, there will continue to be opportunity for collaboration in accelerating the move to next generation switch components as the industry addresses:

- the convergence and interworking of narrowband /broadband networks;
- the future of edge-of-network devices (which will include local exchanges, access nodes for data and mobile access (including intelligent terminals); and

- the opening up of monolithic proprietary switch functionality through the introduction of appropriate APIs.

Conclusions

For a company such as BT, any strategy for the network of the future must incorporate an architecture sufficiently flexible to ensure the delivery of the whole range of narrowband and broadband services. Much is written of the opportunities presented by edge-of-network IP based technology and the relative merits of packet- and circuit-switched solutions. However, the global investment in switch-based technology is such that any network vision of the future must cost-effectively encompass a variety of technology, both legacy and future. In particular, while the current trend towards edge-of-network solutions is important, ultimately it will be the solutions which have sufficient flexibility to deal with the continuing shift in markets and technology which will endure.

To ensure that the switching function within the network of the future has the price performance characteristics necessary to retain competitive advantage, it is necessary to adopt leading-edge technology and ensure appropriate economies of scope and scale. In the new world order, given the longer investment lead times of the global telecommunications network, the adoption of standard interfaces and open APIs is seen as an essential step to coping with the uncertainties of future market volatility.

This article has provided an introduction to this topic. In BT's recent competitive procurement of core-network technology in the trunk network, a satisfactory evolution path to meet the key characteristics described in this article has been an important consideration. A future article is planned to cover the opportunities that the introduction of such technology provides for significant network restructuring.

Acknowledgements

The author would like in particular to acknowledge the role of Brian Down and Mel Bale in the early phases of this work. Andy Valdar is thanked for his support and encouragement with this article.

References

- 1 FELBECKER-JANHO, HEIKE. New IN-Services in Liberalised Telecommunications Markets: Potential and Technical Trends. *Br. Telecommun. Eng.*, Special August 1998 issue, **17**, p. 76.
- 2 BUTTERTY, STEVE; MERRETT, BOB; POTE, ADRIAN; and SAGO, ANDY. Local Area Mobility. *Br. Telecommun. Eng.*, Jan. 1998, **16**, p. 290.
- 3 STRETCH, RICHARD. Exploring the Use of APIs in IN CS4, ITU—Telecommunications Standards Sector, Study Group 11, Geneva, 4–22 May 1998.

Biography



Roger Ward
Networks and
Information Services,
BT UK

Roger Ward heads up the Future Switched Network Strategy group within BT's Technology Directorate. He joined Post Office Telecommunications in 1974 after graduating with an honours degree in Electrical Engineering from Cambridge University. After an initial assignment in International Services in London, he returned to study, and in 1978, he was awarded an M.Sc. in Telecommunications Systems from Essex University. In 1980, he joined the System X Joint System Team where, working alongside key players from the switching industry, he had design authority responsibility for the System X digital subscriber switching subsystem (DSSS). Later responsibilities included subscriber line interface aspects of the local exchange competitive procurement programme (System Y), the initial deployment of Centrex in the UK and the introduction of intelligent network databases into the derived services network (DDSN). In 1988, he was selected by BT to participate in Warwick Business School's Integrated MBA programme. He is a Chartered Engineer and a Member of the Institute of Electrical Engineers.

World-Class Network Power— Could That Ever Mean Us?

Part of BT's vision is to be deemed world-class by its customers. However, defining your minimum performance level before you can begin to ask your customers 'How are we doing?' is not always obvious. This is perhaps particularly so when your customers are mainly internal and perhaps unaware of your contribution to the value-chain. This article describes, in a practical way, how BT's Network Power Operations Unit in London set about starting down the world-class road by defining and then following an initial route-map.

Introduction

In 1996, the then BT Director of Networks, John Davies, launched a quality initiative in the access network called *The Battleplan*. This initiative was aimed at reducing the fault rate and improving customer service. The Battleplan was based on seven critical success factors, namely:

- clear strategy;
- responsive and reliable systems;
- production management;
- people capability;
- productivity and cost reduction;
- reliability of the network and service delivery; and
- leadership.

The Battleplan, though predominantly access oriented, was equally applicable to BT's core power activities. Although the engineering issues were different the aims were the same. It was also about this time that the vision of becoming world-class was increasingly being communicated around the business. It was decided that the contribution needed to be as robust as any of the other links in BT's service delivery chain and that the Battleplan philosophy was the way to start.

This provided a vehicle for the aspirations but it still required a destination, way points and a route-map. At this point a conscious decision was made to link all of the unit's activities to the Battleplan. If

an activity did not directly contribute to one of the success factors, it was no longer carried out. At that time the unit's responsibilities comprised: operations, mechanical aids, planning, construction, clerk of works, safety, statutory compliance and related accountabilities.

Having a vision is all very fine but without solid, pragmatic and measurable actions it remains a thing without substance. Using the seven critical success factors, ten measures or goals were devised. Unlike the vision (to be world-class), it was possible to arrive at these goals, fix the position relative to customers, redefine the vision and set off again. In effect this is a continuous process, but it is one which can be packaged into bite-sized chunks and shared with the entire team of 600 people and which, it was believed, would still produce best practice whatever happened.

Therefore the unit's management objectives for the 40 or so managers were linked directly to the Battleplan. In this way everybody and every activity contributed to getting the unit to where it was considered it needed to be in terms of the performance giving the clear strategy required.

Ten measurable goals were defined constituting what are believed to be the entry conditions to being considered near to achieving world-class status:

- To have zero, avoidable, power-related major service failures on the network for which the unit is accountable.
This means that the unit causes no network failures through acts

or omissions, whether through poor planning, construction quality, operation or maintenance intervention. Action by others or force majeure still pose risks but they will be acceptably small from a business perspective.

- To have no call-out attendance for failures which are systematic. *This refers to not being caught in a self-fulfilling prophecy. The unit should never get to a point of spotting potential failures through detecting operating anomalies (for example, wear-out and end-of-life failure etc.), and then doing nothing proactive to prevent the failure but trying to fix it reactively when it does eventually fail in service.*
- To have no failures in capacity. *A good capacity management system under proper change-control is an essential activity. It forms the basis for replenishment and replacement. It is the means of being able to respond quickly to customer demands or being able to suggest alternatives.*
- Limit the production variance on standard, identifiable tasks to 10% or less. *Limiting the production variance enables consistent quality and service delivery. Having a slim distribution really does promote the quality ethic of right first time, every time. The people, the process and the product can be relied upon.*
- Where there is an identifiable correspondence with the next best commercial alternative to the unit's services it equals or betters the industry benchmark. *This is at the heart of any bid to be considered world-class. It also underpins the unit's survival in a competitive environment. All of the ingredients of time, cost and quality are contained in this goal. If Network Power does not provide*

the service then somebody else has to. By definition if this cannot be achieved then the unit cannot be world-class. The pivotal words are 'identifiable correspondence'. This does not mean 'cheaper', or 'riskier' or 'expedient'.

- To identify each year selected systems which have a significant cost of failure and progressively eliminate that cost by 10% each year thereafter, noting that failure rates in world-class manufacturing are a lot tighter than this.
- To always have sufficient numbers of skilled people to meet a short-term shift of two standard deviations (2 SD) in the fault-rate or the planned maintenance workflow. *The work of the unit is specialist covering mechanical, electrical and electronic skills. Here we are seeking to identify the level and mix of multi-skilling balanced against the need to retain and practice the specialist skills (for example, high-voltage, refrigeration and control systems.). The busiest time is the summer for maintaining the cooling systems. Refrigeration skills are a scarce resource, demand for which peaks for just two months.*
- To have a zero sick-leave accident incident rate. *A safe environment is a productive, effective and efficient one.*
- To achieve a minimum CARE line-index of 90 and continuously improve and maintain the gains each year. *CARE is BT's employee attitude survey about the working environment and the way BT people are treated, developed and empowered. The line-index is a figure of merit scored out of 100. It describes how people feel about their managers and the way they are treated as people. It is about mutual respect*

for each other's skills. Customers notice attitude and motivation. The work quality reflects it.

- To operate our systems at an availability of at least 0.9999 *The raison d'être of the job is to provide volts, amps and cooling for the network 24 hours per day 365 days per year. Individual equipment faults should not cause system failures. This level of availability equates to around 50 minutes of downtime per year for each of the system types operated. This places it within the battery reserve for AC failures and within the majority of room thermal capacities in any temperature rise to 40 degrees Celsius in the worst single case if all the downtime were concentrated into one event. (Unlikely but it is the worst level of performance that can be tolerated against the criteria set.)*

The above describes the practical targets set nearly three years ago. At that time some of the measures were known—but a significant area of the performance was unmeasured. Similarly auditable processes, due to the decision to change the way of working, did not wholly cover the unit's activities. In the remaining part of this article there is a description of perhaps the two most radical innovations within the power engineering discipline which are being pioneered in collaboration with colleagues in Group Finance. These are value-chain analysis and activity-based costing (ABC).

Value-Chain Analysis

Although value-chain analysis² is not a new concept, it was the first time the technique had been applied to BT's network power operation. To become truly world-class it was necessary to understand fully all elements of the work including the links and value-chains of the suppliers. By reviewing all the links

the performance could be maximised in terms of reducing costs and improving customer service. The overall aim was to ensure that only those activities, which contribute and add value to customer service are carried out.

Each of the main activities associated with the operation, both internally and externally, was reviewed.

Information was taken from three sources:

- **Activities** Main activities were identified, their individual cost determined, together with the resources they consumed.
- **Activity Analysis** Each process was analysed to determine whether it was a value-add or a non value-add activity.
- **Cost Drivers** By examining our cost drivers, improvements were made by eliminating non-essential activities. Cost drivers are factors that determine the workload and effort required to perform an activity².

The value-chain was defined in terms of primary and support activities. This is shown in Figure 1.

Provision work entry and programming

Work comes into a single point of entry where a team identifies

whether it is to be completed by in-house construction teams, external contractors or in-house project groups. The single-point-of-entry team programme and plan the work which is to be carried out by the in-house construction groups.

Detailed planning

Detailed plans are provided for all work carried out by internal construction groups. For work that is carried out by contractors, a schedule of requirements is produced.

Installation

In-house construction teams and external contractors carry out installations.

Electrical records

The Electrical Statutory Compliance group maintain all electrical records. They ensure that the records are up to date, accurate and do not compromise BT or statutory requirements.

Handover to operations

Once the installation of power equipment is completed, the plant is formally handed over to the operational groups to maintain and repair.

Routine maintenance and repair

The London operational power groups carry out routine maintenance and repair. The life of some items of plant, such as engine-alternators, can be in excess of

20 years. Ownership costs are significant.

The value chain consists of six primary and seven support activities. Once the main building blocks of the chain were identified they were then broken down further into sub-activities. Having established the value-chain, the significance of each activity and the cost associated with it can be understood.

Activity Based Costing

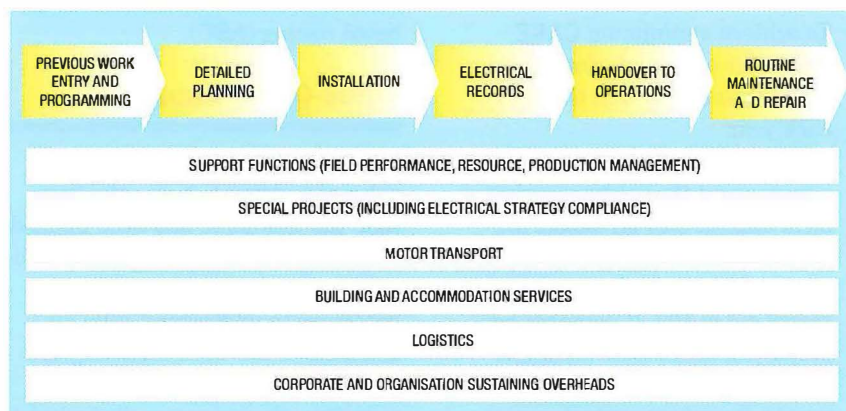
The technique used to analyse the elements of the value-chain was activity-based costing (ABC). ABC was used because conventional costing systems tend to associate costs with departments and not the individual activities carried out. Conventional costing systems determine their overhead costs from the standard labour rate and materials consumed. ABC was found to be superior because it enabled costs to be accurately mapped against each of the activities undertaken and enabled decisions to be taken relating to the allocation of resources.

As Peter Turney² describes: 'Conventional cost information is like the sea that hides dangerous rocks. On the surface, all appears calm and smooth: no inkling of unprofitable products and customers, and no hint of waste in the operations. And like the unwary mariner, the good ship *World Class* sails on, oblivious to the dangers lurking below.' (See Figure 2.)

By working closely with the BT Group Finance unit, data was taken from the BT general ledger systems and assigned costs to the activities based on their consumption of resources. This was found to provide a new source of management information that could be used to make strategic decisions based on activity consumption and not the published (and variable) internal standard labour rate.

To make appropriate decisions, information was gathered on:

Figure 1 – The value-chain defined in terms of primary and support activities



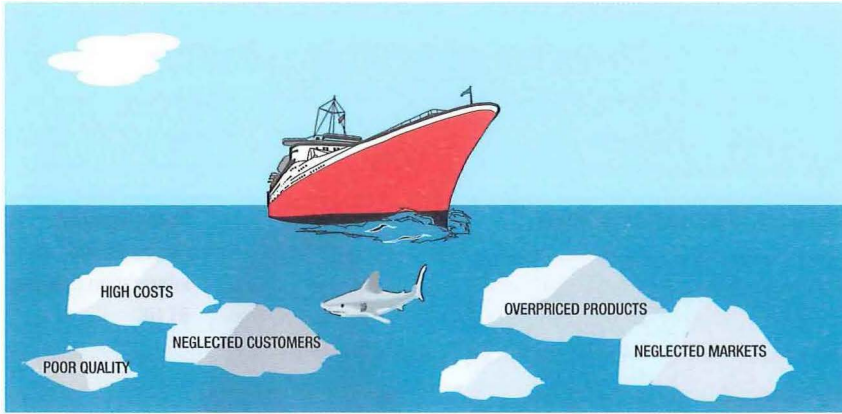


Figure 2—Peter Turney's good ship World Class

- process costs;
- activity costs;
- quality required and quality being delivered;
- quantity;
- customers and
- products

With this knowledge it was possible to identify what needed to be achieved, where value was added, and how cost reductions could be made.

Peter Turney indicates that by adopting ABC it is possible to identify the hidden rocks under the 'good ship *World Class*,² their cost and their importance. The threat that these rocks impose is being gradually eliminated and in turn will improve the overall competitive advantage of the business.

With the support of BT Group Finance ABC has been fully implemented as an ongoing analysis and improvement tool. A specialised software product, HyperABC from Armstrong Laing, is being used to help derive the costs.

Ongoing Work

To become a world-class organisation it is necessary to achieve the highest standards of performance in terms of meeting customers' requirements and be continually looking to improve. Customers must not simply be satisfied, they must be delighted². Customers demand high quality service, flexibility and value for money.

Therefore a production management team has been set up responsi-

ble for continually reviewing all these elements. It was decided these are two of the key areas in which the unit needs to excel. In effect these should form part of the core competence³ to ultimately give a competitive advantage¹.

Core competence is being developed into managing a network of in-house and external suppliers, linked together in close partnership with common goals. By doing this it is expected that costs will be reduced while meeting customers' requirements and progress towards the vision of becoming truly world-class. (See Table 1.)

By using ABC the cost of each activity carried out in Network Power was determined and high-cost low-value-add activities were identified. One particular area, the in-house programming, planning

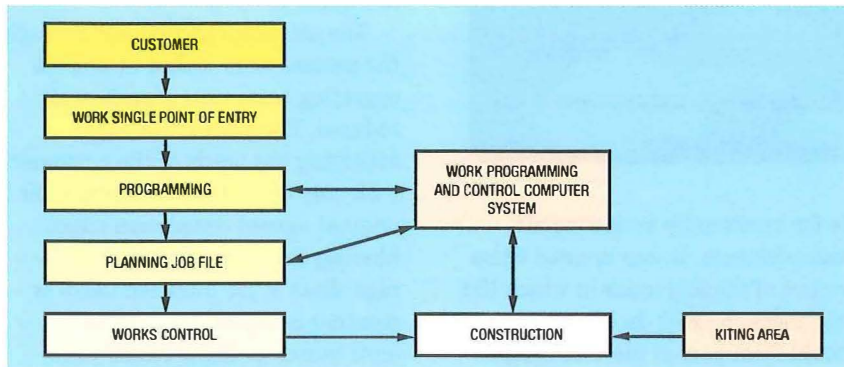
and construction process is shown in Figure 3.

The activities carried out through the process were looked at and its operating costs were significantly reduced. The process was not satisfying the needs of the customer with only 60% of jobs meeting their original agreed completion dates. Missing the date by just one day—regardless of job duration (days or months) is classed as a failure, a *hard measure*. Eight teams were involved in the process, and orders went between the teams without any value-add activity-taking place. With the use of ABC and process modelling in the form of role activity diagrams, the 'hand off' activities were established and a redesign achieved. As part of the route to becoming world class, all the unit's people were encouraged to become involved in making the decision on how to redesign the process. Focus groups were held to listen to the people who actually carry out the work. The result was a resource saving of 10 who were reallocated to areas requiring extra resource and an increase in performance to 95% of

Table 1 Results

Measures	Where we are
Fault rate	Decrease of 17% between 1997 and 1998 (25% since early 1996)
2 × standard deviation for fault-rate	Since January 1998, 2SD is equal to 654 on an average of 1423 faults per month
CARE results	Increase in Line Index from 60 to 85 since 1996
Mean time between faults	Average for the first 10 months of 1998: AC Plant: 0:39 years DC Plant: 1:57 years Fire, Flood, and Gas Detection Plant: 0:22 years UPS Plant: 0:89 years Cooling Plant: 0:90 years Small Power: 0:28 years
Accidents per thousand employees	Sick leave accident incident rate (SLAIR) for year ending March 1998 is equal to 9.12
Faults and maintenance efficiency	Increased from 64% to 94.4% between 1997 and mid-1998
Power related service outages	1 (in first 6 months of 1998)

Figure 3—The in-house programming, planning and construction process



jobs meeting their original completion dates.

Conclusion

World-class is still a vision for the unit. It is now that much closer thanks to the work of the last two years. In that time the Investors in People (IiP) award has been gained and this has further strengthened the resolve. (IiP was particularly beneficial in focusing on the leadership and clear strategy aspects.)

The Battleplan has evolved into the unit's business plan and remains as valid as ever.

The biggest learning points were:

- Trust in, and use the expertise of, the entire team.
- Keep it simple, straightforward and obvious and inject some fun into it.
- Persevere and keep evolving.

References

- 1 PORTER, MICHAEL E. *Competitive Advantage*. Free Press, 1985.
- 2 TURNEY, PETER B. B. *Activity Based Costing, The Performance Breakthrough*.
- 3 PRAHALAD, C. K.; and HAMEL, G. *The Core Competence of the*

Organisation. *Harvard Business Review*, May-June 1990.

- 4 COULSON-THOMAS, COLIN. *Business Process Re-engineering: myth and reality*.

World Class, could it ever mean us? Very probably.

Biographies



Peter Howell
Networks and Information Systems, BT UK

Peter Howell joined BT in 1974 as a graduate and worked for 16 years in power and engineering service system design and maintenance. For the last eight years he has worked in operations for London and International managing power, mechanical-aids, safety and environmental issues for the network. Peter is a chartered engineer and a member of the Institution of Electrical Engineers and the Institute of Building Service Engineers.



Sarah Bennett
Networks and Information Systems, BT UK

Sarah Bennett is a resource and production manager working in the core network. Sarah has been working for BT for two years. She joined BT after graduating from the University of Hull with a B.Sc. honours degree in Engineering Science. She began her career in BT as a field manager working with reactive field engineers responsible for the BT's external network. She is currently involved with a project looking at training and development within BT UK network operations.



Bernie Johnson
Networks and Information Systems, BT UK

Bernie Johnson is a power engineering manager working in the core network division of BT UK. He has been with BT for 25 years and has been involved in all main aspects of telecommunications power engineering. He is currently involved with helping set up the strategy and organisational structure for BT's core network power operations. Bernie has a Diploma in Management Studies (DMS), Master of Business Administration (MBA) and is currently studying for his B.Eng. (honours) degree in Electrical Engineering.

Powering the Internet

Datacommunications Equipment in Telecommunications Facilities: The Need for a DC Powering Option

The Internet industry is seeing an enormous growth in traffic and associated equipment, and it is essential that the supporting power infrastructure is designed with a high level of in-built system reliability.

Introduction

This article was written by a consortium of operators and manufacturers from around the world, who see a strong need for an option to power datacommunications equipment from the telecommunications DC power plant. It was produced at the request of INTELEC (The International Telecommunication Energy Conference), an organisation that is supported by 44 telecommunications operators from 30 countries.

The figures and data used are from a variety of countries each with different standards. Battery reserves, for example, even when coupled with a standby engine, vary from one hour in the UK to three hours and above elsewhere. In the USA, for example, some state legislation requires a minimum of three hours standby. The calculations in many cases are therefore representative of the 'world' and not just the UK, but while the battery times might vary, the underlying argument for using DC is the same.

Background

Telecommunications networks worldwide are evolving rapidly into systems carrying Internet and other data traffic in addition to traditional voice telephony, and operators are seeing a large increase in Internet equipment installed in their networks. These additions consist primarily of datacommunications equipment that typically has not been available with a 48 V DC power supply. Historically, most equipment for Internet service, operator service positions, operations support systems and surveillance

systems operates from AC power and requires an uninterrupted power supply (UPS) to provide service during mains outages.

The new Internet equipment is often collocated in the same room, cabinet or rack as conventional telecommunications equipment. There are differences, however, in terms of battery reserve times, earthing and maintenance. For example:

- Different battery reserve times for system elements, which from operational and maintenance points of view are connected in series, lead to poor availability of service and to maintenance problems. Different system elements along the same chain, on national or global levels, may have differences in battery reserve times. These times can range from four-eight hours or more to as little as five-ten minutes. The lower reserve time sets the limit for the series-connected elements, resulting in unacceptable values for availability of service. The reserve times of the different elements are often unknown, which complicates the analysis of failure situations. Another consequence of multiple battery reserve power supplies is an increase in the number of batteries used, and an increase in the number of maintenance points in the systems. The quality of the batteries in different systems often varies, and maintenance cannot be performed in a systematic manner.
- DC systems are not earthed (grounded) in the same way as AC

systems. The different standards result in practical problems in the implementation of safe and proper earthing of equipment. There are potentially significant safety hazards for both personnel and equipment. There is a greater risk of accidents and fire. Maintenance becomes more difficult and mean-time-to-repair (MTTR) longer, owing to the complexity of the installation and the difficulties in performing preventive maintenance.

- Maintenance costs are higher since two (or more) different power supply systems must be installed and maintained.

Powering Arrangements

Power for traditional telecommunications equipment

Power systems for a traditional telecommunications installation around the world vary but can be characterised as shown in Table 1.

All types of traditional telecommunications installations require highly reliable power systems. Internationally, the typical requirements imposed on the power supply system and the portion of unavailability that may be allocated to the power supply system of a telecommunications installation is 5×10^{-7} , which is 0.26 minutes or 15 seconds of service disruption per year. Table 1 meets this requirement.

Power for data-communications equipment

Power systems for a telecommunications installation with data-communications equipment is characterised in Table 2, which shows that even when a DC source is available, AC equipment is still generally chosen as first choice.

The installations with a UPS do not meet the typical telecommunications unavailability requirement of 5×10^{-7} .

New installations consist mainly of local area networks (LANs) and wide area networks (WANs) in customers'

Table 1 Traditional Telecommunications Power

Power System	Battery	Reserve Autonomy	Comments
1a. Permanent standby engine and 48 V DC power plant (10–20 kW or more)	Typically 1–3 hours	Design standby engine	
1b. Permanent standby engine and 48 V DC power plant (10–20 kW or more)	Typically 1 hour	Dual standby engine	
2. 48 V DC Power Plant only (5–10 kW)	Typically 4–8 hours	Socket for mobile generator set	Back-up for the climate control system must be considered
3. New site (typically <5 kW)	Up to 8 hours depending on services and location	Socket for mobile generator set	Back-up for the climate control system must be considered

Table 2 Datacom Equipment Power

Power System	Battery Autonomy	Reserve Energy	AC Power for Datacom Equipment
1a. Permanent standby engine and 48 V DC power plant	Typically 1–3 hours	Diesel standby engine	UPS or inverter
1b. Permanent standby engine and 48 V DC power plant	Typically 1 hour	Dual single or standby engine	UPS or inverter
2. 48 V DC power plant only	Typically 4–8 hours	Socket outlet for mobile generator set	Inverter or UPS (10 min)
3. New site	Up to 8 hours depending on services and location	Socket outlet for mobile generator set	Inverter or UPS (10 min)

premises. They are being installed in large quantities to provide connections to the Internet and are often located in industrial, office, commercial and residential buildings, that are owned by the property owners or tenants. The telecommunications operators install the interface equipment in the same rooms as the owners and tenants install the network equipment. The types of equipment that might be installed in such rooms are shown in Figure 1.

The power supply equipment in the room should be able to feed both the data and telecommunications equipment. A common standard for the power supply equipment, including

high availability of power for the Internet connection, is therefore very desirable. It is in rooms like these that the meeting of the telecommunications and datacommunications industries, with their different cultures, becomes clearly visible and tangible.

Voice will become an integral part of the Internet and connected LANs, and users' expectations will be the same as for voice service over the public switched telephone network (PSTN), or even higher. Consequently, the traditional high availability of telephony, including operation during mains outages, is likely to be a requirement for data networks.

Figure 1—Equipment rooms

Problems with Present Powering of Data-communications Equipment

Reserve time

Internationally, telecommunications systems powered by a 48 V DC plant typically have battery reserve times of four–eight hours according to a universal standard, unless a standby engine set is provided. Data-communications equipment powered by UPSs usually have reserve times of five–ten minutes. The differences in reserve times will be a particular problem in installations without standby engines, and will entail high costs for operators if the differences are eliminated by increasing the number of batteries in the UPSs or by installing engine-generator sets.

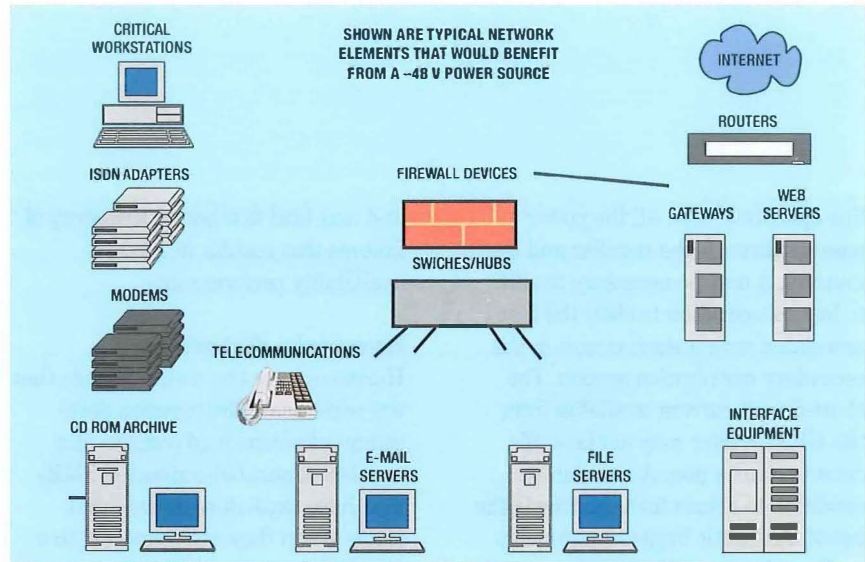
Earthing (grounding)

Standards, rules and regulations for the earthing of telecommunications systems differ among vendors, operators and countries. Different priorities are given to different safety and operational aspects such as protection against injury, lightning strikes and transient overvoltages, and interference in signalling systems.

A major risk of mixing DC- and AC-powered equipment within the same system is that the cross-sections of the distribution conductors are dimensioned in different ways. A short circuit in a DC system can generate large fault currents to earth that may affect the operation of, or even damage, the electronic systems powered from AC feeders. Such failures are often very hard to detect, and their causes and effects are hard both to find and to explain. There is therefore an advantage in powering all the sub-systems from DC.

Distribution

Comparisons of DC and AC systems should consider their ability to protect the connected telecommunications or datacommunications equipment from disturbances on the mains. Examples of such distur-



bances are high switching voltages, transients, lightning strikes, harmonic distortion and interference from other equipment. With the 48 V DC system, the telecommunications and datacommunications equipment is always galvanically isolated from the mains. The system battery also works as a filter against possible transients or harmonics passing through the rectifier. This isolation and filtering yields an almost total elimination of problems with disturbances passing from the mains to the DC distribution system.

Short circuits in a distribution system cause transient voltages with great variations, depending on the resistances and inductances in the distribution system. In a 48 V DC installation, a well-defined and predictable distribution system may be arranged, with control of both resistance and inductance. In such a system, voltage transients caused by short circuits and fuse clearing are under control, and do not spill over to loads served by other feeders. The battery and electrolytic capacitors in secondary distribution panels are

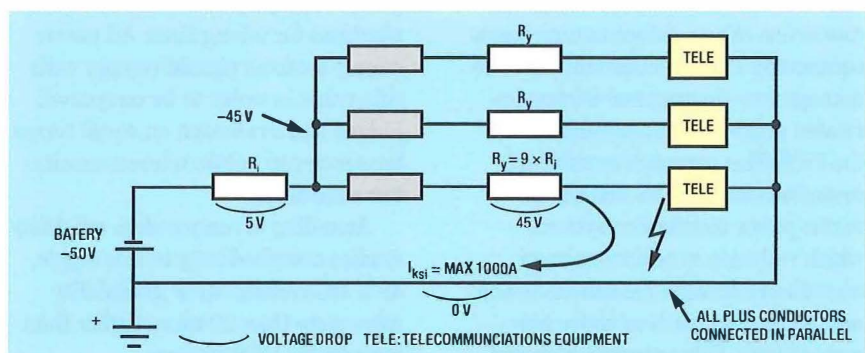
capable of providing enough current to clear the fuses quickly. The fuses and circuit breakers generate an alarm when cleared or tripped, thereby reducing the MTTR.

A unique transient limiting distribution system, shown in Figure 2, has been available in some countries for electronic switching systems for over 25 years.

What happens during the period of short-circuit current is just the ordinary voltage division in the circuit. The voltage drop is distributed between the source resistance R_1 and the resistance R_y of the distribution line with the short circuit. Thus the voltage drops a maximum of 10% or 5 V at the fuse panel. The DC-DC converters and board-mounted power modules in the loads easily handle this change in voltage.

In bypass operation with a UPS, the telecommunications and data-communications equipment is directly exposed to the mains, without any protection against disturbances. Clearly, the bypass line can be fitted with various kinds of protectors and filters, but doing so adds cost. In on-

Figure 2—Transient limiting distribution system



line operation, with all the power passing through the rectifier and the inverter, it may be necessary to shift to bypass operation to clear the fuse associated with a short circuit in the secondary distribution system. The short-circuit current available from the UPS inverter may not be sufficient to clear a fuse. A mandatory condition to obtain fuse clearing is the use of automatic bypass equipment.

For secondary distribution, normal AC distribution practices are followed. With this type of distribution, it is hard to say how other parallel loads will be affected by a short circuit and fuse clearing. There is another problem with distributing AC over long distances in a building; abnormal voltage drops can occur and cause undue stresses on switch-mode rectifiers (SMRs). One way of circumventing this problem is to use decentralised UPSs, one for each item of equipment. This approach, however, can give problems with harmonisation of the battery reserve times and cause higher maintenance costs. There are evidently several factors of uncertainty in the UPS systems architecture that must be taken into account if it is to be used in telecommunications or data-communications installations with high-availability performance requirements.

Another problem in all AC distribution systems is the occurrence of harmonic distortion. Harmonics are a very tangible problem with UPSs since they are not low-impedance power sources. For these problems to be overcome, the UPS and its load must be well known and preferably matched at the planning stage. There is limited room for uncontrolled upgrading and connection of new datacommunications equipment. The action usually taken to manage fuse clearing and harmonics-related problems is to over-dimension the UPS. This approach entails high costs, both for the UPS and for the mains power distribution system, which will have to be dimensioned accordingly. Another measure is to add new equipment, such as active harmonic filters. This approach adds cost,

and may lead to a higher frequency of failures that results in a lower-availability performance.

Harmonic distortion

Harmonics are the result of loads that are non-linear. The common data-communications load consists of a number of parallel-connected SMRs that have capacitors at the input stage. When they are connected to a sinusoidal voltage source, they do not draw a sinusoidal current. The voltage crest factor remains near $\sqrt{2}$, but the current crest factor can exceed 2.5, causing voltage drop in feeders. This voltage drop, however, is typically not a problem in small installations.

In three-phase UPS installations, the most powerful harmonics are the third and fifth. The third harmonic can cause a serious problem in three-phase systems, since these currents do not cancel in the neutral conductor and may therefore overload it. This conductor is often not dimensioned for high currents, since the fundamental currents cancel in the neutral. If it comes to the worst, the overload may cause a fire. However, what usually happens is that circuit breakers trip or fuses clear for no apparent reason, causing unexplained disturbances in operation.

Reliability

As previously stated, the typical annual unavailability allocated to a telecommunications power supply is 15 seconds. This value is generally applicable to all types and sizes of nodes. It is also a value that has been in force particularly in the USA and served as guidance in the design of power supply systems for telecommunications for a long time. All power supply systems should comply with this value in order to be compared, judged and evaluated on equal terms for service in public telecommunications networks.

According to comparative reliability studies described later in this article, 48 V DC systems show availability rates more than 20 times higher than comparable UPS systems.

Batteries of UPS systems generally have four–ten times as many cells connected in series as 48 V DC systems, bringing the fault frequency of the UPS battery strings up to four–ten times that of the batteries in the 48 V DC systems. The result is lower availability and more service-affecting incidents such as fires or post meltdowns. Higher ripple currents and deeper discharges can cause higher maintenance and replacement costs for batteries in UPS installations. Batteries with UPSs typically have a lower life expectancy than batteries used in telecommunications DC power plants and must therefore be replaced more often.

Different power systems with different availability rates affect the overall telecommunications system reliability, and an AC power source will bring the availability of the system to far below the current telecommunications standards.

Safety

A complex mix of power supply systems in telecommunications exchanges is more likely to give rise to faults and accidents. Operating, maintenance and installation organisations face a more complex situation that can be difficult to grasp as a result of this mix. The risk of faults due to human error is clearly increased. Major and minor service disruptions in telecommunications systems occur largely in connection with work by people at the telecommunications exchanges. It is therefore vitally important that the telecommunications systems have simple and uniform designs to reduce the risk of these service disruptions.

UPSs are typically high-impedance sources of AC power and therefore do not have the same capabilities as the AC mains for clearing fuses and suppression of voltage harmonics. This characteristic must be taken into account when dimensioning the power distribution system from a UPS, in order to meet the requirements for coordination of fuses so that only the fuse closest to the fault will clear. There

is the risk that different electricians may install additions or modifications without due regard for the capabilities of a UPS. Overloads in neutral conductors may be hard to detect and, in the extreme, may cause a fire.

Maintenance

The operators of facilities with a mix of 48 V DC power plants and AC UPSs will incur higher costs for the maintenance of two battery systems and two different types of power electronics, double storage of spare parts, and the like. To these costs may be added the need for an increase in competence and/or a new category of maintenance technician, or the costs of double maintenance contracts with two different suppliers, since 48 V DC and UPS installations are usually provided by different suppliers. The two power systems have significantly different mean-time-between-failures (MTBF) rates and may require a different frequency and intensity of maintenance. This situation will complicate the planning and optimisation of the maintenance systems and organisations in large-scale operations.

Advantages of DC Power

Technical simplicity

A 48 V DC power supply system is characterised by its technical simplicity. It consists of a number of paralleled rectifiers that connect to one or more battery strings that are also connected in parallel. In the event of a mains outage or rectifier failure, the load continues to operate from the batteries without switching or interruption. The distribution of power to the loads originates at the point where the battery strings are paralleled, with only fuses or circuit breakers interposed. The electronic equipment in the load has built-in DC-DC converters or board-mounted power modules that interface to the battery.

Modularity and maintenance

It is relatively easy to connect rectifiers and batteries in parallel for reliable load sharing, since the

voltages are low and there is no need to consider phasing. These attributes pave the way for power supply systems with a modular design of the rectifiers, batteries and the conductors that interconnect them. Modular systems are imperative for simple and inexpensive maintenance of installations by persons with limited training on power systems, a need shared by telecommunications operators with other companies and organisations.

Redundancy and battery charging

During operating conditions, the rectifiers provide the current consumed by the load, the float current for the batteries and the additional current for recharging the batteries after a mains outage. If redundant rectifiers are supplied, they fulfil two needs: increased battery recharging after a mains outage and continued operation if one rectifier fails.

Established world norms (standards) and safety

The battery voltage of most telecommunications switching equipment is -48 V. This voltage is a universal standard for telecommunications equipment, and is well defined by both the European Telecommunications Standards Institute (ETSI) and the American National Standards Institute (ANSI).

One of the motives behind the universal 48 V DC standard is that it allows work on a live conductor with minimum risk for personal injury and without special safety measures. Dealing with live circuits is a practical advantage when maintenance engineers work at a distance from the voltage source and therefore cannot disconnect it.

For AC power, there are 14 different voltage and distribution systems defined across the world.

Installed base

In the existing infrastructure of the telecommunications industry, around the world there are millions of highly dependable 48 V DC installations

already in use by the operators. Economically and technically, it would be a mistake to rebuild these installations to equip them with a new system for no-break AC power.

Distribution and isolation

The DC distribution system is well defined and predictable since both resistance and inductance can be specified and controlled. In such a system, voltage transients caused by short circuits and fuses clearing are under control, and do not affect loads served by other feeders. In one system, large capacitors in power distribution frames provide fuse-clearing currents. In the transient-limiting distribution system discussed previously the battery alone provides the current to clear the fuses quickly.


Additionally, the galvanic isolation in the rectifiers and the large battery prevent mains disturbances from reaching the loads.

Case Studies

The following list illustrates the types of problems that some telecommunications operators have encountered with their AC-powered systems.

- An exchange site/data centre suffered an unnecessary loss of service during a mains outage. Due to a minor control problem, the standby engine started but failed to transfer the entire load. The facility was equipped with a card-access security system backed up by a small UPS. By the time an electrician arrived to correct the problem, the small UPS batteries had discharged, and the card-access system had no power to open the doors for entry. Had the card-access system been powered by the 48 V battery there would have been no such problem.
- An exchange site was equipped with a cluster of modems and channel service units used by several banking institutions. The cluster was powered by a small

Figure 3 – Reliability of the mains 240V AC supply in urban and rural areas of the UK



UPS, but small increases in load as the facility grew gradually 'out-grew' the UPS. During a mains outage, the equipment failed after only five minutes of operation.

- An engineer opening a power distribution unit (PDU) to take inventory information knocked a main breaker, disconnecting power to a major subset of an operator service system and causing widespread service delays.
- An engineer was removing a circuit breaker panel faceplate prior to adding a branch circuit to a power distribution cabinet. The faceplate slipped and turned off the main breaker to a 42-pole panel-board, causing the loss of several critical subsystems. Typically, DC secondary distribution systems are designed for orderly growth with very little intrusive work.

- A telecommunications operator's business office, with approximately 250 service operators, had telephones and computer systems that were fed from the AC mains and backed up by a standby engine. Unfortunately, all telephone calls to these operators were disconnected when a mains outage occurred since the engine-generator did not have a 'no-break' power back-up. As a further embarrassment, seemingly no power engineering went into the design for these facilities. As built, a small transformer assembly powered each cluster of 25 telephones. Each transformer was plugged into a multi-outlet strip, daisy-chained from other multi-outlet strips. Ultimately, a socket outlet provided power for the entire operation. This was changed and inverters added to improve system reliability, but during the interim, a major marketing facility was dependent on a wall socket outlet that shared a circuit breaker and branch circuit with other sockets. One cleaner with a floor buffer could easily cripple such a system.

- The failure of a static switch in a 50 kVA UPS took out approximately 100 telephone operator consoles for several hours, resulting in network delays for operator-assisted calls. This event is one of many similar failures over the years. Since most operator service consoles now are based on personal computers (PCs) and the telephony portion is DC-powered, DC-powered PCs would be much more reliable for this application.

The service-affecting incidents described above would be far less likely to occur if the equipment involved was powered from the telecommunications DC power plant.

AC Mains Availability

North America, Europe and Japan

Studies of disturbances on AC mains by IBM¹ and AT&T² in the 1970s provided a view of AC power in computer centres during that decade. A more recent and more comprehensive study by the National Power Laboratory³ (NPL) covered 112 locations over a five-year period. North America has a vast grid of interconnected generation and distribution systems. This grid,

coupled with adequate generating capacity, ensures excellent frequency stability and infrequent outages caused by lack of capacity. The NPL data for total outages (zero volts) has a median value of 1.0 events/month and a mean value of 1.3 events/month. The NPL data for rms voltages $\leq 75\%$ of nominal suggests 2.66 events/year with duration > 5 minutes and 1.52 events/year with duration > 30 minutes. For North America, therefore, it would be reasonable to engineer for one outage/month, with one outage/year that lasts for one hour.

Outage data for AC mains in the UK are shown in Figure 3⁴.

The data for mains availability in the UK shows that one-hour outages can be expected to occur 0.12 times/year in urban areas and 1.11 times/year in rural areas. Values similar to the UK or North America can be expected for most of Western Europe.

The data for commercial AC power outages in Japan are shown in Table 3⁵. This information is for 200 V and 100 V AC mains. It can be seen that the values are not significantly different from those of North America and the UK.

The AC mains in these three parts of the world, and undoubtedly some others, can be characterised as 'good',

Table 3 Commercial AC Power Outages in Japan

Fiscal Year	Number of Outages per User*		Duration of Outages for a User** (minutes)		Mean Duration of Outages (minutes/event)	
	Planned Outages	Unplanned Outages	Planned Outages	Unplanned Outages	Planned Outages	Unplanned Outages
1990	0.09	0.24	11	19	122	79
1991	0.07	0.43	8	158	114	367
1992	0.06	0.13	8	9	133	69
1993	0.06	0.17	6	32	100	188
1994	0.04	0.21	4	38	100	181

* The total numbers of outages divided by the total number of users.
 ** The total duration of outages divided by the total number of users.

and power systems can be engineered as suggested for North America.

Other parts of the world

There are many areas of the world where a rapid build-up of cellular systems is underway and much data on mains power quality has been collected to help determine the causes of problems in telecommunications power systems. Many countries have insufficient generating capacity. The AC mains experiences long outages, frequency instability and large transient disturbances. As a starting point, one can engineer for 30 outages/year with three outages having duration greater than five hours.

Trends

It is probably safe to predict that things will get worse, not better, in those parts of the world where AC mains availability is now 'good'. Increasing demand, concerns over emissions of pollutants and greenhouse gases, finite supplies of fossil fuels and separation of generation, transmission and distribution of electric power pose threats to the AC power systems in countries where high availability is taken for granted. Telecommunications operators will increasingly need to consider co-generation and to insist on lower power consumption in new equipment. The latter has the double

advantage of reducing not only the power consumed in the telecommunications equipment but also the power required for cooling.

Reliability Analyses

System configurations

This section compares the reliability of DC and AC systems, which are considered as the principal options for powering datacommunications

Figure 4—System configuration for DC power supply

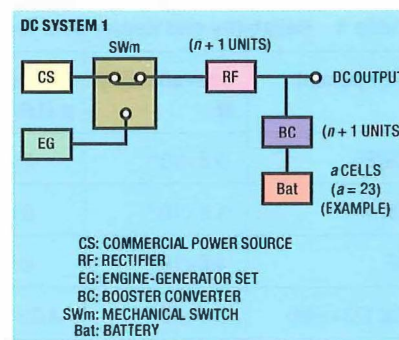
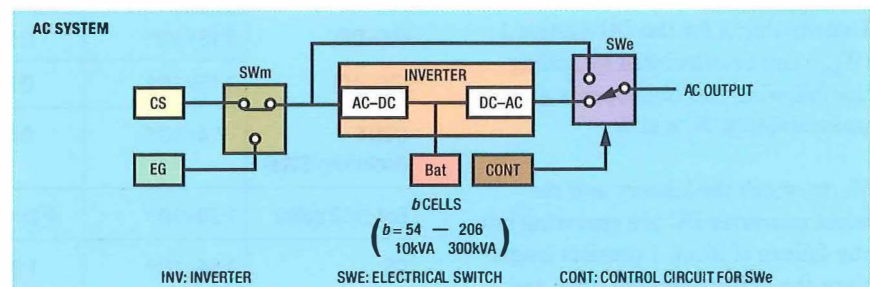


Figure 6—System configuration for AC power supply



equipment. References 6 and 7 provide additional information on reliability and system architectures.


DC power supply systems and an AC power supply system are shown in Figures 4, 5 and 6.

In DC System 1, the batteries and the boost converter back up the rectifier (Figure 4). The battery string typically comprises 23 cells. The reserve time is three hours. The rectifier and the boost converter are composed of (n+1) units, in which the extra unit provides redundancy.

For small plants, a DC system without a standby engine but with eight hours battery reserve is used (Figure 5). This system is called *DC System 2* in this discussion.

The AC system shown in Figure 6 is typically used for powering datacommunications equipment. The AC system omits the boost converters and adds an inverter and an AC bypass circuit. Electrical switch SWe connects the output to the inverter or to the bypass circuit without interruption. The control circuit CONT controls the operation of the switch SWe. The number of cells in the battery string depends on the output power and the supplier's design; it ranges from 54 cells (10 kVA output) to 206 cells (300 kVA output) in NTT, with a maximum reserve time of

Figure 5—System configuration for DC power supply without engine generator set






Figure 7—Reliability model for DC system 1




Figure 8—Reliability model for DC power supply without engine generator set

three hours. In this section, cases for reserve times of three hours and of 10 minutes are calculated. The concepts of unavailability analysis used in this section are applicable to other AC power supply systems, such as those that have redundant sub-systems. The results will, of course, depend on the system configurations.

Models for unavailability analysis

Models for unavailability analysis are shown in Figures 7, 8 and 9.

The model for DC System 1 comprises two blocks, Block 1 and Block 2, which are connected in parallel (Figure 7). The DC System 2 model is composed of only three components (Figure 8). The model for the AC system comprises two series-connected blocks, Block 3 and Block 4 (Figure 9). Block 4, which contains the power path of the switch and the controller for the switch, is the common component in the AC system reliability model, because failures in these components cause system failures.

Unavailability Calculations

DC system 1

Unavailability for the DC System 1 (W_{DC1}) can be calculated by adding the following two components of unavailability, W_1 and W_2 :

W_1 , in which the battery and the boost converter BC are operating but the failure of Block 1 persists longer than the battery reserve time; and

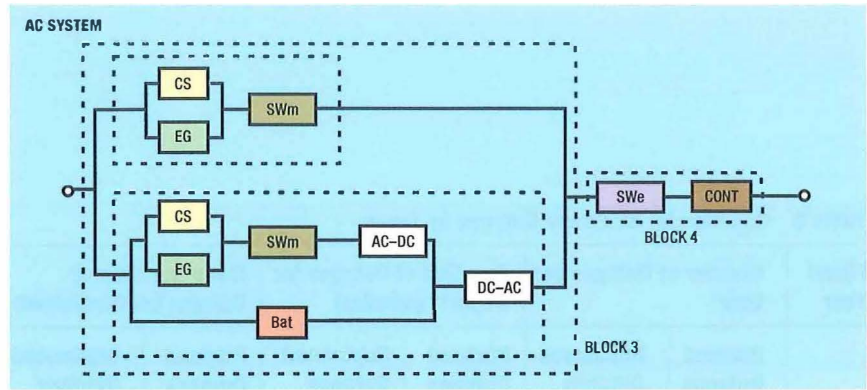


Figure 9—Reliability model for AC power supply

W_2 , in which the battery and the boost converter BC fail during failure of Block 1.

$$W_{DC1} = W_1 + W_2$$

$$W_{DC1} = \{W_{CS/EG} \times \exp(-\mu_{CS/EG} \times T) + W_{SWm} \times \exp(-\mu_{SWm} \times T) + W_{RF} \times \exp(-\mu_{RF} \times T)\} + \{(W_{CS/EG} + W_{SWm} + W_{RF}) \times (W_{Bat} + W_{BC})\}$$

$$W_{DC1} = 3.45 \times 10^{-7} + 6.37 \times 10^{-12}$$

$$W_{DC1} \approx 3.45 \times 10^{-7} \text{ (for } T = 3 \text{ hours)}$$

Reliability values for this calculation are shown in Table 4. From this result, we can understand why battery failures in DC systems have little effect on availability of power because the probability of simultaneous battery and rectifier failures is very small.

DC system 2

The unavailability for DC System 2 is calculated as:

$$W_{DC2} = W_{CS} \times \exp(-\mu_{CS} \times T) + W_{RF} \times \exp(-\mu_{RF} \times T) + (W_{CS} + W_{RF}) \times (W_{Bat})$$

Using values from Figure 8:

$$W_{DC2} = 9 \times 10^{-10} \text{ (for } T = 8 \text{ hours)}$$

AC system

The unavailability for the AC system is calculated as:

$$W_{AC} = (W_{CS/EG} + W_{SWm}) \times \{W_{CS/EG} \times \exp(-\mu_{CS/EG} \times T) + W_{SWm} \times \exp(-\mu_{SWm} \times T) + W_{AC-DC} \times \exp(-\mu_{AC-DC} \times T) + W_{DC-AC}\} + (W_{SWe} + W_{CONT})$$

(1) In the case of $T = 3$ hours (W_{AC1}):

Table 4 Reliability and Repair Data for each Component

Component	Unavailability W	Repair rate μ (1/hour)	Failure rate λ (1/hour)	Reference
CS/EG	6.8×10^{-8}	2.2	1.5×10^{-7}	Field data
SWm	1.5×10^{-6}	0.5	7.5×10^{-7}	Design data
RF	4.6×10^{-8}	0.5	2.3×10^{-8}	Field data ($N+1$ units)
Bat (23 cells)	3.9×10^{-6}	4.2×10^{-2}	1.64×10^{-7}	Field data ⁸
BC	4.6×10^{-8}	0.5	2.3×10^{-8}	Field data ($N+1$ units)
INV	6.3×10^{-5}	0.5	3.13×10^{-5}	Design data
(AC-DC)	3.15×10^{-5}	0.5	1.57×10^{-5}	
(DC-AC)	3.15×10^{-5}	0.5	1.57×10^{-5}	
CONT (including SWe)	7.4×10^{-6}	0.5	3.7×10^{-6}	Design data
Bat (100 cells)	1.70×10^{-5}	4.2×10^{-2}	7.13×10^{-7}	Field data ⁸
CS	1.54×10^{-5}	1.85	2.85×10^{-5}	Field data in 1994

$$W_{AC1} = (1.57 \times 10^{-6}) \times (3.89 \times 10^{-5}) + (7.4 \times 10^{-6}) \approx 7.4 \times 10^{-6}$$

(2) In the case of $T = 10$ minutes (W_{AC2}):

$$W_{AC2} = (1.57 \times 10^{-6}) \times (6.18 \times 10^{-5}) + (7.4 \times 10^{-6}) \approx 7.4 \times 10^{-6}$$

The unavailability of Block 3 is much lower than the unavailability of Block 4 since the AC bypass backs up the inverter sub-system. The electrical switch SWe and its controller CONT can be said to be the 'bottleneck' of the AC system reliability. To increase the reliability, the unavailability of the common part (SWe and CONT) of the AC system should be decreased.

Reliability conclusions

The calculated availability of the DC system is more than 20 times that of the AC system. In the AC system, the failure of a single component, such as the control circuit for the electrical (bypass) switch, can cause loss of power to the loads. Loss of power results in loss of service to customers, which is intolerable for telecommunications systems. If, as expected, AC mains availability worsens, DC systems become even more advantageous for powering all equipment units.

Cost Analyses

One general perception of UPS powering is that it is more economical than 48 V DC powering. In this section, a comparison is made between the cost of deploying UPSs versus 48 V DC powering in two typical telecommunications scenarios.

For simplicity, the costs for upgrading the standby engine, the distribution and cabling from the AC panel to the UPS or 48 V DC power plant, potentially reinforcing the floor for the weight of additional batteries, equipment transportation and travel and living expenses during installation are considered equivalent in both scenarios, and ignored in the analysis. Only the true differential costs, which are the costs of equipment, distribution cabling and installation labour, are compared.

Central exchange

Table 5 considers the addition of a 10 kW load in a central exchange. In this environment, one–three hours of battery backup time are usually required and a standby engine is typically used to provide power during AC mains outages. For an AC load, the cost shown includes the UPS and two large external battery cabinets. A second UPS alternative with only 30 minutes of reserve is also presented, but this option is not recommended since a three-hour battery reserve time has already been established for the telecommunications equipment. For a 48 V DC load, the cost shown is that of adding additional 48 V DC rectifiers, batteries and distribution to the existing power plant in the central office.

The comparison shows that 48 V DC powering is the lower-cost alternative for central exchange powering, when equivalent battery reserve times are deployed. It should also be noted that a long duration UPS often requires an external battery charger unit to maintain a reasonable battery recharge time. This unit was not included in the cost comparison. The 30-minute UPS solution has only minimal savings compared to the DC plant, but as noted above is not recommended.

Remote sites

Table 6 considers powering options for a 5 kW load in a remote location, where four–eight hours of battery reserve time is generally the norm

Table 5 Addition of 10 kW Load in a Central Office

Load Type	AC	AC	48 V DC
Powering	UPS with 3 hour battery	UPS with 30 min battery	DC plant with 3 hour battery
Equipment included	— 10 kW UPS — 2 external battery cabinets — Load cabling	— 10 kW UPS — 1 external battery cabinet — Load cabling	— 200 A rectifier — Battery for 230 A/3 hours — Distribution and load cabling
Cost for equipment and installation	£14k	£10k	£11k

Table 6 Powering a 5 kW Load in a Remote Office

Load Type	AC	AC	48 V DC	48 V DC with DPA
Powering	UPS with 4 hour battery	UPS with 30 min battery	DC plant with 4/8-hour battery	DC plant using distributed powering architecture with 4/8-hour battery
Equipment included	— 5 kW UPS — 3 external battery cabinets — Load cabling	— 5 kW UPS — 1 external battery cabinet — Load cabling	— 150 A rectifiers ($n + 1$) — Battery for 118 A (4 hours and 8 hours) — Distribution and load cabling	— Battery for 118 A (4 hours and 8 hours) — Battery charging, control and shelf — Battery distribution and cabling
Cost for Equipment and Installation	£10k (4 hours)	£6k	<i>With new plant:</i> £11k (4 hours) £16k (8 hours) <i>With existing plant:</i> £6.5k (4 hours) £10.5k (8 hours)	£4k (4 hours) £7k (8 hours)

Figure 10—Global data communications equipment segments

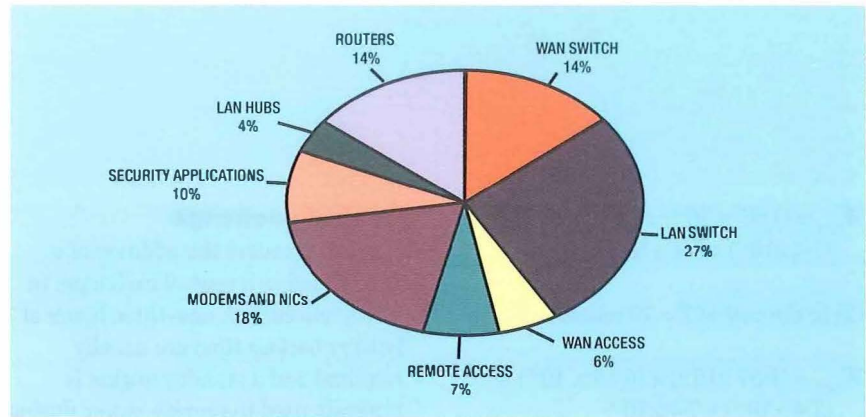
and no standby engine-generator set is available. Again, similar scenarios are considered: a UPS with a long-duration reserve time, a UPS with 30 minutes of reserve time and a DC plant with long-duration reserve time. An additional option, 48 V DC with a distributed power architecture (DPA), is also considered. This option reflects the use of a 48 V distributed powering architecture that is increasingly being used in some countries for medium and higher power datacommunications equipment. In this option, for reasons of modularity and the difficulty in distributing logic-level voltages at high currents, an intermediate 48 V bus is established and DC-DC converters or board-mounted power modules are then used to derive the lower voltages. This option also allows relatively easy addition of battery backup, since only the 48 V intermediate DC bus has to be backed up.

It should be noted that eight hours of battery reserve time are rarely available with most UPSs. In addition, the costs shown in Table 6 do not include the additional battery-charger unit needed to maintain a reasonable recharge time for the four-hour UPS.

For the 48 V DC plant, the top figures shown in the table include the cost of installing a 600 A (ultimate capacity) complete power plant with $n+1$ rectifiers, batteries, distribution and cabling. There is, however, a future benefit from the DC plant alternative. The cost for powering loads beyond the initial requirement will be less, since no additional plant infrastructure or redundant rectification will be needed.

If this scenario considered expanding an existing DC power plant rather than installing a new plant, the costs for the DC powering option would decrease to £6.5k (4 hours) and £10.5k (8 hours).

The 48 V DC with DPA option also offers the opportunity to increase power conversion efficiency when on



battery operation and thereby reduce the battery sizing. The additional costs when using this option are then solely due to the battery and charging/control equipment needed. This option results in a better than two-to-one cost advantage when compared to the four-hour AC UPS option, and is even less expensive than the AC-UPS option with 30 minute reserve.

The UPS with 30 minute battery reserve was included in both scenarios, but not recommended because a longer battery reserve time had already been established for the site. It should be noted that batteries sized for a 30 minute reserve are discharged at a higher rate than batteries sized for three hours or more. The batteries for the 30 minute UPS, therefore, typically have thinner plates than the telecommunications power plant batteries, and have a shorter installed life. The costs of battery replacement were not included in Tables 5 and 6, but would, of course, increase the cost of the 30-minute UPS alternative.

The cost of providing power for four hours of battery reserve time in remote locations is generally no more and may be less with a 48 V DC plant than with a UPS. One must also remember that powering with eight hours of reserve time is often not feasible with a UPS.

Market Estimates

By the turn of the century, the global market for data communications equipment is expected to exceed £40 billion per annum. This market consists of several categories of products intended for LANs, WANs, remote access and network support and services.

Figure 10 illustrates the relative size of each equipment category. While the majority of sales have historically been for private network applications, a growing portion is now deployed into public networks. This deployment comes about because the public network operators (mainly the local exchange carriers, the competitive access providers or newly licensed operators, the long-distance carriers and the Internet service providers) must modernise and expand their network capacity to handle the fast-growing demand for datacommunications traffic. Consequently, by the turn of the century, sales to these operators are expected to represent from 25–35% of the total demand for datacommunications equipment. The major applications for public network operators will be for WAN switching and access (ATM, frame relay and router) and for network servers and data security solutions.

To preserve the quality of service (QoS) and level of reliability that their customers expect, the public network operators will demand 'telecommunications-grade' solutions from the datacommunications equipment vendors. These levels differ markedly from those historically found in private networks, where price has often been the key driving factor. One can also conceive that in many critical applications where the traffic originates and terminates in a private network (LAN or other CPE terminal), the same higher level of performance and reliability will be required. One such requirement is to have the reliable source of DC power deployed in voice telephony networks today.

Most vendors who are already involved in the WAN switching

business, or who have a strong base in voice telephony, already offer telecommunications-grade data-communications products. Other vendors, who have historically been more involved in the LAN, CPE and server segments, still tend to design for AC powering. They will need to make the transition to DC powering if they intend to sell to public network operators.

Available Equipment

Much of the switching equipment associated with Internet service is available with a DC power input as standard. Third-party operators providing value-added services in their own premises have traditionally installed servers and routers in computer room environments and made use of AC-fed equipment.

The growth of the Internet network, coupled with awareness among the major equipment suppliers that telecommunications operators are becoming increasingly interested in providing Internet services, has increased the demand for optional DC-fed equipment.

This customer demand has prompted major industry suppliers to provide routers and servers that can be fed with either a DC or an AC source, and there is clear evidence that new products by other suppliers are in development. The 'pressure' to supply DC-powered equipment varies with the position of the equipment in the communications chain, and, to a lesser extent, cost and assumed overall reliability of the product. The nearer to the central exchange, the more likely that a DC-powered option will be considered. On the other hand, large customers with in-house routers currently see no need for DC-fed equipment, arguing that their in-house PCs are AC-powered and would not operate during mains outages unless supported by UPSs. Similarly, when used in customer premises, terminal adapters are available only with an AC power option. Some network terminating

equipment with both ISDN and conventional telephone inputs will, however, still support the telephone service in case of a mains outage.

Where routers, servers or switching equipment are available with different power options, they come in various formats. Examples are:

- AC as standard with DC as an option;
- DC as standard with AC as an option;
- dual AC or DC supplies; that is, either AC- or DC-fed with a redundant input capability; and
- twin AC or DC supplies, with either a third or a fourth redundant AC or DC supply option.

Routers that normally use AC power supplies typically cost slightly more when they are equipped with an optional DC power supply. Some servers, on the other hand, were initially DC-powered, but now come with AC power as an option. Typical voltage ranges for DC inputs are -48 V to -72 V. For AC inputs, typical ranges are 100-240 V_{rms} and 50-60 Hz.

To show the diversity of equipment, typical examples are shown in Table 7.

To further increase reliability, equipment is currently being designed that uses AC or DC with dual redundancy and dual internal power supplies. This new equipment can be fed from two separate power sources and has a redundant

internal power supply. With this approach, it will not be possible to mix AC and DC power supplies. However, equipment is being developed that has dual DC inputs and a separate AC input-DC output power pack, making it possible to feed the equipment from both an AC and DC source.

The introduction of a distributed powering architecture that uses a 48 V DC bus to distribute power to DC-DC converters and/or board-mounted power modules should be welcomed by operators who must install datacommunications equipment in telecommunications facilities. The increased efficiency of power distribution within the datacommunications equipment is utilised at all times, and battery reserve operation is achieved by powering the DC bus from the telecommunications DC power plant. Before connecting this bus to the DC power plant, the operator needs to verify that the bus converters and power modules can operate over the full range of the power plant output voltage.

As in the PC industry, there is likely to be an increased trend of end users being able to specify and purchase equipment designed to meet their specific site requirements.

Conclusions and Recommendations

This article has looked at the powering arrangements that have typically been chosen for Internet equipment installed in telecommunications facilities. The problems

Table 7 Examples of Power Supplies for Switches, Routers and Servers

Equipment	Power Supplies
Switches	AC only AC and redundant AC option AC or DC with dual redundancy option
Routers	AC or DC, dual redundancy optional
Servers	AC or DC, dual redundancy DC is standard, but with custom rectifier

inherent in AC power systems, the difficulties encountered when both AC and DC power systems are used within a facility, and the advantages of using DC power for all equipment that is added at a facility have also been considered. The entry of telecommunications operators to the business of Internet service provider (ISP) and the anticipated growth of Internet protocol (IP) telephony will lead customers to expect the same level of service on the Internet as they receive on the circuit-switched telephone networks today. This expectation mandates that all telecommunications and datacommunications equipment that provides communications service should have four–eight hours of battery reserve, as a system requirement.

Considerations of availability, equipment and maintenance costs, distribution and safety reveal that operating all equipment from the telecommunications DC power plant is clearly the optimum choice.

Many vendors already offer optional DC–DC converter power supplies in their switches, routers and servers. Although the converters often have a higher price than the SMRs, the size of the market for datacom equipment for telecommunications operators is so large that the demand for converters should be large enough to generate high-volume production prices. The introduction by datacommunications equipment vendors of a distributed power architecture based on a 48 V bus offers operators the opportunity to save money by omitting the bulk rectifier when purchasing datacommunications equipment for telecommunications facilities.

The primary responsibility for achieving the vision of using DC as an energy source lies with the operators. Power engineers need therefore to identify the organisations in their companies that specify datacommunications equipment and convince them to insist on DC power supplies for all their applications in telecommunications facilities wher-

ever a telecommunications grade-of-service is needed

It is hoped that this article has gone some way to explaining the reasoning behind this vision and will encourage system specifiers to use DC as a first choice.

References

- 1 ALLEN, GEORGE W.; and SEGALL, DONALD. Monitoring of Computer Installations for Power Line Disturbances. C74 199-6, IEEE Power Engineering Society, 1974.
- 2 GOLDSTEIN, M.; and SPERANZA, P. D. The Quality of U.S. Commercial Power. *Proceedings of the Fourth IEEE International Telecommunications Energy Conference – INTELEC® 82*, Washington, D.C., Oct. 1982, pp. 28–33.
- 3 DORR, DOUGLAS S. Point of Utilization Power Quality Study Results. Presented at the 1994 IEEE Industrial Applications Society Annual Meeting, Denver, CO, 3 Oct., 1994.
- 4 Data from the Electricity Council of the UK.
- 5 Statistics of Electric Failures in 1994. Electric Power Technology Division, Public Utilities Department, Agency of Natural Resource and Energy of Japan, 1994 (in Japanese).
- 6 YOTSUMOTO, K.; MUROYAMA, S.; and WATANABE, H. Design Considerations of a Distributed Power Supply System for Telecommunications Network Use Based on Reliability Analyses. *The Transactions of the IEICE*, Vol. E 73, No. 3, March 1990, pp. 395–400.
- 7 YOTSUMOTO, K. Research on System Architecture of Telecommunications Power Systems. Kyushu University Publications, Feb. 1992, (in Japanese).
- 8 ISHIYAMA, Y.; FUKUDA, Y.; NOGUCHI, H.; and ASANO, N. Quality Management System for Telecommunications Power and Air-conditioning System. *Proceedings of the Nineteenth IEEE International Telecommunications Energy Conference – INTELEC® 97*, Melbourne, Australia, Oct. 1997, pp. 297–302.

Biographies




John Parsons
Networks and
Information Services,
BT UK

John Parsons is manager of the Network Power and Cooling Systems unit, in Networks and Information Services, BT UK, which has overall responsibility for all aspects of Network Power and Cooling plant, including policy, operational support and capital programme management. He graduated from Loughborough University in 1971 with a double honours degree and has worked in various parts of the business since joining the company in 1972. He was a member of the technical panel that wrote the Guidance Notes to IEE Wiring Regulations and the government Technical Panel that wrote the Electricity at Work Regulations. He is a Fellow of the IEE, and a Member and Technical Advisor to INTELEC. He has presented five international papers covering subjects from battery monitoring systems to the history of power in the telecommunications industry, session chaired six international power conferences and been an invited speaker to a further two international conferences. He is also author of the Power Section in the Telecommunications Engineers Reference Handbook.



John Åkerlund
Telia AB

John Åkerlund graduated with an M.Sc.E.E. from Chalmers University of Technology in Gothenburg, Sweden in 1974. He worked for two years with the Swedish Coast Guard. From 1976–1989, he worked for Ericsson Energy Systems; he was involved in developing power supply systems for AXE and other telecommunications systems and was project manager for Ericsson's alternative energy project Ericsson Sunwind. Since 1989, he has worked in the Telecommunications Energy Department of Telia Network Services in different positions engaged in developing reliable power for telecommunications within Telia.




Chris Riddleberger
Consulting Engineer

Chris Riddleberger holds a master's degree in Electrical Engineering. He is a member of Eta Kappa Nu and Tau Beta Pi, and an associate member of Sigma Pi. He spent over 27 years at Bell Laboratories, engaged in product development. For more than 15 years, he managed development of telecommunications power systems, and is now a consultant on telecommunications power. He is a Senior Member of the IEEE and has worked on the IEEE International Telecommunications Energy Conference since its founding in 1978.



Bo Lindemark
Ericsson Components AB

Bo Lindemark received his master's degree in Power Electronics and Industrial Engineering from Skovde University of Technology, Sweden, in 1968. He has over 28 years experience with Ericsson Energy and its predecessors. He has held various positions in sales, design, market organisation and international. He has presented a number of INTELEC papers. Since 1996, he has been product manager for AC back-up power.



Dan McMenamin
Bell Atlantic

Dan McMenamin is a telecommunications engineer for Bell Atlantic Corporation with over 30 years' experience in central office energy, switching, and transport systems. He is a member of the Lightning Protection Institute and the IEEE Communications Society. Mr. McMenamin is author of approximately 20 technical papers and numerous other published works.



Seiichi Moroyama
NTT

Seiichi Muroyama received the B.E. degree in Electrical Engineering from Tokyo Metropolitan University in 1972. Since joining NTT Laboratories in 1972, he has been engaged mainly in developmental research on uninterruptible power supplies and DC power supply systems for telecommunications systems. He is currently Senior Engineer, Supervisor of Research and

Development Department in NTT Power and Building Facilities Inc. He is a member of the IEEE Power Electronics Society, the IEICE of Japan and the IEE of Japan.



Katsuichi Yotsomoto
NTT

Katsuichi Yotsomoto is Executive Vice-President, NTT Building Technology Institute (NTT-BTI), and is responsible for research on telecommunications energy systems. For 31 years, he was engaged in research and development on telecommunications power systems at NTT Laboratories, Nippon Telegraph and Telephone Corporation and NTT Power and Building Facilities Inc. He received the B.S. degree in Electrical Engineering from Kagoshima University and the Ph.D. degree in Electronic Engineering from Kyushu University. He is a Fellow of the IEEE and a member of the INTELEC Advisory Committee.



Jacques Poulin
Nortel

Jacques Poulin holds a degree in Electrical Engineering and a Master in Business Administration from McGill University. He worked for 15 years in telecommunications in the design and marketing of fibre optic transmission, access and power equipment. His most recent positions have been Senior Product Manager and Senior Marketing Manager for the Power Division of Nortel. He is presently the Director of Marketing for Astec Advanced Power Systems.

A People and Information Finder for Virtual Construction Project Teams

As virtual organisations become more abundant, the need to provide a means of locating people within them becomes more important. This paper identifies an intranet tool designed to support interaction – the People and Information Finder. Although developed for use in the construction industry, where virtual organisations are common, it is applicable to any industry employing them.

Introduction

The construction industry makes wide use of virtual project teams. These teams, made up of many different organisations, exist for the duration of a construction project. Each part of the organisation carries out a defined part of the construction process depending upon their area of speciality. The size of the task they perform and its position in the construction process will often determine the length of time that they are members of the virtual team. Communications between the different parts of the virtual team can be problematic. The transient nature of the organisation can mean that the visibility of others and their work areas is poor. This can lead to ineffective decision-making and costly delays. The problems are compounded by the sheer complexity of modern buildings and the fact that many buildings are unique (each with a unique set of requirements to satisfy).

The Collaborative Integrated Communications for Construction (CICC) project is looking at improving communication in the construction industry¹. It is an ACTS project, funded by the European Commission, involving telecommunications companies, construction companies and research establishments. The project is developing technologies and trialling them on real construction projects. One of these technologies is the People and Information Finder (PIF) which is designed to aid communication

between individuals on a project and to aid access to project information. The ad hoc meetings (such as meeting someone in a corridor) that can offer an effective form of communication between individuals in the same locality are lost when a project team is a virtual one. Technologies such as the PIF are attempting to allow such meetings to occur between members of the virtual team, who are often geographically distributed.

This article details the development of the PIF within the CICC project and within the context of concurrent life-cycle design and construction.

The People and Information Finder (PIF)

Architecture

The PIF is an Internet-based technology. By using a browser, users can access information stored on a server over a network. The protocol used is hypertext transfer protocol (HTTP) and the data is generally in the form of hypertext mark-up language (HTML). Since the PIF is all about finding information, users should have the facility to search through the data in order to find what they are looking for. The implementation of the PIF discussed here allows users to search through a database of people information. Other systems developed in CICC have employed search engines to enable text searches of Web pages containing project information.

Figure 1 – PIF architecture

Figure 1 shows the PIF architecture. Users browse with an HTTP client (most commonly Netscape Navigator or Microsoft Internet Explorer) over a network. The network will generally be an intranet (closed corporate network) or an extranet (closed virtual organisation network). The HTTP server employed is Microsoft Internet Information Server running on Windows NT. This serves data, in the form of Web pages, to the clients. The PIF Web pages are a mixture of forms, that allow users to specify search criteria or to enter data (as shown in Figure 2), or dynamic pages, containing information from the database, as a result of searches (Figure 3).

The PIF employs WebDBC (a commercial product) to achieve connectivity between the Web server and databases. This software converts the contents of Web forms into database queries and permits the contents of databases to be published on Web pages. This type of technology is becoming more important, since the Web, as a means of providing heterogeneous access to distributed data, is growing in popularity.

The queries generated by WebDBC are ODBC (Open Database Connectivity Protocol) compatible. ODBC is a Microsoft supplied, application programming interface (API) that provides a standard for accessing different kinds of database. The intention is to provide a level of abstraction such that the calling application does not need to know about the workings of the underlying database and can treat all databases in the same manner.

BT has developed AP-IA, the Availability of Person—Intelligent Agent, within the CICC project. AP-IA is used in the PIF to augment the information in the database with estimates of the current availability of people. AP-IA is explained in detail below but it essentially works by monitoring the activity of equipment that people use frequently in their work (that is, the telephone and the

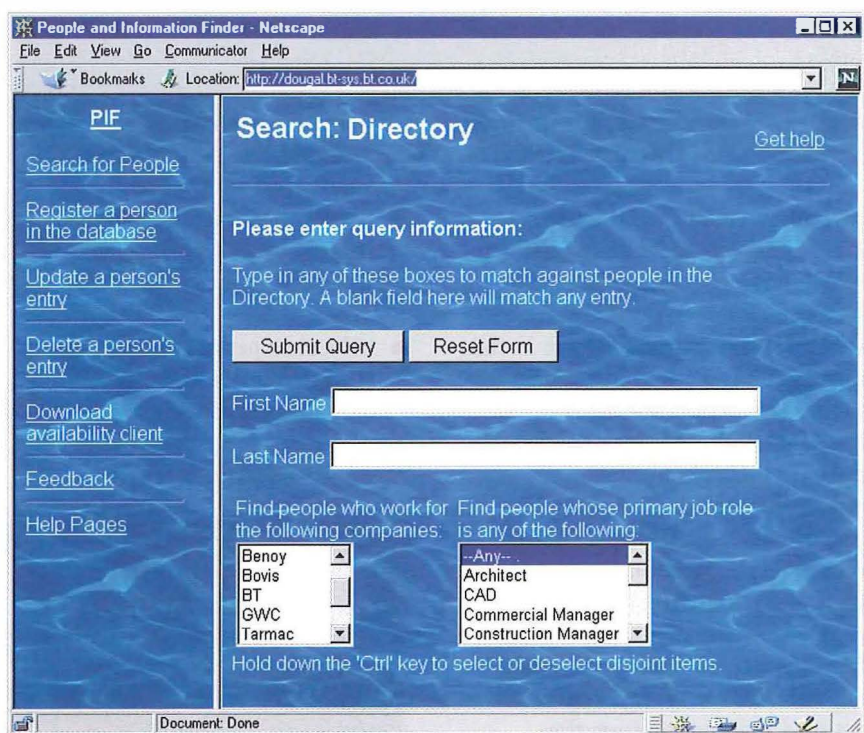
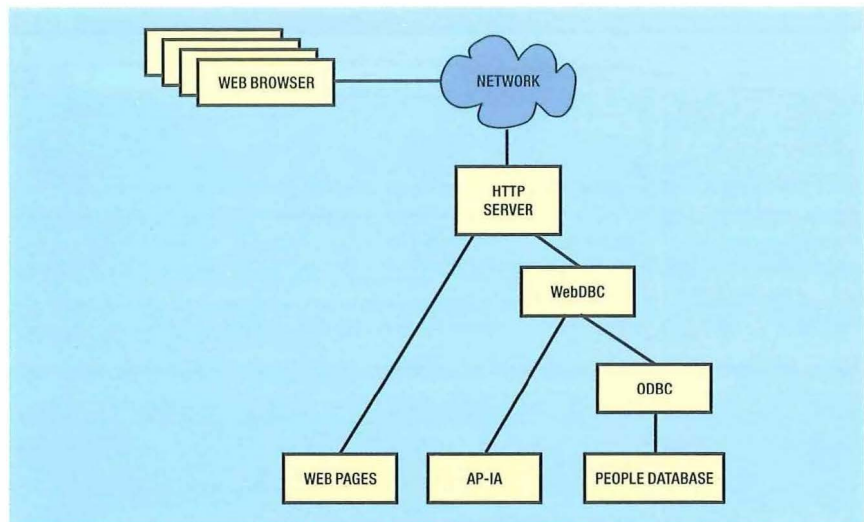


Figure 2 – WWW form

computer). It is a low-level agent that works best when used in conjunction with other agents or information systems. Agents such as Grapevine² which uses the commonality in different people's expressed interests to promote personal contact could use AP-IA to inform users of the current availability of others with similar interests.

Operational context

The following description is an example of a user session. It explains the processes involved. A member of the design team wishes to alter the dimensions of a particular area of the building. The team member sees that this impacts upon

the available corridor space and feels he/she should check that no safety regulations covering the size of corridors are contravened. The team member wishes to contact an individual in the virtual organisation who will be able to help them with the issue. Not knowing the name of such a person, he/she can specify 'safety' in the job role section of the search page (see Figure 2). Upon receiving this information, the Web server calls WebDBC to construct a database query. This is then passed to ODBC, which carries out the search on the appropriate database. For each matching record, WebDBC calls a program that interrogates AP-IA and returns the

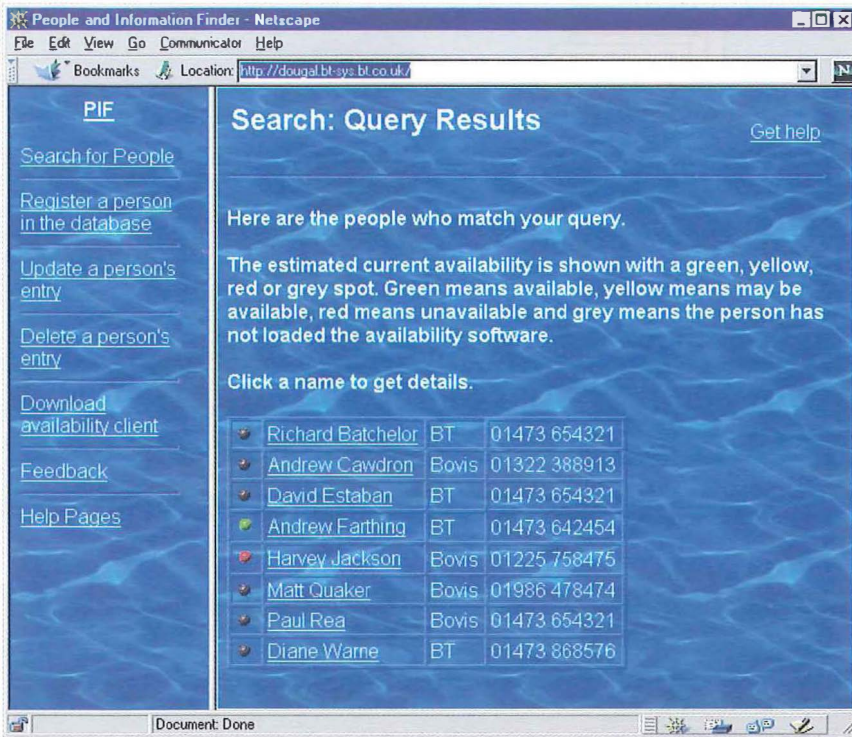


Figure 3—Results page

current availability of the person. The results are added to the database information and displayed on the user's browser via the Web server (Figure 3). Each person in the result set has a coloured dot next to their name to indicate the result of the AP-IA look-up (green means they are probably available, red means they are unavailable and grey means that AP-IA has no record of that person). The user can then see, at a glance, which of the people is available and select them from the list. This selection (again via WebDBC etc.) returns more detailed information from the database, as shown in Figure 4.

Figure 4 also includes images from a 'Webcam'. The image from a camera pointing at the desk area of the individual is periodically written to a file. This is then added to the Web page. Although giving a very good indication of availability, it is more intrusive than AP-IA and perhaps more intrusive than some people would prefer:

Using the PIF

It is envisaged that people might use the PIF for a number of different reasons:

- to find out how to contact colleagues, whom they already know, be it to find their telephone numbers or their e-mail addresses;

- to monitor the availability of a person or a group of people to see when it is best to contact them;
- to locate a person who performs a specific role within a part of the organisation with a view to

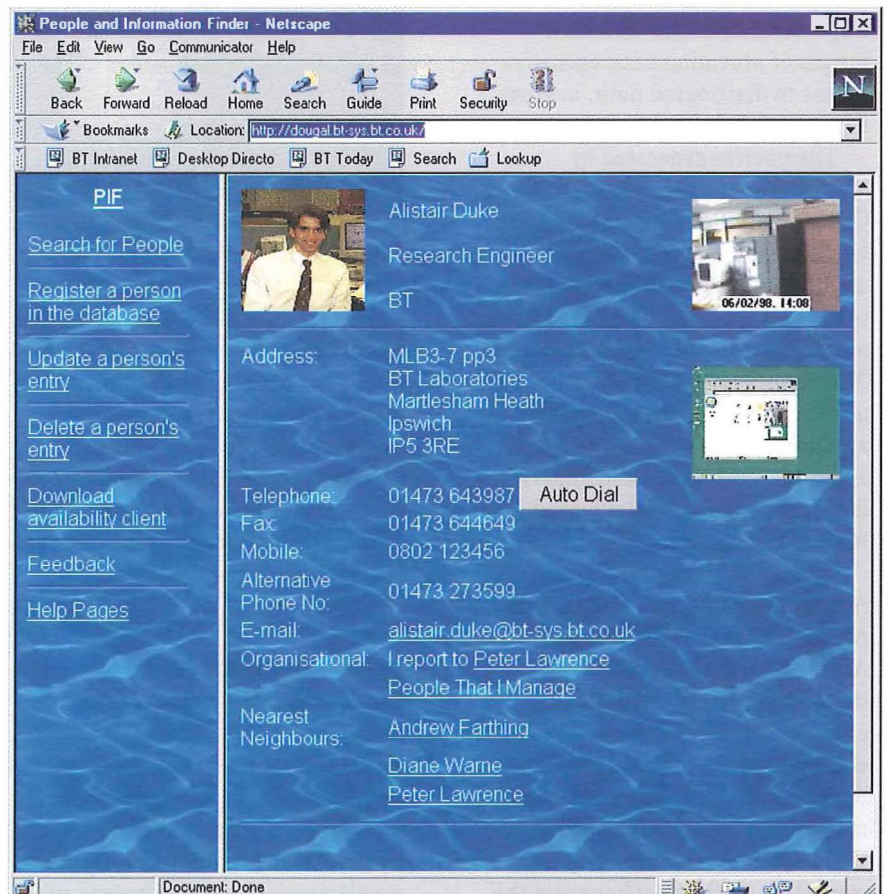
contacting that person to discuss an issue; and

- to browse around in order to increase their general visibility and understanding of the virtual organisation—fitting names to faces and so on.

Specifying search criteria on the Web form can fulfil all these requirements. Specifying a name will return information about an individual, while specifying a job role or a company will return a group of people who may be appropriate to deal with an issue. In all cases, the availability of information will enable people to decide whether to try to make contact with an individual or to try a colleague or alternative contact.

Once the required person has been located in the database, a desirable feature would be to enable users to initiate an appropriate communica-

Figure 4—Personal details page



tion channel with them directly from the Web page. This is easy to implement for e-mail since most browsers support this facility. It is more problematic for the telephone since some form of computer telephony integration is required. The system currently employed uses a single PC with a telephony card³. The 'Auto Dial' button shown in Figure 4 sends a request to the PC with the source and destination telephone numbers. The PC then uses its telephony card to set up a two-way call between these two telephones. The length of time it takes to set up the call is dependent upon network traffic and the load on the telephony PC. Currently, the process is generally longer than it would take to dial the call manually; however, it is expected that this will improve.

Besides searching, the other main user operation of the PIF is to enable users to maintain the information held about themselves. The level of access given to each person is an important issue. Should users be able to change data directly or merely through an administrator who ensures integrity of the data? The type of organisation and the skill level of the users will be important factors in determining this.

Availability of Person—Intelligent Agent

Architecture

AP-IA is a system designed to determine the likelihood of a person being available. The intention is to aid the initiation of communication between people at remote sites. AP-IA is a low-level system. Its use in systems such as the PIF is intended to allow ad hoc meetings to occur between remote colleagues—the sort of meetings that people who work in the same place find very useful and often take for granted (for example, discussions across the office or chance meetings in a corridor). AP-IA works by monitoring the equipment that people use in the working environment (for example, telephones

and computers). Transponders residing on, or associated with, these pieces of equipment generate event messages and send them to the AP-IA server. They also listen for information requests from the server, which they respond to. The AP-IA server holds this information in a database.

The server responds to requests made to it with the data stored on a particular person, giving an indication of the availability of that person. It is envisaged that requests will be generated by:

- users browsing a people database with a view to establishing communication,
- an automatically updating WWW page giving the latest estimate of a person's availability, and
- a virtual world with representations of people, which change depending upon their estimated availability.

The AP-IA client

The AP-IA client is known as the *PC transponder*. This resides upon the PC of each person that AP-IA is to monitor and collects information about the usage of that PC. The following are monitored by the transponder:

- changes in mouse position,
- changes in caret position (keyboard cursor),
- current Schedule+ appointment,
- number of unread mail messages (MS Mail or Exchange),
- user message, and
- do-not-disturb selector.

The transponder is written in Microsoft Visual Basic. This rapid application software allows lightweight clients with supporting




Figure 5—Transponder GUI

graphical user interfaces (GUIs) to be developed. All the CPU usage required for this monitoring process is carried out using idle time. This means the user should not notice any degradation in performance as a result of the transponder. The transponder is designed to run as minimised on start-up. User interaction is achieved via the GUI illustrated in Figure 5. This allows the user to enter a message or specify whether they wish to be disturbed or not.

The mail and Schedule information is determined with the use of MAPI—Microsoft's Messaging Application Programming Interface. This allows the status of messaging applications to be interrogated by other processes on the same machine.

PIF Pilot Trial

The PIF is currently being trialled as part of a CICC pilot on a large construction project. The project involves a large number of organisations distributed across many locations. The site itself is very large, so even those located on site are distributed among the site offices and the management team office. In such an environment, technologies such as the PIF can be beneficial.

The trial includes both technical evaluation and usability studies. This is helping to determine how useful such systems are in the workplace and how they impact upon the working environment. It is also highlighting technical issues with the system and steering the technical development process. One such example of this was the personal address book facility. Upon exposing the PIF to construction managers, they revealed that although the search facility is vital during the

Generally, the results have shown that the PIF will be a useful tool for overcoming many of the communication problems that exist in large virtual organisations.

earlier stages of the project, at later stages the need for it is diminished since most communications tended to take place with a few key individuals. A facility was required that could display information about the key people that each user contacted most often. The personal address book was added that allows users to select people to add to their individual pages. The page (available within one mouse-click after starting the PIF) not only shows the telephone numbers of people, but also their current availability. Thus by keeping the page open, users can see at a glance those who are available for contact.

The PIF has been installed at the project management office, on-site, and at the offices of the client in London. New users were introduced to the technology and observed while performing a series of tasks using it. The intention was to evaluate the usability of the system and its interface. They were then asked a series of questions related to how they find and contact people currently and whether a system such as the PIF would be beneficial on construction projects.

Generally, the results have shown that the PIF will be a useful tool for overcoming many of the communication problems that exist in large virtual organisations. However, there is quite a large investment involved in setting up and maintaining the information. This is a potential barrier to the successful adoption of the technology. It quickly became clear from the trial that the 'bottom-up' approach of allowing people to add their own information was problematic. Users are more than willing to add data about themselves but have high expectations of the data available about other people. If that data is not yet available then the likelihood of them using the system is diminished. A 'top-down' approach, with information being set-up and maintained from a central resource, may well have been more successful.

The results of the pilot will be used in further evaluation activities

involving the PIF and related technologies, including those outlined below.

Telepresence in Concurrent Life-Cycle Design and Construction

The PIF is being further developed as part of a collaborative project between BT and Loughborough University, entitled *Telepresence in Concurrent Life-Cycle Design and Construction*⁴. The aim of the project is to produce an information landscape for virtual construction project teams. This would be a multi-user, navigable environment containing visual representations of the construction and of members of the project team. These representations will provide access to underlying project data (drawings, schedules, rationale, etc.) and to the people themselves via appropriate communication channels.

Such a system would be implemented using Internet/intranet technology. This would allow users access to information stored in the network. They would be able to examine, in three dimensions, specific parts of the project. Upon selecting an individual component, the user would gain access to all the relevant project information about it such as dimensions, cost, scheduling information, and with whom responsibility for it lies. This information (which could be stored in one or several databases or information sources) would be referenced from the object of interest. The fact that the information is being pooled from potentially disparate information sources is not apparent to the user.

Significant research within the industry is into the concept of a single project model. This will enable sharing of all the information related to the construction by representing it as objects within a networked project repository. Developments in this area will be utilised by the project in order to enable this important area of interaction.

In addition to representations of the project, there would also be those of

project personnel. These avatars could give an indication of an individual's availability, rather as the home page does, with the appearance of the avatar changing over time to reflect this. Individuals navigating around the space would be represented by their avatars thus allowing others to see what their current focus is. The act of approaching individuals in the landscape would automatically open up an appropriate communication channel depending on the capabilities of the individuals involved. This would facilitate ad hoc project discussions. Individuals inquiring about a common area of interest would be able to interact with each other while using the relevant part of the model as a reference for the discussion.

Conclusion

The need for an effective communications infrastructure for virtual construction project teams cannot be over-emphasised. This is even more so in a concurrent engineering environment where it is essential to closely co-ordinate several concurrent processes and activities. The People and Information Finder provides an effective mechanism for project team members to have direct access to information about other project personnel. PIF will improve the visibility of other people on the same project. It will also increase the likelihood of finding the right person to talk to and facilitate the initiation of communication with someone who is currently available. Due to the improved visibility, the number of people any individual has contact with could be increased. This will enable the flow of knowledge between different parts of the virtual construction organisation and will reduce the level of uninformed decisions.

Development of the PIF will continue. It is an important building block for the Telepresent Information Landscape which aims to further increase the effectiveness of communication for virtual construction project teams.

Reference

- 1 ANUMBA, C. J.; and DUKE, A. K. Internet and Intranet Usage in a Communications Infrastructure for Virtual Construction Project Teams. Proceedings, IEEE 6th Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises, 18–20 June 1997, Boston, USA.
- 2 CRABTREE, B.; SOLTYSIAK, S.; and THINT, M. Adaptive Personal Agents. *Personal Technology Journal*, Special Issue on Personal Information Agents, Dec. 1998.
- 3 BROCKBANK, R.; PEIRCE, S.; and AMANATIDOU, Z. Computer Telephony Integration and Java. *Br. Telecommun. Eng.*, April 1998, **17**, pp. 10–17.
- 6 ANUMBA, C. J.; and DUKE, A. K. Telepresence in Virtual Project Team Communications. Proceedings, ECCE Symposium on Computers in the Practice of Building and Civil Engineering, Lahti, Finland, 3–5 Sept. 1997, pp. 80–84.

Biography



Alistair Duke
Networks and
Information Services,
BT UK

Alistair Duke graduated from Aston University in 1994 with an M.Eng. in Electronic Systems Engineering. Since joining BT he has worked on the collaborative projects Virtuosi and CICC. Both of these have been concerned with using virtual reality and collaborative communications to support group working. He is a member of the Affective Communication and Computing Centre at BT Laboratories and the manager of the CICC project.

Journal Awards for 1997/98

Introduction

IBTE's *British Telecommunications Engineering* journal is an important record by which the membership and others can keep abreast of various items of interest in telecommunications.

To encourage readers in furthering the role of the *Journal*, and to give authors due recognition for outstanding contributions, the Board of Editors operates an annual award scheme. Prizes are awarded to the authors of articles in *British Telecommunications Engineering* and its accompanying *Structured Information Programme* which, in the opinion of the Board, demonstrate excellence in content and presentation and which enhance the quality and range of contributions published.

The prizes for 1997/98 were presented in October 1998 at the IBTE Awards Dinner and the IBTE London Colloquium.

Top Award for 'Fighting the Information Terrorists'

The prize for the best article from Volume 16 of the *Journal* went to Simon Phoenix and John Regnault for their article 'Fighting the Information Terrorists', published in the January 1998 edition. They received

Winners of the Best Journal Article Award John Regnault (left) and Simon Phoenix (centre) with IBTE President Chris Earnshaw



crystal bowls, inscribed with the IBTE's insignia, a cash prize and certificates.

Information technologies and new services like e-commerce are fundamentally changing our lives. However, the very aspects of this technology that are making it so successful, also bring new threats to the network and its systems. As a 'New Wave' company, BT's future will critically depend on people having trust in the services we offer. But how secure are these new services against malicious activities? The excellent article by Simon Phoenix and John Regnault gives us a glimpse of another world—a world in which there is a continuing struggle to protect our network and stay a few steps ahead of those with ill intent. As telecommunications engineers, we cannot be ignorant of these threats, indeed it is important that we understand this new dimension to our work. We can sleep soundly at night aware that Simon and John are burning the midnight oil finding ways to protect us from the information terrorist.

Journal Runners-Up

The Board of Editors awarded two runner-up prizes for the *Journal*. The authors received crystal goblets, cash prizes and certificates.

Iris Recognition for User Validation

The first runner-up prize went to Chris Seal, Maurice Gifford and David McCartney for their article 'Iris Recognition for User Validation', published in the July 1997 edition.

This article is unusual in that, given the number of photographs of eyes, it appears to be reading you! The authors explain the need for reliable user identification without the need for cards, badges or keys



Runners-up in the IBTE Journal Awards (left-right) David McCartney, Chris Seal and Maurice Gifford receiving their award from Chris Earnshaw

and without being reliant on memorised data such as passwords and personal identification numbers (PINs). The solution, they tell us, is in the unique pattern in the eye's iris. The pattern of the iris is turned into an *IrisCode*[™] using a video camera and computing power.

The article goes on to explain how these codes for individuals are stored and compared each time an individual wants access to something protected by the iris recognition system. Statistically, the odds of being misidentified are enormous, ranging from one in a billion when only 70% of the bits in the codes match, to one in 10⁶⁸ when 99% of the bits match. Given this reliability, the recognition system works even when the eye is wearing contact lenses or glasses (provided they are not excessively scratched).

The authors explain the system and the complex statistical theory in an understandable way and make the subject readable and interesting.

Telecommunications and the Disabled—Friend or Foe?

The second *Journal* runner-up prize was awarded to Paul Tomlinson for his article 'Telecommunications and the Disabled—Friend or Foe?',



John Griffiths (left), winner of the Best Structured Information Programme Unit Award, and Paul Tomlinson (right), runner-up in the Journal Awards, with Paul Nichols, IBTE Office and Publications Manager

published in the October 1997 edition.

The preamble for this article reminds us that Alexander Graham Bell did not set out to invent the telephone. He was investigating the electrical analogue of speech in the hope that he could find a method of using this knowledge to help his pupils overcome their handicaps. He would be very disappointed to be told that today his invention is often seen by hard-of-hearing people as a potent symbol of their isolation from society.

Starting from this irony, the writer describes, in a matter-of-fact style, the magnitude and range of problems faced by those with the many forms of disability. How does a person in a wheelchair get into the design-award-winning Scott telephone kiosk? How does a person with coordination difficulties or arthritis operate the small keypads on a modern cellular telephone? Ergonomics these days leave a lot to be desired when all user requirements are taken into account.

However, there are some successes, for example BT's Typetalk. A quote from a lady deaf from birth, informs us 'I have found the service invaluable. The telephone was the biggest source of discrimination, but a whole new world has now opened up for me. Typetalk has given me the independence and confidence to do my own work'.

Other successes are the calling line identification products which offer a deaf person quite effective answering machines. However, it is doubted whether the designers ever had this application in mind!

In all, this is an informative and illuminating article, especially for those without disability, and proposes that all service designers must develop an awareness of disability issues in their bloodstream. They must instinctively ask 'how will people with various forms of disability use this service?' A simple question, but one with far-reaching consequences.

ISDN is a Winner

The prize for the best unit from the units published in issues 20-23 of the *Structured Information Programme* went to Professor John Griffiths for his unit 'The Integrated Services Digital Network', Chapter 9, Unit 4. John received a crystal bowl inscribed with the IBTE's insignia, a cheque and a certificate.

Little did Alec Reeves realise that his patent of pulse-code modulation (PCM) in 1937 would form the basis of today's and, in all probability, tomorrow's integrated services digital network (ISDN). John Griffiths's unit looks at the technology, evolution and application of the ISDN. After a potted history of telephony, the unit moves on to address the initial challenge facing proponents of ISDN, namely to find a practical transmission method allowing 64 kbit/s links to be extended from local exchange to customer. How this challenge was addressed is explained in simple clear terms with the aid of appropriate diagrams.

The unit then looks at the customer interface, and the signalling channels. (Anoraks will enjoy these two sections, as there are numerous references to S-bus and T-bus, which they could erroneously construe as single-decked and trolley.)

Before addressing the application of ISDN, the unit explores the marketing viewpoint, giving both sides of the marketer's rationale.



Tony Mullee (left), winner of the Structured Information Programme runner-up award

Finally the unit looks at the applications from telephony to data, from facsimile to videophone, and explores the new services of high-quality voice and music, leased line backup, and Internet access. Given the ever-converging markets between telephony, data and entertainment this unit is a must for anybody wishing to understand the underlying principles. Never let it be said that ISDN stands for Integration Subscribers Don't Need.

SIP Runner-Up

The runner-up prize for the *Structured Information Programme* went to Tony Mullee for his unit 'Service Surround and Service Management', Chapter 2, Unit 4.

It was not so many years ago that BT was criticised from all quarters for its poor standard of customer service. Those dark days are thankfully gone and BT's customers experience a standard of customer service which is among the best in the world. It is an achievement for which those people involved can be justly proud.


In his unit Tony Mullee provides a highly readable and instructive introduction to the thinking and processes that underpin that success. As Tony indicates you can have the best products, the best systems, but if the quality of the total service experience is poor then you lose your customer. The service surround and the management of that service surround is the battleground of today and it is where BT is differentiating itself and gaining competitive advantage.

37th European Telecommunications Congress

Following the publication in the August 1998 special edition of the Journal of the papers from FITCE's 1998 Congress, held in London, Tim Jeanes, one of IBTE's sponsored delegates, gives his impressions of the event.

Tim Jeanes is Secretary of the IBTE South West Zone committee.

A highlight of the opening day of the Congress was a round-table discussion. Participants were (L-R): Kevin Fogarty, EURESCOM; Prof. Sergio Treves, Alcatel Italia; Chris Earnshaw, BT; Ian Vance, Nortel; Frank Heller, Deutsche Telekom AG; Peter Walker, OFTEL; and Franco Bigi, European Commission, DGXIII



Introduction

During the last week of August almost 400 delegates from 14 European countries, including 20 nominated by the IBTE, gathered at the University of London for the 37th Congress of the Federation of Telecommunications Engineers of the European Community (FITCE).

The IBTE was a major sponsor of the event together with BT, which hosted the event; Accunet/Matra; Ericsson; Fujitsu; the Institution of Electrical Engineers; Lucent Technologies; Marconi Communications; Nortel; Siemens; Motorola and the European Commission.

The theme for this year's Congress was 'Diverging Roles in a Converging Marketplace', which typifies the change of emphasis which has come

about over the last 10 years in the focus of this annual meeting of FITCE. Predominantly technical in the past, the papers selected this year by the Technical Committee chaired by Professor John Griffiths, from Queen Mary and Westfield College, University of London, were far more wide ranging touching on many of the issues facing telecommunications operators in an increasingly liberal marketplace.

Opening Ceremony

At the opening ceremony Chris Wheddon from BT, Chairman of the Congress Executive Committee, noted the world of communications was undergoing a tumultuous change brought about by rapid advances in technology in the areas of computing, fibre optics and digitalisation of content. This was causing a convergence of traditionally separate industries and a huge increase in entrepreneurial operators in the merged marketplace, the result of which was a sometimes bewildering amount of choice for consumers, but with an underlying need to ensure that, what Chris called the 'terror of the 3Is'—Interworking, Interconnection and Integration—was effectively resolved.

Chris Earnshaw, Managing Director of Networks and Information Services, BT UK, who is also IBTE's President and President of FITCE UK, then gave a lively presentation highlighting the changes in our industry and the response required to remain competitive.

Chris noted that, in the past, life had been simple with defined markets and slow technology. Now that was all drastically different with rapidly advancing technology demanding ever-increasing bandwidth meaning the establishment of superhighways that could not only serve fixed, but also mobile, customers. The 1980s had brought the destruction of the vertical structure in the computing industry and this was now happening in telecommunications driven by customer choice and the emerging global electronic economy.

Chris predicted that in future it would not be sufficient to be just a network operator but that, to meet the increasingly complex requirements of customers, multiple packages would be needed to yield seamless, high-integrity, 24-hour service. This would in turn mean cheap bandwidth and more complex relationships with partners, subsidiaries and joint ventures to yield what Chris dubbed a new *Telecosm* for a new millennium.

Chris's final message was that although all of this would require considerable new skills it would also yield an opportunity to change the future—and the future is here now.

After Chris Earnshaw, Signor Franco Bigi of the European Commission, DGXIII gave an overview of the work of the RACE programme which had started to examine the concept of 'one bit stream to carry all' in 1988 with a view to a practical implementation by 1995. The objective had proved more illusive than initially thought but much progress had been made along the way.

Technology has extended the life of the copper into the future by using developments such as xDSL. Mobile, which was important to RACE in its anticipated role for business people, is now used by people doing their shopping!

The future of wireless communications is of particular interest in undeveloped areas where low-earth-

orbit satellite systems such as Iridium and Globalstar can reach individuals who have never made a telephone call in their entire life.

Signor Bigi speculated that electronic commerce would become far more prevalent more quickly than people imagined, driven by the realisation that home computing equipment provided for leisure is equally applicable for business purposes. The important thing was to ensure that we engineer a user-friendly information society through ease of access, use, affordability and interoperability to yield systems and services for the citizen of the future.

In closing the initial session, Dipl.-Ing. Guntram Kraus, President of FITCE, welcomed all the delegates and outlined the issues for FITCE which he stated as (a) Our Conference, (b) The Future of FITCE and (c) The Visionary focus of FITCE. Herr Krauss noted that the concept of globalisation had only really taken off in the last 10 years with its requirement to connect a network of people anywhere in the world. FITCE should be proud that it recognised this in the 1960s and that cooperation should be undertaken with telcos in the USA. FITCE had a good track record and, with care, could develop successfully as long as it remained constructively critical of what it undertakes.

Herr Krauss thanked all those involved in the organisation of the Congress, which was the highlight of the FITCE year, and all the sponsoring companies without whose gracious support the event would not be possible.

The Imperial String Quartet then played FITCE's overture 'An Invitation to European Dialogue' to set the scene for a round-table discussion by a panel of industry experts on the topic 'Diverging Roles in a Converging Marketplace—Disruptive Technologies and their Effects on the Market'. The moderator Ian Vance from Nortel was successful in stimulating a lively and interesting debate.

Technical Sessions

The technical sessions, which constituted the major part of the Congress, occupied 14 sessions over the following five days and in total 42 papers were presented.

The themes of a liberalised environment were apparent in session with titles such as 'The Marketplace' and 'The New Order' where the issues of a multi-provider environment were explored along with the operation of the regulator, created to oversee the liberalisation process. The overall picture was one of rapid change, not only on the technological front, but also in relationships between vendors, operators and customers. New ways of working together will be required to release fully the potential of the convergence of telecommunications, computing and the media industries.

Within all this change customer satisfaction cannot be lost from sight and a number of presentations focused on the need to maintain and grow customer loyalty in this area.

The enablers for liberalisation such as intelligent networks (IN) and mobility were considered and the operational challenges that stemmed from increasingly complex interworking requirements were presented. These were examined in terms of management systems, quality of service and the need for inter-network accounting and billing. There was also a fascinating presentation on the potential for telecommunications fraud in the new environment.

A clear message was that the convergence of all forms of communications is probably the biggest challenge and the biggest opportunity for future expansion of operators, allowing marketing of their services to more users.

The ever-increasing demand for bandwidth and the services and applications driving this were covered in a number of sessions on broadband in the core and access networks. To meet customer demand

it will be necessary to deploy scaleable and flexible multi-service platforms able to support existing and new services enabling customers to be pre-eminent in their own markets.

Additionally, Internet possibilities and the potential applications that cheap bandwidth bring were discussed in several papers ranging from e-commerce to interactive teletraining systems. The Internet was seen as a unique opportunity to provide customers with personalised services moving away from the product-centric approach of the past to a situation where a truly interactive relationship can be established on a mass basis.

Finally a little crystal ball gazing was undertaken in the sessions entitled 'Futures' and 'Into the Millennium' where the findings of the CONVAIR project, a joint undertaking between European telecommunication manufacturers and operators, were considered. It is clear that the coming Millennium offers vast new opportunities not only for the citizens and enterprises of Europe but also for the vendors and operators who are able to fully meet their customer's needs; however, none of this will be possible without a suitable infrastructure to drive societal development.

Congress Awards

Each year FITCE gives awards to two papers at the Congress as voted for by delegates. This year's FITCE awards were as follows:

- *Best Paper Award*: 'Convergence and Divergence in Business Communications' by Hans Meijer and Marcel Geraads (The Netherlands); and
- *Best Young Presenter Award*: 'Data Networks Integration' by Julien de Praetere, Philippe Maricau and Marc van Droogenbroeck (Belgium) (Julien de Praetere was the presenter).



Hans Meijer (centre), receiving the Best Article Award. Also pictured are Prof. John Griffiths (left), chairman of the Technical Committee, and Dipl.-Ing. Guntram Kraus, President of FITCE

In addition, awards were made to two further papers:

- *IBTE Journal Award*: 'The Euro Side of ATM—A Business Case!' by Claire Ahern (Ireland); and
- *FITCE Forum Award*: 'The Inter-Network Call Accounting (INCA) Project' by Eden Phillips and Paul Edwards (UK).

Social Aspects

One of the great strengths of FITCE is the opportunity it yields to meet colleagues from across Europe and to discuss, in an informal atmosphere, current problems, issues and experiences.

In support of this, technical visits were organised to various sponsoring manufacturers and to BT Laboratories at Martlesham Heath. These visits were extremely stimulating and enabled participants to see technology at the cutting-edge of development.

Also, several evening events were organised including a reception at the British Museum, kindly sponsored by BT, and culminating in the Gala Dinner on the Friday evening.

Congress 1999

The 38th European Telecommunications Congress will be held in Utrecht in the Netherlands on

24-28 August 1999. The Congress theme will be 'Networking the Future' and further information is available on the FITCE Internet Web site, <http://www.fitce.org>.

Conclusion

Although the emphasis of the Congress has moved away from the purely technical, the telecommunications industry is, and will remain, an area where technology has a major impact. It was apparent that there are still lively debates around the merits of various technical solutions typified by the suitability of asynchronous transfer mode (ATM) versus Internet protocol (IP) to handle large volumes of mixed broadband traffic, with staunch champions on both sides.

The marketplace will resolve these debates at the end of the day and the critical issues will come down to the rapidly changing and evolving relationships between vendors, operators, customers and society.

Again and again, the theme of change came through in an environment where technology now gives multiple choices rather than one clear way forward. Relationships are the important thing and a paradox was put forward by John Griffiths in his closing speech. John observed that to succeed in a competitive environment operators must delight their customers, to maintain and grow market share they must fight their competitors vigorously, but in many cases for incumbent operators, such as BT, who carry large volumes of interconnect traffic, their competitors are probably their biggest customers!

Telecommunications is converging with computing and media industries, opening opportunities for traditional operators and newcomers to the market alike. This Congress explored the many technical, commercial and social aspects of this change and gave a fascinating look at the many challenges and potential rewards the future holds for us all.

Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his regular column in the Journal. In this issue he looks at the medical world from a telecommunications point-of-view.

The Wrong Diagnosis

Over the past few weeks, I have been inundated by e-mails and World Wide Web promotions advertising drugs, medication and medical services on-line from Mexico and the United States. No need to see a doctor, no need for a prescription, no cross-checks and no safeguards, just do your own self-diagnosis, select your cocktail of medication, and give your credit card number. Almost anything you ask for is delivered to your home within a few days. In most respects this has to be a dangerous step in the wrong direction. A no-holds-barred world of any drugs and medications you choose has the potential to create a massive problem for the future.

Unregulated and uncontrolled drugs and medication on-line probably presents a far bigger threat than that posed by the discovery of antibiotics and the concomitant decision of administrations to cut back on hospital cleaning. Running dirty and fixing direct- and cross-infection later by dosing patients with antibiotics looked cheaper at the time, but as we now know, it was an uneducated decision and a foolish step that is now costing us dearly. With far more managers and administrators than medics, we can safely assume that health care is in great difficulty everywhere.


No nation can cope with the combination of increasing customer/patient demand driven by the demographics of age, a growing service expectation hyped by technological advance and media reporting, all exacerbated by exponentially rising costs. No amount of taxation moneys can correct these trends because the operational model is both old and wrong. Like it or not, health care is driven by the economics of

availability and doctors, nurses, pharmacists, specialists are apparently in short supply. The only thing that appears to be universally available in abundance is administrators and managers. In some parts of the United States this has reached epidemic proportions with one administrator per patient being treated. No other business could survive such a skill imbalance, except perhaps government.

In the UK, antiquated systems see the generation of upwards of 60 pieces of paper for just one visit to an outpatients department. At least 25 per cent of all hospital records are lost at any one time, resulting in multiple X-rays, examinations, and expensive (and sometimes life-threatening) mistakes. The poor handwriting and errors made during the drug prescription process introduces more waste at the point of dispensing. And extra telephone calls and cross-checks exceed 30 per cent of the pharmacist's time.

Contrary to common perception, the UK has more doctors and nurses per capita than many nations. The real problem is nursing staff spend less than 15 per cent of their time with patients, while over 25 per cent is devoted to filling in forms, 25 per cent looking for records, and 25 per cent physically collecting and transporting information from one location to another. Similarly, many general practitioners and specialists spend precious time on administration, as well as being councillors and social workers rather than practising medicine. Operating theatres and specialist equipment costing millions are only in use from 08.00 hours to 17.00 hours on weekdays only. No commercial operation could survive such waste.

A radical change in the processes involved in patient care are long overdue. Information technology (IT) can cure much of this by introducing unified electronic record systems, with the automatic and simultaneously scheduling of appointments and operations across many regional facilities, cross-checks on diagnosis, prescription and aftercare. There will



also have to be a prioritisation on the basis of patient behaviour. Can any system afford to continually try to repair people who persistently try to kill themselves through tobacco, alcohol and drug abuse?

In the future, the maxim may well be—patient heal thyself. With insufficient resources each of us will have to become more responsible. Preventative medicine is far more effective than any curative process. So what can IT do in this arena? The next big advance will be the data mining of patient records. Here lies a wealth of information on the symptoms, diagnosis, treatment and outcome of millions of cases histories. Once in a database the possibilities are endless. No doctor can know everything, and getting to a specialist in time is now more often an opportunity way beyond the average. PC-based doctors already equal their human counterparts in some aspects of diagnosis. A marked feature of PC-patient interaction is the level of honesty in response to detailed questions. This often surpasses the equivalent human-human session and can lead to greater accuracy of diagnosis and speed of treatment.

Specialist Web sites for almost every human ailment are now available. These span the basic facts and figures on diagnosis, treatment and survival, through to support groups. Increasing numbers are advising and assisting in the supply of drugs and medicines by post from far-flung places on the planet where the laws and limitations on the use of medicines are markedly different

from the UK. The advantages and risks to the individual user and the health care system are obvious. In the UK we spend over £40 billion on our National Health Service, and an estimated £40 billion outside the system for alternative treatments and services. So far no one has dimensioned Internet-based spending, but it is growing fast. Arranging an operation in the United States or Europe via this medium is now straightforward, and there are no waiting lists. All you need is a PC, Internet access, and money.

The development of head-mounted and hand-held cameras for surgeons and doctors, remote sensory systems

worn by patients coupled into the telecommunications networks have already demonstrated great potential. Specialists can perform, participate, support and guide operations across the planet. They can conduct fetal scanning, X-ray, MRI, and other diagnostic processes remotely without leaving their home hospital. Patients with heart, breathing, kidney and other ailments can be remotely examined and advised by a paramedic. A doctor need only be brought in for the exceptional and dangerous cases. All of this is under test. Might the patients then choose a doctor in the United States where this is possible, and where money grants

instant access? Again, the advantages and dangers are very clear.

Governments may look to providing more money for future health care of a nation, but the answer lies elsewhere. Electronically networked facilities offer the only hope of maintaining current service levels, let alone improving them. The reason? A paradigm shift which will see the customer/patient in control with freedom of action well beyond and outside the national model. Internet medicine is only the start, and there is much more to come. Administrations are guilty of making the wrong decisions because they made the wrong diagnosis!

field focus

Square Pole from a Square Hole

A little bit of history was made on 17 June 1998 when the last square-section pole in the South-West was recovered from its position in Chippenham, a market town in Wiltshire. Was this the last such pole throughout BT's UK external network?

The pole began life as a traditional, round, 50 ft stout pole, originally sited by the A4 to carry the Bristol-Reading overhead route. In 1957, it was recovered in good condition, following the recovery of the overhead junction route. It was then reduced in length to 40 ft, converted to its unusual square cross-section because of the reconditioning process involved, before being reinstalled as a distribution point in, appropriately, Wood Lane, Chippenham, to serve a nearby school. Because of its size and position, partly in the road and partly on the pavement, it was expected that it would not come out of the ground without a fight. The on-site preparations were eased considerably by a very helpful traffic warden who closed the road for the critical part of the operation, and the local council had people on hand to reinstate the large hole that the pole was expected to leave behind.

A pole erection unit and a heavy-duty hydraulic pole jack were used to remove the pole from its hole. The pole

withstood the first attempt without budging, but after resetting the jack the pole gave a slight shudder and rose by an inch. As the process was repeated, the square pole continued to rise, albeit rather reluctantly. As more of the pole became visible, it was obvious that the 18 inches that had been immediately below ground level were badly decayed. Below the decay the wood was still sound. The pole finally cleared the ground, leaving a hole 7 ft deep and 18 inches square.

Because of its unusual shape, the Chippenham Town Council had placed a restriction order on it. In view of their interest in the pole, it has now been donated to the Council and so it is on the brink of yet another life when a suitable site has been located.


The BT team with the vanquished pole



book reviews

Speech Technology for Telecommunications

Edited by F. A. Westall, R. D. Johnston and A. V. Lewis



This book is one of a number in the BT Telecommunications Series that aim to cover the wide spectrum of telecommunications technologies. This edition focuses specifically on speech technology in all its myriad forms.

As distinct from others in the series, the topic for this particular edition is aimed at a very broad subject area and, with a few exceptions such as the Laureate text-to-speech system, it is not product specific. Contributions are included from a range of experts in the field, both laboratory and university based, resulting in a book which will appeal to readers with varying

levels of knowledge, background and expertise.

The book succeeds in being accessible to all and is likely to prove highly useful to students because of the many and varied areas it covers in depth within one volume, and because it represents current state-of-the-art speech-technology thinking. The reference potential of this volume is likely to be its most useful feature.

It is laid out in a conventional format with early chapters providing a broad overview and explaining much of the terminology for those new to or inexperienced in the area. This is done in preparation for a more detailed treatment of the subject in subsequent sections, where the authors go on to review the basic mechanisms of human speech and to explain the key techniques of analysis central to all speech technologies. It then moves on to 12 chapters which focus on the core enabling technologies. In subsequent chapters the book aims to provide an insight into the challenges of integrating a range of speech technologies into effective commercial solutions aimed at consumers and, not unusually, the final chapter speculates on the future of telephony.

So does the book actually add anything to the vast amount of accumulated written matter on this subject? In my opinion it does. It provides in one volume a summary of speech technology from an introduction to human speech production, hearing and analysis, through lower-rate speech coding, to the performance of existing text-to-speech synthesizers including speech recognition—beyond prompt and response basics.

The book contains contributions from around 40 individuals, each of whom is recognised as an expert in their particular subject area. It retains a genuinely upbeat feel throughout and holds the readers interest by presenting much that is both technically diverse and difficult in concept, in a progressive way. This is important in a book where technical developments mix with theory and is something that is handled well throughout.

Overall the book will appeal to experienced and inexperienced users alike and the subject is handled in a varied and interesting manner. The authors have done an excellent job both in compiling and coordinating the inputs in a wide and varied subject area and have produced a most useful book.


*Published by Chapman & Hall
ISBN 0-412-79080-7*

£80.00. pp. 564 +xii

Reviewed by Tom McGrane

1.1.98 Telecommunications In The New Era: Competing In The Single Market

by Michael Tyler and Carol Joy



This book gives an insight into the opening of the telecommunications market and how the barriers will be brought down to allow a more competitive market. It assesses the implications for the telecommunications business. The book is in a language that the ordinary person can understand, unlike some that are written for the specialist. Some countries have fully opened their markets such as the United Kingdom and New Zealand, whereas in the United States (which has an open market) any company wanting to go in is challenged by the American companies as a delaying tactic. Some countries allow a telecommunications company to come in as a joint venture with a local company.

The book consists of five parts looking at what is meant by a single market, how the new environment will work, the rules of the game in a single market (not the finer points of regula-

tion but the key elements), the scenarios that operators can adopt as they prepare for the new environment and the main conclusions.

It talks about telephone calls, the Internet, data services, interconnection, how each company charges for calls going through their network etc.

I felt that the only problem the book has is that there is no glossary including all the abbreviations and their meanings, although they are spelt out when they are first used.

I would recommend people to read this book if they are interested in the changing market within the telecommunication market.


*Published by Multiplex Press
ISBN 0-9532323-0-1*

£24.00. pp. 248 + iii

Reviewed by Richard Craig

The Complete Idiot's Guide to the Internet

by Peter Kent and Rob Young



Being a complete idiot (on the Internet) the title of this book immediately appealed to me. On receipt of the book, my first impression was: This is exactly like the Internet, very colourful, packed full of detail (it is 2.6 cm thick), and it is probably going to be impossible to find what I want in all that detail.

However, the book is actually very well structured. The four main sections headed 'The Least You Need to Know', 'There's Plenty More', 'Getting Things Done', and 'Resources' are well written in a light-hearted way. They start from the assumption that you know nothing about the

book reviews

Internet (although it does assume you know Windows 95 etc.) and discusses the topics logically, using simple everyday language.

The sections are set out to give a new Internet user enough information to get going quickly without giving exhaustive detail all at once. When you feel confident with the basics, you can find the detail in the later sections of the book. These cover, buying goods, virus protection, protecting your children from pornography, searching

news groups, creating your own web page and much more.

The structure of the book is convenient and logical for the new user, but it can be annoying to experienced users who are looking for a certain piece of information, (and find it is scattered across two or three chapters). However, the index and cross-referencing are very thorough and there are many links to useful sites throughout the book, (and copied in the Resources section).

So if you are the sort of Internet user that could find out about President Clinton's difficulties before the BBC, there is probably nothing in this book to interest you. Anyone else will almost certainly find something of use, and Internet novices will find it invaluable.

Published by Que Corporation

ISBN 0-13-786823-5

£19.99. pp.364 + xvii

Reviewed by Haydn Singleton

telecom focus

Automatic Details for 999 Services

BT has unveiled a new service to help police, fire and ambulance authorities reduce response times to emergency incidents.

Currently, call information is passed orally to the emergency authority; the BT operator passes on the caller's telephone number, but the caller is required to give details of the location.

The new information service, however, will allow details of both the calling number and the address from which a 999 call has been made to be transferred automatically to the emergency authority operator's screen. This enhancement prevents misunderstandings caused by uncertainty over the precise location, unusual spellings, panic, local accents or language difficulties, and improves call handling and vehicle despatch times by an average of 30 seconds.

Extensive trials of the enhanced information service have been carried out with the West Midlands Ambulance NHS Trust.

Barry Johns, Chief Executive of the West Midlands Ambulance Service (WMAS), said: 'We are delighted to have been involved in the trials of this new BT technology. Where speed of response is essential, the ability to save vital seconds represents a real enhancement in the quality of service delivered to the public.'

'Using the technological capability of WMAS and BT, the new facility will ensure improved service to life-threatening incidents.'

Home Office Minister George Howarth welcomed the initiative and said: 'I am pleased to be part of the launch of BT's enhanced information service for emergency calls.'

'Delivering public services in a more efficient and effective way is central to the Government's "Best Value" approach. BT's new service will be a valuable resource in assisting the emergency services to handle 999 calls more effectively and provide a quicker response to incidents.'

'I am also particularly pleased to see that the service will help identify hoax callers more quickly. Malicious calls pose a serious risk to life and waste valuable resources. The Government fully appreciates the commitment of BT to reducing this menace.'

Boom Expected in Internet Shopping

Forty per cent of all the shopping done in Europe will be transacted over the Internet within the next five years, according to research carried out for BT.

There are already 500 000 companies selling by e-commerce, and that number will rise to 8 000 000 by the year 2003. About six per cent of homes in the UK already shop on the Internet; the figure in Germany is seven per cent

and in the United States—said to be four years ahead of Europe—a quarter of all homes carry out such transactions.

Sue Davidson, BT's Director of UK Strategy and Business Development, said: 'The research shows that companies cannot ignore e-commerce and remain competitive. Despite the fact that European take-up of the Internet has been slower than in the United States, we believe pay-per-use Internet services and Internet access via television is likely to attract 3.5 million subscribers within the next four years.'

Caution Advised When Surfing the Web

When it comes to exploring the Web—and reading information on-line—Internet users should always use caution to make sure they are not unwittingly getting caught up in fraud, according to AT&T.

These days, the Internet is a haven for news and information on offering stock information or lottery draws. In addition, it seems that the same types of fraud that is committed through direct mail or telemarketing can now be found increasingly on the Internet as well.

The first thing any smart surfer should do before making a decision based on information found on the Web is to know the source. This means checking into a Web site before handing over any personal

information—especially personal financial information. This does not mean you should not surf at sites unknown, but it does mean you should have done your homework before exchanging information.

Deal only with reputable organisations, and do not judge reliability by how nice or sophisticated the website may seem.

Also, read your mail carefully. Many fraudulent offers come in the form of e-mails because the Internet makes it possible to send thousands at a relatively low cost.

Use a mail program that allows you to screen out these mass mailings, and you will spend less time with your finger on the delete key.

AT&T and BT Global Venture

AT&T and BT have announced that they will create a \$10 billion global venture to serve the complete communications needs of multinational companies and the international calling needs of individuals and businesses around the world. The venture will offer communications services of an unprecedented scale, scope and level of quality.

The venture will combine the trans-border assets and operations of each company, including their existing international networks, all of their international traffic, all of their international products for business customers—including an expanding set of Concert services—and AT&T and BT's multinational accounts in selected industry sectors.

Further, the two companies will develop an intelligent, managed Internet protocol (IP)-based global network to be implemented by the venture, its parents and their partners. This network—unparalleled in capabilities and reach—will support services such as global electronic commerce, global call centres and new Internet-based solutions to support global organisations and executives on the move.

Owned equally by AT&T and BT, the venture in its first full year of operation is expected to have rev-

enues of more than \$10 billion, growing thereafter in excess of 15 per cent a year. The venture is expected to contribute positively to the earnings of both parent companies from day one. Operating profits are expected to be around \$1 billion in the venture's first full year, growing at 15 per cent to 20 per cent a year. Capital expenditures are expected to be around \$1 billion.

Through the venture, AT&T and BT aim to be the undisputed leader in the fast-growing global communications services market. The venture, together with partners around the world, will provide an outstanding range of global services far greater than either AT&T or BT could provide alone or with their current alliances.

The venture will:

- offer customers the widest range of seamless advanced global products and services to meet their end-to-end communications needs;
- give customers unequalled levels of integrated service and support; and
- carry trans-border traffic on a greater scale and more efficiently than ever before.

BT Chairman Sir Iain Vallance said: 'The creation of this new venture will be excellent news for our customers; they know that we can—and will—deliver. By positioning ourselves at the forefront of the global information age, we will maximize opportunities to grow revenues and enhance profits for our shareholders.'

AT&T Chairman and CEO C. Michael Armstrong said: 'This announcement is a key part of the overall 'facilities-based' growth strategy AT&T has been aggressively rolling out since January. The moves we have made will enable our shareowners to benefit from growth opportunities in the local communications markets of the United States, and now, in the exploding global market.'

The venture, which will be named later, will be freestanding with its own chief executive officer and management team. Its Board of Directors will include executives from both parent companies. Sir Iain Vallance will be the venture's first chairman.

With its operational headquarters in the eastern United States, the venture will employ initially about 5000 people worldwide. It will have its own sales force to serve directly corporate customers in selected industry sectors around the world.

The venture will stimulate competition in recently liberalised markets by supporting new competitive operators around the world. Many of them will be distributors for the new venture and all will be potential customers for its carrier services. All of this will accelerate efforts to bring accounting rates down, lowering prices, stimulating international calling and creating new businesses and new jobs.

Both companies expect that, while the regulatory approval process in Europe and the United States will be thorough, completion is achievable within about 12 months.

Consultations Over Marine Radio Closure

BT has announced that it is setting up a nationwide consultation on proposals to close its commercial maritime radio services in 1999. Demand from seafarers for long, medium and short-range terrestrial radio services has dropped by 80 per cent in the past five years as satellite and cellular-based services have developed.

Fishermen, tanker operators and other users will be invited to comment on the proposals and to seek advice from BT about alternative means of communications. It is proposed that existing commercial terrestrial services will close entirely, or in stages, from the end of June, 1999.

Safety-related communications also provided by BT, such as MF broadcasts and Navtex, will continue, but under the stewardship of the Maritime and Coastguard Agency (MCA).

BT is advising customers about alternative means of communica-

tions, such as C-Sat and Mobiq, in order to provide a smooth migration to cellular and satellite services. More than 500 fishermen around the UK already use BT C-Sat, which provides secure messaging and e-mail facilities as well as being able to receive information such as weather forecasts.

Mike Wilton, BT Project Manager, said: 'Maritime radio services around the world have experienced a significant reduction in use over the past 10 years due to the advent of new technologies such as GSM and satellite communication.'

'We appreciate that the proposed closure may affect a number of customers but we will do all we can to minimise any disruption by providing world-class satellite communications at attractive rates.'

'We want to ensure a smooth transition to alternative services which already deliver much greater functionality and value to customers.'

BT Highway

BT has introduced a new mass market digital communications service which for the first time allows UK customers to surf the Internet at high speeds and use the telephone simultaneously over their existing telephone line.

BT's Highway service transforms a customer's existing telephone line into a new fast reliable and multi-functional digital Internet and multimedia connection. It emphasises ease of use, flexibility, comprehensive customer support and speed—taking the wait out of the World Wide Web.

The new service uses the latest digital access technology supplied by Marconi Communications and Ericsson, and operates over existing BT telephone lines at speeds several times faster than the fastest modem.

BT Highway offers the customer a choice of using a combination of analogue or digital connections over a single telephone line. These are:

- two digital data channels (each with a speed of 64 kbit/s per second);

- one analogue voice channel and one digital data channel; or

- two analogue voice channels.

Or the customer could combine the two digital data channels to give one line with a speed of 128 kbit/s.

This means families or small businesses can surf the Internet or send e-mails and make or receive telephone calls simultaneously. Currently, users with a single analogue line have to disconnect from the Internet to use the telephone.

Afshin Mohebbi, Managing Director of BT's Business Division, said: 'BT Highway is a major step towards mass digitalisation of the UK. Though Internet growth has been phenomenal, only 1.8 million UK households are connected. BT Highway will help bridge that gap.'

BT's aim is to connect the majority of UK households with the latest high-quality high-capacity digital communications by the year 2003.

BT Highway also means that customers can benefit from a fast and clear digital service without having to give up their existing telephone line and equipment. In addition, customers will have three numbers—one for each analogue line, and one covering both of the digital connections. In most cases one of the analogue lines can keep the customer's existing number.

Digital Interactive In-Flight Entertainment

Airlines and their passengers will soon enjoy the world's first digital interactive in-flight entertainment (DIFE) service, with the signing of a Memorandum of Understanding between two of the world's leading IFE solutions providers—BT and Sony Trans Com.

BT Airline Interactive Service (ALIS) is already at the forefront of in-flight entertainment thanks to its advanced software platform and two-way interactive capability making possible a range of enhanced entertainment, information, and communication services available to passengers

via the Skyphone network. Sony Trans Com has developed P@ssport, an advanced digital in-seat video-on-demand system, whose open architecture supports BT's ALIS system. The two companies have signed a Memorandum of Understanding, heralding the introduction of the world's fastest and most advanced in-flight entertainment service. DIFE delivers the passengers' selections much faster, with higher-resolution graphics and CD-quality audio. Interactivity is also more responsive, so passengers spend more time enjoying the entertainment and using the services, and less time waiting for the system to respond to commands.

Alam Gill, Head of Aeronautical Services, said: 'BT's new working relationship with Sony Trans Com takes IFE into new realms of sophistication. Our customers can now offer their passengers the very best in in-flight entertainment to enhance the passenger experience.'

Doug Cline, President and Chief Operating Officer of Sony Trans Com said: 'The combined efforts of BT and Sony will bring to airlines a comprehensive IFE package that offers enhanced services to passengers while strengthening internal operations via BT's Skyphone communication network.'

With an ever-increasing choice of content available covering games, entertainment, on-line shopping and information, the new digital services are expected to attract greater use by passengers, thereby increasing revenue-generating opportunities for airlines.

BT is unique in providing its airline customers a 'one-stop shop' for IFE and now DIFE. ALIS offers advanced technology, operational and reliable functionality, and customer-focused content, deliverable in minimum timescales with global maintenance and support.

European Commission Decision on International Calls

BT has welcomed the decision by the European Commission to exclude it from its continuing investigations

into the high level of accounting rates which affect the price of international calls between European countries.

A BT spokesman said that the EC's decision to concentrate its investigations on arrangements involving seven other European operators recognised the highly-competitive nature of the UK market for international telecommunications which is far more advanced than in other EC states.

BT is just one of many international telecommunications operators in the UK. Licences have been granted to 80 other companies allowing them to build their own facilities for carrying international calls, and more than 120 other companies are involved in 'reselling' international calls from within the UK.

BT has successfully negotiated some of the lowest international accounting rates in the European Union. Since 1992, it has achieved significantly lower rates on major routes. Accounting rates with United States operators have been brought down by as much as 85 per cent during the same period.

Iridium Goes Live

The Iridium providing global satellite telephone and paging went live on 1 November 1998. With Iridium telephones and pagers, people will have the ability to communicate virtually anywhere on the face of the planet—using just one telephone.

Using its constellation of 66 low earth-orbit satellites, the Iridium system provides reliable communications nearly anywhere—from ships at sea to the highest mountains and remote locations. For people travelling in urban areas in the developed world, Iridium offers a cellular roaming service featuring dual-mode telephones that can be switched to operate with terrestrial wireless services.

Whether their telephone is in satellite or cellular mode, an Iridium customer will be reachable at a single telephone number, anywhere they happen to be.

Interactive Services via Television

Digital television (DTV) is quickly emerging as a second interactive platform in Europe, competing with the Internet for consumer interactivity, and cannot be ignored by commerce, advertising and broadcasting companies.

European consumers' use of DTV is rising faster than its counterpart in the United States. According to a report by Jupiter Communications, DTV-based interactive services will reach over 19 per cent of households in the United Kingdom, 28 per cent in Sweden, and 12 per cent in France by 2002. In addition, a Jupiter/NFO survey showed that 33 per cent of UK households and 29 per cent of French households are willing to pay for interactivity on their television sets, services that many European broadcasters are expecting to offer for free.

While consumer demand exists, broadcast, commerce, and advertising companies face a new series of challenges, including the high cost of entry, possible limited revenues and the lack of experience with the platform. All players will incur some costs to develop interactive services for the TV platform, while the revenue streams to support these ventures may come from current Web budgets. In addition, companies will need to create a strategy that would work on this untested interactive medium.

Pay-As-You-Go Internet Service

Yahoo! Europe and BT have announced an alliance which will simplify the use of the Internet for the UK's web users.

The announcement was timed to precede the launch of Yahoo! Click in the UK, a combined pay-as-you-go Internet access and Web guide. The service needs no monthly subscription or registration process. Click users will be able to take advantage of straightforward access, powered by BT, linking directly to Yahoo!'s Web directory and its additional search, communication and information services. Users will be charged for

the time spent using the Internet on their telephone bill—at 1p per minute above local call rates.

John Swingewood, BT's Director of Internet and Multimedia Services said, 'Ease of access through BT Click coupled with Yahoo!'s highly successful Web guide service makes Yahoo! Click a winning combination. BT launched its innovative pay-as-you-go Internet access service to help grow the Internet market; this announcement will do much to encourage this expansion and it is a firm endorsement from Yahoo! of this pay-as-you-go concept.'

New Look for Oftel

David Edmonds, Director General of Telecommunications, has unveiled the details of the restructured Oftel. The new structure makes Oftel more flexible and better focused so that it can respond quicker to new challenges in the communications world.

In the new organisation, Oftel's 10 branch structure is replaced by a streamlined operations system built around just two operational Directorates—Regulatory Policy and Compliance—backed up by a Business Support Directorate.

The Regulatory Policy Directorate will be responsible for the development of the UK regulatory framework in order to protect the interests of consumers and to encourage greater competition. Within the Directorate a new Customers and Markets Unit will be established. This will bring together Oftel's work on market and technological development with that on analysis and tracking of customer benefits from competition. The unit will contribute to the strategic backdrop against which the work of the Directorate will be planned. The Compliance Directorate is the focal point for all Oftel's work on enforcing licence conditions and handling complaints. It comprises Oftel's consumer representation section and all staff involved in investigating complaints and enforcing licence conditions. The Compliance Directorate will also investigate complaints under the new Competition Act when this comes into force in March 2000.

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by The Institution of British Telecommunications Engineers, Post Point 2D05, The Angel Centre, 403 St. John St., London, EC1V 4PL.

Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in *British Telecommunications Engineering* or the *Structured Information Programme* unless any such statement is made specifically by the Board.

© 1998: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Annual subscriptions (including postage)—companies, universities, libraries and other bodies £60.00 UK, £65.00 overseas; private individuals £36.00 UK, £41.00 overseas.

Single copies and back numbers—companies etc. £15.00 including postage UK, £16.25 including postage overseas; private individuals £9.00 including postage UK, £10.25 including postage overseas.

Overseas customers paying by cheque should use a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London.

Orders can be paid by credit card by calling +44 (0)171 843 7626.

Orders by post should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL.

Cheques should be made payable to 'IBTE' and should be crossed '& Co'.

An index of recent back issues is on <http://www.ibte.org/ibtepubs.htm>

IBTE Members requiring back issues should contact the IBTE Office on +44 (0)171 843 7626.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 (0)171 843 7626. Fax: +44 (0)171 843 7888.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 (0)171 356 8008. Fax: +44 (0)171 843 7888.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and the Board of Editors' permission is necessary for reprinting extracts.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by The Institution of British Telecommunications Engineers for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/98 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone / Centre	Contact	Telephone Number
London	Carl Alderman	(0181) 666 5956
Martlesham Heath	Howard King	(01473) 642583
Midlands	Dave Plumb	(01203) 556034
North East	Geoff Jenkinson	(01226) 248141
North West	David White (in Manchester)	(0161) 600 2468
Northern Home Counties	Kevin Rees	0802 195219 (mobile)
Northern Ireland	David White (in Belfast)	(01232) 215354
Paris	IBTE Office	+44 171 356 8008
Scotland	Cameron Scott	0800 212993
South West	Tim Jeanes	(0117) 9205481
Southern Home Counties	John Dymott	(01202) 554001
Wales	Mike Thomas	(01492) 538103

IBTE International HelpLine: +44 171 356 8008

IBTE On-Line (Internet): <http://www.ibte.org>

(BT intranet): http://networks.intra.bt.com/ibte_nat/ibtehome.htm

BRITISH TELECOMMUNICATIONS ENGINEERING

Journal Subscription Order Form

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager
Post Point 2D05, The Angel Centre,
403 St. John St.,
London EC1V 4PL
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£60.00 UK; £65.00 Overseas**
Private individuals (to be paid for by personal cheque): **£36.00 UK; £41.00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', should be enclosed with the order. Overseas customers should pay a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London. Alternatively, customers may wish to use a local subscription agent.)

Title (Mr/Mrs/Ms etc) _____ Initials _____ Surname _____

Mailing Address _____

Postcode (Zip Code) _____ Country _____

Signature _____

Date _____

Please state the issue you wish your subscription to commence: _____
(April, July, October or January)

VAT Registration Number (zero rate): 243 5917 54 *Photocopies of this form are acceptable.*

38th European Telecommunications Congress
Utrecht, The Netherlands,
24-28 August 1999

FT
CE

CALL FOR PAPERS

NETWORKING THE FUTURE

Public Telecom Operators, Suppliers and Customers are entering a new Competitive World. Fast growth of Data, usage of Inter-/Intra-/Extranets, mobile communication and new technologies are the key and buzz words in this new environment. Bridging between and transformation from the Traditional Telecom World and the New Telecom World are absolute requirements to be successful in the market.

The FITCE Utrecht 1999 congress will create the ideal opportunity to translate "Networking the Future" for the participants into working together in the technical sessions and meeting each other in the scenic, cultural and academic city Utrecht during off-business hours.

Papers are invited exploring the technical and commercial aspects of this theme. Possible topics include, but are not restricted to:

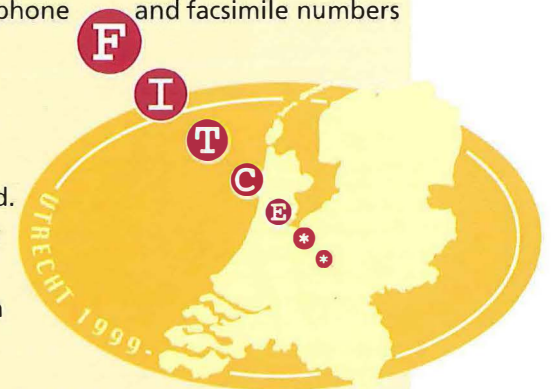
- new technologies (ATM, TCP/IP 4-6, IP Switching, Voice/Fax/Video over IP, Multimedia)
- internet as the new converged platform for electronic communication and commerce
- service development (customer driven, technology controlled)
- purchase, design and manage global networks (new suppliers, outsourcing)
- service- and businessmanagement (customer controlled, end-to-end)
- legal opportunities and bottlenecks (DG IV and/vs DG XIII)
- liberalisation, deregulation, re-regulation after 1.5 years of open market.

1999

UTRECHT

More information

- Updated information about the congress is available by visiting:
<http://www.fitce.org/fitce99>
- Authors are requested to submit an abstract of 200-400 words. Abstracts must be submitted in the English language and should be previously unpublished. Abstracts must be sent by 1 march 1999 to KIVI, Congresbureau, Postbus 30424, 2500GK Den Haag, The Netherlands. (Telephone: +31 70 391 9890; Facsimile: +31 70 391 9840; e-mail: congres@kivi.nl (plain text only)). Please include your name, date of birth, address, telephone and facsimile numbers as well as e-mail address, and affiliation.
- The abstracts will be reviewed by the International Paper Selection Committee for relevance, technical content and originality. Authors will be informed by 31 March 1999 whether their paper has been selected.
- The full text is required by **1 june 1999**. Authors may wish to consult with their respective national representative of the 'FITCE Comité de Direction' with regard to coordinating national contributions.



<http://www.fitce.org/fitce99>

Contents

Guest Editorial: Customer Service Delivery	113
Customer Service—The Key to Satisfaction and Loyalty	114
Strategic Resource Planning for Optimum Service Quality	120
Implementing Software Tools for Planning Automation	126
Micromachining Future Networks	135
Broadband Multicast on BT's Futures Testbed	139
Engendering Trust in Electronic Commerce	150
Whither Video?—Television and Telepresence	158
The Future World of Telecommunications, 2005?	163
Future Switched Network Strategy—Profiting from an Uncertain Future	172
World-Class Network Power—Could That Ever Mean Us?	180
Powering the Internet	185
Datacommunications Equipment in Telecommunications Facilities: The Need for a DC Powering Option	
A People and Information Finder for Virtual Construction Project Teams	198
37th European Telecommunications Congress	204
Journal Awards for 1997/98	207

