

VOLUME 15 PART 1 APRIL 1996

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

Process Re-engineering

*Cashless Services Replacement
Project*

*Key Technologies for the
Information Industry*

Internet



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 15 ■ PART 1 ■ APRIL 1996

Guest Editorial: Making Breakthroughs 2
Lowry Stangne

Process Management

Process Re-engineering: Measures and Analysis in BT 4
Laurie Fineman

Architecture

TRIADS: Planning Changes to BT's Operational Support Systems Infrastructure 13
Peter Skevington, Ian Videlo, John Wittgreffe, Paul Putland, Donald Sloan,
Dan Creswell and Alan Smith

Making it Happen

Cashless Services Replacement System Project 21
A network perspective
Bill Hobbs

Telecommunications in the 21st Century

The Future Isn't What It Used To Be 28
Bruce Bond

Education for Changing Times 32
Chris Fowler, Terry Mayes and Bernard Bowles

The Information Industry and its Key Technologies

Key Technologies for the Information Industry 38
Bonnie Ralph

Internet 39
Paul Jenkins

Distributed Processing—Managing the Future 46
David Freestone, Tony Richardson and Ben Whittle

TINA—A Collaborative Way Forward 54
Tom Rowbotham and Martin Yates

Software Agent Technology 59
Robin Smith

Cover picture:
Checking aerials at Portishead
Radio Station (Highbridge).
BT Corporate Pictures; A BT
Photograph.

Board of Editors

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

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

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

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

The Information Industry and its Key Technologies (continued)

Navigating On-line Service Environments Jonathan Legh-Smith	66
From Books to Bytes—Managing Information in the Information Age Keith Preston	72
Security in the Information Age Chris Gibbings	78
The Information Needs of Network Communities Andrew Hockley	85
Interactive Visualisation and Virtual Environments on the Internet Graham Walker, Jason Morphett, Marco Fauth and Paul Rea	91

Moving to a Paperless Quality Management System Aardvark to Zulu Roderick Macmillan	100
---	-----

Regular Features

pcochrane@btlabs	105
Book Reviews	106
BT News	107
Industry News	109
IBTE Contacts	Inside Back Cover

Theme Editors

Process Management	Lowry Stanage	Architecture	Andy Valdar	Making it Happen	Alan Bealby
Telecommunications in the 21st Century	Peter Cochrane and Bill Whyte				
The Information Industry and its Key Technologies	Bonnie Ralph				

Lowry Stanage

Making Breakthroughs



Breakthrough process design takes courage, not just for the individuals involved but also for the leaders and finally the whole company that has to back them

Making breakthroughs and quickly capitalising on them in the market place have become de rigueur for industry—and none more so than UK telecommunications where the fiercest competition in the world continues to test our mettle. For this reason, it is perhaps timely to consider various forms of process management and evaluate their prospects in this vital emerging field. Such has been the interest that a new national institution has been formed aimed at rapidly exploiting the learning and opportunities across the whole spectrum of industry and commerce. My few words in this introductory piece herald a series of articles that form a theme over the next two years in the *Journal* where many aspects of process management, re-engineering, benchmarking etc. will be explored.

Processes have been around since people realised there is considerable efficiency to be gained through systematising any volume activity. Although pioneering credit goes to Henry Ford, in some respects manufacturing/production industries continue to provide much of practical ingenuity. With the advent of sophisticated computer packages, there is now plenty of scope for everyone to bring a spark of originality to their local activities and, in many instances, see their best ideas adopted nationally. Trying to obtain a single national process has always proved difficult, and sometimes just impractical. However, there is increasing recognition of the benefits of this approach although one of the key issues is the need for a slick process to be in place to enable good ideas and components to be quickly identified and developed for national implementation.

Of all the processes where competitive pressures are most keenly felt, service delivery processes,

including the introduction of new products and services, are at the cutting edge. It is staggering to note that in the finance sector some companies can modify national processes and their associated systems in 48 hours and introduce completely new products in a few weeks, complete with national training. Telecommunications has some way to go to match this, but is more than capable of meeting the challenge.

Quite unlike Total Quality, process management has deliberately not been rolled out in a top down and, its critics would say, sheep-dip fashion. This may be partly due to the perceived need to explore a wide range of approaches and support tools before settling on a preferred methodology for the industry. Stumbles have occurred on the way and it would be true to say that the drive and commitment demonstrated by process leaders have been major determinants of success. Process is a key part of the European Quality Award model; in fact it attracts the lion's share of the 'enabler' criteria points and thus is pivotal for companies such as our own that aim to win that award.

Whatever terminology is used to describe the mechanism for improving processes, the need to continue down this path is now regarded as fundamental to the future success and survival of any medium or large company. The race for breakthrough solutions to some intractable manually based processes is hotting up. Harnessing the best ideas from any company and any employee is recognised as a key issue. That has led to our active participation in numerous approaches and techniques, including benchmarking. Benchmarking of whole and, quite usually, part processes is increasingly seen as of high value now that the

phase of industrial tourism is drawing to a close. To be most effective, it is recognised that information has to be swapped with data in areas where we already have a lead position.

Moving from pedestrian, but often vital, improvements known as *continuous improvement*, to breakthrough process design, often referred to as *re-engineering* or *radical change*, remains one of our biggest challenges. Our company is tentatively entering this field and often has to catch up with others who strike gold early on. Breakthrough process design takes courage, not just for the individuals involved but also for the leaders and finally the whole company that has to back them. In some cases, breakthroughs might require the company to rethink its strategy and investment programme. This requires clear thinking and a responsive senior management. The evidence within BT is good, as the need to create a culture of learning is well appreciated.

The key issue for all process work, which inevitably results in some form of change project or programme, is the need for successful implementa-

tion. All too often the failure of process re-engineering projects is a consequence, not of any flawed concept or proposition, but of an organisation's inability to implement successfully. This is because most, if not all, projects involve the need for people to do something different. Successful change requires the understanding of the 'soft' issues. Indeed, experience in BT is very much that 'the soft stuff is the hard stuff'.

This series of articles opens with a broad overview of process management biased towards measurement and analysis (see p. 4). Later issues cover the design of customer-facing processes, the selection of software support tools, advanced decision software aimed at guaranteed on-demand national customer services, benchmarking processes and costs, and, perhaps most vitally, the mechanism for gate-keeping the industrial roll-out of new process designs.

Lowry Stanage

Director, Process Control and Measurement
BT Networks and Systems

Laurie Fineman

Process Re-engineering: Measures and Analysis in BT

This article focuses on steps taken by BT in over seven years of experience with process management. It considers learning points and areas of success, and provides a unique insight towards a process re-engineering vision by discussing practical steps being taken to computerise key elements of design, activity based management (modelling), process re-engineering, decision making and instant measurement and analysis.

General

Process management, whether it be dramatically improving existing processes or re-engineering them from scratch, is fast becoming a way of life in major enterprises across the world. In some cases, emphasis has been placed on selecting key company activities that are customer facing, while in others even bigger challenges have been mounted by re-engineering the whole business. The first approach can become a driver for the second; however, the scope and objectives of change in a company have to be carefully considered before starting on such a fundamental adventure.

Some industries appear to have adopted process management more readily than others. Not surprisingly, manufacturers of modern technology equipment were early on the scene; benchmark examples include Motorola and Xerox. In both cases, process activities have followed on from an extensive period of internal cultural change in the direction of total quality. While there is much to be learnt from these companies, it is in the financial services sector that the most unexpected and perhaps radical changes have been made. A good example is the National and Provincial Building Society who focused on customer-facing processes to the extent that they can now make changes to nationally available financial products within 48 hours and introduce new ones, including software upgrades to their national computer systems, within 6 weeks. Rapid launch procedures are quickly becoming a vital commercial determinant in many commodity driven markets, including, and perhaps particularly, the telecommunications market.

Several practical points have been absorbed into the chronicle of BT business process-improvement experience. In the early 1990s, there was a search for the holy grail of the best methodology; now there is a realisation that it is best to construct the methodology to suit the problem being tackled. This can now be done by using one of a growing repertoire of process tools and proven methods of each stage of the improvement programme.

Gaining the authority to take a fundamental look at key business processes right across the company has taken some time to gain ground, even though nominal end-to-end process ownership has been around for some years. The advent of BT's corporate BreakoutT project brought this into focus and has, just by the existence of a centrally-orchestrated team comprising directors and senior managers, had a significant impact on change culture within the company. Projects such as dial-up private circuits repair, that have been attempted many times in the past, have now found the necessary backing to operate quickly, concurrently and effectively with up to 100 specialists in virtually every major department of the company.

One issue that has been circulating for five years is the choice between component-led and end-to-end owned process-improvement activity. Latest experience is beginning to demonstrate that both have significant added value if used as drivers, at the appropriate time in the improvement cycle. Put simply, this would mean that an end-to-end process owner would be responsible for defining the overall objective for a process while keeping a grip on the

Process management and re-engineering are about taking advantage of the latest ideas and opportunities to move significantly, and sometimes dramatically, ahead of other players.

component changes. Within the development phase, the component owner will advise on the best approach to upgrade one or more systems and related local task level procedures to achieve the overall end-to-end objective. Both owners will then be in a position to make a valued and effective contribution. However, the component owner will be very conscious of the spill-over effect of proposed changes on other processes, be they system-defined or mainly task-level work carried out in the office or field.

Most process teams are able to redesign a process successfully to reveal substantial improvement opportunities. Early and effective communications to all those affected has had mixed success and requires extensive planning in its own right. Perhaps even more significant is the failure to introduce project management controls at an early stage. Delivery techniques are often overlooked or disregarded in the excitement to finish the project, often leading to disappointment.

Process management and re-engineering are therefore about taking advantage of the latest ideas and opportunities to move significantly, and sometimes dramatically, ahead of other players. Where it is possible to do this by exploiting previous infrastructure investment so much the better. Examples in BT's case might be the ubiquitous nature of the network and world-beating systems such as customer service systems (CSS) handling real-time customer information ranging from requests to services and repairs through to billing.

Benchmarking has become another activity for gaining and maintaining up-to-date knowledge of precisely what others are doing. This requires astute information gathering and exchange across the globe in the same and other industries. The process components of other industries have become an important source; the market place of the same and emerging similar industries will be another. In the early days of

benchmarking, people focused on comparison of overall targets and measures. While this is a good starting point we have come to realise that the task-level procedures developed by others in any industry are even more useful as they open doors to fresh thinking and ideas. This can be compared with the early days of linking physics with chemistry, biology etc.

Having gained clarity of the agreed target for a process being re-engineered and various benchmarked achievements and detailed components for the new process, it is time to get to work. In the early days of process improvement, teams were selected from people working and managing the current process. These people have the advantage of knowing all the problems but often have difficulty in breaking the mould and perhaps vested interests in the shape of the emerging process. In many cases, teams spent an enormous amount of time capturing the existing process rather than moving swiftly onto fresh ground. Arising directly from BT's BreakoutT project came the concept of learning laboratories, in which those currently performing all customer-facing tasks of a given process are brought together. They are invited to design the new process directly in a series of short workshops. Details from the workshops are captured on a database and quickly played back to the mixed disciplined team for checking. A form of iteration takes place that enables the team to move forward confidently to alpha trials.

As part of the learning laboratory activity, process specialists are able to use purpose-built process development software to evaluate the time and cost of the emerging process. In addition, there is a focus on selecting a set of key activity points along the process that are critical to its success. These measurement points are entered into new software, called *online process analysis* (OPA), developed by BT Network and Systems for this specific purpose. Essentially, OPA extracts data at key points as a job

progresses along a process, formats the information and graphically displays it to line managers and those working on the process for instant action, if problems occur. Similar problems occurring on numerous jobs can then be collated and presented for discussion leading to modifications of the basic process that can be immediately applied throughout the company. See later section on 'Analysis'.

Once the new process begins to firm up, it will be opportune to capture task-level procedures in national documentation; called *ISIS* in BT. By doing this concurrently with the process design, it will be easy to check out the detail and be confident *ISIS* is based on facts that can be replicated throughout the country. The move to quality management systems is then a simple one. BreakoutT also took a fundamental look at the need to relate processes, data and systems. This led to a new operations support systems (OSS) framework that focuses on data, information and functional systems in addition to business processes. Three key elements are selling and billing, maintaining service quality and managing portfolio and platform. As will be seen in articles in future editions of the *Journal*, OSS introduces a practical range of tools that enables generic business processes to relate in a practical way with logical groupings of computer systems throughout the company. The ideas emerging can be expected by the late-1990s to influence industry at large as well as external customers and suppliers.

Looking towards the late 1990s reveals some fascinating competitive opportunities for the company as it exploits its strong cultural, system and network infrastructure capabilities through process-management techniques. Among the most exciting projects, there is one being carried out by the department of Advanced Applications and Technology at Martlesham Heath called ADEPT. The project is partly funded by the Department of Trade and Industry in

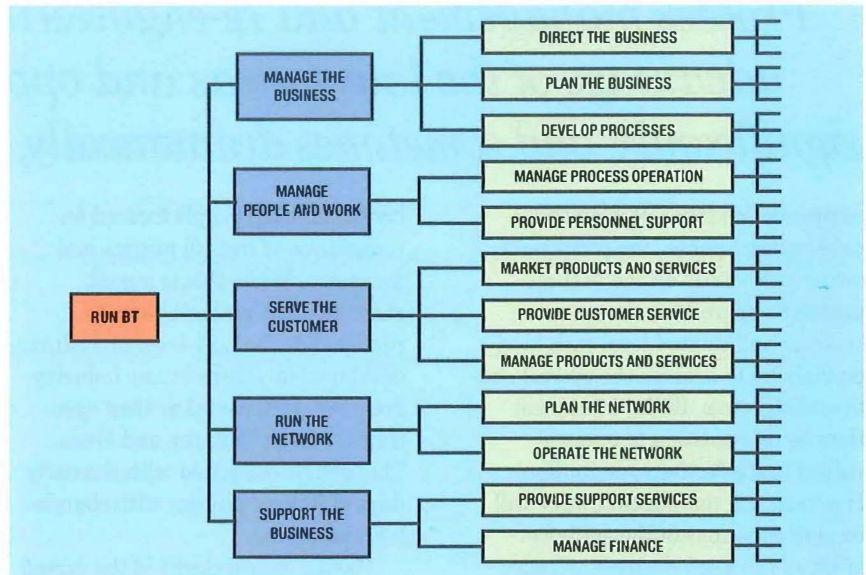
Figure 1—Business activity structure and decomposition

association with a number of universities. A new software negotiating language is being developed that enables agreements to be struck between computers, working through computerised agents, by bidding for resources and skills at specified places and times to carry out a process activity, usually for a paying customer. ADEPT will probably open the door for a number of processes to be re-engineered—see later section on ‘Research’.

Approach

Process management was formally introduced in BT in 1990 as a fundamental part of a major reorganisation project called *Sovereign*. Its history within BT can be easily traced back to 1988 when a joint study contract was placed with the New York telephone company NYNEX. Both companies had recognised the potential value of process management. In particular, they considered it was time service industries caught up with manufacturing companies who, since the early part of the century, recognised the competitive value of streamlining their basic productive activities.

BT and NYNEX approached the subject by identifying all the processes in the company and placed them in a large reference model. The model took on two forms, one part focused on how processes are arranged today and the other identified the future processes. All the basic functions within the model were arranged in a *decomposition chart*. This simply placed ‘Run BT’ at the top layer of the model and then decomposed it downwards towards the precise tasks being performed. At the time, five layers of detail were produced (see Figure 1) for an extract taken from the model. However, this later proved to be some three or four layers away from the tasks being performed by people in the field. Other difficulties arose with keeping the model up to date with the vision of the company, trying to achieve



improvement on all processes simultaneously and recognising that processes actually migrate in time or can even be suspended. For all these reasons, use of the model has been partly replaced by other approaches.

Among the useful concepts retained from this early approach to process management was that concerning the definition of different types of process.

BT’s project *Sovereign* rebuilt the business by placing external customers at the top of the organisation chart. This led to an inverted triangle chart of the business and enabled end-to-end processes to be identified that directly focus on customer requests for action. These processes were named the *business* processes. Examples include provision of various services, such as new telephone lines and the repair of them. To ensure these basic processes work BT decided to name all processes that are vital to their *successful delivery* support processes. These processes can be viewed as being those that take place in anticipation of the

business processes being used. To make matters even more interesting, or complicated, BT decided to mark out those parts of business and support processes that are common to more than one process. These were called *component* processes and often focused on computer systems used for many processes and common tasks that people undertake.

Process owners were appointed to business, support and component processes. Owners were at director level and, in the case of business processes, tended to be those directors who managed operations people working directly with customers. In charting existing business processes, an approach was used that highlighted the naturally occurring organisational interfaces. At these interfaces, service level agreements (SLAs) were written following interdepartmental discussion, to record the quality of service required across the interface to meet the overall process performance targets signed off in the business plan. (See Figure 2.) SLAs also included

Figure 2—Measuring linked activities

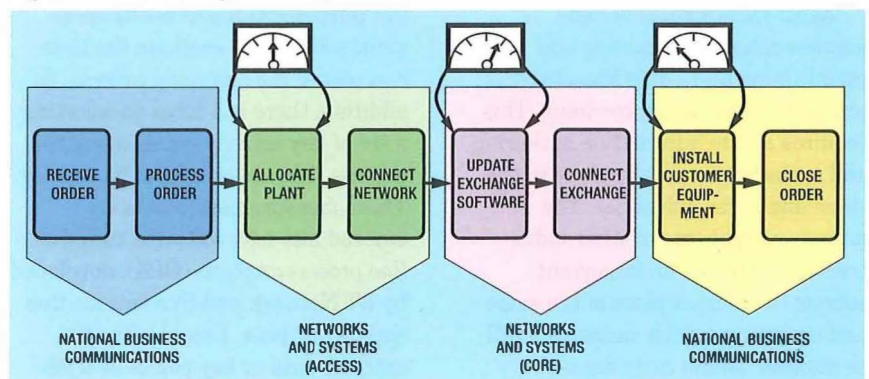


Figure 3—Measuring process success—2 Mbit/s wideband delivery


information of the measures to be used at vital points along the process and who precisely would own them. SLAs have had a mixed reception because they took considerable time to gain sign off at director level and align with the many existing measurement systems and reporting mechanisms already in place around the business. BT has thousands of imbedded measures and targets.

More recently, in late-1993, a fresh approach was taken by inviting an external consultancy to set up a new central change management function that became known as *BreakoutT*. This project took process re-engineering as one of its major workstreams and drew in high-profile directors and other managers to tackle a number of key process problems that had not achieved sufficient improvement. The Group Managing Director sponsored BreakoutT and the project received widespread support and authority to act right across the business.

Drawing on my own experience of involvement with BreakoutT it has been the mix of senior level support, ready access to any key player in the business and a capability to draw on specific skills—from research to wiring frames—that has been excellent. In addition, techniques such as learning laboratories have enabled new processes to be rapidly developed and captured in auditable documentation.

Learning laboratories have brought together the necessary skills sets required to deliver a process and encourage those involved to design new processes to meet new stringent requirements. In most cases, managers of existing processes are not directly involved in the design; however, they have a key role in setting measurement points along the process and sponsoring system upgrades.

Partly as a result of ongoing discussions with OFTEL, BT is moving towards a commercial refinement of the SLAs described above. In the near future, BT will



introduce a range of internal service provision agreements (SPAs) and complementary trading agreements. SPAs will sweep away most SLAs by focusing on products and services supplied through the major internal divisions. Trading agreements will maintain clarity of the end-to-end business process deliverables that can supply a number of products and services. Each major customer-facing division will have separate accounting structures that, through SPAs, will be able to ensure standards of service, volumes and costs for their specific BT network service requirements. BT's aim is to maintain and enhance the cohesion of its processes by continuing their refinement and automation in response to market-led requirements. The process link with financial systems is seen as a vital factor in the company's future development, see the latter paragraphs under 'Modelling' for future discussion of finance-based tools.

Improvement

One of the first and significantly improved BT processes has been 2 Mbit/s digital private circuits—MegaStream. In late-1991, it was quite clear that competition was making quick and easy inroads into this important business product. In addition, some difficulty was being experienced in gaining sign-off at director level for the SLA. Discussion

centred on the target times for one department to complete a series of tasks and the level of support required to undertake the forecast volume of orders.

A meeting was set up to bring the directors together and find a way forward. The outcome was that a small cross-company process improvement team was set up with an objective to design the process that met the customer-facing department lead-time requirements. Within three months, the team had established what was happening (the 'As Is' process) and prepared their design for a new process that could be introduced with minimal delay for system changes and retraining.

Full design of the new process (known as the 'To Be' process) was fleshed out, recorded and training workshops set up for those who would be affected. Minor changes to existing software were required; however, the team was so enthusiastic that they found temporary ways around these potential delays.

By mid-1992, the first effects of the new process could be seen by customers across the country (see Figure 3) and ongoing improvements made by the team has resulted in the provision time being reduced by over 70%. In cases where basic equipment is already installed, the team reduced the lead time from 72 to 5 days. All of this may in fact be viewed as incremental improvement activity owing to

the lack of really fundamental equipment or skill changes being necessary.

Quite separately from this activity, another team has looked at the feasibility of re-engineering BT's approach to providing and subsequently repairing lower-speed private circuits towards on-demand provision and 2 hour repair. Within the company, this radical approach is known as *dial-up private circuits*.

Modelling

A wide variety of software modelling tools are available in the market place. BT has assessed many of them to identify how they would have practical applications in the company's environment. Those attracting most attraction have PC-based features with excellent colour graphics, exploiting widely used databases with easy import/export facilities. An example would be Windows-based packages as this introduces low start-up cost and a wide range of data capture/analysis features including such things as information flow diagrams, process mapping, data flow diagrams, and time/cost overlays. Information flow diagrams record all the source/recipients of information affecting the process under study. This can be used as a starting point for detailed process study as the data will provide a useful insight into the underlying process.

Next, process mapping creates a picture of the components of a process as it threads its way across a company or organisation. This process is known as the *end-to-end* process and has particular relevance when applied to processes that include external customers and suppliers. The power of process mapping comes home when the process is drawn in such a way that it depicts the different departments through which it passes. It is quite realistic to draw a process at any level of detail at which a study is required to take place. Modern process software enables this by allowing the operator

to concatenate process detail—if necessary down to task level. Space does not permit a description of how BT is attempting to exploit all types of facilities and in many cases link them together.

Advanced software packages include 'As Is' charts, choices of 'To Be' process maps and 'To Be Implemented' capabilities. 'As Is' process charts capture the current most common successful route taken through a process. In the case of a 'To Be' process, a vision of the long-term requirement of a process is drawn, usually at a relatively high, and therefore simplified, level. The 'To Be Implemented' process shows what is selected for implementation at this time.

As confidence has grown, BT has been active in trying out various tools to find which combinations enable the company to reduce the time to capture any given process and then analyse it. There was perhaps a tendency to 'run before learning how to walk' through basic processes. Early adventures with software demonstrated that most of the available tools were off-shoots of existing computer-industry systems development tools that had little in common with direct applications in live operational areas. Process teams like the freedom given by straightforward tools such as Post-it-Notes on white boards that encourage changes to take place in the heat of wide-ranging debate.

However, once a process team has exhausted one round of discussion, the process can be captured and analysed off-line to throw up new questions and possibilities in preparation for the next meeting. Overnight analysis that delivers to a hard-working team the following morning the fruits of their effort can be instructive and lead to innovative ideas.

For some time, a part of BT has worked with GTE in Dallas, exploring how PC-based tools can be exploited and functionally improved. Working with another organisation in the same industry has enabled

similar problems to be discussed although both companies understand that there is a limit to the learning if discussion is kept within the confines of telecommunications.

With the growing complexity of OSS, it has become vital for the company to use sophisticated support tools that enable systems, data and information to be brought together for synthesis and analysis. As systems are being enhanced at differing rates, together with new ones being developed for fresh applications, planners and architects are faced with the bewildering problem of synthesising selected elements to meet new requirements, especially for emerging products and services. TRIADS† (telecommunications requirements impact decision support) is a response to this problem that has introduced a world-class on-line dynamically-maintained access capability for disparate areas of the business. Typically, TRIADS enables planners to bring together data, processes and information flows to build current intermediate and speculative models, complete with a drill-down capability.

As will be seen in the reference article, TRIADS provides a remote access facility, together with a data transposition for local off-line analysis by project teams. As TRIADS continues to be developed, it will be directly linked with live systems that dramatically enhance the integrity of the data being processed.

Systems such as TRIADS have a natural counterpart with on-line operations process tools discussed above.

Activity-Based Management

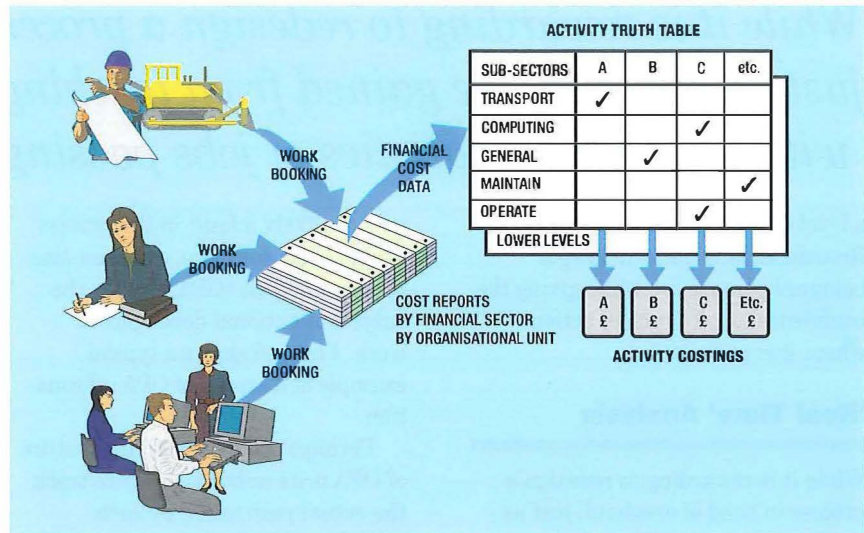
Although process modelling and improvement methodologies are vital

† SKEVINGTON, PETER, *et. al.* TRIADS: Planning Changes to BT's Operational Support Systems Infrastructure. *Br. Telecommun. Eng.*, April 1996, 15 (this issue).

Figure 4—Producing activity costs

tools once a decision has been taken to improve a chosen process, there is a clear need to find a reliable mechanism for selecting key processes. With this in mind, BT has developed an activity-based model that captures its engineering cost structure. This model has been built by taking costs spent on over 200 different basic work-related activities and assigning them to named cost sub-sectors (see Figure 4). As there are far more than 200 different tasks against which work is booked, BT has made decisions about the grouping of tasks into sub-sector activities (see activity truth table in the figure). For any group of activities that can be associated with a process, decisions are made in the truth table of precisely which tasks will be included. Where activities contain cost elements that are inappropriate to a particular process, an analysis is first carried out at task level and then adjustments made to the activity cost definitions. As can be seen in Figure 5, the selected activity costs are then grouped together to form process costs.

It is at this point the vital analysis takes place. Intensive study takes place to assess the cost structure of similar processes in the same or other industries to understand the underlying productivity achieved within their processes. This is naturally going to be easier if comparative statistical data is publicly available. In telecommunications, this happens to be the case, in as far as detailed publicly available financial information of the US telecommunications industry is



Activity cost lessons learnt

- Using surveys to apportion financial report costs to activities suffers from:
 - being based upon anecdotal rather than factual data,
 - apportionments that only last until next financial report,
 - a need for future costly re-surveys to keep up to date.
- Using basic-level financial data to create activity costing via ‘truth tables’ benefits from:
 - a consistent approach that can be easily verified,
 - being robust and future proof against organisational change,
 - being easy to change/correct in the light of experience.

published regularly by the Federal Communications Commission. By gaining agreements to swap data and process techniques, further clarity of the best area for focused improvement activity can be readily achieved.

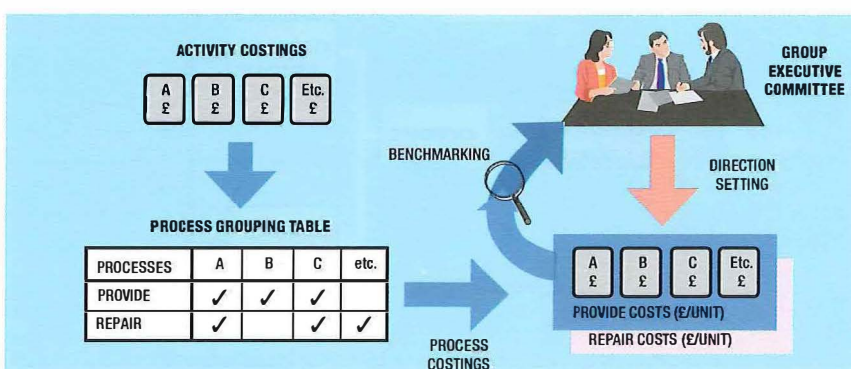
With large models, the main problem is keeping them up to date and relevant to the problem in hand. To this end, BT has quickly recognised that apportionment of cost data is easy to undertake but difficult to justify, especially when those carrying out the apportionment have moved on. A better approach is to create a truth table that draws upon

a mix of basic-level financial data that can be regularly tested for reasonableness (see panel for a summary).

Buy-in from senior management is vital to gain the necessary support and resources to carry out process design, system development and introduce new processes in a consistent manner throughout the company—in BT’s case across the whole UK as a minimum. Benchmark data along with clear activity-based costings that reveal productivity comparators appears likely to make the necessary impact.

As cost information becomes increasingly mechanised and visible at the time resources are used, the speed of process cost assessment can be expected to increase. This will encourage operational quality measures and results to be brought into the picture along with volume drivers to gain a balanced picture of processes. It is becoming clear that there are limits to reducing costs in some areas, especially when some other processes go untouched and could stretch a company well beyond

Figure 5—Activity cost improvement



While it is rewarding to redesign a process in need of overhaul, just as much can be gained from tracking the progress together with completion statistics of jobs passing through the process.

a level of commercial viability. Benchmarking helps to keep a balanced perspective while giving the confidence to take radical action where it is necessary.

‘Real Time’ Analysis

While it is rewarding to redesign a process in need of overhaul, just as much can be gained from tracking the progress together with completion statistics of jobs passing through the process. If this can be achieved in real time, then the rewards might be even greater. In this regard, the idea was to provide on-line information for those operating the process and for the managers directly responsible for the results.

Demonstration software was built that could gain access to existing databases holding current information about selected jobs either in progress or recently completed. In the first case, it was relatively easy to identify those jobs that had already failed to meet certain time milestones. Unless additional effort was put in place to rescue them, they were doomed to failure and BT would be paying compensation to the customer. To ensure the information could be read at a glance, a colour graphics package was developed that depicted each set of activities—component processes—in separate boxes. Different colours were used to represent the status of each box. For example, a red box indicates that the job is ‘dead-on-arrival’ with the next person down the chain, while green means the job is within target time at that stage and orange reflects the job is in jeopardy. Boxes were linked together to show precisely how the activity dependencies related to one another.

By drawing together statistics of many jobs, it is possible to identify those parts of the process that are regularly failing to meeting their target. The cause of the failure can be locally analysed to reveal issues such as the availability of sufficient resource, perhaps operator efficiency

or most likely a fault in the process design. This software, called *on-line process analysis* (OPA), is now the subject of national development work. Figure 6 shows a typical example of a screen of OPA information.

Perhaps the most striking feature of OPA is its in-built ability to track the actual path taken by each individual job through a process. This quickly reveals inherent problems within a process. Most processes are drawn to reflect a success model of the required process. This avoids showing the locally developed rescue techniques to correct its unforeseen failings. By using OPA, process teams can analyse what is actually happening within a process to enable either minor corrections to be made or call for major surgery. Once the key processes of the company have OPA in place, it will be relatively easy for new process improvement teams to move swiftly to the ‘As Is’ stage. As mentioned earlier in this article, the tendency is for teams to spend an inordinate amount of time at this stage—3 to 6 months was not uncommon in the early days.

Opportunities for using OPA as a driver for quickly testing the linkage

of existing or new components of a process are in mind for the future. In an increasingly competitive world, process tools will be vital if the company is to become successful in the global market place.

Research

As computer processors take giant steps in power and their speed and capabilities are exploited, new opportunities for process management are expanding dramatically. At BT’s Advanced Applications and Technology department, a three year project is underway in collaboration with Government, industry and a number of universities. The project, called *ADEPT* (advance decision environment for process tasks), is aimed at creating an environment that will provide a concurrent engineering capability to exploit distributed and varied management information systems on a common platform.

The project will explore many directions; however, for the purposes of this article, just one will be described. A recurring management problem is that of collecting data from many differently constructed

Figure 6—On-line process analysis

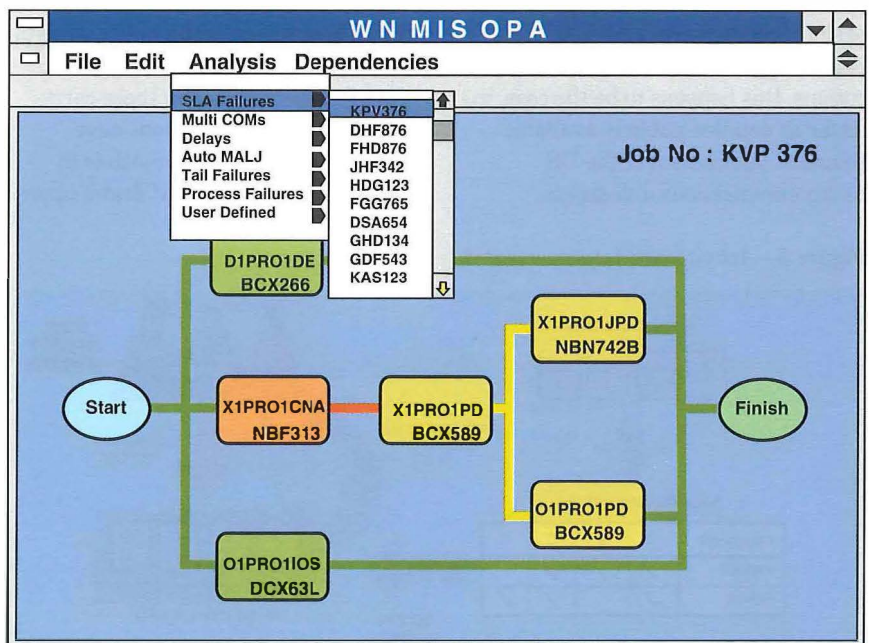


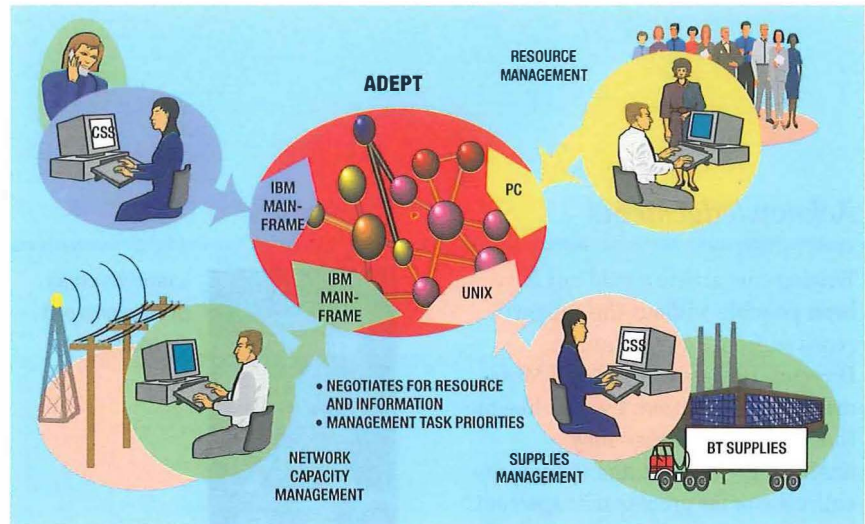
Figure 7—Front-end information

databases in order to make an informed decision. In the case of processes that directly service a customer request for quick and decisive action, this information is vital, provided it can be relied on for accuracy and is available on-demand. As can be seen in Figure 7, one possible application for ADEPT is to send out an electronic request for data concerning, say, the provision of a new telephone line.

A sales representative speaking directly to the customer interprets the customer's request for service. While the customer is holding on the line, the sales person triggers a request via the ADEPT computer agent. This local agent is linked to other agents anywhere in the country that have access to all the necessary databases holding vital data affecting such things as the availability of transmission plant, suitable equipment at the telephone exchange(s), people to do the job at every part of the process and the precise time at which the work can be booked and carried out. These computer agents literally negotiate with each other using a specially developed negotiating language to ensure that all the relevant information is collated concurrently, evaluated and acted upon.

To deliver this capability, the ADEPT project is focusing on three main streams of development: information infrastructure, information management and information presentation. This requires researchers to investigate what they call *information fusion computerised negotiating agents, open distributed processing, heterogeneous information sources and shared ontologies* (information models).

As can be seen, the sales person is insulated from the agents and problems associated with extracting information on previously unrelated databases. ADEPT enables responses for the customer to be achieved within seconds. Taking a customer order can therefore be done with confidence. This will result in new



customer orders being provided at the agreed time and at a price that reflects the competitive advantage gained by removing costs associated with failure, or in many cases, the high cost of success.

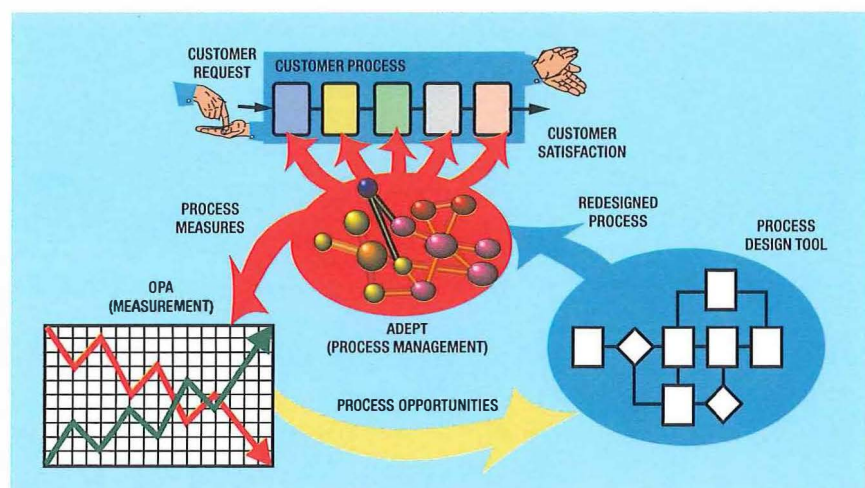
A Vision

Based on discussions in the above three sections it is possible to envisage a kind of virtuous triangle of process improvement activity. Figure 8 attempts to bring these concepts together in a manner outlined in a sketch produced by Jon Martin—see 'Acknowledgements' below. The concept portrayed takes an imaginative step forward by assuming that, by the late 1990s, it will be possible to demonstrate an almost self-perpetuating process operation and improvement cycle. By breaking into the triangle where ADEPT is in real time use, it should be possible to track the performance of a process under

study by exploiting OPA. Output from OPA will flag up difficulties and problems with a process that can be funnelled into a specially designed process modelling tool for analysis.

A skilled process team will then quickly be able to investigate the problem within the process and take in other data—such as the latest benchmarking information from the same or other industries—before deciding on the best path to take for improvement action. It is worth noting that some process tools complement their pictorial data with related text information that define the process in English. This is extremely valuable for checking back over proposed changes to a process as it will quickly reveal if vital steps are left out or are perhaps unnecessarily duplicated. Further, it is not too far-fetched to suggest that such text could form the starting point for creating the information required to establish an ISO 9001 quality management system.

Figure 8—An integrated process



Acknowledgements

Writing this article would not have been possible without the support received from Lowry Stanage, Director, Process Control and Measurement. In addition, I would like to thank all of my team—they unstintingly support and indulge my enthusiasm for process management. In particular, special thanks go to Jon Martin for his contribution to the 'Analysis' and 'Vision' sections of this article and Paul Allen who produced the illustrations. I would also like to thank Robin Smith and Mark Wiegand for permission to use information on the ADEPT project.

Please note that all comments made in this article are those of the author and do not reflect the policy or future intentions of British Telecommunications plc.

Biography



Laurie Fineman
BT Networks and
Systems

Almost unusually these days, Laurie started work for the GPO as a Trainee Technician Apprentice. An interest in telex led him to become a maintenance technician on international telex prior to attending the University of Bath. In the early-1970s, he worked in the new Network Strategy department writing the first national radio paging and radiophone network strategies. This led to competitive network scenario studies, the early BT cellular network strategy and the introduction of local area network products. By the 1980s, Laurie was working in Reading leading a planning and operations team focused on the modernisation of the network. During this time, he became fascinated by the role of processes and the capability of radical processes potentially providing an outlet for innovation and commercial leadership.

*Peter Skevington, Ian Videlo, John Wittgreffe, Paul Putland,
Donald Sloan, Dan Creswell and Alan Smith*

TRIADS: Planning Changes to BT's Operational Support Systems Infrastructure

The TRIADS system supports architectural design and migration planning for BT's operational support systems infrastructure. It provides on-line, dynamically-maintained information to architects, planners and designers working in disparate parts of the business. This article outlines the business model underpinning TRIADS and describes some of the system's functionality.

Introduction

To thrive in today's rapidly changing and highly competitive business environment, BT must be able to deliver new and improved products and services to the market place more quickly and at lower cost than ever before. In this environment, competitive advantage goes to companies who can plan and manage changes to their operations and infrastructure most effectively¹.

The scale and complexity of BT's operational support systems (OSS) infrastructure potentially produce an enormous inertia to change. BT's future success is dependent on the development of tools and techniques which will enable both the tactical changes to existing operational support systems and the strategic migration towards a more adaptable OSS infrastructure² to be effectively planned and managed.

The telecommunications requirements, impact analysis and decision support (TRIADS) system has been developed to support architectural design and migration planning for the OSS infrastructure. It has addressed the need to provide on-line, dynamically-maintained information to architects, planners and designers working in disparate parts of the business. The information is accessed through a powerful front-end application which enables data to be displayed and edited from multiple, customised viewpoints within a user-friendly graphical environment. The application adds significant value to the data by providing a range of facilities for navigation, information

processing and impact analysis, together with electronic tool support for the migration planning process.

Modelling the OSS Infrastructure

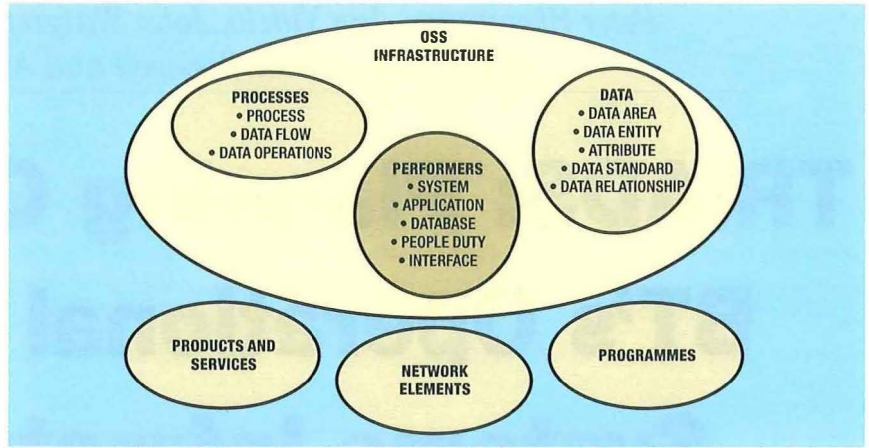
Planning and managing change to BT's OSS infrastructure is a complex, multidimensional task^{2,3}. A comprehensive model of the business is required to handle the key dimensions of the problem domain and enable the impacts of change to be analysed.

The scope of the TRIADS business model is shown schematically in Figure 1. The focus of the model is the OSS infrastructure, modelled using three major domains:

- *Performers* The physical framework of interconnected computer systems, databases and people which together provide the operational support functions for the business.
- *Processes* The logical framework of workstrings which describe the operational support processes carried out by the business.
- *Data* The framework of corporate data resources, entity relationships and data standards used by OSS performers and processes.

The model contains a detailed set of relationships between these domains. For example, it is possible to identify which performers implement which business processes (for a

Figure 1—Schematic of the TRIADS business model



particular service), which processes operate on which data entities, or which data entities are stored on which databases.

In addition, the model contains relationships to a number of other domains which either drive or are impacted by changes to the OSS infrastructure:

- *Products and services* BT's portfolio of products and services.
- *Network Elements* The switches, transmission equipment, etc. making up BT's network.
- *Programmes* The programmes and projects managing change within the business.

The limited data currently populated in these domains has been copied from external sources. Future developments may include the

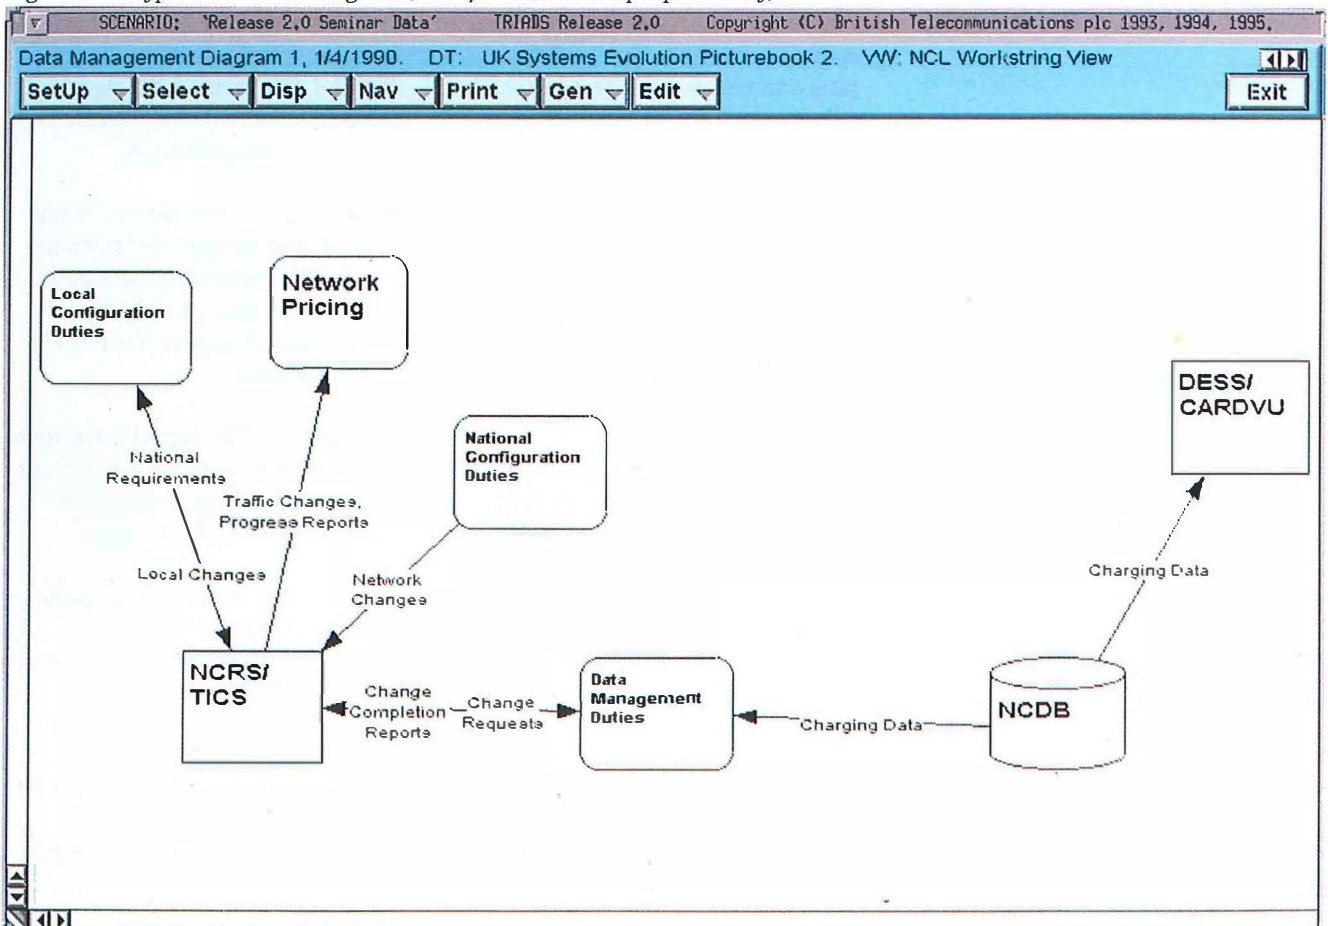
building of electronic links to primary data sources for these domains. For example, products and services data may be sourced from the products and services life-cycle management system (PSLM) and network element data for the public switched telephone network (PSTN) may be sourced from the network capability register (NCR).

The TRIADS model principally deals with relationships between versions of performers, network elements etc. TRIADS is not usually concerned with modelling deployed instances at specific sites.

Viewing the OSS Infrastructure

TRIADS information is typically viewed using topological diagrams in which each node and link is bound to an instance of one of the entity types in the business model (see Figure 2). These diagrams provide a flexible means of displaying a wide range of architectural viewpoints. For example, a workstring may be represented by a set of process nodes connected by data flows, or a systems view may be represented by a set of OSS

Figure 2—A typical TRIADS diagram (data for illustrative purposes only)



TRIADS is facilitating a major cultural change within the business from diagram-based work practices to a much more powerful model-based approach.

databases, people etc. connected by manual or electronic interfaces.

Basic attributes of the diagram contents (for example, name) are shown on the TRIADS diagram. Full details of a content instance may be obtained by 'clicking' on the appropriate node or link and displaying the instance form. The form displays both the attributes of the instance and its relationships to other instances.

Generating diagrams

TRIADS supports both manual and automatic methods for generating diagrams. The manual method enables users to create, save and retrieve diagrams via a graphical front-end to the TRIADS database. The approach adds rigour to existing diagram-based work practices (for example, those developed for drawing packages) by capturing diagrams as part of a structured model. The benefits derived include support for diagram-to-diagram navigation and impact analysis.

TRIADS diagrams are never saved to the database as images. Instead, the 'save' operation creates a set of relationships linking each instance appearing in the diagram to a diagram header. Sufficient information is stored to enable the TRIADS application to reconstruct the diagram each time it is 'loaded' from the database. This approach ensures that any changes to the diagram contents are reflected in the diagram whenever it is loaded. For example, checks are applied during diagram loading to ensure that the date range associated with each content instance is consistent with the diagram date. Inconsistencies are indicated by 'greying out' of nodes and links.

Automatic diagram-generation methods go much further than the manual method, enabling users to auto-generate transient, customised viewpoints from raw data in the underlying model. Users are increasingly recognising the benefits of this approach and the emphasis of data population is gradually shifting away

from the recording of diagrams towards the population of fundamental relationships. TRIADS is thus facilitating a major cultural change within the business from diagram-based work practices to a much more powerful model-based approach.

TRIADS provides facilities for automatically generating both standard and fully-customised viewpoints. Standard, commonly required viewpoints are supported by dedicated tools which provide specialised display and editing functionality⁴. Supported viewpoints include:

- *Decomposition hierarchy*—a viewpoint showing the parent/child hierarchy relationships for a selected instance.
- *Instance connectivity diagram*—a viewpoint showing all information flows for a selected instance. The diagram is 'cartwheel-like' in appearance, with the instance under investigation appearing as the hub, the information flows appearing as spokes, and the connected instances occurring on the rim.
- *Timeline*—a viewpoint displaying the critical path of predecessors and successors for one or more selected instances (see below).
- *Logical-to-physical mapping*—a viewpoint showing relationships between selected logical instances (for example, processes) and the physical instances (for example, performers) which implement them.
- *Physical-to-logical mapping*—a viewpoint showing relationships between selected physical instances (for example, performers) and the logical instances (for example, processes) which they implement.
- *Message sequence diagram*—a viewpoint showing the chronologi-

cal sequence of information flows for a particular business operation.

In cases where a dedicated tool is not available, fully-customised viewpoints can be generated using the auto-layout routines⁵ provided by the TRIADS analyser (see below). By providing both standard viewpoint generation tools and the analyser, TRIADS is able to support a broader, more flexible range of viewpoints than any other known architectural modelling tool.

Navigating around the information

The construction of TRIADS diagrams facilitates diagram-to-diagram navigation using both cross-reference and drill-down methods.

Each real-world entity (for example, a particular version of an OSS) is typically modelled by a single instance in the TRIADS database and this instance is associated with each diagram node representing the entity. A simple search of diagram contents can therefore be used to find all diagrammatic occurrences of the instance and enable the user to cross-reference appropriate diagrams. Using this approach, an information trail can be followed through a succession of related diagrams. On navigating to a new diagram, the node or link used for cross-reference is highlighted to indicate the entry point for the trail.

Navigation through complex infrastructures is often further facilitated by the use of top-level diagrams and drill-down. A top-level diagram provides an abstraction of the problem domain which can be readily assimilated. Any particular aspect of the problem domain can be expanded by 'clicking' on the appropriate node and 'drilling down' to a lower level diagram. Multiple drill-down routes are supported. Hence, drilling down on a process node may provide access to both a lower level

Figure 3—Schematic showing operation of the TRIADS analyser

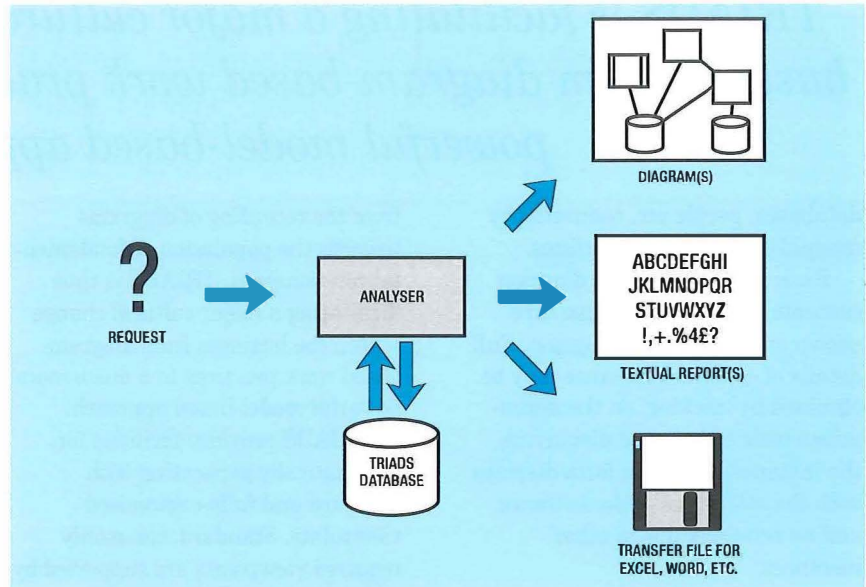
workstring diagram and a set of systems diagrams showing implementation of the process for various products. The drill-down method is fully recursive, enabling a problem domain to be broken down into multiple levels of abstraction.

Customised information processing and visualisation

The most powerful features of the TRIADS application are its abilities to perform customised information processing and data visualisation. These features are provided by the *analyser*. The analyser enables users to gain maximum value from the data stored in TRIADS by fully exploiting the relationships in the business model. The features of the analyser may be grouped into three main categories:

- 1 *Error checking* Facilities range from identification of duplicate data instances to consistency checking of whole architectures.
- 2 *Impact analysis* Customised searches enable all instances impacted by a change to be identified.
- 3 *Viewpoint auto-generation* Results of a database search may be viewed as a diagram, enabling users to create customised viewpoints on the business.

The principles governing the operation of the analyser are summarised in Figure 3. Requests are submitted to the analyser in the form of pre-defined scripts. These scripts may be created by the user or selected from a script library. The scripts are written in a high-level macro language which hides the complexity of the database schema and employs terms which are familiar to business architects and planners. Output from the analyser may consist of a list of data items, a formatted report or a TRIADS diagram. This output may be directed to a screen, to a printer, or to a file for



export to other software packages (for example, word processors and spreadsheets).

Planning OSS Migration

The primary role of the TRIADS business model is to provide a framework for planning and managing change. Consequently, all major components of TRIADS are modelled with a time dependency, and extensive facilities have been developed to support the migration planning process.

Timelining

One of the most common approaches to migration planning is to produce *timelines* (also known as *evolution routemaps*) for individual systems

The primary role of the TRIADS business model is to provide a framework for planning and managing change.

(see Figure 4). Timeline diagrams are similar in appearance to Gantt charts, with instances plotted as horizontal bars against time. The bar for each instance represents the period for which the instance is intended to be operational within the business.

The TRIADS timeline tool enables timeline diagrams to be generated, edited and visualised. Each timeline displays a critical path of predecessors and successors for a single, highlighted instance. Multiple timelines may be plotted on the same

diagram to enable planned dates for various OSS releases to be compared and the impact of any slippages to be assessed. Such knowledge is essential to the successful planning, design and delivery of complex integrated solutions.

Analysing a migration sequence

Each TRIADS diagram (with the exception of timelines) is dated such that it represents a snapshot of the business at a single point in time. A migration sequence may be produced for any business viewpoint by creating a series of diagrams with different dates.

TRIADS provides facilities for viewing and navigating through migration sequences. However, in

general, migration planners are much more interested in comparing diagrams in a migration sequence to identify planned changes over time. For a set of complex diagrams this can be an extremely difficult task. The TRIADS diagram comparison tool has greatly enhanced the value of migration planning data by enabling similarities and differences between diagrams to be displayed graphically or summarised in an auto-generated report.

The diagram comparison tool can be used for any pair of TRIADS




Figure 4—Example of a timeline (data for illustrative purposes only)

diagrams, regardless of whether they are manually or automatically generated. The diagrams do not have to form part of a migration sequence. The tool analyses each node and link in the two diagrams and categorises it as either 'identical', 'modified' or 'different':

- 'Identical' items are those which have a counterpart in the other diagram which represents the same database instance.
- 'Modified' items are those which have a counterpart in the other diagram which relates to a similar real-world entity, but which is modelled by a different database instance (for example, a different version of the same OSS, or a different link between the same two processes).
- 'Different' items appear in only one of the two diagrams and may therefore relate to real-world entities which have been removed or introduced during the period between two migration steps.

The comparison data is visualised by highlighting nodes and links for the various categories in different colours. In addition, comparison

reports may be generated for either whole diagrams or individual instances.

Interpolating between migration steps

A migration planner will often have knowledge of a current situation and a distant vision and be given the task of creating a series of migration steps between the two. The TRIADS interpolation tool has been provided to assist in this task.

Two existing diagrams from the same migration sequence are chosen to act as the initial and final states for the interpolation. A new intermediate diagram is then partially auto-generated by comparing the two end diagrams and automatically copying populated nodes and links to the intermediate diagram. The interpolation criteria are based on diagram comparison and timeline data. For example, if an instance exists in both end diagrams it is assumed that it should also appear in the intermediate diagram. If an instance exists in modified form in the two diagrams, timelines are examined to determine which version of the object should be added to the intermediate diagram.

Having auto-generated part of the intermediate diagram, the interpolation tool displays the initial and final diagrams with colour coding to

indicate which items have been automatically copied to the intermediate diagram. The remainder of the intermediate diagram can then be created by 'drag and drop' of populated symbols between diagrams or by use of the standard diagram editing facilities. Colour coding of the initial and final diagrams is dynamically maintained during drag and drop to assist with housekeeping.

The interpolation tool may be used repeatedly to build up multiple steps in a migration sequence.

Use of the interpolation tool enables most routine aspects of the migration planning process to be automated. This reduces the risk of inconsistencies and enables the migration planner to concentrate on the higher-value aspects of the task.

Managing TRIADS Information

In today's rapidly changing business environment, OSS migration plans must be continuously reviewed and updated to keep pace with changing business needs. TRIADS therefore requires a flexible information management mechanism which enables architects and planners to create and modify data dynamically. The mechanism must provide adequate access control and configu-

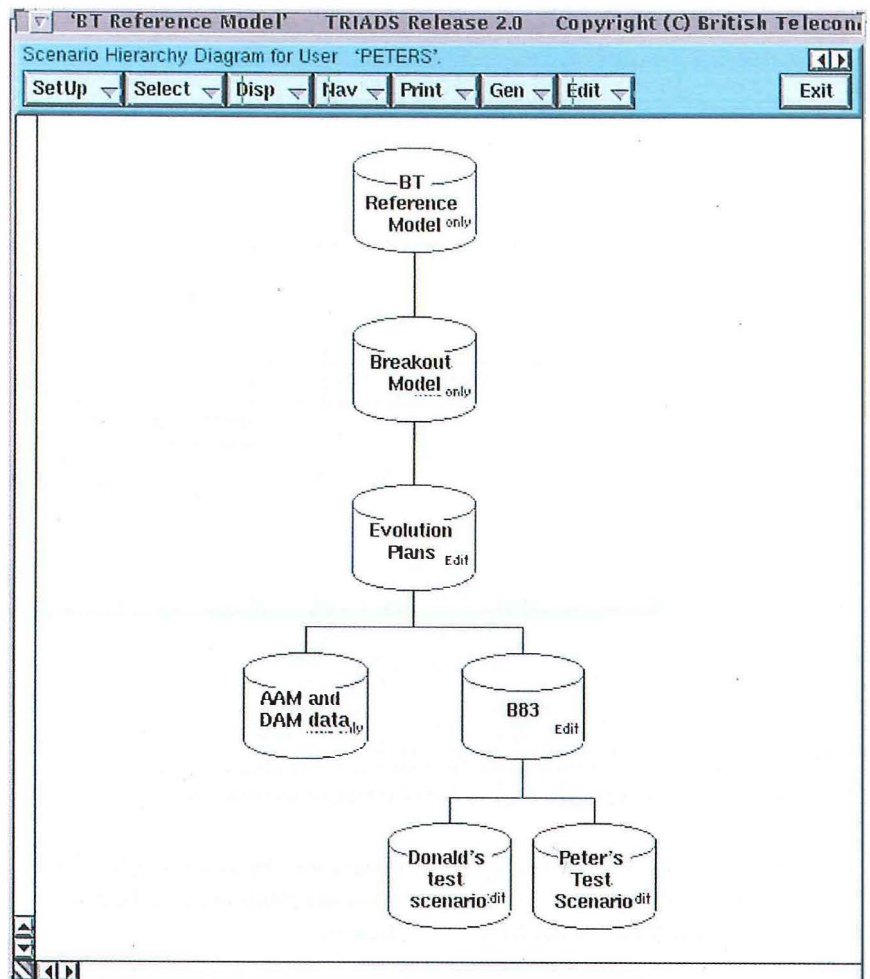
Figure 5—The TRIADS front screen, showing a typical scenario hierarchy

ration management, while at the same time providing the flexibility to enable alternative planning options to be modelled and assessed. These requirements have been addressed using a novel scheme in which the TRIADS database is partitioned into units called *scenarios*.

Each scenario represents a unique, multidimensional perspective of the business. Consequently, scenarios can be used for the development of new architectural models (for example, process and systems architectures for a new service) or for carrying out speculative 'what if...?' studies. Alternative architectural models may be developed in different scenarios for comparison or impact assessment. The scenario provides a suitable partition for analysis and consistency checking of a new model, prior to submission for approval by appropriate design authorities. Authorised scenarios may be made available to a wider community of users or incorporated into higher level models of the business.

The scenarios are arranged in a hierarchy, with the most stable, highly approved model of the business at the top of the structure (see Figure 5). On descending the hierarchy, the stability and approval status of the scenarios tend to decrease, but the data contained in the scenarios is more representative of latest ideas. The lowest-level scenarios typically contain working models or data from speculative studies.

Each instance of data in the TRIADS database is owned by a single scenario. Visibility of instances is inherited down the hierarchy, so parent-owned instances can be viewed and reused by a child scenario, but not vice versa. Hence, each scenario is essentially a perturbation of its parent, maximising the reuse of data via inheritance. This greatly simplifies data management by avoiding the need to create multiple copies of data instances which could subsequently get out of step. The methodology ensures that all models which refer to a particular real-world entity in a



higher scenario (for example, a particular version of an OSS) refer to the same instance in the database. It is thereby guaranteed that if any attribute of the instance is modified (for example, the date at which the new OSS version is to be deployed), the impact will immediately be knocked on to all scenarios which make use of the instance.

Any number of users may have read-only or edit access to a scenario. Consequently, scenarios can function as either private or group work areas. Typically, the high-level, approved scenarios will have a large number of read-only users, but very few users with edit permissions. Low-level working scenarios will tend to be private to a single user or a small group of users, but the proportion of users with edit access is likely to be higher. Each scenario has an owner who is responsible for managing the data instances owned by the scenario. The scenario owner may assign read-only or edit access to other users.

On login to the TRIADS database, a user is presented with a personalised front screen showing the

hierarchy of scenarios accessible from his or her account and the permission level (read-only, edit or owner) for each scenario. The user selects a 'current scenario' to work in by clicking on one of the data-can symbols on the front screen. The current scenario then scopes the data available for viewing and editing. Data from different scenarios may be compared using multiple windows working in different current scenarios.

TRIADS Implementation

The TRIADS application was developed using a suite of tools from ILOG based on the LeLisp™ variant of the Lisp programming language. The principle tools employed were Aida™, a library of graphical components, and Masai™, a graphical user interface builder. The ILOG toolset was chosen to support the evolutionary prototyping methodology used for TRIADS development.

The TRIADS application runs on a UNIX server and is typically accessed via an X-windows emulator on client

PCs. The TRIADS database is implemented on a UNIX platform using an Oracle relational database management system.

Users and Uses

TRIADS is currently being used by more than 100 architects and planners within BT for high-level architectural design and migration planning. Databases are currently deployed at BT Laboratories and three sites in London. Plans are in hand to create a single logical database spanning all user sites.

TRIADS has already contributed to a large number of major programmes within BT, including Breakout OSS design², CSS migration strategy, Billing 90s, Featurenet, SMDS and Cashless. TRIADS is now the master repository for the Network Control Layer Picturebook and will shortly become the master repository for the Systems Encyclopaedia and the Data Architecture Manual.

Summary

The TRIADS system is adding significant benefit to BT in several ways:

- The TRIADS database is providing dynamically-maintained information on the OSS infrastructure and facilitating interworking between disparate groups across the business.
- Data visualisation and navigation facilities are adding value to the information by enabling it to be viewed from multiple perspectives and at various levels of abstraction.
- Advanced information processing and impact analysis capabilities are enabling users to exploit fully the wealth of information in the TRIADS database.
- Electronic tool support is transforming the migration planning process.

As a result, BT is now better equipped to plan both tactical changes to existing operational support systems and strategic migration towards future OSS infrastructures.

Acknowledgements

The authors wish to acknowledge their many colleagues throughout BT who have contributed to the work reported in this paper. The close collaboration between developers and customers has been critical to the success of the TRIADS project.

References

1. FINEMANN, LAURIE. Process Re-engineering: Measures and Analysis in BT. *Br. Telecommun. Eng.*, April 1996, **15** (this issue).
2. EVANS, MIKE. OSS Architecture. Scheduled for *Br. Telecommun. Eng.*, July 1996, **15**.
3. STEWART, HILL. Building on the Architecture Framework, *Br. Telecommun. Eng.*, Oct. 1994, **13**, p. 186.
4. SKEVINGTON, P. J., VIDELO, I. D. E., WITTEGREFFE, J. P., PUTLAND, P. A., SLOAN, D. F., CRESWELL, D. C., and SMITH, A. P. TRIADS: Managing Change in Complex World, submitted to *BT Technol. J.*, 1995.
5. FRUCHTERMAN, T. M. J., and REINGOLD, E. M. Graph Drawing by Force-Directed Placement. *Software—Practice and Experience*, 1991, **21**, p. 1129.

Glossary

CSS Customer service systems
OSS Operational support systems
NCR Network capability register
PSLM Products and services life-cycle management
PSTN Public switched telephone network

SMDS Switched multi-megabit data service

TRIADS Telecommunications requirements, impact analysis and decision support

Biographies



Peter Skevington
BT Networks and Systems

Peter Skevington joined BT Laboratories in 1985

after graduating in Physics from the University of Cambridge. For the next seven years, he worked on the growth and analysis of semiconductor devices for optical transmission systems. His work on the development of chemical beam epitaxy resulted in more than 25 publications and the award of a Ph.D. from the University of Wales, College of Cardiff. In 1993, he joined the Applications Unit within Advanced Applications and Technologies to work on TRIADS. He is currently working on several projects associated with the planning and management of change in complex enterprises. He is a Member of the Institute of Physics and a Chartered Physicist.



Ian Videlo
BT Networks and Systems

Ian Videlo graduated from the University of

Southampton with a B.Sc. in Physics. He joined BT Laboratories in 1981 where he spent several years working on the reliability of electronic and opto-electronic components. In 1992, he joined the Applications Unit within Advanced Applications and Technologies to work on the TRIADS project.



John Wittgreffe
BT Networks and
Systems

John Wittgreffe graduated from York University in 1989 with an honours degree in Physics. He joined BT in 1989 and worked for three years in materials and device analysis using SIMS, auger, X-ray, and electrical characterisation techniques. In 1992, he transferred to the Applications Unit within Advanced Applications and Technologies to work on the TRIADS project. His roles within the TRIADS team have included system design, leading the software implementation team, providing software demonstrations and training users. He is a member of the Institute of Physics and a Chartered Physicist.



Paul Putland
BT Networks and
Systems

Paul Putland joined BT Laboratories in 1985 after graduating in Physics from the University of Exeter. For the next seven years, he worked on the reliability and analysis of semiconductor devices for optical transmission systems. He gained an M.Sc. on the Physics of Laser Communications from Essex University in 1992. In the same year, he joined the Applications Unit within Advanced Applications and Technologies to work on TRIADS. He is currently working on projects involving Internet-based services.



Donald Sloan
BT Networks and
Systems

Donald Sloan graduated from the University of Strathclyde in 1984 with an honours degree in Mining and Petroleum Engineering. He went on to gain an M.Sc. degree in the Design and Manufacture of Microelectronic Systems in 1985 at the University of Edinburgh. He joined BT Laboratories in 1985 and worked for the next six years on projects concerning the reliability of semiconductor integrated circuits and various opto-electronic components. In 1992, he joined the Applications Unit within Advanced Applications and Technologies to work on TRIADS. He is currently working on projects that support the ongoing maintenance and development of TRIADS and is looking at the application of artificial intelligence to the management of change in complex enterprises.



Dan Creswell
BT Networks and
Systems

Dan Creswell graduated in Computer Science from Loughborough University in 1992. After a brief period with Cegelec in Rugby, he joined BT Laboratories in 1993 to work on the TRIADS project within the Applications Unit of Advanced Applications and Technologies. He is currently working on projects covering data visualisation and the adoption and impact of object technology throughout the business.



Alan Smith
BT Networks and
Systems

Alan Smith worked for HM Customs and Excise upon leaving school. He subsequently studied Computer Science at Lancaster University, graduating in 1992. He joined the Applications Unit within Advanced Applications and Technologies to work on TRIADS. He is currently working on quality of service and dependability in distributed systems.

Bill Hobbs

Cashless Services Replacement System Project

A network perspective

Since the introduction of the BT Chargecard product in 1988, customer demand and expectations have increased. A complete review of this service initiated the cashless services replacement system project. This article describes the impact of the changes to the cashless services network, improvements to customer service and some of the techniques used to manage a large team across a variety of companies working in different countries and time zones.

Introduction

In 1988, after extensive field trials with customers, BT opened its automatic cashless services network. Since then, growing customer expectations and requirements, synergistic opportunities created by company initiatives and the desire to offer the best service have been the main drivers for enhancing the cashless services network.

The cashless services portfolio includes BT's Chargecard, Creditcall and several other services. Cashless services allow customers to access the fixed network from any telephone connected to BT's network and have the call fees debited to their nominated telephone account or credit card. The call-making element of a mobile telephone is conveniently replaced by a plastic card that will easily fit into a purse or a wallet, and customers only pay for the service when they use it; that is, there are no rental charges or annual card fees.

In 1993, a BT Chargecard sponsored project initiated a complete review of the whole of the cashless service system. This review covered the end-to-end delivery of the services from customer order handling through to the revenue collection methods. Sandwiched between customer orders and revenue collection are the network components, billing systems, fault repair and in-service support systems and processes.

The outcome of the review was the cashless services replacement system (CSRS) project. The main objectives

of the project were to deliver a platform that would improve customer service, to provide a basis for developing the BT Chargecard product and to simplify growth aspects.

This article has been written from the author's personal perspective and experiences from his involvement with the network aspects of the project. The article is biased towards the network, but describes some of the improvements gained within the service management, customer service support and billing areas, and goes on to some of the project management techniques used across several diverse functions and external companies.

It must be borne in mind that the network aspect was only a part of the project and the significant contributions by other BT Departments, MCI and external suppliers have not been fully reflected owing to the limited space available.

BT Chargecard Service Review Project

The BT Chargecard service review project, referred to in the introduction, identified the following disadvantages and constraints of the old system†:

- Very little 'customer cuddle' existed. This means that when a customer got into difficulties using

† POPE, N. The Cashless Services System. *Br. Telecommun. Eng.* July 1990, **9**, p. 112.

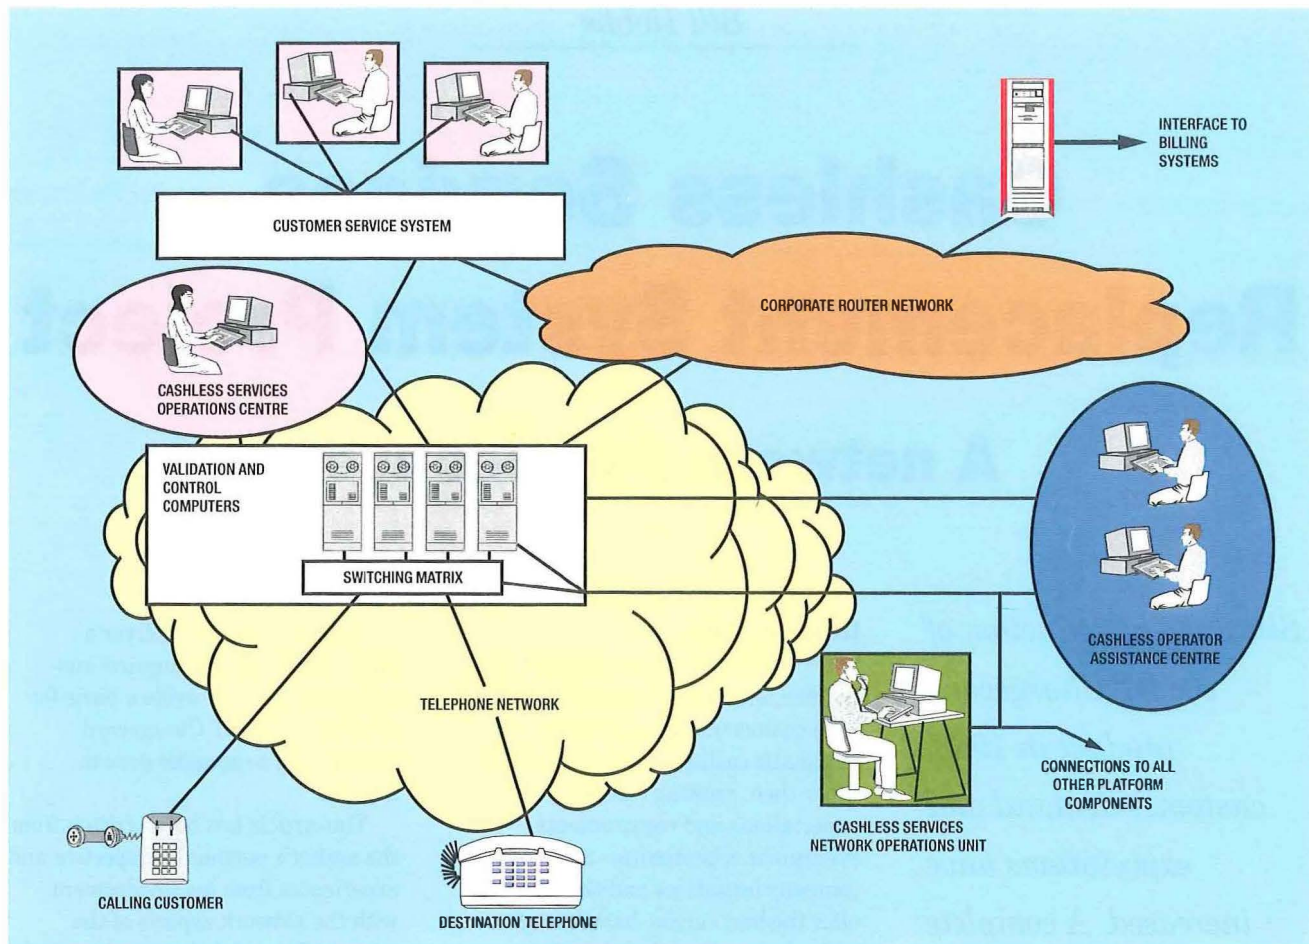


Figure 1—Cashless services system

the automatic system, they either abandoned the call or rang 100 for operator assistance or rang the help line where a member of the customer services operations centre (CSOC) would attempt to explain how to use the system.

- A centralised customer ordering process prevented BT from maximising possible benefits of using the standard 150 and 152 dial-up sales channels. Also, there was potential for the CSOC support element to be subsumed by the work involved with order processing.
- Customers with loop-disconnect telephones (still a large part of the UK telephone population) could not use the automatic system.
- The pricing and billing methodology was inconsistent with BT's Billing 90s strategy.
- Growth of the customer base was restricted by inherent capacity limitations.

- Development potential was limited.
- The network-based equipment used outdated signalling systems, which also contributed to unacceptable post-dialling delays. Also, the discrete components used within the hardware were rapidly becoming obsolete.
- The number of call handling sites (36) made the roll-out of enhancements and extensions resource intensive, complex and time consuming.
- The old network prevented exploitation of other network (PSTN) components or new developments; for example, calling line identity.

The Cashless Services Replacement System Project

BT Group Systems Engineering (now part of Network and Systems, Design and Build) recommended a replacement system that combined the MCI

Star card network and a variety of BT systems to produce the physical components of the CSRS. The end-to-end system included process re-engineering and a broadening of the network support function to include all aspects up to the billing interfaces.

General principles of a cashless system

A comprehensive cashless services network requires customer account details, interfaces to the telephone network and billing/revenue collection mechanisms, network management, fault reporting systems that cover all types of service problems, systems to support the customer when difficulties are encountered, ubiquitous access and an easy-to-use system.

Figure 1 shows the functional blocks of the cashless services replacement system; the following describes the major changes that have resulted from the project.

Call network

In simplistic terms, the call network consists of the validation and control

The CSRS created a major development programme that brings together the billing requirements of BT Chargecard into one cohesive, but distributed, network.

computer (VCC), which holds customer details and service access data, and a switching matrix to handle the telephone call element. The old network and CSRS configurations can be likened to an intelligent network where the service and feature availability reside on the computer platform but the functionality and call-carrying capacity reside on the switching matrix.

The BT Chargecard customer dials the network access code (144) and is routed to the switching matrix interface. The switching matrix is a sophisticated three-way conference bridge which bridges the calling customer to the VCC. The VCC checks the customer account details before a call can be set up to the destination telephone. If the customer account is accepted, the VCC instructs the switching matrix to connect the calling customer to the PSTN (the type of service the customer has requested), and then to request destination telephone number details. The VCC drops out from the call set-up process and is reused only if the customer wishes to make another call or wishes to use another feature/service.

The other components of CSRS are explained below.

Service management

Service management can be regarded as the end-to-end treatment of the customer from provision of service, pricing, revenue collection, in-service support through to fault management. The following paragraphs explain some of the improvements that have emanated from the CSRS project.

Service provision

Ease of obtaining service is a prime consideration of both the service supplier and customer. Prior to CSRS, all orders for BT Chargecard and service enquiries were centrally processed by a specialist unit at Liverpool, known as the *cashless services operations centre* (CSOC). All customer orders had to be entered

twice, once into a cashless services database and repeated into BT's customer service system.

Software enhancements to BT's customer service system now allow customers to order BT Chargecard service by dialling the 150 or 152 sales channels or the CSOC. The customer service system downloads the information direct to the replacement system and updates any other associated systems, such as the billing suites. The role of the CSOC has been maintained and further improved by allowing the CSOC personnel, known as customer service agents, to set up a call after a service enquiry from a BT Chargecard customer.

In-service support

One of the prime responsibilities of the CSOC is to provide guidance and support to BT Chargecard customers. Before, CSRS customers rang or were referred to an 0800 helpline number. With the CSRS, cashless operator assistance centre (COAC) personnel, known as *cashless operators*, can automatically transfer customers with service enquiries to a customer service agent.

The customer service agent can then try to answer the query or resolve the problem. If the customer then requires a call to be set up to the destination number, the customer service agent can provide this service from a desk-mounted terminal connected to the CSRS.

Operator assistance

On the old network, once customers dialled 144 to gain access to the cashless services automatic network, they had no fallback to any system if they encountered difficulties, unless they cleared down and dialled BT's operator assistance on 100 or 155. This inconvenience has been removed.

When the CSRS call network detects a customer not completing certain actions within a given time, entering incorrect account information or using a loop-disconnect type

telephone, the computer part of the platform requests the switching matrix to transfer the call to a free cashless operator's position.

The cashless operator is able to assist the customer by either completing the call set-up to the requested destination telephone number, after successful account validation, or, if the customer requires help beyond this, the call can be transferred to a customer service agent.

The customer can also obtain COAC services by entering *0 after dialling the 144 access code.

Billing

The cashless services computer on the old network, referred to above in the section on 'Service Management', used very sophisticated software that provided the service management, billing and pricing, and validated BT Chargecard accounts against the stored customer details. To ensure that calls were priced correctly, tariffing data had to be kept in step with BT's pricing systems and the output had to go through a protracted process before it appeared on a customer's bill.

The CSRS created a major development programme that brings together the billing requirements of BT Chargecard into one cohesive, but distributed, network. Benefits have been gained by utilising, wherever possible, existing billing capability and adding new systems where they were needed.

A single entry of the BT Chargecard customer details onto the customer service system ripples through the billing systems. In a similar manner, any changes to prices can be reflected immediately by the CSRS referring to the appropriate national tariffing database.

Fault management and support

A cashless services network operations unit (CSNOU) was established in 1994 within the network operations unit at Worthing. The network management centre (NMC) at

Oswestry still manages the PSTN elements of the service. The CSNOU monitored 36 switching matrices (in 33 locations), liaised with other field units and equipment suppliers and became the focal point for all fault reports on the old network. As can be seen from Figure 2, the logistics of managing this network across the UK were complex.

The CSNOU's original responsibilities were geared towards maintaining and clearing problems on the switching matrix part of the old network. This role has been expanded to include the management and control of the entire CSRS network up to the Billing 90s and PSTN interfaces and managing all customer fault reports from the CSOC or via the 151 fault reporting channel. This includes hardware, software, capacity/performance management, managing all fault reports, all data build, wide area networks, local area networks and any other links within CSRS.

The CSNOU has connections to all of the major sites and can monitor faults as well as localise them to individual pieces of equipment. A wide variety of monitoring systems, skills and intensive training has been devoted to the establishment of this unit. It is interesting to note that the CSNOU is unique within BT as it is the first operations unit that combines telephony and computing operational management for a national network within one unit. Figure 3 shows the impact of concentrating CSRS into fewer sites; however, the skills, interdependencies and complexity of network management have increased.

More information on the role and achievements of the CSNOU will appear in a future article. The CSRS Network Design Authority will also be providing an article covering aspects of this challenging work area.

Project Structure

The CSRS project has been a large undertaking which has depended on

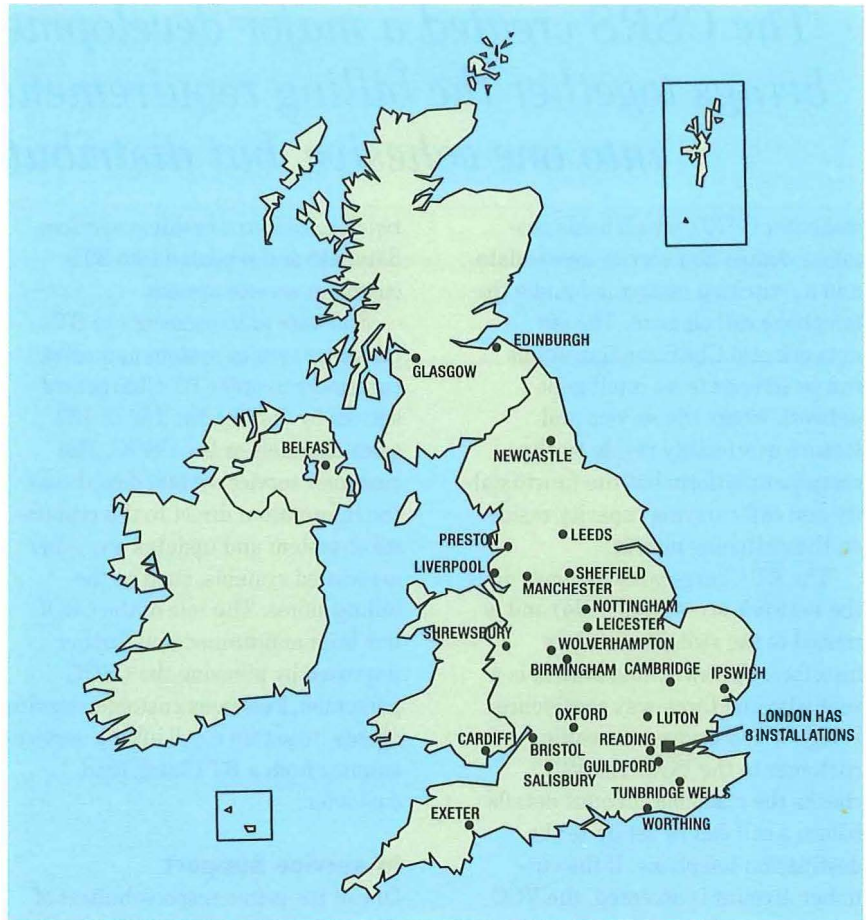
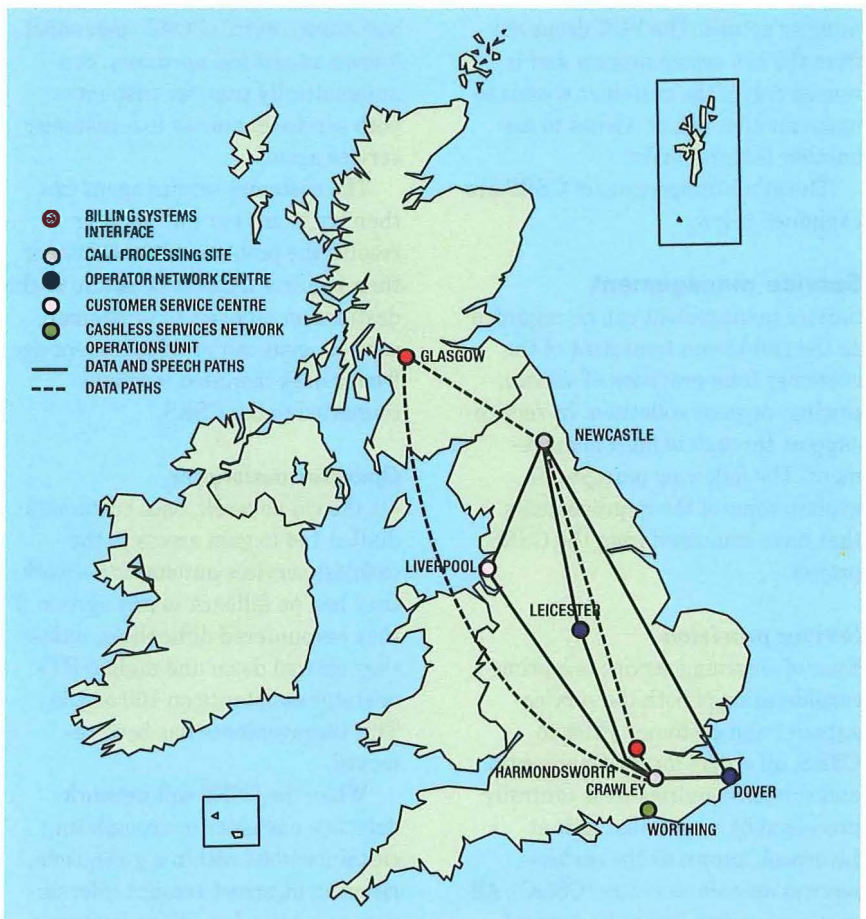


Figure 2—Old network locations

Figure 3—CSRS call processing network



The geographical separation of people involved and the added complexity of matrix management across companies and divisional boundaries demanded a good communications system.

delivery from a variety of functions and involved personnel from BT, MCI, two switch suppliers, four major computer companies and other companies covering the provision of specialised equipment or expertise.

The emphasis was that the **team** meant **everyone**: BT, MCI, equipment suppliers, etc. It would be impossible to arrange the management of the different companies or divisional units or time zones into one functional block. Matrix management has been used extensively.

It is not unusual for matrix management practices to create a plethora of paper systems and bureaucratic structures. These aspects were minimised—meetings and conference calls kept the need to produce written information to a minimum.

The project manager, Charles Conley, had a small team of direct reportees and delegated the functional responsibilities to the appropriate areas of expertise. The project organisational structure was divided into seven main units, see Figure 4:

Requirements ensured that the customer requirements were being met. This function was headed by the BT Chargecard customer champion (the Card Services Product Team) who actively participated throughout the project by being an important part of the feedback loop.

Development produced the low-level design and then coded and 'unit' tested the billing, service management, management information system and field operations system software.

System design coordinated the end-to-end design and broke down the overall system into manageable building blocks for each of the directly reporting functional units.

Implementation project managed the planning and implementation of the installation of all CSRS components up to the Billing 90s interface and the supporting PSTN infrastructure.

Integration ensured that the designed system components worked together and defined/completed the end-to-end testing of all of the hardware and software that makes up the CSRS platform, excluding the PSTN.

Operations re-engineered operational processes, established the CSNOU, established the Glasgow command centre (Computer Service Operation's management of the data network and associated computer hardware components not covered by the CSNOU) and COACs, enhanced the CSOC, implemented operational training programmes, support system development and ongoing support for the CSRS.

Migration managed the transfer of all data/traffic between the old systems and CSRS, migration of all service support, and the redesign of the BT Chargecard card, and ensured continuity for any residual services remaining on the old network and all customer communications.

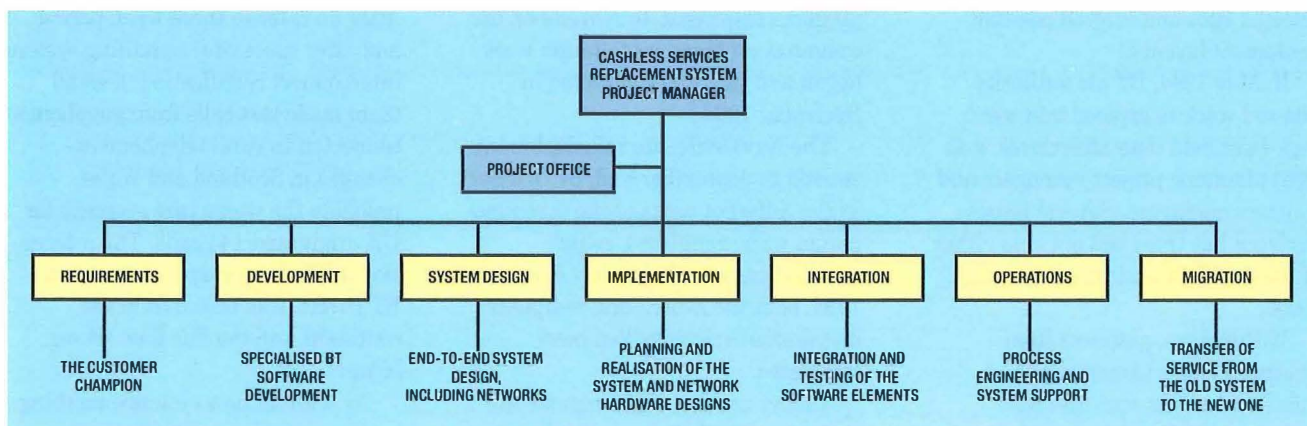
Communications

The geographical separation of people involved and the added complexity of matrix management across companies and divisional boundaries demanded a good communications system.

Regular joint team meetings (JTM) were held involving representatives from the major functions, external suppliers and any other project participants who were playing an important role at the time. No subject was sacred and the taboos of them and us were eradicated; this enabled the 'joint' part of JTM to come to the fore. Each of the project structure blocks had their own meetings which discussed topics relevant to that area of expertise. Once a week, a project team meeting was held to monitor progress, set new objectives and resolve issues.

Conference calls between BT personnel based in the United States and UK became commonplace. Also, meetings within divisional units were held to ensure that local and project objectives/deliveries complemented each other.

Figure 4—Project structure



With the CSRS moving to an in-service platform, the meeting formats and terms of reference are changing to meet the challenges ahead. New forums have been established to deal with product development, in-service support aspects, change control etc.

Team Working

The interdependencies between the different functions and complementary skills between the functions cannot be very easily shown in a diagram or drawing. The project structure diagram gives an idea of the work flows through from customer requirements to traffic migration (note: the operations unit also has post-migration responsibilities). As an example of how some of the network activities were realised, a brief synopsis of the installation of the call processing network follows.

In anticipation of project authority, the Planning Department gathered information from field planners about suitable and well positioned (in network access terms) locations to site the new equipment. In March 1994, accommodation was reserved by field accommodation tactical planners in Southern Home Counties (Crawley) and the North East (Newcastle).

In parallel to this activity, the design of the site configuration began. This exercise included MCI, equipment suppliers, BT's CSO, Network Planning and the more detailed design skills of the Network Design Authority (NDA). This first iteration ratified the selected sites and evolved possible equipment layouts.

In May 1994, BT plc authority allowed work to proceed and meetings were held soon afterwards with local planners, project managers and contract managers. Several issues surfaced but these did not stop either of the two field units from starting work.

Within days, planners from Southern Home Counties had initiated various activities at Crawley, which resulted in the

recovery of redundant Strowger hardware, and site surveys by the various teams involved with refurbishment, power provision, hardware suppliers, accommodation planners etc. Switch installation began in July.

Timescales for the Newcastle site were more generous, but 1200 racks of redundant crossbar equipment needed to be recovered. Prior to the recovery, it was discovered that the ceiling tiles contained asbestos. After a 12 week site quarantine period for asbestos removal, site clearance and accommodation refurbishment began.

The diversity of equipment types, power requirements, resilience factors and variations in operational environments required manpower resources that were not readily available. To overcome this problem, a coalescence of local planning expertise, supplier support and centrally coordinated information gathering/distribution techniques emerged. The level of support supplied by the local power, heating, ventilation and environment control planners allowed efforts to be focused on the important requirements and reduced the planning time needed to a few months.

During September 1994, additional requirements arose. A review of the dimensioning and network configuration aspects took place. By redefining some parts of the designed platform, very little change was needed. By October, the Crawley switching matrix was successfully passing test calls to a simulator that replicated some of the computer platform responses. In November, the computer platform installation work began and this was completed in December 1994.

The Newcastle site refurbishment started in September and, by October, as the different parts of the accommodation were completed, switch installation work began. By February 1995, both the switch and computer installation activities had been completed.

To give an idea of the equipment that was installed, the major hard-

ware components installed at just one site included four switches, 15 voice guidance systems, 14 computers, a variety of bridges, routers, Ethernets, token rings and several hundred connections to the PSTN.

By February 1995, both of the call processing sites were ready for the integration test team to begin work. The testing methodologies were proven on a CSRS network model at Martlesham Heath. When this was completed, testing using direct links between the Martlesham Heath test bed and the Crawley call processing site took place. Once the basic system proving was over, the integration team relocated themselves closer to the equipment by moving to Crawley.

Once on site, the team undertook more intensive testing, system (hardware and software) performance measurements and various call simulations, including provocative testing under failure conditions. Gradually, the CSRS interfaces to the billing systems, COACs, CSOC and cross-network connections to the second site became available. Each of these required 'within platform' tests. When all of the integration testing was completed, the CSRS network was handed back to the implementation and operations teams in July 1995.

The task now was to prove that the CSRS call processing platform worked with the PSTN. Before this could be done, 28 major elements had to be connected together to form the core of the CSRS network, which then had to be connected to the PSTN. After an intense three-week period and after successful signalling system interconnect certification, a small team made test calls from payphones connected to rural telephone exchanges in Scotland and Wales, probably the worse case scenario for UK originated 144 calls. The international BT Chargecard service, Auto BT Direct, was tested from the continent and the Far East (Hong Kong).

As with all new systems, teething problems, both on and off the plat-

form, appeared. The CSNOU training was soon put to good use, as well as the experience and network knowledge of others. Similarly, the planning, Network Design Authority, implementation and integration teams, Operator Services, CSOC, MCI and the central operations unit (COU) all became active participants in identifying and clearing the teething problems. Some were relatively minor data-build errors and others were more complex problems such as satellite downlink interworking with CSRS. Fortunately, the number of problems reduced to such an extent that the decision to introduce 'live' traffic onto the platform was taken.

By 28 August 1995, the migration manager was able to start pre-beta trials using a selection of BT service account holders as the guinea pigs. In October, the trials were extended into full beta trials and included real customers from the Liverpool and Severnside areas of the country. Special attention has been paid to the feedback from customers—this enables the success of the CSRS to be determined. On 4 February 1996, all BT Chargecard traffic was migrated from the old network to CSRS. This involved the transfer of large data files between the computer platforms and changing the 144 traffic routing data in BT's telephone exchanges. The exercise involved personnel from MCI, CSNOU, CSOC, COACs, NMC, COU, field data managers, Networks and Systems, BT Syntegra and a large number of people closely involved and associated with the project.

Conclusions

Since 1988, new network capabilities have been added to BT's infrastructure which BT Chargecard customers have not been able to utilise owing to the constraints of the old network. The CSRS project has delivered a platform that has extended the boundaries of BT's Chargecard service management, enhanced customer support elements while at the same

time providing a launch pad for innovative products that will be able to take advantage of the 'telecomputer' infrastructure that has been provided. Additionally, CSRS has the potential for acting as a gateway to a variety of existing services and the time to market of new services will be greatly reduced as will the roll-out timescales for future projects.

The broadened CSRS project definition of the team certainly eased the planning and implementation aspects. Within the planning area, benefits of this approach are still being realised, especially when the impact of changes or new services is discussed with the suppliers.

The cooperative collaboration between MCI and BT augurs well for the future of the alliance and further joint projects.

The old system is not entirely redundant as it has attributes that, for the short term, suit the Payphone Creditcall Service (another one of BT's cashless services).

Glossary

- COAC** Cashless operator assistance centre
COU Central operations unit
CSNOU Cashless services network operations unit
CSOC Cashless services operations centre
CSRS Cashless services replacement system
JTM Joint team meeting
MCI Microwave Communications Incorporated
NDA Network Design Authority
NMC Network management centre
PSTN Public switched telephone network
VCC Validation and control computer

Acknowledgements

This article gives a very brief overview of some of the network aspects of the CSRS project. It does not cover the extensive work of the various units involved. However, the author

would like to make special mention of some of the 'back room' people who have contributed very much to the realisation and success of the project.

The patient and tolerant Southern Home Counties and North East planners and project managers, the central operations unit and local circuit provision teams who provided the transmission infrastructure, the CSO router network teams, field data managers and the network management centre.

Similarly this article does not reflect the work of the implementation or migration teams or the skill and professionalism of the clerk of works and CSNOU who were very much involved in bringing the CSRS into service.

Biography



Bill Hobbs
BT Networks and
Systems

Bill Hobbs joined the Post Office in 1969 as a Trainee Technician Improver on customer apparatus maintenance. The next 12 years were spent on various maintenance activities, culminating with implementation and support to one of London's measurement analysis centres. He became a manager in the early-1980s, initially providing support to exchange planning and construction teams. Since then, he has helped with the development and support of the Works and Stores Programme, been closely involved with project and programme implementation on overlay networks connected to the PSTN and, in 1994, moved into cashless planning in support of the CSRS project. In 1991, Bill obtained Diplomas in Marketing and Management Studies to complement his technical background and technical qualifications.

Bruce Bond

The Future Isn't What It Used To Be

Technology developments, coupled with easy access to a wealth of information, are driving tremendous change in the ways companies do business, people work and societies organise. In this article, Bruce Bond, Managing Director of BT's National Business Communications, examines these changes and discusses meeting customers' needs in the new knowledge age. He sees BT playing a special role in making sure this country adapts to the new environment.

A Time of Change

We are all currently involved in an incredible amount of turbulent change. It is almost as if we are crossing a great divide, beyond which everything will be different—so different that the next generation will have difficulty grasping how people lived and worked before the 1990s.

There is a children's story called 'Where the Sidewalk Ends' by Shel Silverstein. In it, he talks about the 'Tesseract' or a wrinkle in time where people on one side are incapable of conceiving how those on the other side lived and worked.

This is what we are experiencing now. It is sometimes hard to see when it is happening, but in retrospect we will all remember that we lived through it and we will see that the effects of this period of change will be considerably more profound than those of the industrial revolution. The changes are so dramatic that the future is not only unknown, it is virtually unknowable.

Technology Drivers

The main drivers of this change are developments in technology. The power of computers is doubling every year; at the same time, the cost is falling dramatically.

In 1980, the cost of a modest 4.5 MIPS computer was \$4.5M, which was equivalent to 210 people. By 1988, the cost was \$100 000, and the people equivalent had reduced from 210 to two people. At the time, forecasters predicted that by the year 2000 this 4.5 MIPS would cost about \$10 000 which would be equivalent to a quarter of a person.

Table 1 Computing Price/Performance

	1980	1990	1995*
Functionality (MIPS)	4.5	4.5	4.5
Cost (\$)	4.5m	100 000	1.35
People	210	2	0.000016

* Sony Playstation

Last year, Sony launching its new Playstation, which priced 4.5 MIPS at around \$1.35. \$4.5 million to \$1.35 in 15 years (see Table 1). In effect, you now have a toy with more processing power than the American Department of Defense obtained from its most powerful computer, a Cray, five years ago.

Developments in telecommunications such as fibre optics mean that we will have almost unlimited capacity to deliver information wherever we choose.

The current world record for the number of simultaneous telephone calls carried by a single fibre in a fibre optic cable is around 31 000. This compares to the 30 carried by a twisted pair of copper cable. Today, BT Laboratories can transmit a million calls over a single fibre the size of a human hair. Within the next couple of years, it will demonstrate a billion calls over a single fibre. This is very difficult to comprehend, but the implications for future networks are clear. We can confidently expect to do more with less—almost ad infinitum: more traffic carried by less physical material, less energy and at an exponentially falling cost. We are developing almost unlimited capacity to deliver information wherever we choose.

The future for businesses will be understanding how to gather information and apply knowledge to serve their customers better than any other organisation.

A New Era for Business

This combination of extremely powerful and affordable processing power, coupled with easy access to a wealth of information, is shaping the way companies do business, people work and societies organise.

The new era we are entering can be characterised as the *knowledge age* (Figure 1), where information and understanding about how to create value for customers are the most precious resources.

Companies are beginning to see that their core competencies are not products, but skills and knowledge, and an understanding of the market and the customers, and the processes and behaviour needed to respond to their needs.

The future for businesses will be understanding how to gather information and apply knowledge to serve their customers better than any other organisation. It will be about bringing together the right people and skills and using technology to gain a greater understanding of customers and deploy resources to meet their needs.

Almost without exception, the customers we talk to are considering how they can begin to use communications and information technology to implement dramatically different business structures.

These new structures are around re-engineered value chains; for

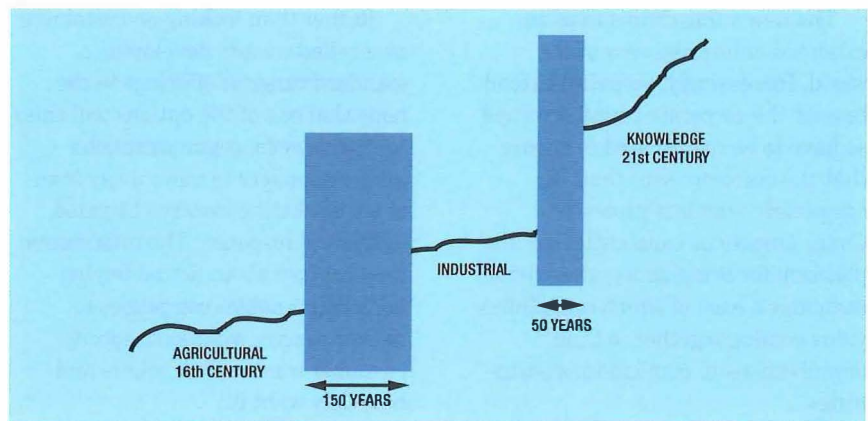


Figure 1—Tesseract


example, the retailer's value chain, which comprises consumers, retailers, manufacturers, wholesalers, delivery functions, and the raw materials of the suppliers for example. Goods and services move in one direction and money in the other. (See Figure 2.) Today, virtually all companies and organisations are struggling with how to make their value chain create more value for the end user than that of their competitors. They are looking at ways to slim down, become more efficient, do things in new ways, and reduce the amount of inventory in their value chain. The Japanese have a concept of 'muda' or waste—how to eliminate time wasting, how to eliminate inventory, how to pull out whole functions—not to just compress processes but to stop doing them if possible. Technology is providing

opportunities for our customers to start thinking about removing unnecessary processes.

In the new streamlined structures, companies focus on the specific value they add and look for relationships with complementary companies to deliver a complete service to the customer. In this environment, they are beginning to outsource more and more of the other activities. They are using more part-time employees and more contractors to help them do their business, resulting in a real shift away from doing things in-house.

In the United States, the Efficient Customer Response (ECR) looked at how to eliminate muda in retail businesses. The main source of waste in retail is inventory, and the programme is predicting that, by using communications capabilities, around 41 per cent of the inventory tied up in the retail channel can be eliminated. Over \$30 billion of hard cash can be saved, before taking into account time wasting, and making sure that you have, in stock, what the customer wants when the customer wants it. This saving is based around more direct shipments to customers, eliminating part of the retail piece, allowing customers to use call centres to come directly to the manufacturer. Freefone and lo-call services, coupled with computer-integrated telephony capability and EDI/EFTPOS, mean that the request and fulfilment can be completed entirely over the telephone.

Figure 2—The value chain



The new value chains take an extended enterprise view of the world. Increasingly, activities extend beyond the corporate boundaries and so have to be coordinated to ensure that the customer sees them as completely seamless processes. Communications capabilities are the platform for bringing together virtual companies, each of which contributes value coming together in fluid organisations to respond to opportunities.

The concept of a virtual corporation at work can be seen, for example, in BT. We outsource the manufacture of switches and even ask our suppliers to perform some of the installation and maintenance so that we can deliver value to the customer without bearing all the cost involved in the work and developing skills in people.

This model of virtual business extends to individuals who are increasingly working as part of virtual teams. More and more, the workforce will comprise networked individuals with responsibility for applying their particular skills and knowledge to deliver value to customer.

Meeting Customer Needs in the New Age

So we see that technology is driving the period of change, but the new capabilities also offer the way forward. In times of dramatic change, the safest place to be is holding hands with your customers. Examples of companies which have failed to do that are legion; they have had excellent technology but have not understood the needs of their customers and driven forward in that direction.

All businesses need to become learning organisations, and develop learning relationships with their customers based on direct contact and a detailed understanding of the actual needs and preferences of the customer (whether they are the end consumer or other organisations in the value chain).

Rather than looking on customers as a collective and developing a standard range of offerings in the hope that one of the options will cater for their needs, communications allows suppliers to move away from mass marketing towards targeted, individual response. The information they capture about actual buying behaviour enables companies to deliver exactly what each specific customer wants when, where and how they want it.

In this type of learning relationship, the supplier becomes increasingly able to meet customer needs, and customers become extremely loyal as they have invested time teaching the supplier about their preferences and they benefit from excellent service.

Our customers are coming to us to help them find customers with new routes to market, develop learning relationships with them, bring together virtual communities of complementary companies to respond to their needs and increase the value of their offering. We are working with them to apply capabilities such as groupware, Internet and Intranet type services, distributed processing and sophisticated databases which allow them to manage the business in radically different ways.

One of the clearest examples exists in the financial services community, where we are working with banks and building societies to help them to see the total requirements of each customer and how they develop over time. This enables them to understand the life cycle of their customers and allows them to cross-sell service based on their understanding of the complete picture. Many other companies in industries such as utilities and in retail are also asking themselves how they can better understand their customers, and what are the most efficient means to deliver value including outsourcing non-core activities and managing surge capacity to respond to high and low points in demand.

Visual communications and groupware applications allow customers to manage teams of people in different geographies, time zones and even different companies. We can eliminate diseconomies of scale and scope which come from distance and lack of accessibility and we can work in international teams with a 24 hour working day, completing projects in days or weeks instead of months.

Multimedia and video-on-demand services are not about movies—not about delivering Jurassic Park or Arnold Schwarzenegger to the home. There is a much more fundamental benefit in that they open up new channels to market. There is a recognition of this in almost every area of society—hospitals, police, transport, education, newspapers, television, radio, publishing and many others. The major retailers, banks and building societies are investigating how to use this new networked capability. What they see is not a TV set, but a total interface to the customer with a single delivery charge.

All these requirements are driving a movement away from people putting in place private networks and more into managed networks with network-based applications. Across virtually all customer segments and all sizes of customer, we are seeing there is a need to build an organic information system which allows them to expand and contract, which flexes to suit their business needs, and new ways of working which mechanistic systems based on traditional communications networks cannot support.

Over the last few years, we have been very successful in working with our customers to implement managed data networks, managed voice network services, and applications such as call centres. Now we see that the call centres are not enough. Customers want computer telephony integration which gives the person on the telephone the information they need about the customer and the business to be able to fulfil the request over the telephone.




Figure 3—Society dimensions

As we enter the knowledge age, all these capabilities are about increasing the understanding and wisdom of the organisation and its people. The simple words 'people are our most important asset' have never been more true. This whole game is much more about behaviour than it is about technology. It is about how we apply and work with the new capabilities. It is not the tools which make a successful craftsman, but his skill in using them.

Future Organisations

Organisations, like all organisms, must adapt to changes in the environment or die. Creatures whose behaviour fails to match the environment cease to exist; witness the dinosaur and the dodo.

The successful organisations of the future will look very different, and the working practices in the new age will be in sharp contrast to those we have lived with in the post-war years.

They will be based on networks of companies and individuals in organic structures, each one bringing diversity, different insights and knowledge. Instead of searching out people who think like you, successful behaviour will value difference and pluralism. We will be dealing with such different problems that the best ideas on how to move forward could come from anyone. All talents will have to be valued, and we will have to transcend traditional boundaries of geography, time, culture, race, background. Everyone will work in organic organisations based on co-operation rather than competition, the environment will be centred on customers not bosses, participation will be open to everyone and the role will be defined by the value they bring. (Figure 3.)

This new paradigm is a serious threat to the status quo. It means a

radical rethink at the highest level in every single company and organisation and by every individual about power. Creating an environment where everyone contributes value, where employees lead and customers teach is a real challenge to the old ways.

BT's Fundamental Role

BT has a very special role to play in making sure that this country adapts to the new environment. The ability of societies and people to adapt behaviour and embrace change is actually more important than wealth, strength and intelligence. Other cultures have historically been more ready to embrace change; we need to help the UK get over innate conservatism.

Customers and citizens are relying on us to show the way to the future. Our challenge is to find ways to use the new tools to deliver a better life, to develop new structures that are appropriate for the knowledge age and to encourage people to grow to their full potential and provide a confident environment.

The race to the future is a struggle between societies. The essence of society is a shared view of the future and how to approach it; we must develop a shared set of assumptions about the future that we will create together, a sort of compelling wire frame model that works for everybody who wants to play. We have to create that view and drive it forward. The stakes are too high for us to risk not coming first in this race; it's about jobs, quality of life and the future success of the country.

Biography



Bruce Bond
Managing Director
BT National Business
Communications

Bruce Bond has a Master's degree in Business Administration from the University of Dayton and a Sloan Management degree from the Massachusetts Institute of Technology. He joined the Ohio Bell Telephone Company in 1969 and has held a wide range of posts in the North American communications industry. He joined BT as Director Corporate Strategy in 1989 after leaving his post as Corporate Vice-President, Strategic Planning, in US West, where he was responsible for coordinating the long-range plans of 43 US West subsidiaries.

Chris Fowler, Terry Mayes and Bernard Bowles

Education for Changing Times

A resurrected Archimedes would be staggered by the technological advancements made by the human race since his death. At the same time, he would be amazed to find how little has changed in the teaching methods used in our classrooms and lecture halls. He might well ask the obvious question: why are you not using all of this wonderful technology to help you teach and enlighten your students?

Peter Cochrane, IEE, 1994

Introduction

Technology is now at a point where it can help learners become less reliant on teachers through self-directed interactive multimedia tutoring. Children can access and communicate with experts from all over the world and collaborate with other cultures. The effective dismantling of geographical barriers and harmful stereotypes will follow as telecommunications effectively shrinks physical distance, time and information space. It is therefore necessary to anticipate a future where more and more learning will take place in the home, office and place of work throughout a lifetime.

Technology can not only change the way education is structured but also accelerate a paradigm shift in thinking about learning and teaching. This article explores the existing paradigm, explains the new technologies and demonstrates how they can change the pupil–teacher perspective, particularly concentrating on how change required can be managed to ensure a universal acceptance of the contribution that new technologies can make to education.

Teaching and Learning Paradigms: Old and New

In the ancient world, informal education often took the form of discussion or dialogue between mentors and their protégés. The protégés learnt either by directly questioning the mentor or by observing and noting the mentor's responses to other people's questions. This didactic or instructional approach assumes that the teacher possesses the knowledge and this

knowledge can be accessed by detailed and critical dialogue. This model of learning is effective as long as the group sizes are small, the knowledge domain is restricted, and the student sufficiently mature and confident to both know and be able to articulate the right questions.

Over the centuries, the approach has become distorted. Until a few decades ago, the emphasis shifted away from dialogue to a 'deficit' model¹ that assumes knowledge is in some way transmitted from someone who 'knows' to someone who doesn't. Dialogue was no longer such an important component. In particular, large class sizes and the belief that most children lacked the cognitive capabilities for entering into meaningful dialogue led to an emphasis on the learner's lack of knowledge and an adoption of a passive rather than an active role in the learning process. However, children are not naturally passive, and the deficit model often marginalised the active or easily distracted child, or those with different learning styles and abilities. More recently, there has been considerable progress in trying to anchor teaching objectives to a real-world setting and to attempt to make learning more interesting and more of an active 'discovery' process. This change has been most noticeable within primary schools, whereas secondary and higher education still retain a heavy reliance on 'chalk and talk'.

Today's learning paradigm should be learner-centred, collaborative and constructional, with the emphasis on the learners constructing or building their own understanding through an active process². The learner can then take more control of the learning and learn through not just listening but

Figure 1—Concepts: teaching paradigms

by working collaboratively with other learners, or the medium.

The active and social nature of the 'construction' approach to learning requires and encourages the development of a new set of learning and teaching skills and competences. In particular, exploration, communication, negotiation and other personal or social skills have to be developed to exploit the cooperative or collaborative element successfully. Likewise, the teacher's role changes from being the source of all knowledge to a 'facilitator' that enables students to develop and take control of their own learning. The differences between the two paradigms are illustrated in Figure 1.

Although there has been a gradual paradigm shift towards the 'constructive' approach to learning, the new technological advances, which could support and accelerate this process, have not been universally applied and assimilated within educational institutions. Now learning can take place well beyond the classroom through adopting a 'telematics' approach.

The Telematics Approach to Learning

The range of technology and associated functionality covered by the term *telematics* is broad, including computer conferencing, e-mail, video and audio conferencing, sharing of files, chalk-boards, control of remotely located applications, scanners, Internet access, and so on. Its main applications have been in open and distance learning, and the combination of the technology and its functionality is commonly referred to as the *distance learning platform* (DLP).

DLPs promise to deliver life-long learning which can be:

- anytime—when you want it;
- anyplace—in your home, office, place of work, hotel, car; and
- anywhere—global availability.

	INSTRUCTION	CONSTRUCTION
CLASSROOM ACTIVITY	TEACHER-CENTRED DIDACTIC	LEARNER-CENTRED INTERACTIVE
TEACHER'S ROLE	KNOWLEDGE PROVIDER ALWAYS EXPERT	KNOWLEDGE RECEIVER SOMETIMES LEARNER
LEARNER'S ROLE	LISTENER ALWAYS LEARNER	COLLABORATOR SOMETIMES EXPERT
EDUCATIONAL EMPHASIS	FACTS MEMORISATION	UNDERSTANDING DISCOVERY
CONCEPT OF KNOWLEDGE	ACCUMULATION OF FACTS	TRANSFORMATION OF FACTS
SUCCESS CRITERIA	QUANTITY RETAINED	QUALITY OF COMPREHENSION
ASSESSMENT	NORM-REFERENCED WRITTEN EXAMINATIONS	CRITERION-REFERENCED PROJECT PORTFOLIOS
USE OF TECHNOLOGY	DRILL AND PRACTICE	COMMUNICATION, COLLABORATION, PRESENTATION INFORMATION ACCESS AND RETRIEVAL

(Adapted from Dwyer, 1994)

The issue is not one of technology—the platforms (videoconferencing and Internet) exist and are well understood. The problems arise more from application and the lack of the information tools that allow the exploitation of the platform to achieve educational benefits. Some of these can be explored by gaining a better understanding of the telematic perspective on 'distance' and 'learning'.

Concepts of Distance

Telematics can remove the physical distance between teachers and

learners with education transported to the people rather than the people going somewhere to be educated. In the future, more people will be educated in their own homes, their work places or in any place of their choosing, and through a variety of just-in-time (JIT) methods. Such conditions could be essential in realising the benefits from life-long learning for a future of even faster technological change.

Physical distance is also critical if one considers learning as a social process (Figure 2). Indeed, one of the main reasons people register for

Figure 2—Learning as a social activity using a PC-based videoconferencing system (VC8000)



Figure 3—Substituting computers for teachers is not always desirable

traditional courses run out of educational establishments is to meet other people. For children, this is an essential part of socialisation; they have to learn their social skills, norms and more from interacting with others. Any telematic solutions need to support this social aspect. There is a real and genuine concern that technology will encourage physical and psychological isolation. The image of a single child sitting in front of a computer for most of its waking hours is one that any telematic solution must dispel (Figure 3).

The DLP paradigm has been very effectively applied in schools located in remote regions or where universities have a number of different and distant campuses. In the former case, the need is to bring together small numbers of students into a 'virtual' cohort so that they can collaborate and share experiences with other remotely located children of the same age or ability. Equally, it allows teachers to be trained or specialist teachers to be 'brought' in without having the inconvenience and expense of physical travel.

The university experience of telematics is often driven by the needs of the teacher. It is designed to allow student access to tutors by a more economic and efficient means. The same lecture can be presented to a number of sites simultaneously and at least it provides the appearance of a normal lecture (that is, live and interactive). New technology also facilitates independent research and study, and student communication beyond the institution is increased.

Temporal distance refers to the time relationship between the learner and what needs to be learnt. It assumes that most learning has to be timely in that, to maximise learning, the learner has to be ready to receive the learning experience and, increasingly, has a need. For example, the learner may lack sufficient prior learning or cognitive ability to understand or make sense of the new material. Equally, readi-



ness could reflect the appropriateness of the context or situation. The learner may not be in the right situation to either receive or benefit from the learning. The concept of *JIT learning* arises from this situational constraint. There is a right moment to learn when the problem is very salient and the requirement is very immediate.

Psychological distance refers to more social and cultural differences. For example, neighbours may be physically very close, but may be 'miles apart' in terms of interests, attitudes or other social factors. Allowing, for example, children in an inner city school access through videoconferencing to a rural school must surely broaden cultural horizons and undermine the development of harmful stereotypes.

Recent work has concentrated on networking schools to support collaborative learning, with Suffolk schools working on joint projects with pupils in Wales. Not only do they learn about the contrasting geographical differences (both physical and economic) but they also gain an opportunity to make new friends with different backgrounds. It is intended to expand these links with schools in Scotland, California and Canada.

One of the interesting findings arising from observing children in their use of videoconferencing technology is the role of gender in the

children's friendship patterns. Within their own schools, most friendships are of the same sex, but across schools most contacts are, surprisingly, made with children of the opposite sex. Children appear less inhibited in forming cross-gender friendships over a distance. Although these are the results from one pair of schools, it does suggest that by breaking down this polarisation of friendship at such early ages may help also break down some of the gender stereotypes prevalent in society.

Other findings show considerable improvements in children's self-confidence and self-esteem. In many cases, the children understand the technology as well as their teachers. Indeed, teaching the teachers is critical in building self-confidence and seeing the pupil-teacher relationship as a partnership built on mutual respect and shared experiences. In terms of collaboration, the schools soon identify clusters of expertise held by different learners and these are quickly communicated and shared. The pupils recognise themselves as being part of a learning organisation where skills and information can be exchanged for mutual benefit.

Teachers, too, are beginning to see the power and benefits of this new technology. In one set of schools which had no previous contact or

Figure 4—A telematics-based model of learning

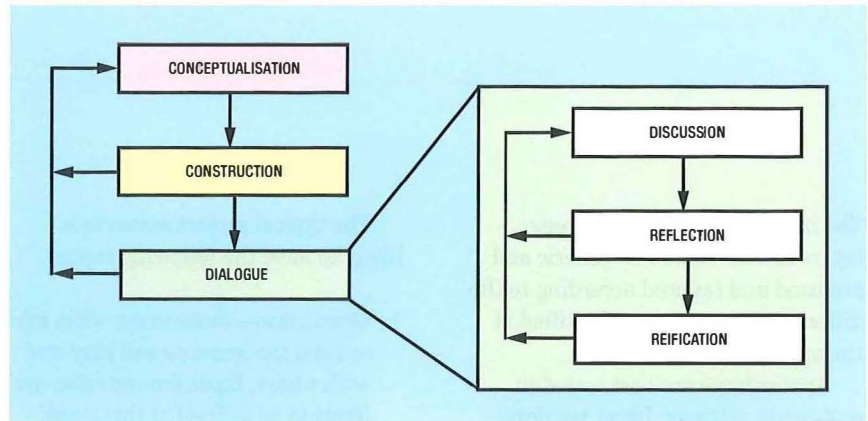
exposure to this type of technology, clear management-of-change issues have emerged. There was initially a significant and understandable element of technophobia. However, this fear was founded in a lack of technical awareness rather than an adverse attitude to technology. It took time to recognise what the technology was capable of doing and to position it in the right educational context to exploit its power fully.

The Learning Model

The learning model which underpins our Education and Training research programme was developed in conjunction with Heriot-Watt University. The major components of this model are summarised in Figure 4 and a description of the various elements in relation to secondary and tertiary education follows.

Conceptualisation refers to the users' initial contact with other people's concepts. The learners' contacts are traditionally through text books or seminal readings—these are referred to as the *primary courseware*. Primary courseware is resilient; it does not need frequent updating, and paper-based media are an adequate way of publishing, though electronic media provides tools to help search, browse, and otherwise enhance the information.

Construction refers to the process of interpreting and combining concepts through their application to some meaningful task such as laboratory work, writing, preparing lectures and presentations etc. The results of such a process are products like essays, notes, handouts, laboratory reports and so on. Such material, or *secondary courseware*, is volatile and of varied quality. The Internet in particular lends itself to publishing such courseware. The tools available are also important for supporting aspects of the construction process. Generally these are search, retrieval and authoring tools, but there is a need for a set of 'mind tools' (for example, an 'outliner' on a



word-processing package) which can help learners to reflect about their concepts at a deeper level than they might otherwise do.

It is in the *dialogue* phase that the telematics approach comes into its own. Dialogue refers to the debate, discussion and conversations that take place between learners, or between learners and teachers. Learning can result from the type of question and from reflecting on the answers. The roles of *discussion* and *reflection* are paramount to deep learning, but so is the concept of *reification*. This literally means 'making an object of' and in this learning context refers to the output of reflection and discussion in terms of the construction of a new understanding or conceptualisation. The learner is now ready to come into contact with new ideas and concepts based on the new understanding, and so through a process of iteration the learner returns to the conceptualisation stage.

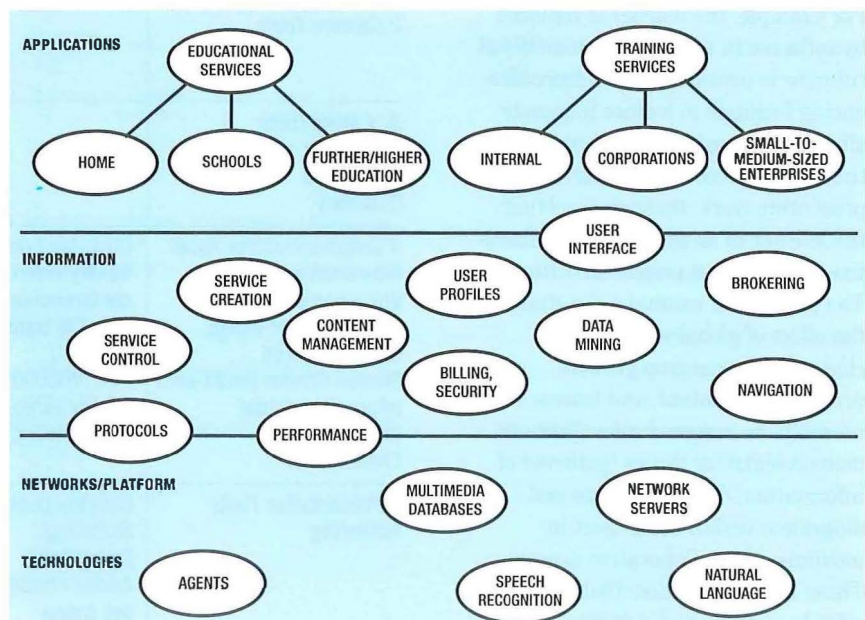
The products of dialogue (the conversations, questions and answers etc.) are also important sources of material for other learners. If the material could be captured and processed, it could also be reused so that new learners can benefit from other peoples' learning experiences. Evolving or organic databases like 'answer gardens', or more simple notions like 'most frequently asked questions', should be part of the *tertiary courseware* in any telematics solution.

From such a model, it should be quite clear that the power of any telematics solution increases as one goes from conceptualisation, through construction, to dialogue. The real strength of telematics is therefore through supporting the interactive components of learning.

The Platform Requirements

Figure 5 shows a representation of the key technical layers of the DLPs.

Figure 5—A layered model of the technical requirements



The information and technology layers can be viewed as generic and are used and tailored according to the different requirements identified at the application level.

Applications are best tested in real-world settings. Ideas are developed and refined in a laboratory and tested **in action** with the appropriate groups in the natural environment on a test bed. The effectiveness of the interventions is then evaluated, problems identified and solved, and new solutions implemented. The process is, of necessity, dynamic and iterative. A test bed is currently based on integrated services digital network (ISDN) service with servers linked to the Internet. Currently, four schools in East Anglia are on the network, with a further two in Wales. Within the next year, it is planned to add Suffolk College to the educational test bed and to build *integrated project environments* (IPEs) in the four Suffolk schools.

Integrated project environments

IPEs are dedicated learning environments rather than electronic classrooms. This is an important distinction as electronic classrooms all too often use technology to emulate poor practices in existing classrooms. For example, the teacher is replaced by software in the form of an artificial tutor, or is provided with videoconferencing facilities to lecture to remote sites. In contrast, the IPE emphasises the learner discovering, sharing and presenting work, through involving the learner in an active and collaborative role within a project structure. The project (for example, the study of the effect of global warming on climate and thus crop growth) provides the context, and learners will normally be assigned roles (farmers, meteorologist) or duties (gatherer of information, drawer of maps and diagrams) within the project to maximise the collaborative aspects. These roles can be distributed across schools, nations and cultures.

The typical project scenario is likely to have the following stages:

- 1 *Orientation*—discovering what role or roles the learners will play and with whom. Basic ground rules are likely to be defined at this stage.
- 2 *Information gathering*—the learner will search and retrieve information relevant to their role. This may involve collecting new materials using a variety of recording devices (for example, camcorders, cameras, tape recorders).
- 3 *Conceptualisation*—the learner will synthesise, interpret and formulate key concepts for the project from the gathered information.
- 4 *Refinement*—the learner will test out his or her concepts through discussion, questions and answers

and expert review. The concepts may well be refined as a result of sharing his or her understanding with a range of different people.

- 5 *Presentation*—ideally the results could be published on the Internet and become itself a resource for other users.

The role of the teacher in such a scenario is to facilitate and guide the learner. The technical requirements to support an IPE are given in Table 1.

These IPEs are not just important for exploring new ideas and testing this new paradigm. They are also our 'Trojan horse'. Their very existence will cause change. The issue is to accelerate that change. We see the IPEs becoming personal development centres for teachers from other schools. They can visit the schools on our test bed, gain hands-on experience of using the technology and

Table 1 Sample Requirements for an Integrated Project Environment

Tools	Applications	Content/ Courseware	Hardware/Network
1, 2 <i>Search and Retrieval</i> Browsers Navigation aids Text summariser Agents	Netscape (browser)	<i>Primary</i> Text Videos Stills Voice Music	Multimedia PC/Mac with Internet access ISDN2
2 <i>Capture Tools</i>		<i>Primary</i>	Camcorders Digital cameras Tape recorders
3, 4 <i>Mind Tools</i> Simulators Visualisers Outliners			
4 <i>Communications Tools</i> Videoconferencing Voice conferencing E-mail (text or video) Bulletin boards Shared Spaces (multi-user interactive virtual environments) Chalkboards	CU-C-Me? (low- fidelity video- conferencing) FTP (file transfer protocol) PCC (VC8000) (higher fidelity video- conferencing)	<i>Tertiary</i>	
5 <i>Presentation Tools</i> Authoring	Desktop publishing Scanning PowerPoint Adobe Photoshop MS Office	<i>Secondary</i>	Laser printing Scanner Video-editing suite

make the all-important links with other more experienced teachers. It should be a window onto the future and, as such, sow the seeds of change.

A Word About the Internet

Too many people believe that the solution to all educational problems lies in connecting schools to the Internet. Connection is not enough. Learners must be provided with a lens through which they can filter out distracting and unhelpful material, and a set of guidance tools which allows them to locate and use what information they need for the learning task in question. The Internet itself needs a new layer which genuinely provides a supportive environment for learning. Many of these and other innovative ideas are being implemented on BT's CampusWorld, which will give schools cheap, safe, and beneficial access to the Internet.

Conclusions

Current classes are too big; books are in short supply; the information world is expanding exponentially, teachers, learners and parents are becoming disenchanted—a change is needed. The educational technology and pedagogy to bring about that change are now available. The question is do we have the money and the will to change? Undoubtedly, the hardware described is not cheap, but costs will reduce over time, and the benefits are very real. Educational institutions need also to look at the cost saving from using these new environments (for example, travel and subsistence, sharing of resources), as well as revenue generation opportunities (for example, the use by local business out of school hours and terms).

A number of areas have already justified the cost of installing some of this hardware (in Wales and Scotland) and are making significant savings. But by far the greatest

challenge is to win the hearts and minds of the teachers. This is best achieved by showing them the educational benefits resulting from the use of the technology. Let us also not forget that learning and teaching should be fun, and this technology is fun to use—as our children already know!

References

- 1 DWYER, D. *Apple Classroom of Tomorrow: What we've learned*. Educational Leadership, Apple Incorporation, 1994.
- 2 PAPERT, S. *An Introduction to the 5th Anniversary Collection*. In Harel, I. (ed.) *Constructionist Learning*, Cambridge, MA, MIT Media Laboratory, 1990.

Biographies



Chris Fowler
BT Networks and
Systems

Chris Fowler is a psychologist who works as a Technical Group

Leader in the Human Factors Unit at BT Laboratories, Martlesham Heath. Before joining BT in 1990, he completed a Ph.D in Cognitive Psychology at London University and then lectured for 12 years in the higher education sector. He now manages BT's Education and Training Research Programme.



Terry Mayes
Heriot-Watt University

Terry Mayes is Director of Research in the Institute for Computer-Based

Learning a position he has held since its formation in 1990. Previously, he was Deputy Director of the Scottish

Human-Computer Interaction Centre at Strathclyde University from 1986-1990. He studied psychology at Bristol University and gained a Ph.D. in Human Memory at the University of Newcastle-upon-Tyne. As a lecturer in psychology at Strathclyde, he published widely in cognitive aspects of learning and in the experimental psychology of human memory. He has worked extensively on the development of interactive learning from hypermedia. Current work centres on videoconferencing and on using computer-mediated communications for distance learning, on cognitive mindtools, and on organisational issues concerning the large-scale development of advanced learning technology. His commercial experience began with the formation of a company to manufacture teaching machines. He has consulted widely in industry. Currently he is working on projects with the Defence Research Agency and BT. He is a member of BT's advisory panel for education.



Bernard Bowles
BT National Business
Communications

Bernard Bowles graduated from Manchester University in

1971 with an M.Sc. He held senior positions in a number of private companies specialising in data transmission in the oil and mining industries before joining BT in 1987. Initially, he managed BT's managed router service development and is currently leading the development of broadband local community services within ATM Marketing in National Business Communications such as for linking higher education, further education and schools. He has recently presented a paper for Telecom 95 and written a number of articles on enterprise opportunities in the health and manufacturing sectors.

Bonnie Ralph

Key Technologies for the Information Industry

This short editorial note introduces the final section of the theme 'The Information Industry and its Key Technologies'. In the October 1995 issue, we opened with an introductory article on the information society. In the January 1996 issue, we examined the emerging characteristics of the converged information industry. We saw how new roles and types of player are emerging, and how, in many parts of the industry, new ways of doing business are emerging.

This issue considers the key technologies on which the success of information services rests. By no means all the pieces are in place. You could say their lack constitutes a barrier to the development of an information-based society, but I think it is truer to say that industry developments take place only in response to the availability of new technologies. It is by no means the case that new techniques always emerge in response to a clear commercial need. Many of the key strands of the new technical environment came about in response to a private problem (World Wide Web is a prime example—its creator at CERN wanted to use the Internet to correspond easily with co-workers around the globe) or even as the result of simply 'mucking about'. Very, very little seems to have developed as an explicit answer to a fully-articulated customer requirement.

It may very well also be true that for every major technology or concept which has made its mark with the user community, and has been recognised in the following pages of the *Journal*, many have died for want of a useful application, or have been left to languish in obscurity on a few university workstations.

Exceptions to this 'natural selection' are the concept of the 'distributed processing environment' (DPE) and TINA (telecommunication information network architecture). A thorough-going implementation of neither of these concepts actually exists. It is more that people believe they **ought** to exist, and that benefit would be derived if they did. DPE and TINA are concepts with particular attraction for the telecommunications component of information industry culture (although, of course, DPE is an IT concept, born out of IT client-server thinking). They are **systems** (in the same way that the 'intelligent network' is a system) and are all-embracing. Many of the key components of information networking, such as the 'trader' concept in navigation, owe their acceptance to DPE or TINA, but, by and large, the builders of today's techniques seem prone to finding the loose end of a problem and pulling at it, rather than trying to build a plan that solves **all** the problems.

The articles in this issue are divided into groups. The first group deals with these overarching systems. First is the Internet. The Internet is, of course, actually existent and in use. The word 'Internet' has to be interpreted according to the context, and Paul Jenkins does that for us. We then have the DPE vision, explained for us by David Freestone and colleagues, and the TINA vision, articulated by Tom Rowbotham and Martin Yates.

We then have a group of technologies which enable **access** to information over the information network. Key to future manageability of information resources is the

concept of the agent (carefully unwrapped and set before us by Robin Smith). Navigation, and information management, key areas in which BT has highly-developed skills and has built useful capabilities already, are covered by Jonathan Legh-Smith and Keith Preston respectively.

Information networks will not work without navigation, and they won't work without security systems either. Chris Gibbings tackles the issues for us.

Finally we have two articles about human needs and perception, and about how technologies are developing to underpin interface mechanisms which allow human users a more satisfactory relationship with cyber-environments of many kinds. Andrew Hockley explains the needs in his article 'The Information Needs of Network Communities'. This leads into our final article by Graham Walker and colleagues on 'Interactive Visualisation and Virtual Environments on the Internet'.

I hope you have found this theme helpful and thought-provoking. It is about real commercial and technical developments which are taking place **now** all around us, and in which BT is participating energetically. You can see that it takes hard work to keep up with the industry. Fortunately, it's also fun!



Bonnie Ralph
BT Networks and
Systems

Paul Jenkins

Internet

The Internet is enabling information held on computer systems around the world to be accessed by terminals with suitable access rights. This article provides an introduction to the Internet. It describes where it came from, how it works and what it may become.

What is the Internet?

The Internet means many different things to many people. As with all technology that receives a lot of attention in the media, the range of reactions to the Internet is very broad, from a place to meet friends to a new thing to be feared. This article looks through some of the more extreme descriptions to describe what the Internet is and may become.

At its simplest, the Internet can be described as a number of independent computers that are linked together using transmission control protocol/Internet protocol (TCP/IP), and other related protocols. A terminal can access all of the computers once it has accessed one of them (subject to security mechanisms). On top of the TCP/IP protocol is the application protocol. The application protocol is chosen according to the application in use (for example, e-mail, hypertext browsing, etc.). It operates directly between the terminal and the system on which the application resides. While the terminal is connected to the Internet, it can be involved in communicating with a number of different applications in different systems. Indeed, the hypertext browsers like Mosaic and Netscape operate so that the terminal system communicates with a series of systems as part of the application.

The Internet is not administered as a whole system, although overall administration of certain things is necessary (like the allocation of names and addresses). Indeed, the independence of system providers and the freedom of speech enjoyed by the users is held as a major benefit by the users. Many of the current systems, and the communication links between them, that are used to create the Internet are paid for by academic and government bodies.

This is a hangover from the early Internet use, when the costs would be borne centrally as support for academic activities as a whole. With the increase in commercial use of the Internet, the infrastructure is increasingly being provided by commercial concerns.

How Did It Start?

The protocols used for the Internet started to be developed in the late-1960s as part of a US Department of Defense project to interconnect the facilities undertaking research projects for them. After some early starts, by 1973 the ARPAnet had become a network that was available across America. By the mid 1970s, the network included many sites outside America and was experiencing increasing traffic. Many of the design assumptions made during these early stages of development remain embedded in the current protocols.

Internet origins

- Started as a US defence research project
- Purpose to link scientists and researchers together
- By 1973, a national network
- In 1983, the ARPAnet split into the Internet and Milnet
- Early 1990s, over 2 million users, primarily funded by government and universities
- Today, estimated 25 million users
- Commercial companies providing access

The IP is used to link networks and leased lines together to form a platform of connected systems that can reach beyond the scope of any single network technology.

The fundamental issue during these early years was to interconnect islands of computer resource that consisted of one or more mainframes serving a large number of terminals. In such an environment, only the mainframes needed to operate the protocols and hence only they had addresses. As a result the early protocols had only 1 octet of address space (allowing 256 systems to be identified in the whole world) and addresses were allocated as a flat address space rather than as part of a hierarchical address structure to simplify routing (as is done in public network numbering plans). As time progressed, the number of users increased and the computers moved away from a mainframe architecture to a client-server one that causes the terminals to require IP addresses as well as the servers. This increased demand for addresses created problems and their length was increased from 8 bits to the current length of 32 bits (able to identify almost 4 300 000 000 systems). However, IP addresses are still allocated as a flat address space, a practice that is now causing difficulty with routing systems.

In addition, the mainframes were either connected by leased lines or local connections using manufacturers' proprietary protocols. Owing to the variety of proprietary protocols used to transport it, the IP protocol made minimal assumptions about the capabilities of the network technologies it was being carried over. The IP assumes that the underlying protocol when asked to transfer a sequence of packets may deliver the packets with some packets missing (lost packets), with other packets delivered more than once (duplicated packets) and with no relationship to the order of the packets that were sent (mis-ordered packets). Although the underlying network technologies have become far more reliable, complex error-recovery mechanisms are still embedded in the protocols.

How Does It Work?

To understand how the Internet works it is necessary to consider three protocols and three types of system that are important to the Internet's connectivity. The connectivity between systems is achieved by using consistent protocols at key layers of the communications architecture, specifically TCP, IP and user datagram protocol (UDP) (Figure 1). The IP provides connectivity across dissimilar networks, and TCP and UDP operate over it to provide the application with either a reliable or unreliable communication path, respectively. The key Internet systems are routers, hosts and domain name servers (DNSs). Routers act on the IP protocol to pass information (messages) between hosts. Hosts are either a computer where the desired application resides, or the terminal trying to access it, and hence the source, or destination, of an IP packet. DNSs perform the mapping from user-friendly names like archstan.agw.bt.co.uk to the corresponding IP address (for example, 151.147.136.70) which is used by the IP protocol.

IP is used to link networks and leased lines together to form a platform of connected systems that can reach beyond the scope of any single network technology. This is achieved by placing routers at the boundary between the networks being interconnected. These routers operate IP protocol between them to convey all the information required to




Figure 1—Protocol stack

deliver the communication to the desired destination. The IP protocol itself is so flexible, in terms of the networks over which it can operate, because it requires very little from the underlying network technology.

The IP is a packet-based protocol (Figure 2), which is to say it conveys blocks of data from a source point to a destination point, and it operates in a connectionless mode. Connectionless operation means that the protocol does not establish a route across the networks for a series of packets to follow, but each packet has all the addressing information required to reach the destination and is routed independently of the others (in other words, it is more analogous to a series of post cards than a telephone conversation). This means that a user-to-application communication, which is likely to involve multiple packets in both directions, is not constrained to follow the same route through the network because each packet will be routed independently of the others. This provides a significant degree of network resilience in that, if a router fails, the communications can continue unaffected as none of them

Figure 2—Packet-based protocol

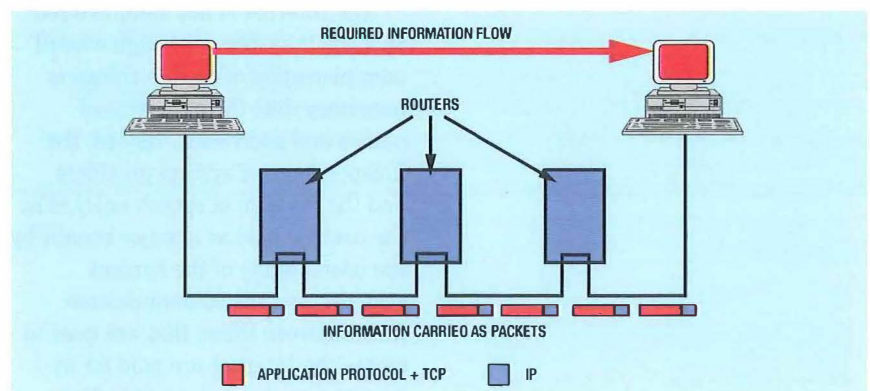


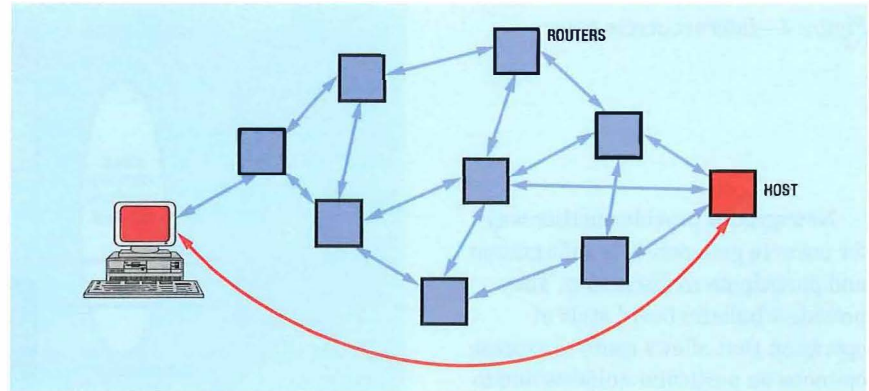
Figure 3—The Internet network

is intimately bound to that router. However, the down-side is that the router has to undertake a route-determination process for every packet it receives and it cannot control the flow of packets directed towards it by exerting flow control on the sending systems. If a router does become congested, it discards excess packets until it has recovered.

The routers forward IP packets towards their destination based upon information contained in a local routing table. This table provides a mapping between destinations and where to go next in order to reach it. It should be noted that this information is based solely on the ability to reach a destination, and there are no quality-of-service considerations in selecting a next hop on the path. The routers operate a variety of related protocols between themselves, and with hosts, in order to maintain the accuracy of these routing tables.

In this way, the IP provides a means of carrying blocks of information across an arbitrary set of networks, but it does not try to recover from any lost, duplicated, or mis-sequenced packets. It delivers such packets as they arrive. It is up to the hosts to operate a higher-layer protocol between them if error recovery is required. TCP is a complex protocol that recovers from the possible errors generated by IP and the underlying network technologies to provide the applications with a reliable communication path. UDP takes no recovery actions and just passes the information on to the application as it arrives. Recently, UDP has become more important for applications (like video) where the delay involved in recovering the lost information is more damaging to the application than failure to deliver it.

The hosts operate application-specific protocol end-to-end over TCP, UDP or both, in order to provide the desired application to the user (Figure 3). Before initiating a communication, the sending host often needs to obtain the destination address from the name information



provided by the user. The host does this by undertaking a protocol exchange with a DNS server. There are three possible results from such an enquiry:

- the DNS has the information and returns the required IP address;
- the DNS does not have the information but contacts another DNS server to get it and return it to the host; or
- the DNS does not have the information but returns the location of another DNS server that may have it.

The latter two options allow this directory look-up process to operate as a distributed service, but they do have response-time implications.

What Does It Do?

In addition to the basic transport protocols, which provide a global interconnected set of systems, capable of relaying electronic information around the world, there are a number of protocols, developed by the Internet Engineering Task Force (IETF)[†], to support the access and retrieval of information on remote systems. As a result, the information held on computer systems around the world can be accessed by a terminal with suitable access rights. Because of the academic background of this

[†] The Internet Engineering Task Force (IETF) is a committee that meets regularly to agree developments and enhancements to protocols used on the Internet. The body has representation drawn from technical experts from around the world.

interconnected set of systems, a vast number of systems allow free access to the information stored by any remote system. The result is an environment that offers the user access to an enormous amount of free information.

The remote information was traditionally accessed using *Telnet* (a terminal emulation application) or a file transfer protocol (FTP). However, these required the user to be familiar with UNIX (a computer operating system which is as obscure as DOS, but different) and hence were not user-friendly. The recent development of user-friendly browser applications (like Mosaic and Netscape) has led to the availability of many World Wide Web (WWW) sites that can be browsed by using these applications. Thus the free information resource has become much easier to use. This new-found ease of use has brought with it additional benefits; specifically it has raised the level of awareness of the Internet of an entirely new group of people who are not, and never will be, conversant with the detail of computer operating systems.

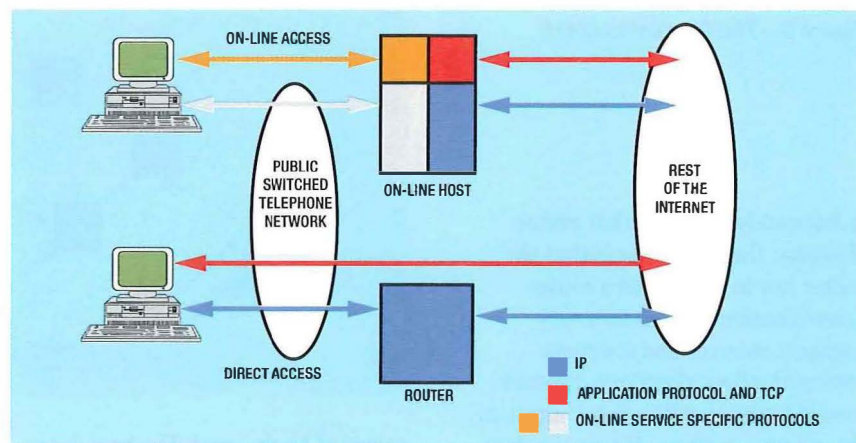
This growing Internet population has had another effect, which is to promote the use of Internet-based e-mail as the primary form of electronic mail, especially between companies. Addresses like fred@bloggs.bt.co.uk are appearing on business cards and it is the ubiquity of availability that is the Internet's major advantage. Although the service itself is primitive and can transfer only plain text messages (not binary files like word processor documents), a series of translation programs has been developed to turn a binary file into a text file and back again. Most mail reading programmes will invoke the translation automatically and hence give the user effective binary mail transfer.

Figure 4—Internet access types

Newsgroups provide another way for users to gain access to information and participate in discussion. They provide a bulletin board style of operation that allows users to express opinions on particular subjects and to respond to others' opinions. There are many such newsgroups and these can generate very large amounts of data. The groups are themed so that the messages are related to a specific topic, and some are controlled by moderators who will sift contributions for relevance before they are added to the group. Newsgroups are at once very useful and very controversial. Groups exist for discussing software problems, religion, sex, politics, and many other topics. Much of the controversy about the Internet, and many of the calls for censorship, are related to newsgroups.

The basic problem in using the Internet as a source of information is finding a way through the vast quantities of spurious and irrelevant information to find what one is looking for. Searching for information on a particular subject can make finding a needle in a haystack a very trivial problem indeed. Some help is at hand with a more complex set of programs, called *search engines*, that are designed to search the Internet for the location of information on a particular topic.

The engines vary in the level to which they search and whether they search the Internet itself, or reference sites of information on what is available through the Internet. Reference sites contain lists of locations on the Internet and the types of information they contain. The information could be gleaned from the file names, the directory structures, or a keyword search of the files themselves. Some search engines then use these sites to undertake a search and locate potential sources of information. Another class of search engines (*crawlers*) search through the Internet to find matching information. Crawlers are often used to establish and maintain the information held in reference locations.



Even with all this help, it can be very difficult to find the information being sought.

Myths

The Internet is free!

As has been mentioned earlier, the Internet's history has been to provide network services (at no cost to the user) that provide access to systems containing information (for which there is no access or copying charge). However, just because the user does not pay for the service, it does not mean that there is no cost associated with it. Historically, the costs of providing the service would be borne as an element of a larger project (for example, academic research) or through some central funding for the common good (for example, governmental support for communications between academic sites). Thus the costs associated with the communication were borne by third parties that were not directly involved in the communication.

This history still has an effect on the way the Internet recovers costs today. Since the costs were not associated with the use of the resources, none of the functionality that would be necessary to provide usage information from which to derive billing is available. As a result the vast majority of Internet access providers (as opposed to on-line service providers) generate their revenue from a subscription fee (monthly, or annual, is usual). This is not unreasonable, since the vast majority of a service provider's costs are fixed, and there is little additional cost whether the service is fully occupied all the time, or never used.

Related Activities

On-line services

On-line service providers (like CompuServe) also offer their customers electronic access, via a dial-in line, to information services (see Figure 4). However, these have one crucial difference compared to Internet access in that the terminal is not connected to the Internet but is connected to a host computer and interacts with the applications on the host. The access protocols used by the terminal are defined by the host's requirements and are not necessarily anything like the protocols used by the Internet. This form of approach has many advantages stemming from the contractual relationship between the on-line service provider and the user. For example, since all the contact information is known in order to undertake billing, it is also possible to use this information to support commercial services such as shopping.

However, the similarity of service offered to that of the Internet has resulted in a convergence of these markets. In order to compete, the on-line service providers are offering gateways into the wider Internet for their customers. In the early days the host acted as a proxy. In other words, the host communicated with the Internet using TCP/IP and the Internet application protocols while the terminal controlled the host end of the Internet application by a service provider specific protocol. Unfortunately, this prevented the on-line customers being able to use some of the Internet applications (like browsers) and hence was not as flexible as a direct Internet access.

The Internet is moving away from a free, academic, uncontrolled environment, and towards a means to a commercial end.

Today, most on-line service providers offer direct Internet access in addition to their specific on-line services.

Intranets

The links between the approval process for Internet specification and the availability of interoperable implementations mean that there are cheap implementations of the Internet protocols widely available. Some computer systems are sold with the software already installed; other systems require a small additional cost. In any case, the specifications are open and their development is outside the control of any one manufacturer. As a result, many companies see the use of Internet protocols as a practical way of addressing their internal communication requirements. The resulting private networks that are based on Internet protocols are termed *intranets* to indicate their closed nature.

This use of Internet protocols is very appropriate because the origins of the Internet were as a private network to support US defence research activities. Hence the underlying design of the protocols is more suited to this use than to their use in providing public network services. The real difficulty arises because most companies do not want their intranet totally isolated from the public Internet, but seek to have some controlled access. The point of connection between the Internet and the intranet is a gateway that usually includes security features intended to prevent unauthorised users gaining access to the intranet through it. Such a gateway is often called a *firewall*.

Where Is It Going?

The Internet is undergoing a massive change in its philosophy. It is moving away from a free, academic, uncontrolled environment, and towards a means to a commercial end. By 1990, the number of users had increased to an estimated 2 million. These users were generating a significant level of

traffic and the usage profile had begun moving away from the original academic/research focus envisaged for the network. Today the number of users is estimated at 25 million and the focus of use has clearly migrated away from its academic origins. The growth has been so great that in the near future the access protocols will be revised to increase the number of addresses available as the current 32 bit address space will run out in a few years.

There is a growing desire to explore the commercial opportunities presented by this fast-growing customer base that can be reached with relatively little cost. The range of commercial activities being explored is wide and varied. The scope includes providing a value-added information service at a cost, using the Internet to make contact with customers and/or sell goods and services that are not inherently reliant on the Internet (for example, car hire), and offering telecommunications services via a TCP/IP access. In order to support such services, there are a wide range of protocol developments under way to produce a technical framework that is capable of supporting the requirements of the proposed commercial services.

Trading

The obvious development to support trading on the Internet is the body of work to add security capabilities for commercial transactions that can protect the purchaser from the risk of fraud through the reuse of the information (for example, credit card numbers) provided for a transaction. For internal communications, or prearranged communications between companies, it is clearly possible to agree on the security mechanisms to be used. However, in this new selling situation, the transfer of funds is being agreed upon between two parties which do not have a prior agreement on the security mechanisms to be used.

Another aspect that is coming to the fore is service quality. The traffic

levels on the Internet have been increasing, but, as there were no direct user costs involved, the slow, and sometimes absent, service has been tolerated by a user community that understood how the network worked and hence the sort of things that can go wrong. The population of Internet users now includes people who do not understand the complexities of networking, and hence are less tolerant of the visible effects of these Internet errors. If the Internet is being used to sell goods and services, then the Internet is being used to support the image that the company is presenting to the user. Good or bad Internet links could be the difference between making or losing the sale. Companies entering into this market will increasingly demand higher-quality service levels to ensure the quality of their corporate image.

Real-time services

The use of the Internet to convey real-time service is relatively recent. The challenge for such services is that the existing Internet infrastructure has developed around the requirements of text-oriented messages and is not well suited to the demands of other media types. Having said this, there are several developments to try and use the Internet to convey such services. Examples include Netphone (supporting telephony), and CUseeme (a program that aims at providing real-time videoconferencing).

The major perceived benefit of using the Internet to support these services is the cost benefits to the user of the Internet compared with traditional telecommunications networks. These benefits accrue because the Internet's infrastructure is very effective at getting the most use from its internal links by sharing the capacity between all the communications on a per-packet basis. However, this form of operation means that the network resources are not reserved for the use of a communication but are grabbed for use and released as each packet goes by. This

causes an effect called *jitter* which is of no consequence to data traffic but has significant effects on real-time services. While the infrastructure can cope with a certain amount of this type of traffic without generating too much jitter, as levels increase the effect of competition for resources will also increase and the overall service will become less acceptable.

In order to address this problem, there is a body of work that may lead to a change in the Internet infrastructure in order to better support the jitter-sensitive services. The proposal is to operate a bandwidth reservation protocol (called *RSVP*) so that jitter-sensitive applications are not competing for resources throughout the communication. If such a proposal is to succeed it will require the willingness of the Internet service providers to modify their elements of the infrastructure, in concert. The major application currently envisaged for this reservation protocol is in conjunction with a multi-point service where there is a single source of information (for example, a video feed from a conference) that is being distributed to many recipients.

In addition to the opportunities offered by such development, their effect on the simpler data traffic that the Internet supports has to be considered. A real-time communication will occupy a significant set of communications resources, whether by explicit use of the RSVP protocol or by continuous competition for resources by a continuous stream of packets. The existence of several such communications will remove large chunks of resource from the pool available to support other communications which do not have such stringent requirements. Thus the support of real-time communications could have a significant impact on the perceived quality of the service being provided for current applications. The coexistence of these two different types of traffic is leading to the consideration of support for a range of qualities of service in the future Internet.

Innovation

One of the big reasons for the growth of the Internet and its rapid evolution is due to its culture and the innovative nature of contributors to it. Historically, the network was intended for helping government researchers and, later, university researchers, collaborate and share computing resources around the world. In the Internet culture, people who created software found that the easiest way to distribute their software was to give it away for free.

The Internet has created/stimulated a large number of different pieces of application software. The software is focused on how people use it. It may appear technology driven, but software on the Internet succeeds only if a large community uses it; hence most software products from the Internet are actually customer-requirements driven. Some of these pieces of software will survive in the Internet and some will vanish from disuse. There is a high attrition ratio between the survivors and those that disappear. However, as the Internet is such a large market, those that survive can have an enormous influence on the environment.

An example of the strength of this way of doing things is the story of Netscape Communications. Recently, someone wrote a piece of software which, combining a number of network protocols, allowed its users to easily navigate through many kinds of information out on the Internet. Within two years, a company was floated and instantly valued at about \$2.7 billion (on total sales worth \$12 million). A typical software development project within an organisation such as BT (or Microsoft) generally takes over two years to reach a first solid implementation. One of the reasons for this success is that the software was distributed for free to begin with, and was rapidly modified to suit the needs of its users in its early stages.

For every success such as Netscape's, there are tens, possibly hundreds of similar ideas which have

not quite succeeded. In a population the size of the Internet this does not matter as there will be a large number of successes. In a population the size of a typical software house, the chance of success is much smaller.

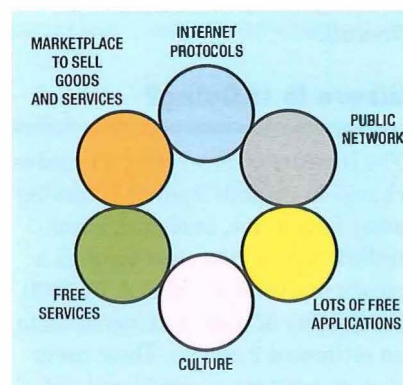
The Internet Engineering Task Force (IETF), which controls the evolution of the Internet protocol suite, has a culture of its own. While the meetings are open to any participation, the influence that can be exerted in a meeting will be based primarily on the participant's knowledge and standing rather than his or her company's market influence. As a result there is a strong resistance to new ideas being introduced from outside the core of participants and a strong 'not invented here' attitude to such contributions.

Conclusion

As has been mentioned several times, the Internet is undergoing a period of significant change. In order to consider the Internet as a whole, it is necessary to consider its different aspect; its protocols; its public network offerings; its culture; its ability to provide innovative solutions; a continuing source of free information, and a potential commercial opportunity. (See Figure 5.)

The infrastructure of the Internet is changing away from 'national' resources paid for by central funds, to an infrastructure provided by commercial service providers as a

Figure 5—Aspects of the Internet



business proposition. This has caused a move from indirect charging to direct charging, albeit on a subscription basis. In the future, the use of reservation protocols (which are particularly greedy with resources) will provide a useful means to identify the start and end of a communication. Such information could be used for per-use charging!

The user community is growing significantly and as a result the demographics are also changing. The user base is becoming less knowledgeable of the technologies involved, and hence less tolerant of the difficulties and failures that may occur in achieving such a difficult task (in other words they will become less excited by the fact that the dog can talk and will be prepared to criticise its diction).

Wider access to the free information resources will generate pressures. The first will be the need to provide high-quality support for resources that are presenting corporate image. The second is the pressure to conform to the local mores that tend to be based on national lines via a technology that is not constructed along geographic boundaries. The third will be the possible withdrawal of current resources that are being provided as a low-priority task on a large system, as the number of visitors increases to unacceptable levels.

With the growth of user population, the technology is providing a new means to market for many companies (either to sell goods and services, or to provide corporate image) and hence the provision of services to these companies based on the Internet becomes a commercial opportunity.

Yesterday's Internet was different from today's Internet, which is different from the Internet that will be tomorrow. The innovative culture of the Internet will ensure a continuous stream of new opportunities, some of which will result in large commercial returns. The only thing that is likely to be constant on the Internet is change.

Biography



Paul Jenkins
BT Networks and
Systems

Paul Jenkins works in Global Engineering, BT Networks and Systems. He worked for over 10 years in data communications, and computer communications standards before becoming active in determining the requirements of multimedia systems. Paul is currently involved in developing BT's approach to the information age and works to determine how BT can most effectively use its resources in this area.

Distributed Processing — Managing the Future

The convergence of computing and telecommunications has yielded a coming of age for distributed processing in terms of market, technology and standards. Distribution is not a panacea, and there remain a number of barriers to its uptake, including technology maturity, and the need to handle existing systems. However, the benefits from distribution can be massive. Distributed processing is fundamental to the way the information industry will do business; exploiting technology wisely is a key activity for the next decade.

Building Systems for the Future

A dictionary-style definition of distributed processing might be: *a combination of heterogeneous hardware, software and network components, in which the software components reside and execute on two or more of the hardware components, communicating and interacting using a network.*

Key points to note from the definition are:

- the software, hardware and network components can come from different vendors and/or use different technologies;
- software execution is not a simple download and run—it involves run-time interactions between components over the network; and
- data is (usually) distributed and processing is executed in parallel across the network.

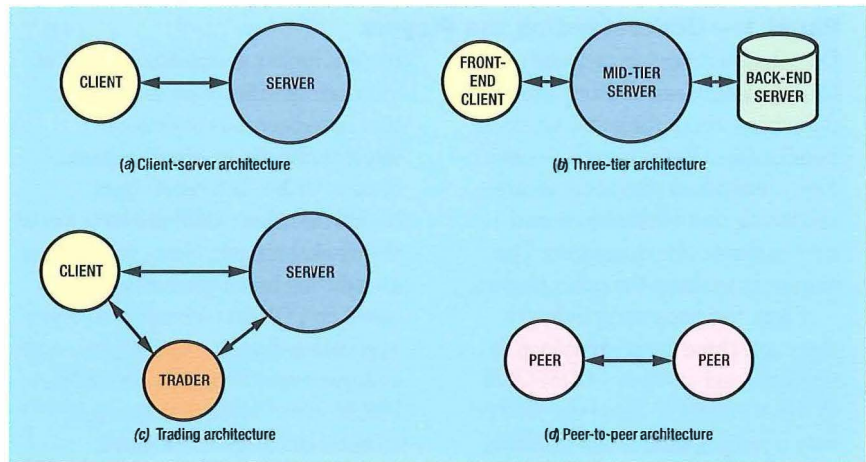
The key benefits of distributed processing come from the modularity of distributed components. Each module can be separately developed, extended and maintained, to provide a system which is potentially highly configurable. Modules are also less location specific, although often it is useful to reduce network traffic by locating modules with high interaction close to each other. All this leads to greater system flexibility and robustness. In particular, this offers a manageable way of interfacing between systems from different companies; for example, in the alliances which characterise the

information industry. However, distributed systems are inherently more complex, and therefore harder to develop.

Three forces are driving distributed processing technology forward. These forces are all aspects of the convergence of telecommunications and computing. First, the converging markets are demanding ever more complex systems. Second, the convergence of the technologies is making distributed processing technically feasible. Third, standards are emerging which allow open distributed processing. This article explores the technological convergence aspect, highlighting standardisation efforts as appropriate. Within the information industry, there are three domains where distributed processing is essential for future growth:

- distributed information services, which are intrinsically distributed across a network (the Internet is a useful conceptual model for this domain);
- business systems to deliver information services commercially (for example, pricing, billing, provisioning, helpdesks, fault handling), and the systems which assist the management of these business systems, each of which will need to mirror the distributed nature of the services they support; and
- advanced network services (also known as *network intelligence services*), where distributed processing is required within the telecommunications network itself.

Figure 1—Distributed processing architectures



Beneath the Surface

Distributed processing is not a new idea, but it has been impossible to achieve on an industrial scale until now. Various underlying technologies now make it possible, or at least more manageable. This section briefly describes the most important of these.

Architecture

There are several common architectures for distributed processing environments.

Client-server architecture

In a client-server architecture (Figure 1(a)), a client process (typically a front-office PC) sends a *remote procedure call* (RPC) to a server process (typically a mainframe system) in a different location, to carry out some remote processing (for example, a database transaction). Parameters are passed with the RPC and the result or an exception message is sent back. This architecture is generally used to partition processing into presentation (on the client) and information processing (on the server). Clients may link to many servers and many clients may link to each server. This has become the most popular distributed processing architecture in industry so far.

Three-tier architecture

A three-tier architecture (Figure 1(b)) is essentially a double client-server architecture. The mid-tier is both a client and a server: it acts as server to the front-end client and as client to the back-end. Typically, the front-end client handles information presentation, the mid-tier contains functional logic (for example, transaction logic) and the back end is a large data store. This architecture is growing in importance, especially for enterprise-wide applications, where there are large amounts of enterprise data to be stored and made available to many front ends. Three-tier systems can involve multiple instances of mid- and back-end servers, as well as

many clients. When multiple instances of servers exist, mid-tier peer-to-peer communication is used to ensure consistency and to balance workload.

Trading architecture

In a trading architecture (Figure 1(c)), clients use the services of a broker or trader to select an appropriate server. The trader's function is to know about all the servers available (a real-time, on-line directory enquiry service). It passes back the location of a suitable server, and the client then makes direct contact with that server. This is the basis of the open distributed processing (ODP) architecture.

Peer-to-peer architecture

In a peer-to-peer architecture (Figure 1(d)), each processor interacts with many other processors in complex ways, and without any master-slave relationships. At different times, a processor can act as either client or server. Because of the flexibility that the concept embodies, peer-to-peer is thought by many researchers to be the ultimate goal in distributed processing. For example, a search request on World Wide Web expands to many sub-requests on many different nodes of the Internet, with information from the original node being part of the answer.

Object technology

Object technology¹ (OT) is not essential for distributed processing, but it is a natural technology to use because it simplifies the design process.

In OT, *objects* are self-contained system components which act as black boxes. They can be accessed

only via well-defined operations (*methods*) at their interface. All data is internal and private — there is no global data in an OT system. Reading, writing and manipulating all have to be carried out via appropriate interface operations. This property is known as *encapsulation*. It means one object cannot corrupt data in another object. It also means objects can be distributed, without any impact on the user's perception of the logic (*location transparency*²). Objects can be used at many different granularities; that is, objects can be composed from smaller objects. Objects can also be used to represent processes. A business can be modelled using entity objects (for example, products) and process objects (for example, order handling) which relate events and actions to entities. These entity and process objects are collectively becoming known as *business objects*.

OT for distributed processing is increasingly supported by industry standards and commercial products. The CORBA standards for middleware (see below) are based fully on OT. So are the standards for OpenStep, the development and runtime platform, and standards from the Object Database Management Group (ODMG). Products are now becoming available conforming to each of these standards.

Middleware

Middleware is the technology which connects different processors together and enables them to interwork. Different middleware services are needed for a variety of different purposes.

For synchronous interaction between client and server where a

Panel 1 — Understanding the Players

Distributed processing stands at the boundary between two converging industries: computing and telecommunications. Both industries (and many individual players in it) are producing new technologies and new variants of technologies. This summary outlines the main players.

From the computing industry, there are three main groupings of products and vendors—COM/OLE, CORBA and DCE. However, this is only a rough guide. Some products straddle different groups, and there are signs of convergence in some areas. From telecommunications, the highest profile activity is TINA. A research consortium, ANSA, has been included because it has had (and continues to have) a significant impact on DPE technology.

COMIOLE (Microsoft)

Microsoft's software architecture is based on technologies called *object linking and embedding* (OLE) and *component object model* (COM). COM/OLE³ is the underlying architecture for the Windows 3.1, Windows 95, and Windows NT operating systems. OLE provides mechanisms for compounding documents while COM provides the communication mechanisms for OLE objects. Microsoft is planning a distributed version of its COM technology, namely Network OLE, for release in the summer of 1996.

CORBA (OMG)

OMG is a consortium of over 600 software vendors, developers and end users, including telecommunications companies. The goal is a common architectural framework for object-oriented applications based on widely available interface specifications. CORBA is the common object

request broker architecture⁴, a set of interface specifications designed so that heterogeneous objects can interoperate via an object request broker (ORB)—a kind of object backplane. Many ORB products are on the market already. Some vendors are introducing broad product families based on CORBA technology to cover application development, deployment and operation. Examples are IBM's DSOM (distributed system object models) and SUN's NEO/Java.

DCE (OSF)

The Open Software Foundation Distributed Computing Environment (OSF/DCE)⁵ supports a remote procedure call (RPC) service. It allows applications in one execution environment (that is, combination of hardware and operating system) to call procedures which run in a different execution environment. The internal architecture of OSF/DCE is shown in Figure 2.

Apart from the basic RPC service, various vendors are also using OSF/DCE to provide foundation services for further levels of value-added products, examples include: IBM/Transarc and Novell/Tuxedo (distributed transaction processing monitors), Oracle (distributed database access), Digital/TeMIP (distribution

mechanism for Version 3 of the management platform).

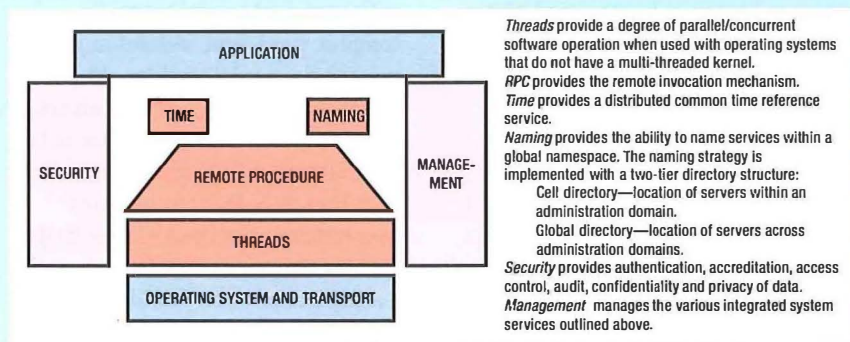
TINA

The Telecommunications Information Networking Architecture Consortium (TINA-C) is a consortium of major telecommunications companies. Its goal is to develop a new global distributed architecture and a service rich environment for deploying and managing advanced services in an ever increasing competitive worldwide marketplace. It is defining a consistent reference architecture^{6,7,8,9} using open distributed processing and CORBA specifications as components for DPE.

ANSA

ANSA is the technology from a research consortium of computer companies, telecommunications companies, systems integrators and end-users. Earlier ANSA work was influential in shaping the ODP reference model and CORBA. Current work is extending this technology to building coherent multivendor, multiservice heterogeneous networks, where much of the information is multimedia and where there are complex requirements for transactions, dependability, federation and real time.

Figure 2—Structure of DCE



response is expected (because the server is always on-line), an RPC service is suitable (for example, OSF/DCE). For asynchronous interaction where the server cannot be guaranteed to be on-line to receive the call, a reliable message queue service is suitable. (Note, here synchronous and asynchronous refer to the availability of the server to respond not to the blocking nature of the client.)

For atomic operation supporting the 'ACID' properties of interaction, a transaction processing service may be needed. For information sharing and access, various forms of data distribution and management may be needed (for example, distributed databases). For interactions based on object technology (OT), an object request broker (ORB) would be most appropriate.

Distributed databases

Distributed databases have data stored at more than one node of a distributed system and present a shared schema to their applications. The common schema approach makes distributed databases different from systems with any distributed objects, each containing private data.

Distributed databases are not essential for an effective distributed

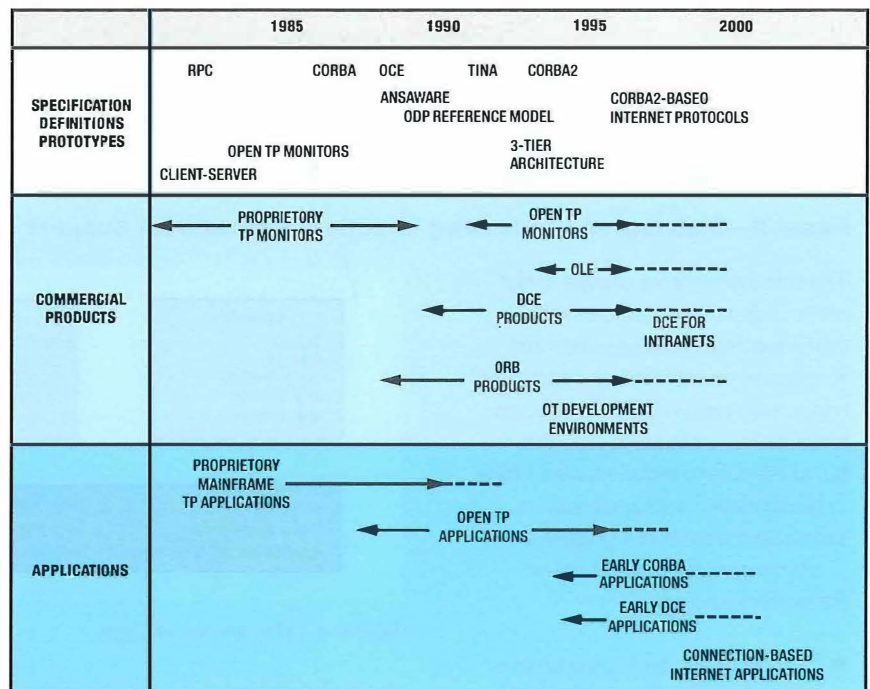
Figure 3—Roll-through of DPE technologies

processing environment. Indeed, they significantly increase system complexity, because of the need to keep component nodes synchronised. This is especially difficult for 'distributed write' where records have to be updated in more than one node. Most database vendors are beginning to market distributed versions of their products. Many are involved in the Object Database Management Group (ODMG), which is developing some standards in this area.

Showing the Pedigree

Distributed processing has not arrived from nowhere. Although it is radically different from centralised computing systems, its introduction has been progressive and incremental. Some component technologies have been common for several years; for example, on-line transaction processing monitors (which ensure that transaction atomicity is not violated by remote processing) and, more recently, client-server computing. The difference now is that distributed systems architecture and design are more comprehensively understood² and commercial-grade distributed-processing products are coming on to the market.

Figure 3 indicates how technologies have been developing, the current state of the art, and what can be expected over the next few years. Typically, each component technology progresses from concept (specification, definition, prototype) through to commercial products and then applications. In the table, the sweep of development moves broadly top left to bottom right. So, for example, on-line transaction processing started as concept in the early 1980s, moved to a commercial product in the mid-late 1980s and is now in widespread use in business applications. Secure middleware (DCE) is now beginning to be used widely, and early CORBA applications are emerging. OT development and run-time platforms (for example, OpenStep) are just being marketed and it is reasonable



to believe their use will grow over the next few years.

For the future, Java is emerging from definition and holds the promise of dramatically reducing deployment and operational costs of major distributed processing systems. CORBA-based protocols for the Internet are about to be defined, and could lead to another revolution in major systems design.

Making the Choice

At present, distributed processing technology is in a transition phase. The concepts, and increasingly the products, are now available for application, but there are still areas where it is immature. Distributed processing is the vision for many (or even most) enterprise-wide systems. However, since much of the technology needed is still at the leading edge, much of the current technical debate is about timing and deployment plans and the balance between benefits and cost.

Whether distributed processing is right now for a particular application (or which form is right) depends on a number of factors, including: application domain, whether it is a completely new system or a change to an existing system, throughput volumes, and the importance of key distribution characteristics.

As mentioned earlier, there are three domains within the information industry where distributed processing is essential for future growth:

distributed information services, advanced network services, and business support systems.

For distributed information services, distributed processing is often intrinsic (for example, to access real-time information from a remote source). However, the hard distribution work is often done by the network platform on which the service runs (for example, the Internet). This provides the great wealth of opportunities for novel information services, but it also acts as a constraint. It is interesting to speculate whether the current Internet protocols and design will be scalable, or whether approaches from more classical distributed systems will need to be used to cope with congestion.

For advanced network services, where distributed processing is required within the telecommunications network, the network designers have to do the hard work of creating a flexible and robust distributed computing environment (DPE). However, in some ways the domain is a green-field site, because adding intelligence within telecommunications networks is still relatively new. This offers the possibility of using leading-edge technology.

For business systems, the main question is usually how to handle existing (legacy) systems (see Panel 2 for a case study). The economics of upgrade versus swap-over versus migration has to be investigated very carefully on an individual basis.

Panel 2—Distributed Processing in action: Operational Support Systems in BT

The three drivers for change in the move to distributed processing (business/ market, technology and standards) are at work on operational support systems (OSSs) for telecommunications companies. The combined effect is to require future OSSs to be developed within a common distributed framework (Figure 4).

Particular features of this framework include:

- built from reusable components to meet the business drivers for cost reduction, speed of OSS delivery, improved flexibility etc.;
- developed using client-server design and employing open standards in line with major technology drivers; and
- appropriate use of standards (managed diversity) and vendor products to deliver cost-effective framework-based systems.

This panel reviews the impact of the distributed framework on OSS design, looking first at the way in which telecommunications companies' OSSs have been designed previously, as represented in Figure 5.

Present-day OSS designs tend to be produced as stand-alone 'islands of processing'. Each OSS (be it of network or service management form) is designed as a self-contained item. This means that it has its own internal form of data representation and storage, design of processing and form of user access. These OSSs are often built upon a variety of operating systems and hardware platforms. Also OSS functionality (for example, customer handling) tends to be reproduced when services are offered over different transmission media (for example, plain old telephony service, private circuits, intelligent networks etc.). For the above reasons, many

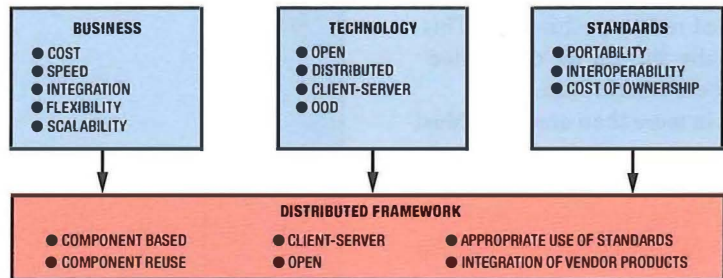


Figure 4—Drivers for change

telecommunications companies have a large number of OSSs (often the order of several hundreds) which are often functionally very similar and therefore should be rationalised. Additionally, the planned levels of interoperability between these systems is often quite weak.

The alternative approach to OSS design, which is gaining prominence in telecommunications companies, is to develop OSSs according to a three-tier architecture. Figure 6 shows the logical structure for OSSs (user, business logic, information layers) in the context of the overall distributed infrastructure.


The distributed infrastructure is shown as a number of layers, with the relative placement of middleware and management services, namely:

- At the lowest level is the data network, which is used to join the various distributed computing resources together. This network

must be capable of supporting the full range of client-server interactions; for example, single-shot (RPC-like) calls, stream (file transfer) etc.

- The next layer consists of the various execution environments which support the distributed infrastructure. Each execution environment is a combination of appropriate hardware and operating systems. Examples include PCs, UNIX and mainframe environments from suppliers such as HP, SUN, Digital and IBM etc.
- The next layer consists of the range of middleware services. These must operate across the various execution environments. The middleware services may themselves be viewed as a number of value-added layers. For example at the lowest-level

Figure 5—Traditional OSS design



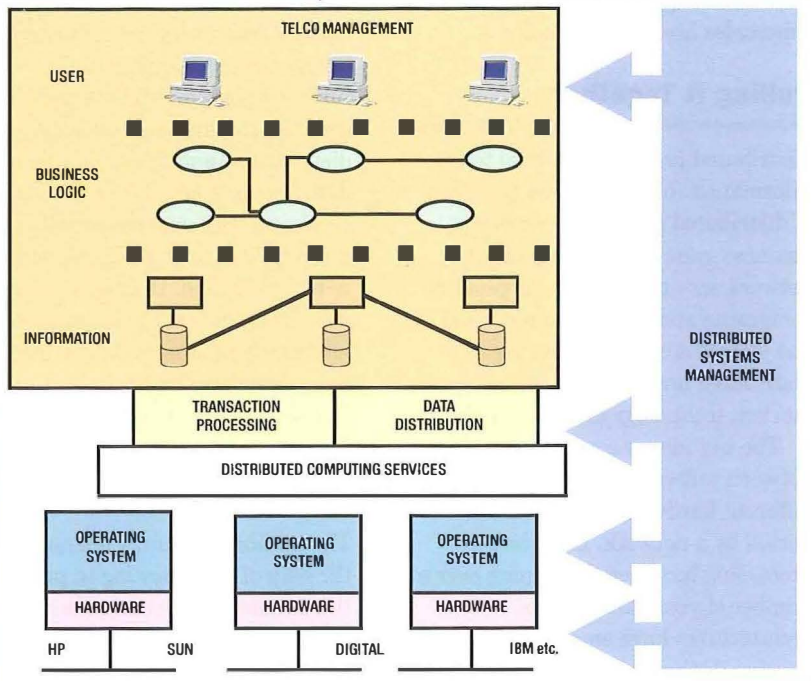


Figure 6—OSS logical architecture

distributed computing services (DCS) which may be based upon RPC or message passing technology (elsewhere in this article the OSF/DCE is discussed as an example of 'standard' RPC middleware). Above this, are services which build upon the DCS; for example, data distribution and transaction processing. The 'stepped' nature to the middleware services (shown in the diagram) indicates that different level interfaces are available to support the applications architecture (that is, TP and data distribution as well as access to the lower DCS services).

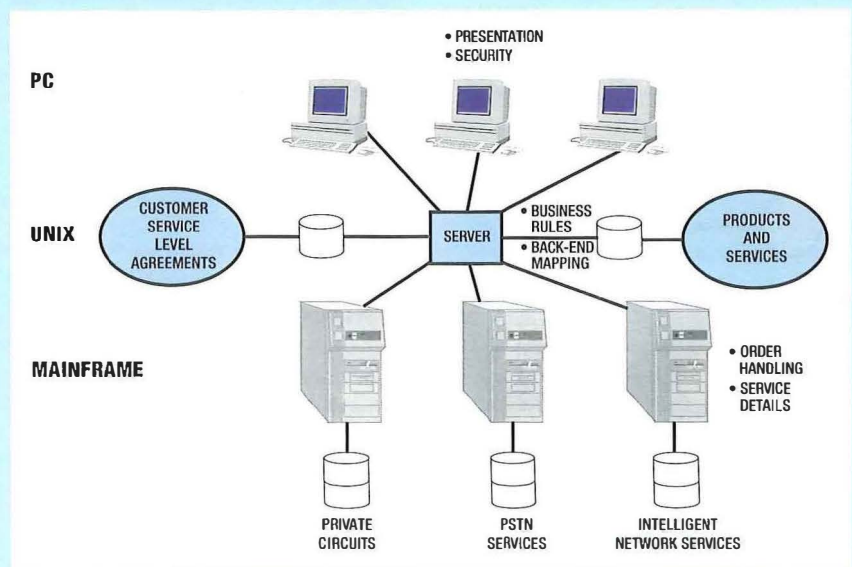
- At the top-most layer is the complete applications architecture infrastructure being supported.
- Figure 6 also shows that each of these layers must have appropriate system management applied. This must be integrated as much as possible in both the standards used and the views presented to the operational staff.

Unified Customer Handling

One example of the way in which the three-tier architecture is being deployed within BT is the Unified Customer Handling (UCH) Programme. Figure 7 shows a future 'typical' customer handling deployment based upon the three-tier approach.

Particular features are worth noting about the distributed processing approach to customer handling.

Figure 7—Typical three-tier deployment—customer handling system



First, the application functionality is no longer completely mounted on a mainframe machine (for example, in the past for public switched telephone network, private circuit or intelligent network service orders).

Second, mainframes may now be viewed as systems which contain elements of business process and information. In this respect they may also be thought of as legacy systems which need to have these business process and information building blocks accessed by other systems (for example, existing ordering process and customer details etc).

Third, mid-range machines (for example, UNIX) are used as intermediate servers which may also be used as containers for new elements of business process or information building blocks (for example, details of products or service level agreements). These servers may also provide some of the mapping mechanisms to access building blocks in the legacy systems.

Finally, PCs will often be used as the containers used for presentation services and the means to provide single logon to a complete security infrastructure (provided in the middleware layers).

Distributed processing is vital to the information industry in the provision of distributed information services, business systems, and advanced network services.

Legacy system migration is the process of integrating existing information systems into a new environment that is based on the new technology. An example is the process of integrating mainframe-based systems into client-server environments. In this process, old technologies are replaced by new ones, and older development techniques are also updated.


The three main options for introducing DPE into an existing (non-distributed) systems environment are interoperating, encapsulation and rebuilding see Figure 8.

Interoperating is replacing one part of the existing system (or building a new extension) using a DPE and building a customised interface between that and the existing system. It is an incremental step which protects the existing asset base, but its flexibility and modularity are restricted.

Encapsulation is building an object wrapper around the existing system, controlling access via new interface methods. It allows the existing system to operate as one node of a DPE. The possible disadvantages are the difficulty of producing an acceptable wrapper, performance degradation and the wrapped system being too big for full flexibility in the DPE.

Rebuilding is the most radical, and probably the most costly, approach. It takes full advantage of the

Figure 8—Ways of handling existing systems



new technology, but the cost and timescales are often prohibitive.

Pulling it Together

Distributed processing is vital to the information industry in the provision of distributed information services, business systems, and advanced network services. The convergence of computing and telecommunications has yielded a coming of age for distributed processing in terms of market, technology and standards.

The key idea is communication between software components on different hardware platforms connected by a network. Distributed processing has been developing over a number of years and several architectures have emerged to support it, the most important being the client-server, three-tier, and trading architectures. Peer-to-peer architectures in which each node can act as both client and server are considered to be the ultimate goal in distributed processing. Other key technologies include object technology, middleware and distributed databases.

Distribution is not a panacea, and several barriers to its uptake remain, including technology maturity, and the need to handle existing systems. However, the benefits from distribution can be massive. Just consider the following as examples:

- collaboration through connectivity and interworking,
- performance through parallel processing,
- reliability and availability through replication,
- scalability and portability through modularity,
- extensibility through dynamic configuration, and
- cost-effectiveness through resource sharing and open systems.

In fact, distributed processing enables many services that would not otherwise be possible. There is further progress still to be made in some areas, but with full-scale distributed applications now being deployed on a large scale for the first time, and with experience still growing, it is an exciting technology to be involved in. Distributed processing is fundamental to the way the information industry will do business: exploiting the technology wisely is a key activity for the next decade.

Acknowledgements

The authors gratefully acknowledge the help of the following in preparing this article: Bob Briscoe, Eurescom Project P517 team, Jonathan Legh-Smith, Steve Proctor, Neil Winton.

References

- 1 CUSACK, E. L., and CORDINGLEY, E. S. Object Orientation in Communications Engineering. *BT Technology Journal*, July 1993, 11(3).
- 2 Open Distributed Processing Reference Model (RM-ODP), ISO/IEC 10746-1 thru 10746-4, ITU-T X.901 thru X.904.
- 3 <http://www.microsoft.com/>
- 4 <http://www.omg.org/>
- 5 <http://www.osf.org/dce/>
- 6 TINA-C Baseline Document No. TB_C3.JB.001_1.0_93, Connection Management Architecture, Dec. 1993.
- 7 TINA-C Baseline Document No. TB_A1.EAC.001_1.0_93, Information Modelling Concepts, Dec. 1993.
- 8 TINA-C Baseline Document No. TB_A2.NAT.002_1.0_93, Computational Modelling Concepts, Dec. 1993.
- 9 TINA-C Baseline Document No. TB_A3.NS.005_1.0_93, Engineering Modelling Concepts, Dec. 1993.

Glossary

Client-server A DPE architecture where the system is divided into client processes (typically front-desk

PCs) and server processes (typically back-end mainframes holding enterprise data).

COM/OLE Common object model/object linking and embedding. Middleware for DPEs, from Microsoft and associated organisations.

CORBA Common object request broker architecture. A set of standards defined by the OMG consortium for DPE middleware, using OT, and products which conform to the standards.

DCE Distributed computing environment. A set of standards defined by the OSF consortium for DPE middleware, and products which confirm to the standards.

DPE Distributed processing environment. A combination of heterogeneous hardware, software and network components, in which the software components reside and execute on two or more of the hardware components, communicating and interacting using the network.

Message queuing A mechanism for passing messages from one systems component to another, based on queues, designed to be used when time to respond does not need to be guaranteed.

Object A self-contained system component which can be accessed only via well-defined operations (methods), and where all data is internal and private (encapsulation).

ODP Open distributed processing. A standardised architecture for distributed processing which has been widely adopted as the right way to model distributed systems.

On-line transaction processing monitors System components which ensure that transaction atomicity is not violated by remote processing.

ORB Object request broker. The component in CORBA which locates and connects objects.

OT Object technology. The set of component technologies which enable systems to be built and deployed using objects, including platforms, design techniques, languages.

Peer-to-peer A DPE architecture where there are no master/slaves.

RPC Remote procedure call. A mechanism for passing messages from one systems component to another, based on near-real-time interactions.

Three-tier architecture A DPE architecture where the system is divided into three layers with client-server operation at each interface.

Transparency Hiding some attribute of the DPE; for example, location transparency where the user does not know where a component is located.

Viewpoint A distinct autonomous scope for the specification of an open distributed system. OP defines five viewpoints: enterprise, information, computational, engineering, and technology.

Biographies



David Freestone
BT Networks and
Systems

David Freestone leads a team of researchers at BT Laboratories into operational

support systems for the future. This includes distributed processing, as well as legacy systems issues, architectural design and advanced prototyping. Among his previous work are requirements engineering research, formal methods in computing, software development environments and transaction processing. David joined BT in 1980 and has degrees in mathematics and computation.



Tony Richardson
BT Networks and
Systems

Tony Richardson joined the then British Post Office in 1965. He has spent most of his career working in various areas associated with systems research and

development. These have included man-computer human factor developments, OSI communication products and office automation systems. Tony's present role is that of Senior Advisor on Management Platforms within the Ipswich Systems Engineering Centre. Here he specialises in the use of vendor management platforms and client-server middleware technologies in the development of BT's operational support systems. He also represents BT at the Open Software Foundation and has carried out work for the Network Management Forum. Tony is a Chartered Engineer and holds a Master of Science Degree in Computer Science from the University of Hertfordshire.



Ben Whittle
BT Networks and
Systems

Ben Whittle graduated from the University of Wales at Aberystwyth in 1989 with a BSc. in Agricultural Economics. Having realised that computers would help him to count sheep, he proceeded to an M.Sc. in Computer Science and was subsequently invited to study for a Ph.D., the subject of his thesis was Software Component Reuse. In 1993, he finally moved away from Aberystwyth to take up a post at the Rolls-Royce sponsored university technology centre at the University of York, where he helped Rolls-Royce to introduce a systematic reuse programme for safety-critical aero-engine control software. He joined BT's SoftLab in November 1995 and has been working on three-tier architecture and reuse related projects. He is the chair of the BCS Software Reuse Specialist Group and has recently been invited to join the IEE Professional Group Committee C (Software Engineering).

Tom Rowbotham and Martin Yates

TINA—A Collaborative Way Forward

Telecommunications Information Networking Architecture Consortium (TINA-C) is an international consortium defining an architectural infrastructure for information services and network resources. It is capable of giving consumers adaptive and universal access to their customised services. The design enables diverse businesses to cooperate, running the components necessary to fulfil a service request. It anticipates advanced software technologies using object orientation, distributed processing environments, intelligent agents and multi-service networks.

A Brief Introduction To TINA

When telecommunications companies contemplated their forthcoming role in the information revolution, a consensus arose that information trading would require radically new protocols among operators—protocols much richer than those existing for basic transport. A vision emerged of operators selecting and using components of information services and computing resources across company and international boundaries. This is a level of interaction and cooperation never previously experienced for basic telecommunications services. Another aspect to this challenge was the exploitation of existing investment in programmes such as intelligent networks (IN) and telecommunications management networks (TMN).

To lay the foundations for the vision, in 1992 over 35 companies formed an international consortium called *TINA-C* that included network operators, computer suppliers and network equipment vendors. All had a mutual interest in establishing a telecommunications information networking architecture (TINA). This article describes the motivations and objectives of the consortium, together with some of the expected benefits. An overview is also given of the technology and commercial world anticipated by the consortium.

Why Now?

Several engineering and commercial trends are causing the well-known convergence of telecommunications and computing. The future is expected to reveal mobile consumers

buying services that manipulate and supply information; the network connections required will be 'incidental' to the customers' view of the products. The service functionality and data will be implemented in software running on a widely dispersed hardware base.

Many related commercial and engineering factors make this decade particularly opportune. Some major ones are:

- the introduction of interactive multimedia products into household markets;
- ergonomic and simple human-computer interfaces (note the boost to Internet use encouraged by WWW hypertext browsers);
- the rapid take-up of portable technology such as notepad PCs and mobile telephones;
- the completion of investment programmes in digital trunk networks;
- the ending of telecommunications monopolies and the imposition of regulation on traditional telephony in comparison to unregulated information markets;
- the aspirations of media and entertainment enterprises to supply consumers electronically; and
- software engineering advances achieved with design tools and language development.

TINA companies are looking toward a new architectural infra-

Figure 1—The enterprise model

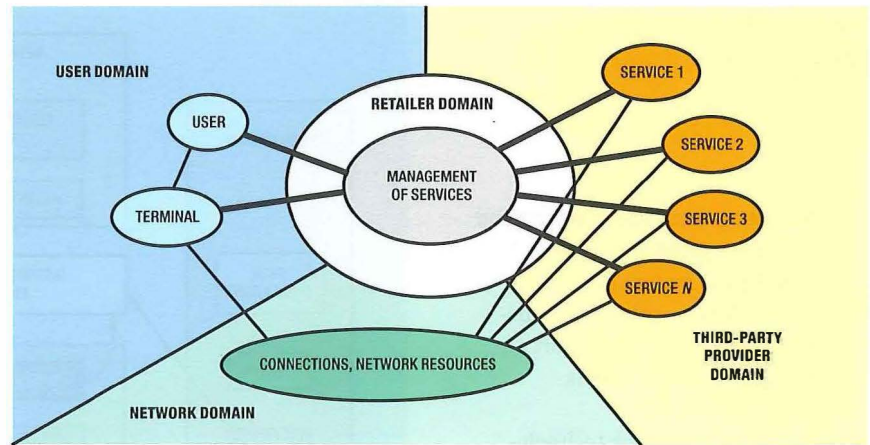
structure on which to deploy the host of multimedia services envisaged. The infrastructure will utilise existing and emerging networks, capture the diversity and dynamism of the Internet, but place the deployment of services on sound operational and management foundations.

Objectives of the TINA Consortium

To hasten progress, member companies seconded about 40 telecommunication architects to form a core team, collocated and hosted at Bellcore, New Jersey, USA. The core team started in earnest at the beginning of 1993 and will run for five years. The first three years have culminated in the definition of an architecture, aspects of which were demonstrated at Telecom 95, Geneva. During 1996, the consortium will move into a rigorous evaluation, validation and refinement phase to complete the architecture by the end of 1997. The consortium will deliver definitions, specifications, methods, tools and validation prototypes. Since the consortium is pre-normative and pre-competitive no products will be produced. It is useful to consider two primary dimensions on which the consortium's output is based—enterprise and technology.

The enterprise context

The TINA consortium aims to define an architecture capable of realising an information network on which consumers can find, request and receive multimedia services. The architecture must recognise and support the different duties of individual enterprises involved in that end-to-end supply chain, and the capacity of vendors to build hardware and software components of the delivery infrastructure. For example, the architecture must be flexible enough to support the operational independence of the different companies involved, but also recognise and enable the federation of companies' operations through open interfaces.



There is widespread acceptance that enterprises supplying or consuming in the information world will have roles that fall into one or more of four domains. These domains, illustrated in Figure 1, are described below.

- *User domain*—where the consumer or user of information services resides. This usage is through a terminal of which there may be many varieties with widely different capabilities (for example, multimedia personal computer or mobile telephone).
 - *Third-party service or content domain*—covering the third parties that produce and license service content and may also supply the applications to view or generate that content. This domain is expected to comprise a large and disparate set of third-party providers who sell their content to consumers through the services of brokers or information retailers.
 - *Retailer domain*—representing brokers or retailers of information services whose activity is analogous to a supermarket. The products or information services are presented and sold to users in a pleasant 'shopping' environment but the information content or the service application could be supplied by a third-party provider.
 - *Network domain*—embracing organisations that run the access and trunk telecommunications networks that can deliver content and interconnection among the domains. In many circumstances, the user will be billed directly by the information retailer and will not have a direct contractual relationship with the network provider.
- Many businesses will operate more than one domain, and the TINA architecture supports the separation and federation of these distinct commercial activities.

Technology foundations

Supporting TINA are four technologies that are either in current use or maturing.

- *Object orientation* is the main software structuring paradigm. This provides strong data and functional encapsulation and promotes interface reuse. TINA objects may have many interfaces each offering a defined set of operations.
- *Distributed processing environment (DPE)* systems are utilised for building cooperating software components that can execute on different computers dispersed throughout a network. A DPE system provides a software backplane or 'bus' that hides the complexity of distribution and allows components to communicate in a way transparent to physical location. This enables the construction of scalable and manageable software systems.
- *Intelligent agents* is a new technology that has emerged from the field of artificial intelligence. Agents are software systems that act on behalf of other parties, such as users, and have been applied to

Figure 2—TINA enterprise and resource model

learning, information filtering and negotiation problems. This technology is being applied in TINA to design software components particularly to support cross-domain interactions.

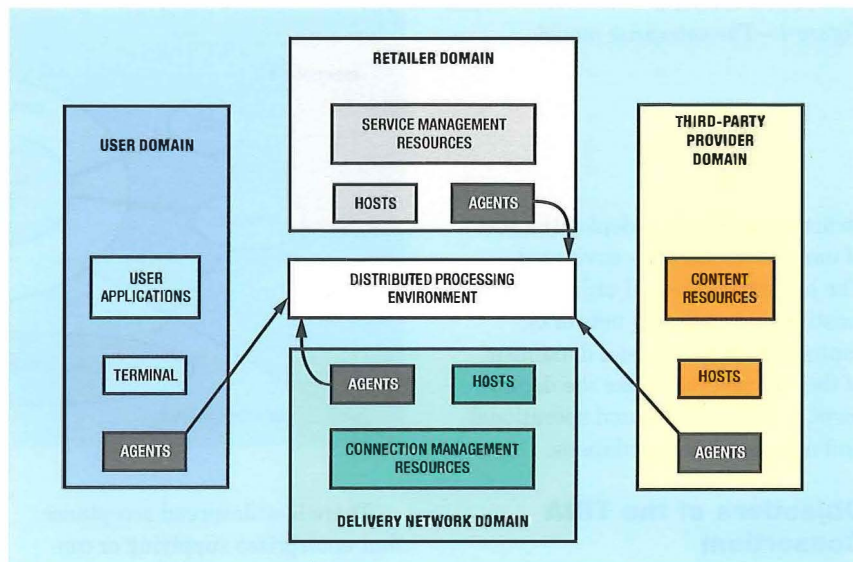
- **Multi-service network** technologies, such as asynchronous transfer mode (ATM), provide the flexible backbone necessary for delivering a diverse range of service contents and are assumed in the architecture. However, TINA is fundamentally independent of network technology and so explicitly accommodates other transport network types, particularly legacy systems such as the intelligent network (IN) and integrated services digital network (ISDN).

A principle of the consortium is to use existing work and create new ideas only when necessary. Thus TINA has reused *de jure* and *de facto* standards from the related technology areas, including TMN, IN, open distributed processing (ODP), Object Management Group (OMG), ATM. TINA has also a relationship with the Digital Audio Visual Council (DAVIC) which is progressing architecture and protocol standards for interactive multimedia.

During the remaining two years, the Consortium intends to submit the maturing architecture to these standards bodies. This will be done through member companies and the core team. It is anticipated that many TINA ideas will form the basis for proceeds of formal standard bodies or be incorporated by vendors through more product oriented fora such as the OMG.

Rationale Of TINA Specification

It is acknowledged that a varied, mobile and geographically dispersed consumer base will only be adequately served by the interoperation or federation of many companies.



TINA aims to reduce the complexity and cost of achieving this breadth of interoperability both for operators and component vendors.

In the past, the telecommunications industry has employed monolithic vertical solutions to provide networks with rudimentary services. Interoperability has been achieved at international standard gateways. For information services, such solutions are unlikely to be feasible because a simple information service request may involve many enterprises using a complex set of resources such as public networks, software and hardware components. These resources may be owned, managed, controlled and used by any combination of enterprises. Hence the interfaces between resources must reflect these relationships and provide interoperability on all aspects such as usage, control and management.

Figure 2 illustrates how resources required to fulfil an information service are mapped onto the enterprise model. It shows that in each domain there are multiple host computers and software resources, such as agents and applications, which are performing duties within the domain. However, it is important to note that DPE technology will support the applications and agents of one domain being run on hosts in other domains. This will be a common circumstance for many commercial and engineering reasons. An example could be video on demand (VoD), where a content provider wants only to supply licensed content files to different retailers who are

operating proprietary VoD services. The TINA consortium is considering such resource-to-enterprise relationships and is specifying a suitable architecture using the design rules discussed below.

Architectural design rules

The TINA architecture is decomposing the complexity shown in Figure 2 into discrete identifiable software components or objects occurring in each domain. These objects are meaningful separations of behaviour and data such that their interfaces coincide with primary commercial, operational and engineering divisions in the overall design. For example, a TINA computational object in the retail domain, called a *communication session*, has an interface with a *connection coordinator* in the network domain over which connections can be negotiated and modified. This interface coincides with a commercial, operational and engineering separation because the objects may be built and managed by different organisations.

Objects are characterised by the operations they offer to other components with which they must interact. Many components will be responsible for allocating or serving resources, and this can lead to complicated negotiation relationships between objects. Furthermore such components may need to be run on any dynamically allocated computer host (for example, for proximity or load sharing). The combination of resource negotiation and mobility of autonomous components is a reason for adopting some features of intelligent

TINA will support the commercial relationships expected for the information superhighway—an environment in which customers can shop around and pick and mix solutions.

agent technology. The technical challenge of object and host distribution is handled by using DPE technology; this enables objects to find and communicate with each other in an environment that is transparent to location.

The TINA Architecture Specifications

It has been explained that the TINA software architecture requires the definition of many types of interacting objects distributed throughout a network. The specification and nature of these objects are described in the three architectural themes illustrated in Figure 3:

- The *computing architecture* describes how software objects should be specified and 'plugged' into the DPE that provides object location and communication facilities. Thus DPE offers important facilities common to all objects using the environment and additionally to the administrations that are responsible for the smooth running of the DPE. For example, a key TINA assumption is that of a DPE kernel transport network that provides objects with location-transparent messaging or invocation. The computing architecture will specify these DPE facilities or services with respect to issues such as security, object management and deployment.
- The *service architecture* provides concepts and principles for service life cycle and multi-party service

federation. The life cycle describes the roles of the computing architecture and service objects through the stages of deployment, consumer subscription, usage and withdrawal. Objects have been identified that provide a flexible, user-oriented access and service session scheme. For example, users are able to subscribe, launch service sessions, invite other users to participate, suspend and resume activities in a session. The session concept is applicable to a wide range of service types, including multimedia conferencing, information retrieval and traditional voice-based services. Multi-party service federation explains the paradigms and object interactions required to federate the various components of a service that run dynamically in different administrations. Federation specifications are essential to ensure a diverse range of providers can contribute to TINA services.

- The *network resource architecture* covers a model of network resources specified at different levels of detail. At the highest level of abstraction are specifications of how service architecture components can request and manipulate end-to-end stream connections independently of network technology. The lower levels of refinement concern specifications relating to switching and transmission equipment and are technology dependent. Network software is responsible for translating a service request into particular manipulations on specific network resources. Interface for the reuse of existing network technologies is described in this architecture.

Functions that are pervasive throughout the three architectural themes are those of fault, configuration, accounting, performance and security—traditional FCAPS. In TINA's object-oriented design, these functions are distributed throughout

the system and part of every component. Thus there are no specific objects or centralised systems dedicated to management. This is important because it isolates the impact of change or upgrade within a large distributed software infrastructure. Nonetheless, these management functions are found explicitly and applied locally in the three TINA themes described above.

What Are The Benefits?


TINA will support the commercial relationships expected for the information superhighway—an environment in which customers can shop around and pick and mix solutions. The integration of third-party service content providers is vital to the variety and growth of the information network. It will allow small and large enterprises to offer and consume services without the need for huge infrastructure investment. The appeal of this type of activity is evident from the Internet world.

Customisation of services is a key TINA feature for users, whether they are domestic customers, or small or large multinational companies. The components specified in the service architecture are designed to support requirements for subscription management and personalising services for users.

Service provision that accommodates heterogeneity of terminals and delivery networks will be demanded by users, particularly to support personal mobility. TINA provides a framework to cope with this variety of access circumstances by separating adaptation from the core capabilities of each component. This aims to reduce duplication which is an economical and technical necessity.

Managing infrastructure upgrade and scaling has always been a problem in large computer and telecommunications systems. TINA has adopted important constructs and methods to make this problem tractable; for example, the use of location transparency in the DPE and

Figure 3—Components of the TINA architecture



the very modular design of elements providing service resources.

Low cost is imperative to stimulate market uptake of services. Software development is predicted to be a major proportion of service deployment cost and so features to minimise costs are included in the architecture. The use of object and DPE technologies and emphasis on interface definition will encourage reuse and reduce integration costs. The capacity of this architecture to reuse legacy investment will also be a major advantage.

Summary

The TINA Consortium intends to establish a basic framework that will allow member companies to conduct their business independently, but with the necessary levels of interoperability demanded by regulation and customers. Basic interoperability is a prerequisite for economical supply of infrastructure components and is of keen interest to vendors hoping to supply products to a uniform market.

Historically, telecommunications have built *de jure* standards that are slow to be agreed but have longevity and a widespread usage reflecting extensive interoperability requirements, whereas the computer industry has had varied standards, often *de facto*, with a short life cycle reflecting the rapid change in underlying technology. These trends in standards will need to compromise if the scale of investment is to be supported in a way that is responsive to capabilities offered by the newest consumer technology.

Many stakeholders in the information future have shown their interest in the TINA. Many concepts have already emerged, either informally or through submissions, in other fora working on interactive multimedia networking standards. Nonetheless, there remain the important tasks of refining, testing and validating the architecture through worldwide demonstrations and TINA auxiliary projects conducted within member countries.

The consortium's aims are far reaching and the requirements broad. Furthermore, the trends driving the information revolution will not converge gently but will meet in a serious collision as enterprises battle to define their new businesses and preferred partners. It may be a turbulent and uncertain information revolution but the TINA Consortium is a major collaborative effort to make it happen sooner. Its motto could be 'The only way to predict the future is to invent it'.

Acknowledgements

Parts of this article are reprinted with kind permission of the International Telecommunication Union (ITU), Geneva, Switzerland. The assistance of TINA core-team members in the preparation of this material is gratefully acknowledged.

Further Reading

- 1 BARR, W. J.; BOYD, T.; and INOUE, Y. The TINA Initiative. *IEE Communications Magazine*, March 1993, pp. 70-76.
- 2 NILSSON, G.; DUPUY, F.; and CHAPMAN, M. An overview of the Telecommunications Information Networking Architecture. *Proc. TINA 95 Conference*, Melbourne, Australia, 13 February 1995, p. 1.
- 3 BAGLEY, M.; MARSHALL, I.; YATES, M. J.; and HILL, J. The Information Services Supermarket—A trial TINA design. *Proc. TINA 95 Conference*, Melbourne, Australia, 13 February 1995.
- 4 MARSHALL, I. W.; and BAGLEY, M. The information services supermarket—an information network prototype. *BT Technol. J.*, April 1995, 13(2), p. 132.
- 5 <http://www.tinac.com>
TINA Consortium public World Wide Web pages.

Biographies



Tom Rowbotham
BT Networks and
Systems

Tom Rowbotham was Director of BT's Network Development until taking up an assignment in 1993 as Senior Vice President of Concert, the BT/MCI joint venture, to set up its product development organisation. He returned to BT in April 1995 to be Director of Technology Strategy, including responsibility for all the research and standards activity of the company. He is a Fellow of the Royal Society of Engineering, the founding President of the TINA Consortium, a recent Vice-President of the IEEE Communications Society, a member of the IEE Electronics Divisional Board, and a visiting Professor at Kings College, University of London.



Martin Yates
BT Networks and
Systems

Martin Yates works on future information networks in BT Networks and Systems. Since January 1996, he has been deployed in the TINA core team, New Jersey, following a two-year collaboration with European partners on the design and implementation of a TINA information service infrastructure. He has a Ph.D. from the University of Birmingham in optical network devices and is a Member of the Institute of Physics. In 1988, he joined BT where he has worked on optical network technology and distributed system analysis and design.

Robin Smith

Software Agent Technology

Software agent technology has escaped from research laboratories and is beginning to make its presence felt in the critical environment of commerce and industry. This has allowed techniques developed by the artificial intelligence community to become accessible to a wide audience. The aim of this article is to provide a balanced review of the field. This is a personal perspective of someone who has spent the past 12 years applying artificial intelligence techniques to issues arising in the telecommunications domain. The article commentates briefly on the business drivers, describes the technology and gives examples of agent systems.

Introduction

It is inconceivable that modern industrial society could function efficiently without the capabilities provided by computing and telecommunications. Increasingly, the convergence of telecommunications and computing is forcing companies and individuals to reassess the way they exploit the power of the (global) information platforms formed by the merger. Rich information platforms bring with them the possibility of radical transformations of whole industrial sectors. These transformations may not be welcomed by all players in the sector. Industrial players could well find that within a short space of time a traditional (and profitable) activity has been replaced by a new information-intensive approach that leaves them completely wrong-footed.

Information technology (IT) is now becoming fully integrated with the overall activities of major business undertakings. Business managers are now looking to the IT department to provide new market opportunities and to reduce the costs of commercial undertakings. During the past few years, one particular software technology has been clamouring for the attention of hard-pressed IT professionals. Software agent technology (agents) is promoted by inventors and vendors alike as being the solution to a broad raft of information-related problems and having the capacity to offer new business opportunities. However, this particular branch of computer science has not been greeted with unanimous acclaim. Several critical reviewers¹ have compared the over-ambitious claims of the agent community with the hype that surrounded the field of artificial intelligence (AI) in its early years and its re-emergence in the 1980s. In some regards that compari-

son is misplaced since agent ideas are firmly rooted in the field of AI. In fact, agents have been described as being the Trojan Horse of AI since the software that provides the very functions which distinguish agent technology from more traditional computer programs comes from mainstream AI research.

The reason for this focused interest in (information) agents is that they alone (many would have us believe) hold out the promise of helping people deal with the profusion of data and information that the new information age has let loose on the world. Thomas Wheeler in a recent article² has rightly pointed out the distinction between knowledge and information (or data). This steadily growing data machine will overwhelm us unless we can carry out meaningful filtering and draw knowledge from the electronic babble. Before exploring these uses of agent technology in more depth, it is perhaps appropriate to take a look at the software systems which underpin agents.

Agent Definition

Agent software can be divided into two distinct classes defined by their state of mobility:

Static agents are programs which run on one host processor and interact with the outside world by the exchange of messages in a manner similar to traditional software.

Mobile agents are programs that are transmitted between processors in order to carry out their functions. So, rather than transmit messages to gain the information required they visit the distant site and run on a foreign host computer.

At the present time, static agents systems and architectures have

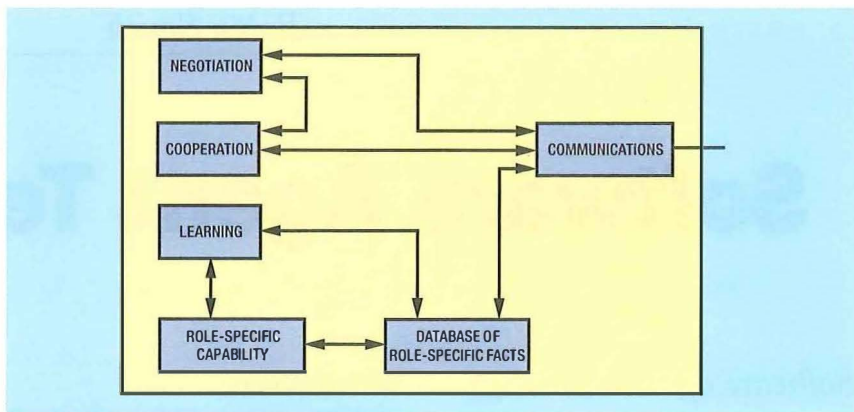
Figure 1—Generic structure for agent software

gained the most attention from the research community and the publishing industry. For this reason, most of the software technology descriptions and future applications will reference static agents. However, acceptance of mobile agents is growing and a brief description is provided at the end of this article to complete the picture.

The characteristics that most distinguish agent technology from more conventional software programs are the capabilities (in varying degrees) of negotiation, learning, cooperation and communication. This eclectic synthesis of ideas from system theory, computer science and artificial intelligence is at the root of the promise held out by agent technology. In its purest form, it offers the ability to view an application area in a completely new light. But, and this is a major problem with the image of the field, the agent paradigm abounds with terms and phrases which have subtle and complex meanings in a human context. The distinguishing characteristics used above (negotiation, learning and cooperation) are ideas central to any description of the human cognitive condition. And as is well understood, at today's state of the art, no computer program exhibits the range and cognitive power of the human mind. This is not to say that agents are unable to perform very useful and smart actions, but it is still very premature to talk about intelligent (agent) systems. This has been the underlying dream of AI and, even though much excellent progress has been made, the goal of intelligent machines remains elusive.

A generic structure for an agent can now be defined (Figure 1).

While the capabilities (of an agent) have been shown as well separated and isolated software modules in Figure 1, in any given implementation many functions will be merged in the interest of providing compact code. However, the agent can be considered to be divided into two complementary parts: the generic



capabilities which define the software as an agent and the role specific part which specialises the agent for the task at hand. In many ways, this description, or way of viewing agent technology, is reminiscent of the descriptions of expert systems in the 1980s. During that epoch the database of facts was considered to be distinct from the rule engine. Such a partition was conceptual rather than actual and a similar holistic approach may overtake agent design as the field matures.

Agent Classes

The simple structure shown in Figure 1 indicates how it will be possible to classify the roles or tasks that agent technology is able to address. All such classifications are to a large extent arbitrary. For instance the classification can be made on the type of function they provide—advice, task execution, prompting etc.; on where they carry out their function—Internet or telecommunications service; or a role specific classification³—management, development,

education, etc. This article uses a practical taxonomy based largely on the way agent technology has evolved at BT Laboratories. The two superclasses of agents are: agents that primarily interact with technical systems and agents that primarily interact with human clients. The second class also interacts with network databases and other technical aspects of the global information platform, but their main role is to assist their human clients (Figure 2).

Agent Descriptions—Interaction with Technical Systems

Any attempt to describe all the types and incarnations of agents indicated in the taxonomy of Figure 2 would lead to data/information overload which is exactly the situation the agent community is trying to contain. Therefore, examples are given of some of those agents which illustrate general properties of their subclass. First, agents that interact primarily with technical systems are described.

Figure 2—Classification of agent types

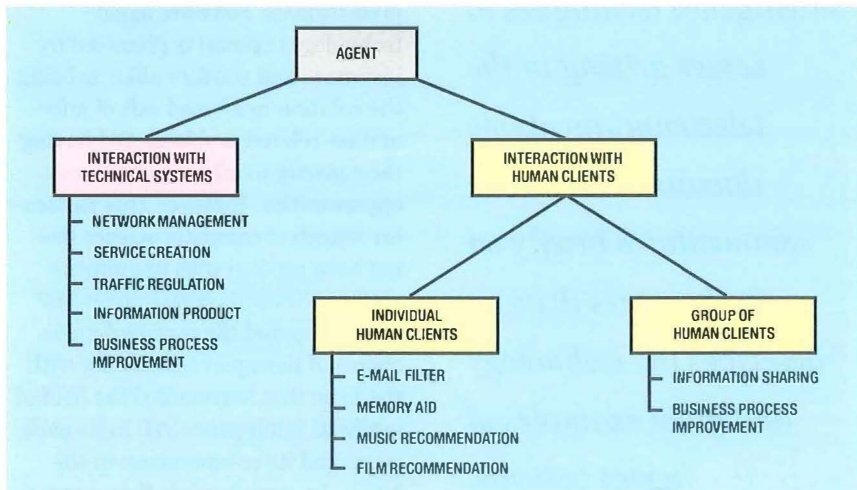
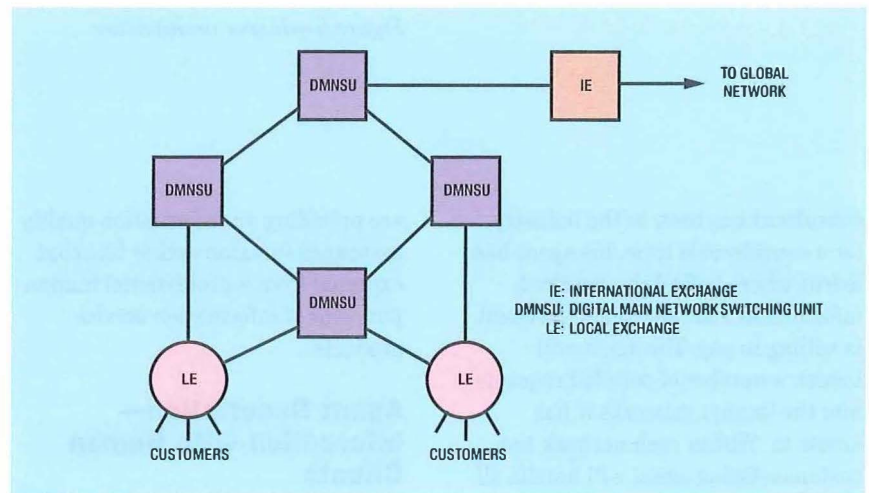


Figure 3—Simplified model of UK telecommunications network

Telecommunications network and service agents

Research at BT Laboratories began by investigating the use of distributed AI systems to manage BT's telecommunications network. This network (Figure 3) is a very large distributed system serving some 27 million customer lines connected to some 5000 local exchanges. These local exchanges are interconnected and provided with access to the global network through the digital main network switching units (otherwise known as trunk exchanges). This network is managed by a central operations unit supported by local network operation units⁴. As has been traditional in telecommunications systems the control regime is centralised and hierarchical. The UK telecommunications network has never experienced a country-wide failure and this excellent record must be maintained. In the future, where a much greater volume and variety of network traffic can be expected, reliance on traditional centralised management solutions could increase the risk of significant network-wide failure.

In order to explore the practicality of new management approaches, a number of agent technology solutions have been explored to the level of detailed computer simulations: fault management in telecommunication networks, certain aspects of service management⁵, multimedia interchange, etc. In an early experiment, a computer simulation demonstrated how, by having an agent associated with each exchange in a network similar to Figure 3, distributed fault management could be achieved. When the fault was located at the agent's exchange, only local corrective action was required. But when faults spanned a portion of the telecommunications network, the agents automatically formed 'management clusters' to localise and rectify the problem. This idea of agents communicating and negotiating to carry out high-level functions is at the core of the agent paradigm and is explored further in the following two examples.



Information brokerage network

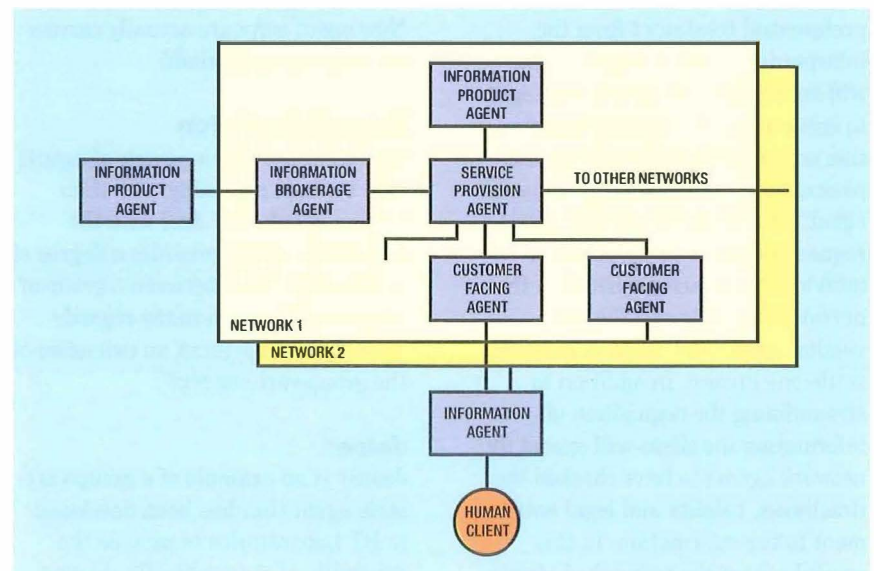
The scenario under investigation concerns a distributed system which seeks out timely information on the behalf of human clients. The vision for this system shares some of the functionality currently embodied in the Internet in the sense that it proposes a set of databases containing the information and a method of searching for specific information products. It is in the latter aspects where the scenario departs from current Internet practice. Figure 4 demonstrates a network of agents involved in managing the location, exchange and costing of the purchase of an information product.

The motivation for this architecture is that the agents located within the networks have detailed information about the products on offer by the product vendors. Additionally, the customer-facing agents will have an

in-depth knowledge of the types of information and style of interaction each customer wishes to transact. Figure 4 shows the customer connected to more than one network. With the increasing competition in communications such a situation is likely to become normal.

A typical information product transaction could conform to the following scenario. The human client is an investment consultant in a major merchant bank and has been briefed to investigate a possible take-over of an electronics company. In order to make a timely and accurate assessment, he needs to know (among other things) the state of the consumer market for the company's products; the health of the competing companies; the stock market value of the company and any underlying trends; the acceptance of the shareholders to a possible take-over, etc. Since our fictitious investment

Figure 4—Information brokerage network



consultant has been in the industry for a considerable time, his agent has learnt where to find the required information and how much the client is willing to pay. The agent will launch a number of detailed requests into the (many) networks it has access to. Within each network the customer-facing agent will handle all negotiations between the client's agent and the network resources. The customer-facing agent will access the request and formulate a request to the (network) service provision agent. The reason for the division of network responsibilities between the two (network) agents is that the customer-facing agent has detailed information regarding each of its customers while the service-provision agent is expert in the many information products and services provided by the vendors connected to the network. These may include constant updates on global stock market prices, insurance and reinsurance options, world commodity prices, etc. The service-provision agent will manage the connection of the vendor agents to the customer-facing agent and thereby effect the flow of timely information including the cost of the various information products required to satisfy the initial request.

Because the networked agents will be handling very many transactions it is likely that they are able to obtain significant discounts and preferential treatment from the information product vendors. This will enable the customer-facing agent to enter into a meaningful negotiation on the price of the information package with the (client) information agent. Also, even though the initial request will have involved many individual interactions between the networked agents and the many vendor agents, the client only has to settle one invoice. In addition to streamlining the acquisition of information the client will expect the network agents to have checked the timeliness, validity and legal entitlement to the information. In this special regard the networked agents

Figure 5—Jasper architecture

are providing an information quality assurance function rather like that expected from a professional human purveyor of information service products.

Agent Description— Interaction with Human Clients

It is interesting to note that the preceding discussions of agents that interact primarily with technical systems (network agents could be an alternative title) led naturally to agents that interact primarily with humans. Perhaps this should not be so surprising considering the present state of machine intelligence (MI). Even at their most powerful (in a cognitive sense) computer systems rarely, if ever, originate requests or actions at their own behest. At this phase of MI, computers are still the slaves of their human clients.

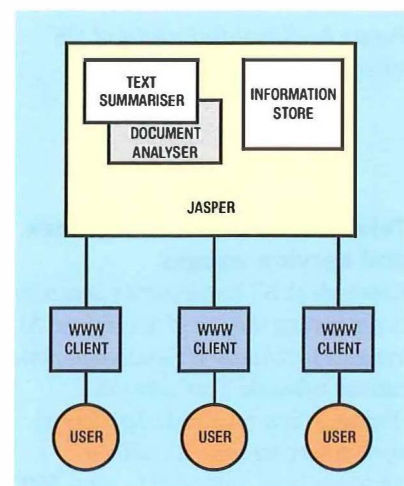
From the viewpoint of the technical literature and especially those more general articles that have popularised the field, agents that interact primarily with people have received the most attention. Perhaps this is to be expected. After all, agents are the first manifestation of computer science artefacts where people have delegated cognitive tasks to software programs. In the past, most aspects of AI were clearly in the role of advisor or analytic assistant. Now agent software actually carries out tasks on our behalf.

Group interaction

The first part of this review of agents that interact primarily with their human clients will deal with the technology which provides a degree of 'information glue' between a group of human workers. In many regards such agents represent an extension of the groupware concept⁶.

Jasper

Jasper is an example of a groupware-style agent that has been developed in BT Laboratories to provide the capability of automatically sharing

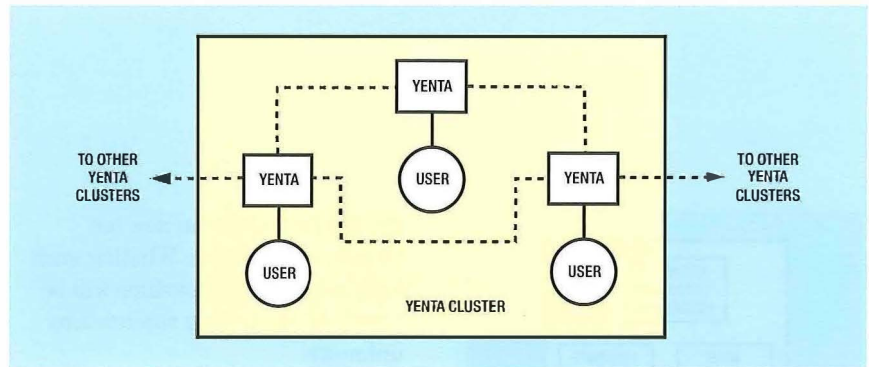


information between a group of research engineers working in a common area⁷. The motivation for developing Jasper was the vast amount of information available on the Internet. Rapid access to appropriate information is vital in the commercial research and development arena.

Before the availability of Jasper-style agents (Figure 5), people would hoard photocopies of significant information or fill the memories of their PCs and workstations with electronic images. Such a duplication of information is not cost-effective and indeed could be considered contrary to the motivating spirit of the Internet. Therefore, a Jasper agent will, upon instruction from its human client, prepare a brief summary of a selected Internet document, generate keywords, record the Internet address (URL) and annotate the entry with date and time stamps. Additionally, the human client can attach his/her own text to improve the subsequent retrieval process.

In order to facilitate the later retrieval of information from Jasper, a user will have prepared a list of key words. By using a keyword matching and scoring algorithm, Jasper will be able to return a list of documents which most closely match a user's declared interest. At this point in the description, Jasper provides a single-user information agent facility. But when other human clients provide a list of their interests, Jasper takes on the groupware role. By assessing the match between every stored document and the declared interest of all the users connected to the system, Jasper ensures that each client benefits from the information found

Figure 6—Yenta architecture



by his/her peers in the working group.

At the present time, when a user of Jasper wishes to change the context of his/her attention, the key word list will need to be modified. This is perceived to be a shortcoming of the current system and research is underway to permit Jasper to track the change of emphasis of its human client community in an automatic fashion. Reference 8 describes a system which provides a different approach to discovering users' interests. Rather than the user entering a set of key words, the agent (called AutoNomy) learns those interests by training an artificial neural network (ANN) on many examples of appropriate documents. This use of ANNs to extract knowledge from text, which is then used to guide search, is a further example of how mainstream AI is influencing agent technology.

Yenta

This desire to make agent technology adapt to the evolving needs of the human clients (and indeed the ever changing external environment) is a research goal at many leading establishments. For instance, workers at the MIT Media Labs have a long-term commitment to making agent interfaces adapt to the needs of people rather than users learning the complexity of (standard) computer interfaces and the rapidly expanding global databases.

Yenta⁹ (Figure 6) is a multi-agent system which has the potential to assist users buying and selling items over the Internet, discovering people of similar professional interests, etc. Yenta employs similar techniques to Jasper but extends the capturing of a human client's interest by including e-mail messages, the content of the user's computer files as well as the Internet text used by Jasper. In one very important way Yenta has extended the (geographical) scope of interaction. Jasper is a centralised system requiring users to be registered with the central store of

information. Yenta, however is a fully distributed agent system. Each Yenta agent will 'discover' other agents who represent clients with overlapping interests. In this way, the agents will build up a set of (agent) clusters representing common interest groups. These agent clusters provide the linkages between people who have opinions, ideas, interests, etc. in common. This capability can be viewed as an example of social information filtering¹⁰.

Morse

This idea of social filtering is at the heart of a number of agent systems which provide recommendations in the area of entertainment. Ringo¹¹ is a system developed at MIT which recommends music, and Morse, developed at BT Laboratories, provides users with recommendations on films to view¹².

These systems share the way the centralised agent combines the opinions of many users to be able to make specific recommendations for a particular user. The prediction of a user's preference is computed by comparing his/her ratings of a set of films with the ratings of a similar set by many other people. As long as the set of films used are compatible (that is, mainstream movies), then it is likely that people will have similar tastes and expectations. Therefore, if a particular user has not viewed a film from the compatible set, the average ratings given by the large number of people who have overall similar tastes should be a good indicator. The results from Morse and Ringo have indeed confirmed that it is possible to predict people's subjective preferences with a satisfactory error performance.

In these modern agent systems, the software builds up a knowledge of

human clients' preferences—for technical information (Jasper), music (Ringo) or movies (Morse)—by inspecting and manipulating lists of simple keywords or numeric ratings. Therefore, even though the learning and inference techniques employed are relatively simple, quite complex (and indeed useful) results can be constructed. It is instructive to contrast these (simple) knowledge representation methods with the more complex rule sets found in expert systems and the deep models which characterise certain knowledge based systems¹³. In many ways, this use of large quantities of simple human preferences represents a departure from mainstream AI research which tries to codify deep understanding. This appeal to a simpler representation of human knowledge may in part explain the greater acceptance of agent technology.

Interacting with individuals

The literature on, and software examples of, agents that interact with individual human clients is growing at a rapid pace. So fast in fact that it is not possible to provide anything like a representative coverage in an article of this size. E-mail agents, meeting scheduling agents, etc. are now an established part of the IT scene, and so the focus of this article is on two agents which illustrate well the leading edge of the field.

Coach

As the need for training and retraining becomes an established fact in commercial and industrial life, many players are looking for ways to automate the process. The days when every student could expect a human tutor to guide him/her through the intricacies of an assignment have passed into history. Information technology is being expected to fill the

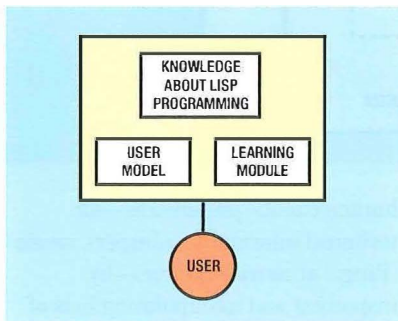


Figure 7—Structure of Coach

vacuum left by the absence of human tutors. Agent-based solutions are being explored in a wide variety of training and educational situations. Fields as diverse as foreign language tuition and the teaching of computer programming are coming under the agent spotlight. Coach¹⁴ is a system developed by IBM which helps students get to grips with the programming language Lisp (Figure 7).

Coach illustrates well how this class of agent fulfils its role. Built into the agent is a model of the user which encodes the proficiency of its students and a set of training rules which help users overcome problems they encounter while working through the structured training session. At the heart of the system is a detailed knowledge base about Lisp programming. This core of knowledge is incremental since the system automatically includes new user-defined functions.

In many regards, systems such as Coach define the novel aspects of agent technology. What we can see is the emergence of a symbiotic relationship between the agent software and its human client(s). When the system is instantiated the agent has a well structured but basic knowledge about the field and a generalised 'understanding' of the human client. As the interaction between the two players unfolds, the agent learns how its particular client wishes to carry out his/her tasks and builds up its knowledge of the specialist domain of activity. In common with all symbiotic relationships, the client will increasingly rely on the agent to

prompt and advise on new (or forgotten) situations. Whether such man-machine relationships will be beneficial in the long run remains unknown.

Remembrance agent

The fact that people do forget is exploited in an agent being developed at MIT Media Lab. The *remembrance agent*¹⁵ is designed to provide an unobtrusive prompting facility which alerts its user to information that may have a bearing on work in progress. This agent inspects text as it is being typed on a workstation or PC and tries to find a match to information stored in the computer memory. Such information could be documents produced by the user, those e-mails that have been retained and of course any Internet text that has been downloaded into the system. In effect, the remembrance agent acts as though it were a colleague looking over the user's shoulder offering helpful pointers to useful information.

Mobile Agents

All of the agents described previously interacted with distant sources of information by exchanging messages over communication channels. A quite different paradigm is evolving in the mobile agent community. Here the idea is that rather than exchanging a sequential set of messages to carry out an information transaction, software and data should be dispatched to the distant site¹⁶. After authentication, the mobile agent software will be able to conduct its information processing tasks on the distant host computer, package the answer and return home.

References 17 and 18 describe the *Telescript* agent system (Figure 8). In order for a user to be able to utilise mobile agents, his/her computer must have a Telescript engine installed. This interpreter converts the Telescript language (which codes the mobile agent) into commands and files understandable by the host computer. This interpreter provides




Figure 8—The Telescript system

two important functions: it imparts computer independence on the system and provides a measure of security for the host site.

If mobile agents are to command a significant place in the information age, platform independence will be essential. General Magic (the company that owns Telescript) describes its system as '... an object-oriented, remote programming language. It is a platform that enables the creation of active, distributed network applications. There are three simple concepts to the language: agents, places and "go". Agents "go" to places, where they interact with other agents to get work done on a user's behalf. Agents are in fact mobile programs capable of transporting themselves from place to place in a Telescript network.... The language is implemented by the Telescript engine. The engine is a multitasking interpreter that integrates onto an operating system through an applications programming interface (Telescript API). The Telescript engine is a server implementation.'

Adequate security measures are central to the ready acceptance of the mobile agent paradigm. These agents have certain characteristics not dissimilar to computer viruses. They enter a foreign computer system, run their own code and interact with the stored information. The Telescript system approaches this important issue on three levels. Because it is an interpreted system, the Telescript engine acts as a gateway between the agent and the host computer. Therefore it is able to prevent illegal instructions from

being executed. Second the Telescript engine limits the amount of computer resource (CPU time, memory, etc.) an agent can consume. The third level of security features encryption. The origin (computer site address) and total contents of the agent are protected by public and private key cryptographic techniques. In this way, a receiving site can verify the authenticity of a mobile agent.

Although Telescript agents are not the only example of mobile agent systems, they have received the major attention in the technical press. The features that these agents can perform have a great deal in common with those provided by their static cousins. Since the functionality of the two styles of agent (mobile and static) are very similar, market forces will determine their relative importance in the years to come.

Conclusions

Agent technology is now at a watershed. It has delivered practical applications that people are beginning to use and rely upon for their day-to-day activities. So where does all the foregoing leave us? For certain, the application of agent technology is growing in scope and depth. Simple agent systems which contain little in the way of 'machine intelligence' are becoming commonplace. These will provide small pockets of computer assistance to their human users.

Admittedly, the sophistication in handling a broad spectrum of information tasks is not up to demanding situations and agents are not yet that slick in their interactions with humans. But the strong indication is that various forms of agent will be with us well into the next century.

Acknowledgements

This article is a shortened version of a keynote address given at the Practical Application of Intelligent

Agents and Multi-agent Technology Conference, London, 22-24 April 1996.

The author is indebted to Professor Abe Mamdani and Dr Hyacinth Nwana for informal discussions which helped to crystallise the ideas presented in this article. The errors, omissions, etc. are of course the author's.

References

- 1 Critical Technology Trends. Feb. 1994.
- 2 WHEELER, T. E. It's the Information Not the Highway. *IEEE Communications Magazine*, Dec. 1995.
- 3 KING, JAMES, A. Intelligent Agents: Bringing Good Things to Life. *AI Expert*, Feb. 1995.
- 4 Network Administration for the 1990s. *Br. Telecommun. Eng.*, Oct 1990, Special Issue.
- 5 SMITH, R.; and GRIFFITHS, D. G. Distributed Management of Future Global Multi-Service Networks. *Br. Telecommun. Eng.*, Oct. 1994, **13**, p. 221.
- 6 GREENBERG SAUL. Computer-Supported Cooperative Work and Groupware. Academic Press.
- 7 DAVIES, J.; and WEEKS, R. Jasper: Communicating Information Agents. Proceedings of the 4th International WWW Conference, Boston, 1995.
- 8 PARTRIDGE, C. The electronic gofer with a mind of its own. *The Times*, 7 February 1996.
- 9 FONER, L. N. A Multi-Agent Referral System for Matchmaking. MIT Media Lab, Oct. 1995.
- 10 Social Information Filtering: Algorithms for Automating 'Word of Mouth'. Proceedings of the CHI '95 Conference, 1995.
- 11 SHARDENARD, U. Social Information Filtering for Music Recommendation. MIT Media Lab, Learning and Common Sense Group Technical Report 94-04.
- 12 FISK, D. Recommending Films Using Social Filtering. Martlesham M.Sc. Dissertation, Nov. 1995.
- 13 SMITH, R.; BUTLER, J.; and AZARM, N. A research and development programme for equipment and network diagnosis. *BT Technology Journal*, Jan. 1993, **11**(1), p. 79.
- 14 Indermaur, K. Baby Steps. *Byte*, March 1995.
- 15 To be published at PAAM 96, London.
- 16 APPLEBY, S.; and STEWARD, S. Software Agents for Control. Modelling Future Telecommunication Systems, Chapman and Hall, 1996.
- 17 WAGNER, P. Agents Away. *Byte*, May 1994.
- 18 WAGNER, P. Free Agents. *Byte*, March 1995.

Biography



Robin Smith
BT Networks and
Systems

For the past 12 years, Robin Smith has led teams at BT Laboratories on the application of AI techniques to the network and service management domain. During the past two years, his research interests have been concerned primarily with issues in machine intelligence. Robin is a Chartered Engineer (FIEE) and a visiting Professor at Imperial College of Science, Technology and Medicine.

Jonathan Legh-Smith

Navigating On-line Service Environments

As the World Wide Web (WWW) evolves towards an environment for commercial on-line service so too must its navigation systems. At present, users can choose from a selection of keyword-based search engines and hand-populated directories. These facilities provide crude but comprehensive navigation services. In the future, users will expect services that can recognise and distinguish between different types of service, and can match the technical capabilities of their systems with those of the services. To do this, navigation systems, indeed the WWW as a whole, must introduce key elements of a service architecture.

Introduction

Search engines are the most commonly used service in the World Wide Web (WWW) today and for many they represent the hub of the Internet. The reason for this is clear: if there was no way to search for information on the WWW, users would be faced with an impenetrable web of links and references. Search engines are a basic form of information retrieval system and are the principal navigation services on the WWW.

The role of a navigation service is simple: to help users find what they are looking for. Navigation services do not have to be perfect; they just have to meet expectations and search engines are a good example of this. What they lack in terms of quality of results they make up for in speed and coverage. However, expectations are increasing and several issues are on the horizon for navigation services. The goal of this article is to raise awareness of these issues.

This article describes the state of affairs as they exist today in the WWW, both from an Internet and intranet (private partitions of the Internet) perspective. It covers the basic requirements of the two domains and assesses how well these are being met. The article then discusses the two most fundamental problems that are likely to impact on navigation in an on-line service environment: distinguishing between different types of service and being able to ensure technical compatibility between the user's system and the service.

Starting With What Is Known—Navigating The World-Wide Web

The Internet

The dominant navigation service in the WWW is the search engine. Services such as Lycos^A and Alta Vista^B provide huge up-to-date databases covering much of the WWW. The service provided is basic but effective: a keyword search facility over more than 90% (according to the search engines) of the WWW pages on the Internet. It does take a little time and practice but most users on the Internet soon learn how to find the sites they are looking for. The coverage of information is excellent, as is the quality of the information (that is, it is up-to-date), the speed is formidable (although many search engines grind to a halt at midday), and, most importantly, the services are free to the user. What is more, the services are evolving rapidly and now provide additional features such as categorised indexes of reviewed pages and searching through news groups. As a result, it seems that most users are generally happy with the quality of service they get and many regard systems such as Lycos as their window onto the Internet.

The domination by the huge databases of Lycos, Alta Vista etc. does have its drawbacks. These services acquire their data by trawling the Internet and indexing all the information they find on their way. They are large centralised services that rely on a lack of resistance rather than cooperation.

Acquiring this data has a great impact both on the network and on WWW servers. Estimates have attributed up to 20% of network traffic to remote indexing. A more cooperative approach would of course be preferable, one where WWW servers index themselves and make this information available to those who want it. This would give webmasters control over when the indexing took place and hence reduce the load on their servers. It would also enable them to ensure that the indexes were up-to-date. The transfer of data would also be more efficient as it would then be a matter transferring preformatted data rather than downloading each individual WWW page.

Systems such as Harvest^{1,c} have been developed with this form of distribution in mind. Harvest combines indexing technology with distribution mechanisms to allow system administrators to balance the two problems of network traffic and processor loading. However, the large search engines are already dominant within the Internet and there are few users who are likely to worry as long as the service they get is free, fast and comprehensive. As a result, the use of distributed solutions has been restricted to intranets.

Intranets

The Internet is, in effect, supporting the entire global community. Intranets, however, exist to support the aspirations of a closed group such as a company or a virtual business, or possibly just a 'club'. Intranets use Internet technology and may be a part of, or separate from, the Internet.

Centralised search engines may appear to be suitable for the Internet but the approach is unlikely to be flexible enough for many large intranets. An intranet navigation scenario is described in the panel opposite. As the scenario illustrates, in order to meet the potential of the WWW, it will be necessary to utilise more sophisticated distribution techniques than just indexing

everything and putting it into a central database. Of course, any navigation service is better than none at all and a basic intranet search facility can always be provided by a centralised search engine. All that this requires is some party within the organisation (or intranet community) agreeing to maintain a complete index of all WWW servers within the intranet.

Organisations, however, are in a position to adopt a more efficient and flexible approach. Developing a

navigation infrastructure rather than just a central index database allows different parties within the organisation to build customised, business-focused navigation services. Such a navigation infrastructure would divide into two main components: the indexing infrastructure and the navigation facilities.

The indexing infrastructure should provide mechanisms for indexing WWW sites and for distributing and replicating the data between sites

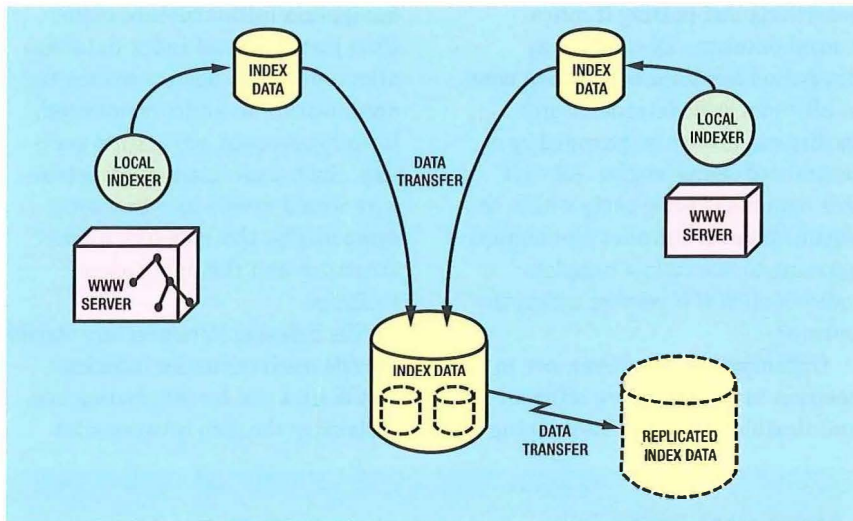
An Intranet Scenario

The most obvious scenario to choose is of course BT. Many parts of the company have independently adopted the WWW as an internal business tool and there are now in the order of three hundred WWW servers within BT. These servers provide access to valuable information such as: project deliverables, quality management systems, internal directories and CVs. In most cases, this information is well presented and structured and there are initiatives underway to ensure the quality and reliability of the information. However, making information accessible to users requires more than well organised home pages. Users need a navigation system to help them find the information in the first place. Such a navigation system obviously needs to ensure coverage of information. If users are to rely on the system as a business information tool they must be confident that they have access to all relevant information.

A central Internet-style index would be able to provide the coverage of information but it would not allow any flexibility in providing more focused services. For example, if a department wished to provide a navigation system that covered only internal information it would probably have to re-index all the relevant WWW servers. This would happen because a central index would be unlikely to understand which WWW servers related to which departments. Similarly, if a cross-business coordination group wanted a service targeted at information of specific interest to their activities they would have to build their own index. Both of these would result in duplication of the indexing process and the resulting database. In order to provide flexibility in configuring different navigation services, it is necessary to be able to break the index down into its constituent components; for example, an index for each organisational unit or project.

Having performed the indexing locally, the index data would then have to be transferred to the system that will provide the query service. (Leaving it on the local system would mean that every system in the organisation would have to support every query on the navigation system. This would have a negative impact on both the network and WWW servers.) The benefit of performing the indexing locally and then distributing the data is that different navigation services can make use of the same data. This would still result in replicated data but would not duplicate the indexing process, nor would it have as great an impact on the network.

Distributing the indexing process and providing mechanisms for transferring the index data provides an infrastructure upon which navigation services can be built.



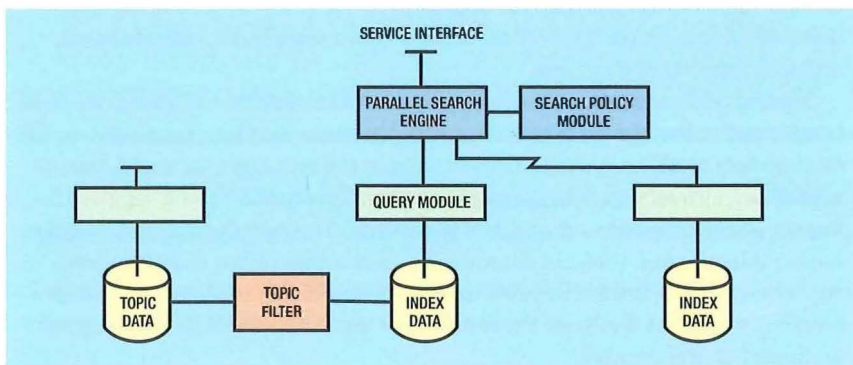
The indexing infrastructure provides mechanisms to gather and distribute the 'raw' data for use in navigation services. Indexing is performed locally (or from nearby) and is then transferred or replicated between appropriate sites.

Figure 1—A distributed indexing structure

(Figure 1). The indexing can always be performed remotely, but this should not be taken as a replacement for a distribution mechanism (which would transfer the resulting data from the local server to the remote server). In a highly distributed environment, such as BT, it is necessary to exploit all opportunities for load-balancing and reducing network traffic. The actual configuration within an intranet will depend on many factors, not least organisational factors such as reorganisations and constantly changing sets of WWW servers.

The navigation facilities would provide a toolkit with which to build the navigation services. This includes the basic facilities to query the index databases and to propagate queries to different systems and collate the results (Figure 2). In addition, the toolkit would be expected to include filtering mechanisms (for example, to select particular types of information automatically) and mechanisms that allow the users to state explicitly the type of information they are making available. Basic user interface modules would also be required to

Figure 2—Functional components of a navigation toolkit



A navigation toolkit would provide the basic functions to configure customised navigation services: parallel search engines enable queries to be performed over multiple indexes and topic filters can be used to select information on particular subjects. The navigation services, of which there may be many, are supported by a common indexing infrastructure.

support the rapid deployment of navigation services.

Much of the technology required for a basic navigation infrastructure already exists. As has been mentioned, Harvest incorporates many of the distribution mechanisms that are required. Parallel search engines exist in a number of forms and technologies are already available for filtering and categorising information such as Usenet news groups. The challenge therefore lies not in developing new technology but in using the available technology to configure services that cross organisational and system administration boundaries.

Summary

To assess progress in WWW navigation systems, it is necessary to distinguish between the Internet and intranets. Within the Internet it seems that the centralised approach of the large search engines is sufficient to meet current expectations. For as long as search engines provide free, fast and comprehensive coverage of the WWW, users are likely to accept the level of service provided. Within intranets however, it is necessary to explore more complex approaches to navigation. A navigation infrastructure, that distributes the process of indexing WWW servers and provides a toolkit to build customised navigation services, will be required if the WWW is to serve as a business information tool. All in all, however, current developments seem to be on course.

Looking Ahead—Navigating On-Line Service Environments

Problems on the horizon

The Internet has more challenging problems to deal with than how to distribute and customise navigation services. The issues that have been discussed so far in this article have been concerned primarily with information retrieval. However,

information is not the only type of resource that users are looking for in the WWW. There is of course the multimedia aspect of information—videos, music, images etc.—but even this is not the complete picture. The commercial model for the Internet is based on companies using the WWW as an access medium for on-line services. This includes selling information—in a sense WWW servers are already providing free information services—but it also includes services such as financial advice, home shopping for flowers, books, music, video etc. It seems only reasonable to assume that users will expect navigation systems to recognise the different types of service available and allow them to search accordingly. There is an obstacle to this, however, and that is that the WWW does not understand the concept of service, never mind how to distinguish between different types of service.

Not having the concept of service is of advantage to the existing search engines. In effect, all WWW pages are equal and hence they need make no distinction between them. Consequently, even if a user found a search engine that could process a natural language query such as 'I would like to buy a pair of shoes' (that is, the commodity is 'shoes' and the action is 'buy'), it would be of limited benefit. The search engine would be just as likely to return references to on-line catalogues (that is, services for browsing a company's product range), as it would references to on-line shopping services that actually allowed users to place an order with the company, which is what the user would have been looking for.

There are a number of existing navigation services that at first sight might appear to address this issue. These are the 'yellow-pages'/Yahoo^D style of services which organise WWW pages into categories. In effect, the type of a service is implicit in its positioning in the category structure. The issue here is that these navigation services are populated by hand

and, with over 20 million pages on the WWW, such services are unlikely to come close to the coverage of the Internet search engines. They are also inflexible: modifications to the category structure are limited by the amount of time and effort required to re-sort the entries in the original categories. To be effective in an on-line service environment, a navigation system must be able to determine what type of service a user is looking for.

Another, possibly more urgent, problem is how to cope with the increasing diversity of the underlying technology. New technologies are always being introduced into the Internet and if navigation systems are to be able to match user requirements with available services they must be able to take into account the technical configuration of both the client and server systems.

The current breed of Internet search engines treat one WWW page just like any other. Even if you ignore the service type issue described above, this is not a valid assumption. Consider HyperText Markup Language (HTML)^E, the 'standard' WWW page description language. A standard exists but this is not the version in common use. Several extensions are available and, while most browsers support most extensions, not all browsers support all extensions. Certainly not all HTML editors support all extensions. While the development of a base-line HTML standard plus standard extension sets may appear to be the answer, this will not help the users much. If their browser does not support the right extension set, they often will not be able to view the page properly.

HTML does not provide the worst compatibility issue. Pages that exploit Java^F and Virtual Reality Modelling Language (VRML)^G, for example, are not visible to all browsers. Technologies such as Vosaic^H and RealAudio^I enable users to listen to music in real-time over the Internet. Not surprisingly, however, the client and server

technologies are not interchangeable. This leaves users with the challenge of finding a site that not only has the music they want but that uses the correct technology for their systems.

The configuration problem becomes harder as the set of available technologies becomes more extensive. Even if the problem of the basic use of technology is overcome (for example, by having a dominant browser or a standard 'plug-in' interface), this will not address the problems associated with user mobility. Using a VRML-enabled page is satisfactory if one is at work and has a high-bandwidth connection, but it is something of a liability if one is at home using a low bandwidth modem.

If navigation systems are to help address this problem they will need some way of determining what technology is available on the user's system and what technologies the service provider is capable of supporting.

These issues and other related topics, such as user profiling, suggest that a service architecture is necessary—one that understands the basic concept of service and recognises the technologies required to access services. This need not be a rigorously defined architecture with standardised components and programming interfaces. It need only be sufficient to provide a coherent framework within which navigation systems can be designed. Architectural coherence has never been the Internet's strongest point however.

Ways forward

The previous section suggested that the Internet must look towards a service architecture if it is to create an environment for commercial on-line services, or at least if it is to support navigation services within such an environment. It is not possible to predict the steps that the Internet will take in this direction. However, looking ahead to future architectures will provide some guidance as to the facilities that will develop. The work on the telecommu-

nications information networking architecture (TINA)^{1,2,3} provides the most appropriate starting point.

TINA provides a service architecture for broadband multimedia services. It has taken a distributed systems perspective and identifies mechanisms to support the replication of components across the system, notification of events between components etc., as well as functions to support naming, trading and type management. Much of this is too advanced for the Internet, either in terms of capability or possibly willingness to cooperate. However, TINA does identify a number of components of direct relevance to navigation systems and that could be used in the Internet. These are the provider agent and the user agent.

The provider agent resides in the user's own system—the client system. The provider agent is involved in establishing and managing connections with remote systems. It is responsible for meeting the user's quality-of-service requirements by matching the client system's capabilities with those of the server system. To support its role, it holds details of the technical configuration of the client system. The user agent is a mobile object (or rather it is not assumed to reside on any particular system) that holds details about the user. It is, in effect, the user's identity within the entire service environment. The user agent holds information such as subscription details, interface preferences, billing details, etc. When the user registers with the system, the user agent is consulted to determine which services the user has access to, which services should be configured immediately, etc. When the user uses a new or non-subscription service, the user agent may be consulted to find billing details. The user agent is not usually held on the client system as that would obstruct mobility. It is possible that when the user registers that he/she is on-line, the user agent is downloaded to the user's system.

The provider agent and the user agent can be used within a TINA-

based navigation system to overcome at least some of the problems identified in the previous section. The provider agent in particular can be used to ensure that the services offered to the user by the navigation system actually support the correct technologies, and, moreover that they are capable of meeting the user's quality-of-service requirements. A restricted form of the provider agent would clearly be of use to Internet-based navigation systems.

An important component in navigation systems is likely to be the use of agent technologies. Their principal use in navigation at present is as personal information-retrieval systems^{4,5}. An agent contains a description of a user's interests and is charged with searching, or maybe monitoring, for information of potential interest to the user. This will be familiar to distributed systems experts in that, from a system perspective, the fundamental property of an agent is that it is an encapsulation of a user's interests; that is, it is an object.

The most important factor governing the use of agents in navigation will be scalability. Consider, for example, an agent which contains a profile of a user's interests and is charged with continually searching the WWW looking for relevant information. If the most powerful WWW crawlers only manage to keep up with 90% of the Internet, what chance does a single agent have? Such an agent could work only if it were designed to exploit the available services, in this case the Internet search engines. Other uses of agents should also identify and develop the supporting services they require for scalability.

As distributed applications, navigation systems will be highly distributed and agents will most likely be found at the periphery of the system. One possibility for the navigation infrastructure described earlier is that it becomes an environment for communities of agents. To achieve this, will require solutions to

problems such as: how different agent technologies interact with each other and how agents find each other. The result is likely to look like a highly distributed object system.

Unfortunately, neither the service architecture of TINA or the use of agents provide a direct solution to the problem of service types. TINA provides a conceptual framework within which to discuss the issue of service types but does not indicate how to represent them. The problem is that the types we are dealing with are rather abstract and very definitely user-oriented. The value, to either the user or the navigation system, of developing a formal type system that believes `pizza_delivery` and `home_shopping(clothes)` are both sub-types of the abstract type `delivery_service` seems questionable. The problem is that users tend to have a rather fuzzy notion of the type of a service whereas the navigation system needs a more formal type system if it is to process queries such as: `((shoes => clothes) + (buy => shopping)) => service_type = home_shopping(clothes)`.

One possibility for resolving this problem is to avoid using a type system at all. If we are to support natural language interfaces then one option might be to represent the type of a service as part of the natural language system itself; for example, using semantic networks. Hence, a query is processed by using the semantic networks from the available services. This is clearly an area for further study.

Summary

Although existing search engines are sufficient to meet current expectations, the commercial concept of service is likely to have a profound effect on navigation services. Initiatives such as TINA provide an indication of the functions that must be introduced into the Internet to support navigation. More generally, it must be recognised that the commercial model of the Internet will require the introduction of a service architecture. The use of agent

technologies is also likely to form a major component in future designs.

Conclusion

This article has discussed how the future of navigation services in the Internet will depend, not on the development of information retrieval techniques, but on the move to a commercial model of on-line services. The implications of providing navigation services in a commercial service environment is that the Internet recognises at least some aspects of a service architecture. Architectures such as that provided by TINA can be used as a target although we cannot assume we know how the Internet will develop.

Within intranets the problem is likely to remain one of information retrieval. Even here we must look towards a distributed navigation infrastructure. Centralised search engines, while sufficient in the Internet, are not flexible enough to promote the use of the WWW as a business information tool. In particular, navigation services must be customised.

The issues associated with navigation in on-line service environments are closely related to the requirements for a service architecture. If navigation systems are to be effective they must be able to understand what a service is and be able to distinguish between different types of service. In addition, navigation systems must begin to take into account the technical requirements of using a particular service. Users should at least be aware of any incompatibilities between their systems and the services they are attempting to use.

In both cases much of the required technology already exists in the contributing fields of distributed systems, information retrieval, natural language processing and human factors. It is therefore important that the opportunity is taken to integrate and experiment with these technologies.

References

- 1 BOWMAN, C. MIC.; DANZIG, PETER B.; HARDY, DARREN R.; MANBER, UDI; SCHWARTZ, MICHAEL F.; and WESSELS, DUANE P. Harvest: A Scalable, Customizable Discovery and Access System. Technical Report CU-CS-732-94, Department of Computer Science, University of Colorado, Boulder, Aug. 1994 (revised March 1995).
- 2 LEGH-SMITH, JONATHAN. Outline Model for Navigation. Subscription and Location Processes. 11 May 1995.
- 3 BERNDT, HENDRIK; and MINERVA, ROBERTO. Definition of [TINA] Service Architecture. TB_MDC.012_2.0_94, Feb 8, 1995.
- 4 DAVIES, N. J.; WEEKS, R.; and REVETT, M. Communicating Agents for WWW. 4th International WWW Conference, Boston, USA, Dec. 1995
- 5 MADSEN, MARK. ANSAweb: A Service Architecture To Enhance The World Wide Web. *First Class*, VI(1).

World Wide Web URLs

- A Lycos. <http://www.lycos.com>
- B Alta Vista. <http://www.altavista.digital.com>
- C Harvest. <http://harvest.cs.colorado.edu/harvest>
- D Yahoo. <http://www.yahoo.com>
- E HTML. http://www.sandia.gov/sci_compute/html_ref.html
- F Java. <http://java.sun.com>
- G VRML. <http://vrml.wired.com/vrml.tech/vrml10-3.html>

H Vosaic. <http://choices.cs.uiuc.edu/Vosaic/Vosaic.html>

I RealAudio. <http://www.realaudio.com>

J TINA. <http://www.tinac.com>

Biography



Jonathan Legh-Smith
BT Networks and Systems

Jonathan leads a multidisciplinary team designing navigation systems for on-line information services. Prior to his involvement in Internet-related activities, Jonathan was involved in the development of distributed computing technologies, in particular the International Standardisation Organisation (ISO) standards for open distributed processing. Jonathan is also BT's representative to the Object Management Group, a computing industry consortium focusing on distributed systems technologies. Jonathan joined BT in 1987 and is based at the BT Laboratories in Suffolk.

From Books to Bytes— Managing Information in the Information Age

With the dawn of the new information age, the risk of information overload becomes ever more apparent. As the volume and complexity of available information rises, people will need increasingly sophisticated tools to help them to locate and manage critical information. Technologies such as BT's text summariser can help by highlighting the most important parts of documents. And systems which can understand text offer even greater potential.

The Information Society

Since the days of Samuel Johnson, the amount of information available has grown enormously, as has the number of places where it can be found. Today's computer systems and telecommunications networks allow information to be generated and stored in unprecedented amounts and allow it to be delivered quickly anywhere in the world. Businesses now rely on accurate, detailed and timely information in order to survive. Indeed, more streamlined information flows within companies are one of the factors which have led to profound changes in the structures of the companies themselves, such as the de-layered, flatter structures which are characteristic of modern business practice. As technology continues to advance and the rate of change in the business world accelerates, companies will increasingly be differentiated by their abilities to cope with large amounts of ever-changing information.

But it is not only at work that effective access to information is needed. At home, we use information to allow us to make decisions, to manage our lives, and for entertainment. And quick and efficient access to information plays a vital role in education. Society is changing from one in which most people remain in the same career throughout their working lives towards one in which people will have many careers in a lifetime. In this environment, 'life-long learning' will become a reality and will be supported by advanced

Knowledge is of two kinds. We know a subject ourselves, or we know where we can find information upon it.

Samuel Johnson (1709–1784)

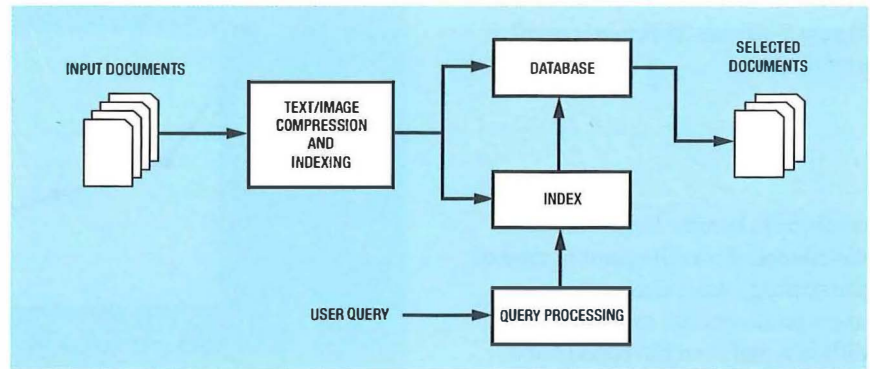
information management technologies to locate, filter and marshal the information people need.

For many years, BT has been investing in the infrastructure needed to handle information, researching and developing technologies with which to build advanced information networks. Of no less importance are the development and provision of appropriate systems to allow users to access and manage the information flows that these new networks make possible, thus providing new value-added services for the future. These sophisticated information management systems will play a key role in the new information industry.

Current Information Retrieval Technologies

Samuel Johnson would have gained much of his information from books to which he had physical access. Through modern telecommunications and computer systems, people are no longer constrained by the physical location of information; they can browse information stored on computers on the other side of the world almost as easily as if it were on a local computer. However, efficient methods of indexing and searching are essential to enable people to find what they require among the sheer volume of information available.

Figure 1—Typical text retrieval system



Information indexing and searching

Since the advent of the Internet and especially the World Wide Web (WWW)¹, many people have become familiar with search engines and information retrieval systems. Figure 1 shows the main elements of a typical text retrieval system. Input documents in electronic form are processed to form a database which contains the document contents and an index which can later be searched in response to users' queries to pick out appropriate documents. If data compression is used, this database and index together will usually take up less disk storage than the original documents!

The index and database do not need to be held on the same computer. For example, a WWW 'spider' such as Lycos² is a program which searches the Web and builds an index to the pages it finds. The spider stores only the index, which contains pointers in the form of hypertext links to the actual pages which remain on their original machines.

Once an index has been constructed, users may pose queries to retrieve the information they are looking for.

Retrieving information

Queries can be divided into two main types, *Boolean* queries and *ranked* queries.

Boolean queries consist of a list of words to be sought, combined using the operators AND, OR and NOT. For example:

text AND retrieval.

This simple query would miss many documents which, though pertinent to the subject, do not include these exact words. A more useful query might be:

(text OR data OR document OR image) AND
(retrieval OR archiving OR storage OR indexing).

Building effective queries requires insight and experience, and many people currently prefer to use the services of librarians and other information professionals rather than attempt to grapple with these problems alone.

Boolean queries are 'all or nothing'; if the requested combination of terms is present in a document, it will be retrieved. This can be very effective where documents are expected to contain specific terms, but is less good where the search is on a more general subject, such as 'information management'.

Ranked queries make use of a heuristic which measures the *similarity* of each document to the query, and returns the most closely matching documents. Queries may consist simply of a list of words which are expected to appear in relevant documents. Steps have to be taken to ensure that rankings are not unduly biased by the lengths of documents, or by how common the query words are. One important class of methods is the *vector space* or *statistical* approach.

In practice, ranked queries can work well, achieving performances which would be difficult to reach with a Boolean query. The penalty is that more complex index information is required than for Boolean queries, and processing costs are higher.

Relevance feedback

Some search engines, once a few documents have been retrieved, allow the user to provide feedback about which ones are most useful, effectively telling the system 'find me more like these'. The documents selected by the user are analysed, the ranking weights adjusted accordingly, and a new search based on the revised weights is carried out. Experimentally, it is found that a

single round of feedback brings substantially better retrieval, and a second round brings a small additional improvement.

As an alternative to analysing relevant documents in their entirety, the system may present the user with a list of important terms from them, allowing the user to add these terms to the query if desired. Although this may be more work for the user, it does give him or her more control over the querying process.

Probabilistic ranking

Another ranking method which can be applied once some documents have been judged relevant is based upon *probabilistic models*. Each term which occurs in a document is considered either as evidence that the document is relevant, or that it is not relevant. Conditional probabilities are calculated using Bayes Theorem and used as weights in the ranking process. One commercial search engine which uses probabilistic ranking is Muscat³.

User interest profiling

Many people have a need to monitor ongoing developments in particular fields of interest. A suitable query, which may be a Boolean string or a set of weights for a vector space model, can act as a user profile and can be run periodically as new information becomes available. This facility is provided by many information systems and is also applied in agent-based systems.

Performance of Current Text Retrieval Technologies

Although currently-available text-retrieval technologies can be very effective, they have their limitations, and these will become increasingly important as information sources

Figure 2—Trade-off between recall and precision

continue to become larger and more distributed. From the point of view of the ordinary user, constructing appropriate queries can be quite difficult, and even the concept of a keyword is unfamiliar to many people. However, these systems have their limitations even in the hands of experts.

Recall and Precision

Text retrieval systems are usually characterised by two measures, *recall* (R) and *precision* (P). Recall measures the retrieval effectiveness of the system: the number of documents which are actually retrieved compared with the total number of relevant documents in the database:

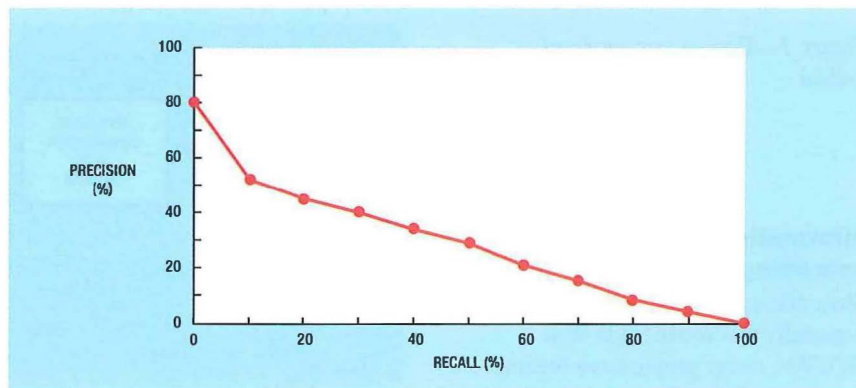
$$R = \frac{\text{number relevant that are retrieved}}{\text{total number relevant}}$$

It has long been realised that the recall of text-retrieval systems can be quite low. One study⁴ concluded that less than 20% of the relevant documents were retrieved from a database containing roughly 350 000 pages of text, using a full-text retrieval system. Interestingly, the subjects involved in the trial believed that they were retrieving a much higher fraction of the relevant documents than they actually were. The authors also claim that evaluations done on small databases do not scale to large databases, since the amount of search effort required to obtain the same recall level increases as the database increases, often at a faster rate than the increase in database size. This has important implications for the very large information sources which are now becoming available.

Precision P measures how relevant the retrieved documents are to the query.

$$P = \frac{\text{number retrieved that are relevant}}{\text{total number retrieved}}$$

A reasonable precision is important if the user is not to be deluged with irrelevant documents. Anyone



who has used text-retrieval systems will be familiar with the effects of poor precision, and while the occasional 'rogue' selection can be amusing, searches which yield large numbers of irrelevant documents can be highly frustrating for the user.

Unfortunately, recall and precision are mutually dependent. The results in Figure 2 were obtained as part of the TREC conference in 1993⁵. This curve shows clearly the trade-off between recall and precision. If a search aims for high precision, then only a small fraction of relevant documents are retrieved. On the other hand, a search with high recall also finds a large number of irrelevant documents which must be discarded. In the limit of 100% recall, the precision has dropped to zero, which means that all of the relevant documents can only be found by retrieving the entire database!

Making Use of Information

Although current information-retrieval techniques have their limitations, they nevertheless play an essential role in helping people to pick out useful material. Imagine for a moment that we had a perfect information-retrieval system, which could find all the documents in the world relevant to our needs, and only those. Would the problems of information overload be behind us? Unfortunately, the answer is no!

Consider a field where information can literally mean the difference between life and death—medicine. An estimate⁶ shows that well over one million items of information are relevant to medical practice. Under ideal conditions, the average student can absorb seven such items per hour. So a full-time student working a 12 hour day for 250 days a year would take approximately 50 years to catch

up with current knowledge! And of course the problems don't end there, since new information is being published all the time.

Problems on this scale, which are by no means unique to medicine, starkly illustrate the growing information overload which we are facing. Quantities of information are now being generated which are beyond the capabilities of humans to handle. In order to cope, people need more than search engines; systems are required to help people to assimilate information—and ultimately systems which understand information for them!

Text Summarisation

In many cases, people find that a whole document is too large a unit of information, and their needs would be better served by a paragraph, or even a few sentences. In other words, the *information granularity* is too large.

NetSumm is a program developed at BT Laboratories which automatically picks out those sentences which it considers make up the 'most important' part of a document. These key sentences can either be highlighted in the original text (giving an effect similar to using a highlighter pen), or can be extracted to produce a summary of the document.

The summarisation process is interactive, allowing the user to choose longer or shorter extracts at will, from a single sentence to the full text. Typically, one begins with a very short abridgement, to see if the article is relevant. If so, the length of the abridgement can be increased to see more details.

NetSumm is currently available on an experimental basis on the BT Laboratories WWW server⁷, and can be used on any WWW document

available on the Internet, as well as on documents provided by users. Figure 3 shows NetSumm about to summarise an article; the summarisation process is controlled by NetSumm's toolbar which can be seen near the top of the screen. Clicking on the SUMMARISE icon (the third from the left) produces the result shown in Figure 4.

How does NetSumm work?

NetSumm uses a combination of natural language and statistical techniques to build up a model of the 'core structure' of the article. Each sentence in the article is then given a score according to how well it is connected to this 'core' model. Summaries are then generated by including all of the sentences with scores above a certain threshold, which is changed interactively as the user clicks on the LONGER/SHORTER buttons on the NetSumm toolbar—a low threshold allows many sentences through, giving a long summary, while a high threshold admits only the most relevant sentences, giving a short summary. When a new article is summarised, the initial threshold is chosen to give roughly a page of text, though this will obviously vary according to the user's screen and font size.

At present, NetSumm's summaries are 'neutral'; they look only at the article itself, though ways are being investigated to take account of the user's known interests to produce a more personalised summary.

The techniques used by NetSumm appear very robust, giving it the ability to work on almost any text, independent of subject. Equally, good results have been achieved with articles on such disparate subjects as semiconductor lasers and red squirrels!

Evaluating NetSumm

The usefulness of a summariser depends on how well it extracts the key information from an article. NetSumm has been evaluated on a number of technical articles with author-written abstracts. On the

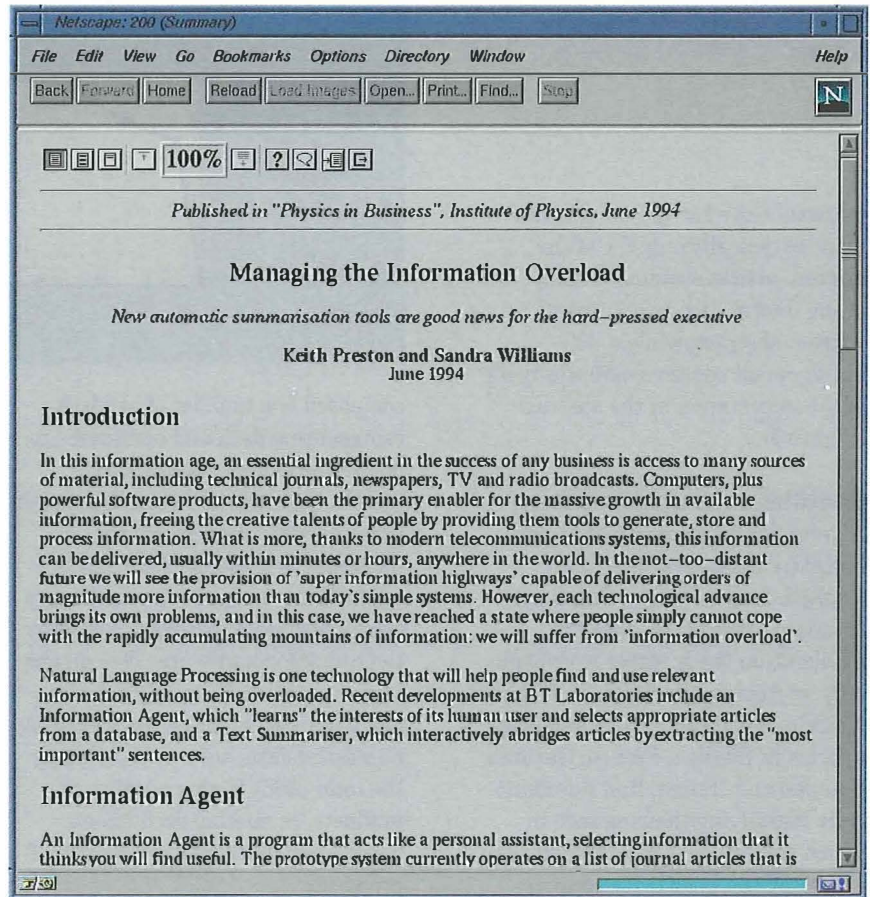


Figure 3—NetSumm in use

Figure 4—19% summary produced by NetSumm

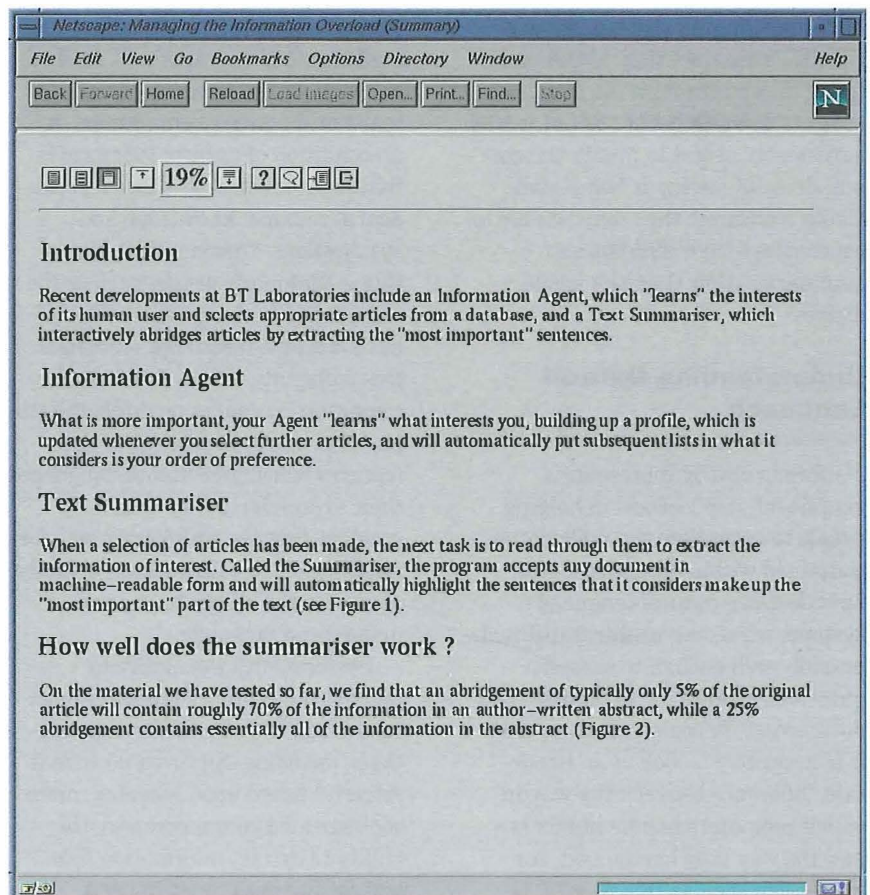


Figure 5—NetSumm evaluation results

material tested so far, an abridgement of typically only 5% of the original article contains roughly 70% of the information in an author-written abstract, while a 25% abridgement contains essentially all of the information in the abstract (Figure 5).

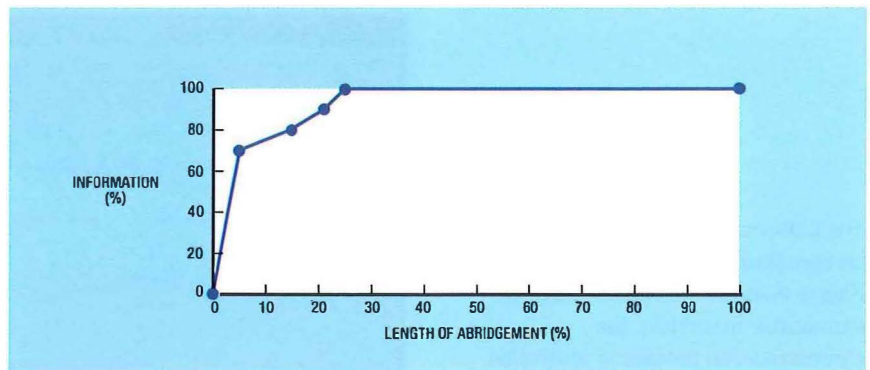
Multilingual summarisation

A great deal of information is published in languages other than English. For example, it has been estimated that of the material available on the Internet worldwide, only approximately half is in English. NetSumm can currently summarise articles in English, French, German, Spanish and Italian, first automatically identifying the language in which the article is written and then accessing the knowledge sources needed for that language.

Even though NetSumm does not translate between languages, its multilingual capability can benefit people who encounter a document written in a language in which they have some familiarity but are not fluent. Ordinarily, they would have to try to read through the article in order to decide whether or not it was sufficiently useful to justify the cost and delay of having it translated. Using NetSumm they may be able to concentrate their efforts on a summary rather than the whole article.

Understanding Natural Language

NetSumm clearly represents a significant step forward in helping people to assimilate the information contained within documents. But how close are natural language systems which can **understand** text, possibly well enough to answer questions about it, or to compare two documents? To begin to answer this, it is necessary to look at an important difference between the way in which computers handle numbers and the way they handle text. An expression such as $(8 + 3) \times 17$ is



composed of a number of symbols representing data and operations to be carried out on that data. It is relatively straightforward to program computers to 'understand' the meanings of these symbols and to apply these operations; in fact many operations are directly supported by the computer hardware. This means that computers can easily calculate new information based upon existing numerical data, such as calculating the total profit from a number of products, or making predictions based upon mathematical models of processes.

A sentence in English or any other language can also be represented as a sequence of symbols (words): 'the man went to the bank'. However, in order to understand this sentence, the computer needs to know far more than a few rules for performing mathematical operations. It needs to know about the words used and the real-world entities behind them. A precondition of natural language is that its users already share a great deal of common knowledge and assumptions. One manifestation of this is that many words, such as the word 'bank' in the sentence above, have multiple meanings. Although this ambiguity causes problems for computers, it causes no difficulties to humans, and is in fact one of the features which give human languages their expressive power. It is no accident that the well-known test for intelligence, the Turing test, actually measures the ability to use and understand language.

Systems with the ability to understand ordinary language open up a range of powerful new applications, including improved document retrieval based upon meaning, more sophisticated summarisation, the ability to extract information from text for further processing and

providing assistance to people authoring new documents. Some level of understanding is also needed in spoken-language systems, which allow people to use speech to access information or to give instructions.

Domain-specific understanding

One way of reducing the problems of ambiguity and need for word knowledge is to build systems which are specific to particular subject areas or *domains*. For example, the spoken-language interface to the BT Business Catalogue⁸ understands a range of questions about telephones and related products and can respond in a reasonable way. But it cannot chat to you about the weather, and in fact would not even recognise the words needed to do so!

To extend such systems we need to give the computer more knowledge; not just about the words used and the ways in which they can be combined, but also about the real world which underlies them. Significant progress is being made in this area worldwide, not just through the development of ever-faster computers to handle the extensive processing required, but also through an improved theoretical understanding of the issues involved and the best ways of realising these systems. The day may not be too far off when people have systems which can read the financial newspapers and tell them about the likely impact of today's news on their investment portfolio, or the possible effects on their companies!

Conclusions

Future information networks have the potential to help people to solve problems and to make decisions, to enhance their creativity and to enrich their lives—at home, at school and at

work. However, unless sophisticated tools to help manage this mass of information are available, these benefits will not be achieved. On the contrary, if people are continually swamped with irrelevant and useless information, they will simply not use these networks.

Advanced information management systems, which not only allow people to find what they want but which can also help them to assimilate and understand it, will be one of the keys which will unlock the full power of the information age.

References

- 1 BERNERS-LEE, T.; CAILLAU, R.; GROFF, J-F.; and POLLERMANN, B. World Wide Web: The Information Universe. *Electronic Networking: Research, Applications and Policy*, Vol. 2, Part 1, pp 52-58, 1992.
- 2 <http://www.lycos.com/>
- 3 <http://www.muscat.co.uk/>
- 4 BLAIR, D. C.; and MARON, M. E. An Evaluation of Retrieval Effectiveness for a Full-Text Document-Retrieval System. *Communications of the ACM*, Mar. 1985, **28**(3).
- 5 MOFFAT, A.; and ZOBELL, J. Compression and Fast Indexing for Multigigabyte Text Databases. *Australian Computer Journal*, 1994, **26**(1).
- 6 BENSON, T.; and NEAME, R. *Healthcare Computing*, Longman, 1994.
- 7 <http://www.labs.bt.com/innovate/informat/netsumm/>
- 8 WYARD, P. J.; SIMONS, A. D.; APPLEBY, S.; KANEEN, E.; WILLIAMS, S. H.; and PRESTON, K. R. Spoken Language Systems—Beyond Prompt and Response. *BT Technol. J.*, Jan. 1996, **14**(1), pp. 187–205.

Biography



Keith Preston
BT Networks and
Systems

Keith Preston gained a B.Sc. in Applied Physics and Electronics from the University of Durham in 1978. After joining BT, he was involved in pioneering work in optical communications, and has many publications and patents in this area. He was also responsible for the design and prototyping of a commercial range of advanced laser transmitter products. In 1992, he moved into the area of advanced software and now heads the Natural Language group at BT Laboratories, investigating a range of futuristic information management and language understanding techniques.

Chris Gibbings

Security in the Information Age

The greatest challenge in future networks will be to produce services that customers can trust. Combined with their needs for mobility, usability and global reach, this represents a difficult goal. Although security building blocks can be very strong individually, it is the way in which they are put together that is important. This article argues that security must become a central part of the design process.

Introduction

Future service and information providers will offer high-value goods and services over the information network. Their confidence in the network provider and the network's underlying security and reliability will be an essential prerequisite to winning their business. Customers too will have legitimate concerns about the security of electronic commerce—will they wake up bankrupt one day after a hacker has impersonated them electronically?

Network abuse has a long history. Worried that operators were passing calls to his rivals in the undertaking world, Strowger was inspired to invent the automatic telephone exchange. If there are inadequacies in the information network, is it too fanciful to suggest that some 21st century Strowger will respond in a similarly revolutionary manner? The stakes are high, as networks will take a more and more central role in commerce and entertainment. Customers will have to have confidence in their network operator. Generating trust and responding to customers' security needs will be areas of competition between network operators.

Despite this importance, the way in which security is perceived may be a major obstacle to progress. Security can appear as a tiresome set of commandments and restrictions that impede real work. It is sometimes thought of as a 'black box' isolated from the rest of the system that will, when developed by the appropriate high priests, somehow prevent any abuse taking place. Where analysis and protection are applied to a limited part of a product, it should

come as no surprise that, while one area is strongly protected, fraudsters simply attack another feature and bypass it.

In this article, it is argued that security should gain a more positive image as an enabling technology that extends the range over which services can be offered, and gives a competitive advantage. The mindset of mechanisms and fixes should be abandoned as security moves to a central position in development.

Firstly, a simple example of electronic commerce is described, followed by a discussion of mobility in both current and future services, to illustrate several aspects of security. These aspects are then considered in more detail in the second half of the article.

Commerce over the Information Network

Shopping

A simple example of an electronic purchase illustrates many of the areas where security must be considered. Dee sits at home browsing through the on-line catalogue for 'Honest Ed's Widget World'. She decides to purchase a new, improved widget for £100. She presses 'RETURN', and waits for Ed to despatch her purchase. There could be a few pitfalls:

- The widget doesn't turn up—Ed claims to have not received the order.
- The widget doesn't turn up—Ed claims to have sent it, and insists on payment.
- Dee dishonestly claims to have not received the widget so that Ed will send another.

- Dee's account is debited, but Ed's isn't credited.
- Ed charges for the expensive HyperWidget, claiming that Dee had ordered it.

This is just the start of a long list of potential problems which have to be taken into account. Services delivered electronically also raise new questions of how to protect copyrights and charge for use. There are several fundamental services that should be provided:

- Guaranteed levels of service—customers and service providers will want to get through to each other first time, every time.
- Some means of identifying the various parties to each other. Is 'Honest Ed' on the line or 'Nasty Nick', his rival?
- Protection for messages in transit, so that nobody else can read or change them.
- The ability to prove that you have sent a particular message ('send me a widget' in the above example).
- A way of confirming that messages were received at the other end.
- Reliable ways of transferring money.
- A method for resolving disputes.

Not all these services may be provided by the network operator, but even services provided by other parties at the edge of the network may place constraints on the network infrastructure.

Premium rate services

Premium rate services represent one of the first steps in information networking. The network operator pays a fraction of the call revenue to a service provider. This arrangement has been a great success for both

parties, and customers have benefited from new services such as televoting. Premium rate services also give pointers to the problems that may arise on a full-scale information network.

- **Control:** Teenage sons used to make calls to sex lines, and disgruntled workers would leave the office telephone off the hook after calling a premium rate number. Call barring and PIN access to 0898 were introduced to deal with these problems, giving the person paying the bill some control over the services purchased.
- **Repudiation:** Customers claim not to have made the calls on their bill.
- **Fraudulent service providers:** With a means of making free calls (for example, an unwitting employer, a cloned mobile telephone, etc.) premium rate services could be a moneymaking opportunity. Fraudsters would set themselves up as a service provider and then make calls to their own number.
- **Fraud migration:** As UK premium rate numbers have become more controlled, sex lines and lines set up by fraudsters are moving to international destinations which are less easy to control.

One important lesson is the need to see problems as a whole, rather than attacking isolated areas. The need to identify customers and avoid repudiation has also been important in mobile telephony, as outlined in the next section.

Mobility

Analogue cellular telephones

Mobile telephones have been an enormous success, changing the way many people do business. Unfortu-


nately criminals have been among the beneficiaries. They would no doubt be prepared to pay a premium for the anonymity of a mobile telephone. It turns out that they don't have to pay anything, thanks to the way in which the system operates.

Cellular mobile telephone companies operate a network of base stations. Telephones transmit their electronic identity to the nearest base station—a series of identification numbers denoted as 'X' in Figure 1. It is relatively easy to make up plausible electronic identities, and this was exploited by crooks. Just cutting off the user after the call doesn't address this problem, as they can just make up another number. The only solution is to check the validity of the number before allowing the call to take place. This fraud was especially significant in the US, necessitating additional signalling messages for inter-network roaming.

As the radio signals are broadcast, fraudsters can, in principle, record another user's identity and use it themselves. This is known as *cloning*, and has become more and more feasible as technology has improved. The equipment is now readily available at low cost. Technological advances have undermined the original security measures.

Clearly, the problem with analogue mobile telephones is that one cannot be confident that the correct telephone is calling. Additional measures such as user PINs or individual radio-frequency profiles for each telephone transmitter have been introduced, with some success. Designing secure

Figure 1—Analogue cellular telephone authentication






Figure 2—Digital cellular telephone authentication

user authentication was a fundamental requirement for the next generation of mobile telephones—the digital cellular system.

Digital cellular telephones

Digital systems such as GSM are beginning to replace analogue telephones. Figure 2 shows how improved authentication has been introduced. Once again, the telephone sends identification data. In reply, the network sends out a random piece of data which is passed to the telephone's security module. It uses its unique cryptographic key to produce a response. A call can be made only if the network receives the correct response.

An eavesdropper gets the response to only one particular challenge. It wouldn't be practical to wait for that challenge to occur again. Deducing the cryptographic key from collected challenges and responses requires a supercomputer and a lot of time. Therefore cloning has become a great deal more difficult.

Digital mobile networks can now be sure that when a call is made, it is from the legitimate customer's telephone and not a clone. However, they cannot be sure that the customer will pay the bill at the end of the month. Subscription fraud—running up large bills and then disappearing—sidesteps the well-engineered security components and focuses instead on poor customer vetting and the delays in passing bills between network operators.

The conclusion to be drawn from this example is that the security

includes procedural and organisational issues as well as technical development.

Future mobility services

Mobility will be an important feature of new services. Today some mobility products use the fixed network, such as BT Chargecard, while mobile telephones and pagers use radio transmission. Both avenues will be used in future. The fixed network has the advantage of high bandwidth. Radio networks reach users in motion and offer high flexibility. *Fixed-mobile integration* will allow seamless use of these mobility products.

Customers will see the benefits of mobility when they can log on to a multimedia terminal and know that their calls will automatically be diverted to their new location. As they walk or drive around there may be links to global positioning technology to direct them to the cinema where they have just bought tickets (over the telephone!). Of course, a knowledge of their whereabouts may well be abused—by terrorists in some contexts or journalists in others. The network operator must make it difficult to track people by means of their mobile telephone.

Other customer requirements will include confidentiality of video, voice and data. Note that some of these services may be provided by the network provider, but others may rely on the applications that run at the edges of the network. For example, current Internet commerce cannot rely on the confidentiality of the network itself, and instead relies on encrypted messages passed to and from security components in user software.

The Need for Security

The above examples illustrate the difficulties faced by the network provider in satisfying the security needs of both users and service providers while safeguarding its revenue stream. Where difficulties

have arisen in existing products, it is often because small parts of the system were considered in isolation, rather than analysing the whole process.

In the remainder of this article, the following areas will be considered:

- authentication,
- access control,
- fraud and abuse detection,
- confidentiality and privacy,
- accountability,
- integrity, and
- availability.

Encryption is an important tool in achieving several of these goals, to the extent where the ends and the means are often confused. It is therefore treated in some detail below.

Who's Calling?

Some agencies want anonymous calls. Confidential helplines are an obvious example. On other occasions, such as the electronic commerce example above, the ability to prove one's identity is useful. Telephone networks have never provided either perfect anonymity or guarantees of identity. Recently, the caller line identity passed between exchanges has been signalled to the called party¹. However, it is important to bear in mind that this only identifies the calling **line** and not the caller himself/herself. This is especially important in the electronic commerce example. 0898 chat lines proved that, when it comes to spending money, the wishes of teenage sons and their parents can be diametrically opposed. If high-value purchases are to be authorised, the vendor has to know that the order comes from the person who will pay the bill, rather than another party.

Authentication using a unique personal number of some kind is relatively cheap and effective. Just as analogue mobile telephones use identification numbers ('X' in Figure 1), people use their credit card or calling card number for authentication. The potential threat is the same in both cases: fraudsters can obtain the number and use it for their own purposes. Credit card numbers are seen by many people, while someone watching a calling card number (and PIN) being entered at a public telephone can memorise the digits and use them. This 'shoulder surfing' is common in US airports and railway stations.

Stolen numbers can be rapidly transmitted to hackers across the world, using bulletin boards or stolen voicemail accounts. As the amounts of money at risk get larger, there will be

ing in a hostile environment where users will go to great lengths to avoid payment. This emphasises the importance of good design at the outset, trying to anticipate future technological advances that might reduce security.

Access Control

Many systems aspire to be the modern-day equivalent of the medieval castle. To enter the castle, it was necessary to cross the draw-bridge and walk under the portcullis. The guard would ask visitors to prove their identity with something they knew (a password) or something they possessed (a letter with the king's seal). Sometimes biometric methods were used (the guard recognised their face). The high walls and defensive measures were designed to convince

people do once inside. In future networks, there will be a wide range of resources shared between the network operator, service providers, information providers and customers. Systems will be operated by a mixture of permanent employees and contractors. In this complex world, there will still be some rigid rules to protect customer privacy and stop other customers reading their data. However, setting read/write privileges in advance and checking them each time a data file or process is accessed is likely to become increasingly difficult. What is legitimate in one context may be rather suspicious in another. A more flexible approach, possibly involving auditing usage after the event, may be necessary.

Fraud and Abuse Detection

There is often a trade-off between security and usability. A telephone calling card that can be used from any telephone is much more usable than one that needs a special smart card reader. On the other hand, the estimated fraud losses on US calling cards of \$1 billion per year owe much to their susceptibility to 'shoulder surfing' and related activities.

Measures to increase security are not acceptable to customers. Some companies even print the PIN on the card to help the customer (and fraudsters). As the product itself cannot be changed to pre-empt fraud, some other way must be found to keep financial losses and customer inconvenience to a minimum. Fraud detection relies on analysing the patterns of use by calling cards. Several high-value international calls in a short time, or simultaneous calls from widely spaced locations are both possible indications of fraud. As the fraud detection algorithms have improved, losses have been reduced².

Reacting to fraud rather than preventing it might seem to be a transitional phase, allowing insecure products to be supported until the proper infrastructure is in place. For example, smart card readers are

Eavesdroppers should be presented with the challenge of solving hard cryptographic problems

a need to move to a more effective means of authentication than just transmitting a personal number. Eavesdroppers should be presented with the challenge of solving hard cryptographic problems to obtain the original key data, rather than reading the data in the clear. This difference is illustrated by the discussion of analogue and digital mobile telephones presented above.

In future, improved authentication will be provided either by security modules incorporated into customer equipment (presumably with a code to lock out teenagers where necessary!) or by personal smart cards. Customers need to be confident that nobody else can impersonate them before they adopt electronic commerce as a major part of their spending.

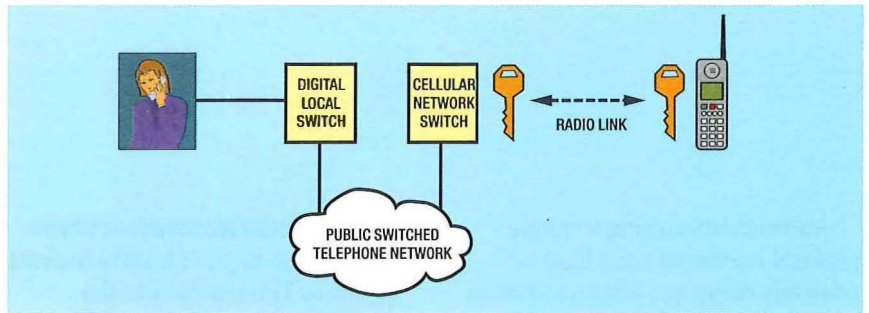
Service providers will also need confidence in the security of customer equipment. Smart cards for satellite television have proved to be operat-

invaders that they would be better off attacking another castle.

Modern-day systems tend to use passwords or some kind of convenient token such as a smart card. The goal is to make the cost of gaining access greater than the potential benefits. However, none of the current means of authentication is without its difficulties. Passwords can be written down or guessed. Tokens such as smart cards are more expensive, and can be lost. Some biometric methods are unacceptable to users; the remaining techniques cannot reliably reject impostors without also rejecting too many legitimate users. The development of effective authentication mechanisms is therefore an important study area.

Modelling computer systems as castles is not appropriate to the information network, and is probably misleading in many present-day systems too. It concentrates on the point of entry, rather than what

Figure 3—Radio link encryption in digital cellular telephony



gradually being introduced in payphones, and more secure calling cards may well be commercially feasible in the medium term. However, it is unlikely that prevention measures and increased security are the whole answer. In a competitive market, cheap, easy-to-use products with acceptable fraud levels may be what customers really want.

In any case, it may be impossible to prevent fraud completely. Many weaknesses only come to light when the product interacts with others. Checking every possible interaction between vast numbers of product features is a difficult problem. Some kind of intelligent auditing of usage data represents a sensible insurance against some new and hitherto unsuspected form of abuse.

Privacy

As well as the fraud factor outlined above, the move from analogue to digital mobile telephones has also been driven by customers' privacy needs. Some of the victims of eavesdropping have been newsworthy, to say the least, even without the sensitive nature of their conversations.

Confidentiality—allowing only authorised parties to access data—is important in many areas:

- Some operational data, such as customer calling card numbers and PINs, would be highly valuable for fraudsters.
- Industrial spies aiming to discover the secrets of competing companies might concentrate their attention on their communications links.
- Many commercial transactions in the information network will require confidentiality

As outlined below, encryption can be an effective way of protecting data, either stored on servers or in transit over networks. Highly

confidential communications may have social costs, as well as benefits. Drug traffickers, terrorists and child pornographers value confidentiality in their dealings. Law enforcement agencies will be impeded in their investigations. Tax evasion may become a problem for society. The debate on government regulation of cryptography, which balances these potential problems against individual liberty, has been vigorous. As well as questioning the desirability of control, many people question its feasibility, given the nature of global networks. From a design point of view, flexibility to respond to differing legal requirements across the world may be desirable.

It should be borne in mind that privacy is not a question of keeping everything secret. It relates to peoples' control over the information held about them. The introduction of Caller Display, where called parties are given the caller's number, illustrates the complexity of the privacy issues involved. The lesson learned from US experience was that the conflicting interests of callers and called parties necessitated a range of barring and unbarring options. Anticipating customer and regulatory demands in a similar way in future will help reduce product whole-life costs.


Encryption—A Magic Wand?

Cheap and effective encryption will revolutionise many products. For example, digital mobile telephones encrypt the speech and signalling sent over the radio link. An eavesdropper can collect and store the transmitted datastream, but without the key it is impractical to decipher the contents in a reasonable time. After decryption by the mobile network the call passes over the public switched telephone network normally (see Figure 3).

In the above example, the protection is concentrated on the radio link because it is so much easier to eavesdrop, whereas the PSTN part of the call proceeds as normal. As the connection between user and exchange is replaced by radio links (mobile and cordless) or shared optical fibres, the encryption of the this link—network level encryption—will become ubiquitous.

Encryption over individual links is not always the solution—it depends upon where the threat lies. Figure 4 shows a laptop computer being connected to an Internet server via a digital cellular telephone. Again the radio link is highly vulnerable to eavesdropping and is protected by

Figure 4—End-to-end encryption over the Internet



encryption. However, the threat to data packets transmitted over the Internet is that of copying by computers en route. There have been reports of 'sniffer' programs placed on insecure sites that look for passwords and credit card numbers passing through the site. Encrypting the data links between computers does nothing to address this problem. End-to-end encryption at the application level between server and laptop computer, as shown below, can effectively protect the data.

Some services (like pay-per-view TV) will rely on encryption to ensure that only those who have paid for the service can receive it. This 'service level' encryption operates on top of the 'network level' encryption that protects the link between user and network provider.

Theoretical Strength

In a hostile environment, where attackers brandish more computing

it is important to consider the whole system, rather than just individual components

power each year, it clearly makes sense to choose an encryption algorithm which is difficult to crack. The key length is very important, as a 'brute force' search through all the possible keys is easier with shorter keys. Even with long keys there may be weaknesses in the algorithm which allow faster attacks. Sometimes suspicions are raised that the designer of the algorithm has put in a 'trap door' for their own use, so that they can decode other people's messages more easily than is officially possible. For these reasons, secret algorithms are often less safe than published algorithms, because their weaknesses have not been probed by the world cryptographic community.

The Data Encryption Standard (DES) is a good example of an algorithm which has been subjected

to intense scrutiny. After over a decade of use in banking and other commercial activities, its 56-bit key is now becoming more vulnerable to a brute force attack, and 'triple-DES', with longer keys, is now being used. IDEA (International Data Encryption Algorithm), developed more recently, uses 128-bit keys.

Given that long keys are better than short keys, and that the computing power to perform encryption is getting cheaper all the time, why not just use long keys and strong encryption algorithms? This is a political rather than a technical question. Encryption products are subject to US export restrictions in much the same way as missiles and tanks. Exportable algorithms often have very limited key lengths. For example, although Netscape can use a 128-bit algorithm in the US, export restrictions limit the key length of the international version to only 40 bits. A French researcher found a 40-

bit key by brute force using spare university computing resources in eight days. As technology advances and short keys become more and more vulnerable, there will be commercial pressures both to ease export controls and to come up with workable solutions to the government regulation of encryption.

Real-Life Weakness

Although there is a natural tendency to concentrate on the strength of cryptographic algorithms, this strength is only occasionally a factor in how effectively the system works. In the case of digital mobile telephones, fraudsters were easily able to find another area to attack, demonstrating again that it is important to consider the whole system, rather than just individual components.

An area which is often difficult is that of key distribution. In the example above, the server and the laptop have the same key. It wouldn't make sense for this to be sent unprotected over the insecure network. It could be transmitted after being encrypted with another key, though this raises the question of how the key-encrypting-key had been distributed! These problems are not insoluble, especially with the use of public key techniques, but care is needed to avoid loopholes.

Encryption is a valuable contribution to security, but not a magic wand to be waved at difficult problems. Future products will benefit from effective cryptography, but applied after a careful and holistic analysis of the system and the threats against it.

Other Customer Needs

Accountability

It should be possible to prove that an action was carried out by a particular person. As well as the negative aspects of tracking down fraud or abuse, this is also an important facility for customers and service providers. Referring to the electronic commerce example above, the ability to 'sign' documents electronically is essential³. Cryptography provides a mechanism for achieving this.

Availability

Quality of service is a differentiator in network services. Customers will not be happy if they cannot reach a service provider. The service provider will be very unhappy. While they are unavailable, it is possible that a rival may gain business. The network operator must protect against abuse (deliberate or accidental) which reduces availability. For example, fraudsters should not be able to overload fraud detection systems by mounting a coordinated attack. Competitors should not be able to damage one another's business.

Systems should be designed to meet required levels of availability.

There is a balance between high availability and cost-effectiveness. A distributed architecture will probably require high levels of reliability only in certain subsystems, and should be able to rely on redundancy in many of the supporting systems.

Integrity

Systems should work in the way that they were intended, and customer data should not be corrupted. Network and service providers must be especially careful to ensure that their billing systems work correctly. This is another area where trust is important.

Services

The network operator may be able to sell some security services to service providers and other customers. For example, service providers may prefer to outsource fraud detection from the network operator and take advantage of the centralised hardware and established expertise.

Conclusions

This article has examined some of the priorities for security in future services. It will be necessary to transform the image of security for many in the development process. In contrast to its current perception as an isolated remedial activity, it should be brought into the development process at an early stage. An influence on high-level priorities and goals will avoid the need to resort to expensive changes and technical fixes later on.

The information network and increasing customer mobility will drive improvements in many areas. Customer privacy, fraud prevention and accountability all require good authentication as a prerequisite, but it has to be made more reliable, cost-effective and user-friendly. Electronic commerce requires several secure components (for example, billing, accountability and confidentiality) but it is not yet clear how these will be divided between the network itself

and the user applications at the edges.

There must be sufficient flexibility and forward planning to cope with a range of regulatory and legal environments across the world, especially in the areas of privacy and cryptography. The strength of security measures should take into account the pace of technological change; what is difficult to crack now may be much easier in a few years time.

Security will be a central component of the information network, enabling flexible and wide-ranging services which customers can trust. Building up this trust will be a difficult but unavoidable goal.

References

- 1 DANGERFIELD, WILLIAM; GARRETT, SIMON; and BOND, MELV. Caller Display and Call Return. *Br. Telecommun. Eng.*, October 1993, **12**, p. 176.
- 2 MESSHAM, JAMES; KELLY, PHIL; and SMITH, DEAN. Telecommunications Fraud Management. *ibid.*, April 1995, **14**, p. 58.
- 3 HICKSON, NIGEL. Digital Signatures: A Solution for Security in the Information Society, *ibid.*, Jan. 1996, **14**, p. 301.

Biography



Chris Gibbings
BT Networks and
Systems

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. In addition to work on security in future networks, he has written guidance notes on telecommunications fraud and has helped define a future fraud management architecture for BT.

Andrew Hockley

The Information Needs of Network Communities

This article examines some of the social issues associated with the construction of network communities. It argues that network communities are an important tool in maintaining social links which are currently being eroded by demographic changes. It looks at the way real societies operate, and defines some of the information that members of the society consume in order to operate effectively in that society. It identifies some of the consequences for the design of interfaces for network communities.

The World is Changing

In 1980, I wrote my first report for British Telecommunications, part of the Post Office. I wrote it on lined A4 paper, in long hand. I got it back from the typing pool three days later. In addition to my report, I received some novel respellings, various typos and a carbon copy for corrections, which I duly marked up. The postman collected my corrections and returned some new misspellings and typos a few days later. The process was repeated a few times. It was labour intensive, and gave at least three people a job. Today, I am typing this article myself and when it is finished, I will print it myself.

In the last 15 years, the world of work has changed, fundamentally and forever. The typing pool is a good illustration of two key trends: the concentration of responsibility onto an increasingly small number of individuals, and the associated disappearance of former roles for an increasingly large number. These changes in the structure of work are reinforced by, and interact with, changes in demographics, the transport infrastructure and geography, further increasing the potential fragmentation of the existing social order.

The result is a world in which there is an increased separation between individual society members, and where a threatening underclass is used as a warning to the rest. A major danger is the alienation and isolation arising from a reduction in opportunities for socialisation. We need to ensure that we take every opportunity to support the **social** effectiveness of technology. If we do not, then it will simply increase the

separation between individuals, creating a new world of sociopaths with unpredictable consequences for us all.

We used to live beside the people we worked beside, watched the same movies and shopped in the same corner shops. Increasingly, we have little in common with our neighbours. In a world where we choose customised services over a network it will be through the network that we meet like minds and groups. It is vital for the health of society as a whole that we generate means of building real network societies, filled with many different types of people, rather than simply an elitist few. These new societies must include those who have been displaced and dispossessed by this same technology.

This article examines some of the issues associated with building social experiences into network communities. It defines some of the things we need to know in order to function as individuals in a real-world community, and briefly describes some implications for the construction of interfaces for network community applications.

Why Community is Important

People are primarily social animals. We live in physical communities with social structures that have evolved over many generations and which have rules defined in rituals and etiquette. Our social context gives meaning to our actions, and provides much of the motivation for them. Isolation from a social context tends to be debilitating, resulting in loneliness, addiction, depression and a sense of purposelessness¹.

Figure 1—Work is a rich social environment

Contact with others is also the mechanism by which we learn society's rules and expectations. If we reduce social contact, then so we reduce the power of these rules to moderate the behaviour of members of the community.

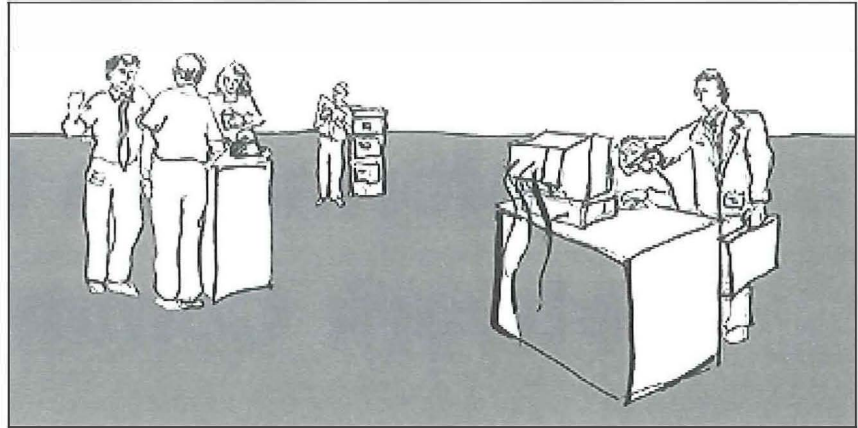
Modern communications networks offer an exciting prospect for countering and tempering the world changes that are currently undermining traditional community experiences.

Communities are Being Redefined by Technology

Network technology means that current views of communities need to be extended beyond simple physical proximity and physical concourse. Numerous electronic communities now exist, and new ones are forming. These range from simple on-line discussion groups, through Internet cybercities such as Virtual Amsterdam² and the YORB³, to sophisticated three-dimensional virtual worlds and microcommunities with spatialised sound such as the Onlive system. Some of the latter will now run over relatively low bandwidth Internet connections. This means that, potentially, people who are geographically dispersed can be part of a community. People with widely differing physical mobility can be part of a community. And the things that can occur within the community can be very different from those experienced in the real world.

What Human Characteristics Does the Technology Need to Support?

Interfaces have traditionally been designed to support the accomplishment of a specific task in a specific social context. Network communities broaden the notion of task context to a richer social experience. So what happens in real communities, and how can we use our knowledge of this when designing an interface for a network community?



Being a community member is a process

An important thing to remember is that being a member of a community is a process rather than event. Our affiliations and loyalties change over time, and the objectives and approaches used by individual members change with them. Five stages in the community membership process have been described⁴:

- *An investigation stage* when both the community and a prospective member vet a prospective relationship, in terms of the degree to which their respective needs are met. A mutually acceptable investigation leads to community **entry**.
- *A socialisation stage* when individual and community share and merge their respective goals and values. Successful socialisation leads to **acceptance** in the community.
- *A maintenance stage* when the new member explores and negotiates an appropriate role, and the community negotiates its relationship with the member. Unsuccessful maintenance leads to **divergence** between individual and community.
- *A resocialisation stage* when the community and individual negotiate any conflicts that may provoke divergence. If this does not succeed, the individual **exits** the community.
- *A remembrance stage* when the individual may reminisce, and the individual's contribution to the old community becomes part of the tradition of the community.

These stages provide some anchors for interface design by giving some basic objectives that our network community interface should support. They also indicate the importance of coping with dynamics in any network product aimed at community. Graham Walker *et al* in this *Journal*⁵ discuss some of the ways of achieving this dynamism through a combination of network technologies such as Java and Virtual Reality Modelling Language (VRML).

Basic information used in the community

In order to operate effectively in a social community, members need to understand things about the community, and their relationship to it. Some of the things members need to understand are:

- *Information about the community's overall goals and objectives.* In physical communities, this is provided formally with tools such as mission and vision statements, but more importantly through the observed actions of key members. It is particularly important that this information is available during the early stages of membership.
- *Community values.* In physical communities, these emerge subtly during contact between community members, and in the consumption of media such as television, print and radio. If we are to explicitly address the needs of a community, we need to identify rich and elegant means of representing existing community values over networks. Currently, community values must largely be experienced through the written

word in network communities, as when a newcomer monitors a topic in a newsgroup without participating. This is effective, but does not approach the richness of value sharing in the real world, where speech, text, dress, and individual success markers all contribute to a rich gestalt of values.

- *Assignment of individual roles.* The assignment of roles to users needs to be clear to the role holder, and to other community members. In physical communities, this is provided formally through task descriptions, ownership and use of titles, uniforms and so on. Network interfaces, particularly, need to support the embodiment of users as role holders, as well as simple individuals. This is an information need which lasts for the lifetime of the community membership.
- *Explication of community and other laws.* All communities have some form of legal system, and an essential component to happiness in a community is to understand what the boundaries are, and what will happen if they are crossed. Codes have arisen on the Internet; for example, 'Netiquette' guides and virtual 'executions' for social transgressions have occurred in network societies⁶. As the concept of networked community begins to take over from the concept of country and location, the legal systems of communities will become more important. Many difficult issues will need to be faced, ranging from simple questions of how best to use multimedia to represent the legal process, to how conflicts between community legal systems will be identified and controlled. Note that this information is frequently hard to access in the real world, and network community information systems could potentially deliver a better-than-real-world experience here.

Note again that all of these types of information are potentially highly dynamic, and also heterogeneous. This implies the need for the integration in real time of information from a large number of different sources into any network platform. For example, we would need to combine information from the workplace, the home, colleagues, friends etc. in order to deliver this information within a community context. This will have implications for future network architectures.

Information to foster community cohesion

It is not enough to support simple information requirements in network communities. For membership of a community to be rewarding, there must be a degree of cohesion between members of the community. Cohesion may be informally thought of as a closeness and a feeling of attachment between community members. Generally, the need for belongingness is a strong social driver. Some key factors generating cohesion between community members are:

- Communities benefit from a strong definition of the community boundary. In the real world, this is often provided by the community name, and very powerfully by physical location and boundaries such as walls and buildings. In addition, the community may invent badges of community membership such as informal dress codes, styles of behaviour and language. A powerful example of such codes is the 'mods and rockers' phenomenon of the 1960s.
- Communities are forced together by the perception of external forces, communities and threats. A strong driver of community cohesion is that members of the community see themselves as a set with a shared agenda, in a world with threats to the community. In terms of networked

community, it is sensible to ensure that the operation of communities is represented in a 'world' context, implying the need for representations of not just the owner's community, but of other defined communities.

- Communities are bound together when the value of individuals' contribution towards the success of the community is made explicit to the community. Feedback allows individual community members to recognise and hence reproduce successful behaviour and to avoid unsuccessful behaviour. The important characteristics of feedback in these circumstances are that the user's performance is reported back in terms of the overall objectives of the community, but that it is identifiable as being uniquely associated with an individual. A good example of such feedback is in BT shops. The sales of individual team members are often indicated on a daily basis in some team location—often the briefing room. Individuals can compare their performance with others, and are aware that their own contribution to the team is being indicated. Networked community multimedia can offer powerful, more timely, cheaper-to-implement versions which can be run over distributed areas. Moreover, the feedback offered can be more engaging and immediate than tallies entered at the end of a day.
- Ties between community members are strengthened by the means of a cooperative reward structure, in the sense that success of one should not automatically mean failure of another. The BT shop example above is an example of one such reward structure: it is possible for all to share in the success of the team.
- Communities need channels and mechanisms for socialisations

external to the satisfaction of the community's main objectives. A key need for the formation of strong teams is that members of the team build effective social as well as working relationships. Networked multimedia provides this possibility through such means as network team games, videotelephony, shared telepresence etc. The provision of areas that encourage incidental communication, should be encouraged. This is an area where current network communities lag far behind the real world experience, and where additional effort needs to be expended.

A key requirement for cohesion is the representation of other members of the society concerned. The BT Laboratories' Portal project, described in this issue of the *Journal*⁵, has implemented one approach to representing other users in the world. Here literal visual embodiments of other users of the space are used to give a feeling of community, and of different interest groups and perspectives.

Information and privacy

The previous sections have focused on communication and information needs. But a successful community in the real world also needs to support personal space, and the ability to withdraw from the community. An important part of any networked multimedia application is that it supports the establishment of an appropriate level of privacy for its users. In other words, it must support intentional retreat from social contact, and measured retreat from contact, rather than disconnecting from the network. In the real world, different types of privacy exist⁷:

- *Solitude*—the degree to which one is physically separate from others. This would correspond to being disconnected from the networked community.

- *Intimacy*—the ability to shield communications between people from others. In a network community, this would correspond to the use of encryption devices.
- *Anonymity*—the ability to hide one's role and identity. In a network community, this would correspond for example to being able to gather information and observe without this necessarily being recorded for future use. The potential for anonymity is important to avoid junk mailings and other 'power communications' which are unwanted.

In the real world, individuals regulate their privacy level along these dimensions using such real-world devices as:

- making a change in the physical environment (for example, open/close a door/barrier/channel);
- performing some social action (for example, an explicit invitation to visit, talk etc., or conversely asking someone to leave);
- moving to a place with associations of privacy or openness (for example, to a personal place or to a public space)—the messages conveyed by presence in a dance hall are very different to those conveyed by presence in one's lounge, for example; and
- manipulate the social communication channel, (for example, moving out of eye contact, moving closer to the other person).

Network community applications need to support equivalent privacy actions.

Information and territory

Humans are territorial. They have areas of their lives over which they need to feel personal control and ownership. In the case of network communities, territory can be

thought of as delineating zones of operation, as well as physical representations of space, interpersonal relationships, and community membership. Territory is important from both practical and motivational perspectives: the need to display and exert territorial right is a key biological driver in its own right. Any doubts may be dispelled by 'borrowing' a colleague's computer without permission and leaving some alien files as a marker.

Four types of territory have been defined as follows⁷:

- *Primary territory*, owned solely by an individual or community, in which any intrusion will invoke territorial behaviour (for example, a person's home, or desk space). One way to support this in an interface is to allow users to personalise their own network community space, to give them a sense of ownership over it.
- *Secondary territory*, shared by many communities or individuals. An example might be a virtual shopping mall concourse or a news group.
- *Public territories*, shared by everyone. Effectively, these are the rest of the world, and behaviour in these territories is regulated by prevailing social norms, rather than local territory norms.
- *Interaction territories*, temporary territories established specifically for the purpose of protecting a communication space. In the context of a face-to-face discussion they may be thought of as a virtual glass bubble around a conversation. The bubble is loosely bounded by the participant's notion of earshot (or other communication limitation) within the world. The need to support interaction territory will be particularly important in network communities, where participants will need to see a representation

of the 'earshot' range of their communication, and where observers will need a representation of the territory from the outside. In the context of a network community, interaction space is equivalent at its simplest to an e-mail creation space, and at its most complex a space in which multimedia communication can occur.

Work on the Portal project, described in this *Journal*⁵, gives some insights into how a physical representation of the environment can be used to indicate territory in virtual communities. Here separate physical islands are used to partition the virtual world into a set of sub territories or contexts.

Information and community rituals

One of the distinguishing characteristics of communities is the presence of ritualised behaviour; that is, stereotyped ways of doing things. Rituals make explicit many of the society's expectations, and illustrate some of its underlying beliefs. Ritualised behaviour is frequently used to mark transitions in the way that individuals relate to the community, and vice versa (for example, a marriage ceremony). Rituals are key components of all communities, and key social drivers, and should therefore be supported in networked communities. Some characteristics of rituals are:

- they have an audience,
- there is a defined and agreed protocol to the ritual,
- there may be a degree of audience participation, and
- failure to follow the agreed protocols leads to social approbation, which reinforces the use of the protocol.

One can imagine that a combination of simple procedure processing

and the dynamic configuration of an appropriate communications infrastructure could be very successful in supporting these basic requirements. As an example of ritual in network communities, the first cyberwedding was recently reported in the press.

Information about success

Humans need to succeed at some level in their life to be happy. The definition of success is clearly relative, but in terms of social drivers may be thought about in terms of their individual contribution towards:

- *Success of the community itself.* That is, success of the unit/tribe/organisation, independent of the individual success of its members. Thus a workforce may feel a bond with a company, and derive satisfaction from seeing the company do well against its rivals. Key to supporting this individual driver in networked multimedia is the publication and explication of basic information about the community's success.
- *Success of other members of the community.* There is much debate about the existence and origins of altruistic behaviour; that is, the tendency of community members to help other members of the same community, even at expense to themselves. For some members of a community at least, expression of altruistic behaviour is a key driver, and one that we should support. This relies on the existence of mechanisms to indicate the success or otherwise of other community members.
- *Success of self.* The key driver for many people is the respect of people whom they themselves respect. This may manifest itself in many ways, depending on the socialisation involved in the community. Public respect requires that badges of success be made apparent to other members of the community, and that

appropriate communication channels be made available within which they can be recognised. Such badges may range from rank within an organisation, wealth, car size, dress sense etc. Supporting such badges of success will be a key challenge for network multimedia.

A simple and powerful way to deliver information about success is through graphic visual representations of other members of the society. These virtual representations are usually called *Avatars*. BT is currently developing the representations used in Portal People to deliver much more information to other community members, including status and activity. The challenge is to extend the view away from literal real-world representations but to keep the representation intuitively understandable.

Conclusions

Network communities are thriving all around us, and hold the promise of providing some of the social glue that is currently being washed away by demographic and employment trends. Network communities can provide islands of structure in an increasingly unstructured world, and hence provide some stability and meaning to people who are losing that stability and meaning as traditional social links are taken away from them.

But in order to be genuinely successful, they need to deliver support to some of the basic human drivers which are expressed through membership of a community. In the real world, people use a vast range of information. Virtual worlds need to represent this information to their members to provide the engagement and structure that real communities provide. The key lessons for the design of network communities are:

- In order to support the process of community, information should be continuously available to members

of the community. Preferably, it should be available to them in ways that do not require them to search for it and request it specifically. Thus they can build their knowledge of their community and its members implicitly and by a process of gradual accretion. This implies the need for network architectures which can combine, in real time, sources of information from many disparate sources.

- Communities should be presented to members in such a way that their community can be seen to exist in a world of other communities. In other words, the community itself should be given a context and a relationship to other communities. This will strengthen and support bonding within a given community. Representing the community as a physical space enables this to be done.
- Communities should show success and failure. Members must be able to see who is doing well, and why, in a continuous fashion. This lets members build role models. This means we need to develop representations of success and failure, and let community members see what other community members are doing. Embodiments of members of the community are a promising way to achieve this.
- Communities should support the concept of retreat from other members of the community, while still preserving the notion of belonging to it. The experience of community is partially defined by being able to be absent from it. In other words, the user's personal workspace should be presented in the context of the total community representation. The workspace is simply part of the community space.

These are some of the issues that are being examined in research being

carried out by BT. The objectives are to develop embodiments of users, representations of territory and space, and seamless communications between community members within the space. A range of specific designs are currently being prototyped. To arrange a demonstration of work in progress, e-mail Andrew McGrath at amcgrath@hfnet.bt.co.uk.

Acknowledgements

Illustration by Amanda Oldroyd. Many of the ideas in this paper have been contributed by Andrew McGrath, Andrew Hockley, Phil Clarke, Matt Shipley, Rob Taylor Hendry, Paul Rea, and Dave Burraston.

References

- 1 RUBENSTEIN, C., and SHAVER, P. The Experience of Loneliness. In PEPLAU, L. A., and PERLMAN, D. Loneliness: A Sourcebook of Current Theory, Research and Therapy. New York, Wiley, 1982.
- 2 Virtual Amsterdam, <http://www.dds.nl/>
- 3 The YORB show. Broadcast weekly on New York Cable.
- 4 DEAUX, K., and WRIGHTSMAN, L. Social Psychology in the 80s. Brooks/Cole, Monterey, CA, 1984.
- 5 WALKER, GRAHAM; MORPHETT, JASON; FAUTH, MARCO; and REA, PAUL. Interactive Visualisation and Virtual Environments on the Internet. *Br. Telecommun. Eng.*, April 1996, **15** (this issue).
- 6 MACKINNON, R. The Social Construction of Rape in Virtual Reality. In SUDWEEKS, F. *et al.* Network and Netplay: Virtual Groups on the Internet. MIT Press, 1996.
- 7 HARRE, R. Social Being. Blackwell, Oxford, 1993.

Biography



Andrew Hockley
BT Networks and
Systems

Andrew Hockley joined BT's Human Factors Unit in 1980, having studied psychology at Sussex University, and human factors at Birmingham University. He initially worked on adaptive interfaces, intelligent tutoring systems, and computer supported cooperative working. After seven years of research work, he moved to the Human Factors Products and Services team, where he worked on the development of interfaces for new network services based on CLASS and ADSI. After five years working on network services, he moved to work on network multimedia products. He runs the Human Factors 'Future Market Opportunities', building demonstrators of needs-driven network products.

Graham Walker, Jason Morphet, Marco Fauth and Paul Rea

Interactive Visualisation and Virtual Environments on the Internet

We are on the threshold of a step-change in Internet services, with implications perhaps greater than those arising from the introduction of the WWW graphical interface. Current developments, such as VRML and Java, promise a future of interactive, multiparticipant, multimedia environments and applications. This article outlines the background to these developments and describes early demonstrators which provide a glimpse of future services.

Introduction

Three-dimensional interactive graphics have the potential to liberate the human-computer interface, and provide intuitive access to data landscapes and immersive applications—an appealing vision, which remains largely unfulfilled. The absence of widely accepted, cross-platform standards for distributed virtual environments and interactive visualisation has restricted commercial applications to niche markets and specialist communities. Now, emerging industry standards such as the Virtual Reality Modelling Language (VRML) and the Java programming language, promise to deliver the vision. These advances are supporting World Wide Web (WWW) applications with universal accessibility, and near-term developments will enable a wide range of Internet service offerings.

This article outlines the background to the emergence of VRML and Java in the context of current Internet development, and describes demonstrators which highlight the early capabilities of these standards: *Portal* is a VRML interface to a range of projects at BT Laboratories, while *Jacaranda* is an interactive Java visualisation. Both demonstrators can be accessed on the BT Laboratories WWW site^{1,2}. The article also considers the potential for future services, and concludes that support for interactive multiparticipant multimedia environments and applications is an inevitable short-term development of the current standards.

Background

Bell et al.³ have identified three stages in the historical development of the Internet, which are summarised in Figure 1. The first stage, *storage*, was the evolution of the transmission control protocol/Internet protocol (TCP/IP) network infrastructure which provides a layer of abstraction between data and physical machines. Stage two, *retrieval*, was the development of the World Wide Web hypermedia system, built on the universal resource locator (URL) addressing scheme and the HyperText Markup Language (HTML) document standard. This made the distributed resources of the Internet more widely accessible, and prompted rapid and sustained growth in network usage. However, the overwhelming majority of material on the WWW remained passive and two-dimensional, comprising, in effect, a vast multimedia database able to send out pictures and text in response to requests.

Figure 1—Stages in Internet development (after reference 3)

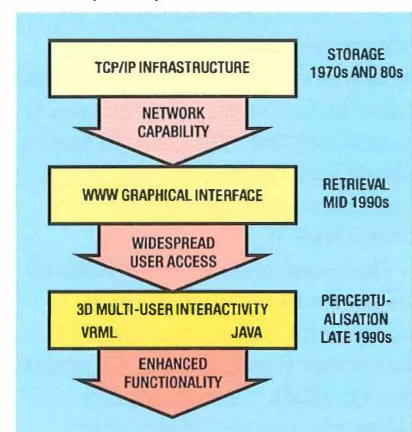


Figure 2—Extract from a VRML 1.0 scene description

This article is concerned with the third stage, *perceptualisation*. We envisage an interface to the Internet which is increasingly built around human interaction skills, and which delivers much more than remote data access. Perceptualisation is the domain of the title of this article: interactive visualisation and virtual environments. Many potential services require greater interface functionality, including high-quality three-dimensional graphics and audio, complex user interaction (possibly tactile), and multi-participant capability. Developments such as VRML and Java are starting to offer the capabilities required for this next stage, and they will provide the building blocks for a range of compelling applications and services.

VRML

VRML was conceived in the spring of 1994 during a special interest group meeting at the first annual World Wide Web Conference in Geneva, Switzerland⁴. Several attendees described projects already underway to build three-dimensional graphical visualisation tools which interoperate with the Web. It was agreed that these tools should have a common language for specifying three-dimensional scene descriptions and WWW hyperlinks—an analogue for virtual reality of the HTML standard. The term *Virtual Reality Markup Language* (VRML) was coined, and the group resolved to begin work on a specification after the conference. The word *Markup* was later changed to *Modelling* to reflect the graphical nature of VRML. Within six months, a draft standard was produced for review, and a few months later the first version was released. It is intended that VRML should become the de facto language for interactive multiparticipant simulation on the World Wide Web.

Released in June 1995, VRML 1.0 allows for the creation of virtual worlds with limited interactive behaviour. Figure 2 shows a segment from a VRML scene description. Although strongly influenced by

```
#VRML V1.0 ascii
# a red cube translated 1 unit along the X axis with a link to a 'test-cube.html' file
Separator {
  WWW Anchor {
    name "http://iron.bt-sys.bt.co.uk:8080/jmorph/test-cube.html"
    DEF redCube Separator {
      Transform {
        translation      1 0 0
      }
      Material {
        diffuseColor     1 0 0
      }
      Cube {
        width            2
        height           2
        depth            2
      }
    }
  }
}
```

Silicon Graphics, and in particular the Open Inventor file format, VRML was designed from the outset as an open standard, with key requirements of platform independence, extensibility and acceptable performance over low-bandwidth connections. VRML worlds can contain objects which have hyperlinks to other worlds, HTML documents or other valid multimedia file types. When the user selects an object with a hyperlink, the appropriate viewer is launched; for example, Netscape. Similarly, when the user selects a link to a VRML document from within a correctly configured WWW browser, a VRML viewer is launched; for example, WebSpace. VRML viewers are therefore a perfect complement to standard WWW browsers for navigating and visualising the diversity of data on the Web.

Just two years on from the initial discussion in Geneva, there is now a plethora of application domains using VRML on the Internet. The following list (taken from the VRML Repository⁵) summarises some of the categories for which VRML has provided a platform independent method of transferring three-dimensional environments:

- architecture
- art
- astronomy
- biomedical sciences
- chemistry
- commercial applications
- computer sciences
- education
- entertainment
- environmental science
- history

- home spaces
- maps and globes
- mathematics
- music
- physics

Discussions on VRML 2.0 are well advanced, with significant contributors including Silicon Graphics, Sony, Apple and Microsoft in addition to a range of universities and individuals⁶. An agreed revision to the standard is likely before the end of 1996, and proposals are focused on extensions to enable richer graphical worlds, interaction, animation and audio. Other enhancements could include explicit support for motion physics and real-time multi-user interaction.

Although VRML browsers are available for a range of hardware platforms, a relatively high-end PC or workstation is required for acceptable interactive performance. VRML models and applications are therefore not currently accessible to the majority of WWW users. Short-term developments in both hardware and software are certain to rectify this situation, and the standard will become increasingly widespread.

Java

Java is a programming language specifically designed for distributed computing environments, and is a second key enabler for increased interest and interactivity of WWW applications^{7,8,9}. Originally developed several years ago by Sun Microsystems as a control language for consumer electronics, Java was relaunched in early 1995. Figure 3 shows a segment of Java code. It is a derivative of today's standard object-

```
import java.applet.Applet
import java.awt.Graphics;

public class HelloWorldApplet extends Applet {

    public void paint(Graphics g) {
        g.drawString("Hello World!", 50, 25);
    } // end of method paint

} // end of class HelloWorldApplet
```

Figure 3—Extract of Java code

oriented language C++, and includes a number of features which make it ideally suited to distributed Internet applications. These characteristics include:

- *simple*, especially in support for the Internet protocols required to access remote network resources. A corollary of simplicity is that Java programs, or applets, can be small and relatively easy to develop.
- *secure*, which is critical for a widely distributed and largely uncontrolled environment such as the Internet. Once downloaded across a network, Java applets are dynamically interpreted within a secure environment with no access to local programs or data.
- *architecture neutral*, enabling the same applets to run on any machine to which the Java interpreter has been ported. This is another essential attribute for a diverse environment such as the Internet, with a range of hardware platforms.

These and other features are not exclusive to Java, and competitors include VBScript¹⁰ and ActiveX¹¹ from Microsoft, and General Magic's Telescript¹². However, Java was relaunched with perfect timing to meet market demand for a programming language for the Internet, and has already achieved widespread commercial acceptance. In the past year, Java has been acclaimed as the language of 'network-centric' computing. A range of Java applications, including spreadsheets and word processors can be summoned over the network as required, and companies such as Sun and Oracle are promising low-cost terminals or 'Internet appliances' optimised for the task. These developments raise important



Figure 4—Portal interface

issues for the future shape of the entire computer industry, but this article is particularly concerned with applications to distributed visualisation and virtual environments. In this context, it is probable that VRML 2.0 will support Java as a preferred scripting language, thereby complementing VRML graphics with the power and flexibility of Java programming.

Netscape, the most popular WWW browser already supports Java applets, highlighting the remarkable pace of development in Internet products and standards. Moreover, several comprehensive commercial Java development environments are available, with programmer support increasing rapidly. The interactivity of Java is therefore already familiar to many Internet users, and the range of applications is growing daily.

Application Demonstrators

This article has introduced the key features and current status of VRML and Java, suggesting that they provide the functionality required for new interfaces and services which will result in perceptualisation over the Internet. This section describes demonstrators which highlight early capabilities of the standards and point to future service opportunities.

Portal

Portal provides an introduction to some of the work and facilities at BT

Laboratories, and is an example of an Internet virtual environment. It is accessible from the BT Laboratories WWW server¹, and is both an evolving and an involving interface to projects and other visitors (Figure 4). As described in this article, it has been implemented using the VRML 1.0 standard.

Portal is prefaced by a three-dimensional model of BT Laboratories (Figure 10), which illustrates the ability to present physical environments using VRML and which points to future applications in planning and design. The user is then invited to select a personality icon, which provides the user with a choice of persona within Portal:

- business,
- human factors,
- marketing,
- education, or
- technology.

Once a persona is adopted, all subsequent information can be tailored to the individual. For example, a user who selects the human factors personality icon might be informed of the psychological, sociological and ergonomic aspects of the projects they visited.

The projects are displayed as islands (Figure 5) and are currently:

- *The London Model*—an area of Central London, with potential radio planning applications;

Figure 5—Portal project islands

- *Fly the Network*—a three-dimensional network management interface;
- *The Electronic Agora*—a user-centred video-conferencing system;
- *CamNet*—a mobile ISDN telepresence system;
- *VISA*—an interface to home services; and
- *Workspace 2000*—a physical desktop of the future.

From knowledge of their personality icon, individuals are represented by automatically placing a colour-coded embodiment of them on the project islands they visit—*Portal people*. This aids in focusing information for subsequent visitors. They will see a presence on the islands, which indicates both the popularity of an island and the personality of other visitors, from the number and colour of the embodiments, respectively. In addition to colour and placement, transparency is used as a means of temporal presence, so that embodiments 'fade away' over time. This use of embodiments provides a sense of sharing and association within the space of Portal, in contrast with the normal sterile isolation of WWW information. For the Portal visitor, there is a sense of association which extends beyond simple multimedia data retrieval.



This feature is achieved through combining Perl scripts with the VRML language, an extremely powerful combination of Internet resources. Perl is a language designed for manipulation of text, files and processes, and it has been used extensively to manage and build intelligence into WWW sites. It has been suggested that Perl could act as the behaviour language for VRML¹³. However, using a script in this manner does not provide truly interactive behaviour, since the VRML environment remains static once loaded into the browser. While this latency is acceptable in applications such as Portal people, future applications will inevitably demand full interactivity and synchronous communication. This could be provided, as already noted, through the ability of Java to manipulate the content of VRML scenes.

Within the Portal islands, the user can fly to pre-set viewpoints, using a 'Viewpoints' menu option within the browser. This provides a macro level of navigation, enabling the world developer to assign cameras to strategically important or interesting

views of the world. Once a viewpoint is selected, the user is flown to that position and is then free to navigate the world at a micro level, using standard VRML browser controls. By maintaining a sense of direction and travel, flying to information reduces the problem of getting lost in jumps between hyperlinked data. Current browsers also support limited interaction within the VRML scene: when the cursor passes over a hyperlink, the object is highlighted, and if the mouse button is clicked, the associated URL is retrieved.

In the current version of Portal, the links from the VRML worlds are to HTML pages relevant to the selected object. The pages give a textual outline of the projects, and are, in turn, hyperlinked to further three-dimensional VRML models (Figures 6 and 7). It has also been proposed that such links could provide the basis of asynchronous communications (such as e-mail) within the Portal virtual environment¹⁴.

Currently, six projects are within the Portal environment. More projects will be introduced, and there are plans to automate aspects of this

Figure 6—London model—a 3D environment for line-of-sight radio planning





Figure 7—Workspace 2000—flexible support for remote collaborative work



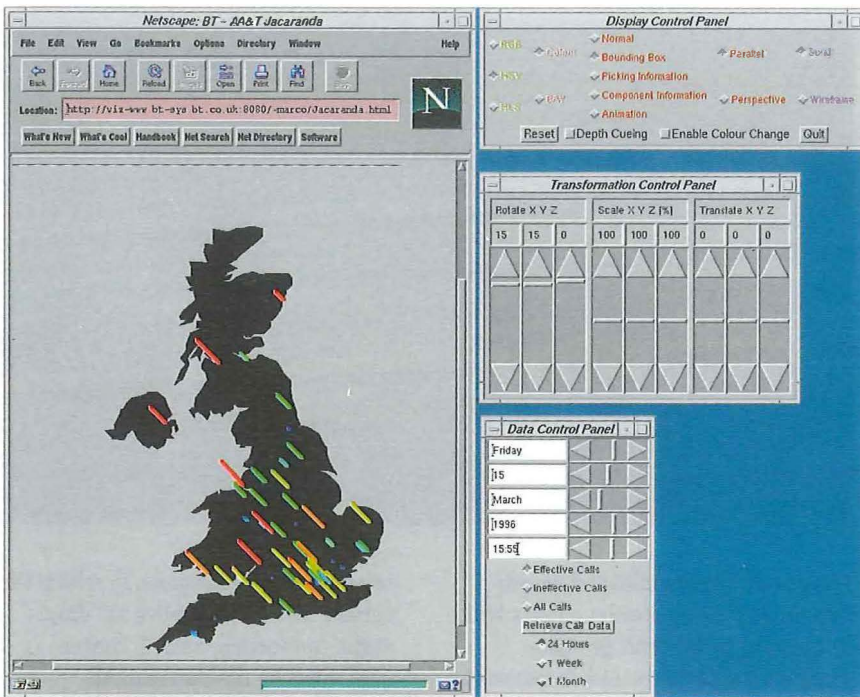


Figure 8—Jacaranda interface—overview

process. For example, the size and spatial arrangement of project islands could be automated to reflect popularity and relevance to the personality icon, and the date of creation might be mapped to transparency.

In the first six months following the release of Portal in July 1995, almost 100 000 access requests were registered. Allowing for the limited availability of VRML browsers, such statistics emphasise the power of the Internet in reaching a wide audience. The wide diversity of the user community is also confirmed, with requests ranging from Brighton to Bolivia.

Jacaranda

Jacaranda (*J*AVA *C*ALL *R*eporting *A*ND *A*nalysis) illustrates the potential for Java to provide interactivity within Internet-based visualisation applications². The Jacaranda demonstrator conveys a vision of future developments in on-line service reporting and management, in which Java is used to deliver the results of a database query together with an interactive visualisation application. The end-user is able to explore the underlying data, and to interact with a visual display of telephone call record statistics.

Jacaranda is an interactive visualisation of telephone calls made from within the UK to a single

enquiry point; for example, a product support-line or customer enquiry number. The geographical source of calls is represented by towers located at the main centres of population. The height and colour of a tower is mapped to the number of calls originating from that town in a single hour (Figure 8).

The user can select the date and time of interest, and initiate animation sequences which might reveal interesting patterns and trends. The ability to link from selected towers to supplementary information is demonstrated, and it is also possible to rotate and scale the three-dimensional visual display. The application comprises four primary windows (Figure 8): the main visualisation, the display control panel, the data control panel, and the transformation control panel.

The display control panel is used to select the overall mode, while the transformation control panel allows the user to scale, rotate and translate the visual output. The data control panel selects the call records to be displayed and controls the animation function. A fifth window, the 'pick information' panel provides more detailed data at a selected location (Figure 9). An extended description of the application can be found on the WWW pages².

The functionality of Jacaranda is almost identical to a visualisation concept demonstrator which was implemented over two years ago



Figure 9—Jacaranda interface—'pick information' panel

using a commercial software development environment for stand-alone visualisation applications¹⁶. However, the ability to deliver this functionality over a network to any terminal with a Java-enabled browser offers important practical advantages, including:

- wide base of installed and configured WWW clients—the application developer need not be concerned with either the hardware or software of the user environment, and hence incurs greatly reduced client set-up and support costs;
- platform independence eliminates problems and costs of porting—there is no need to write and maintain separate versions for Macintosh, Unix, NT, Windows 95 etc.; and
- single current version on the central server—reduced costs of version control and distribution, and immediate, universal upgrade and bug-fix capability.

Jacaranda was first developed in the spring of 1995 using the alpha release of Java at a time when there was minimal developer support. It was therefore necessary to code all the three-dimensional graphics and interface classes, and the final application was only compatible with a limited range of WWW browsers. Jacaranda has recently been updated to Java 1.0, making it more widely accessible. The figures in this article are taken from this latest implementation.

Future Services

Our initial concept demonstrators, Portal and Jacaranda, are only scraping the surface of the service functionality that will shortly be available as standards and support tools continue to develop. This section

Figure 10—BT Laboratories site model

outlines several potential applications in interactive visualisation and virtual environments, and includes brief consideration of Internet audio developments. The applications are grouped into categories requiring increasing functionality.

Three-dimensional models


The current VRML standard permits the construction of three-dimensional models, and provides pre-set viewpoints in addition to user fly-through. It is suggested that standard libraries might be used to distribute textures and other bandwidth hungry details which help to provide photo-realistic scenes. Specific services might include:

- on-line product catalogues or virtual shopping mall;
- architectural walk-throughs (Figure 10), and public review and comment on proposed urban developments; and
- inclusion of three-dimensional 'figures' within electronic publications and on-line manuals.

Modelling and user interaction

The inclusion of support for interpreted scripts is a certain near-term development in the VRML standard. This will enable the functionality and interactivity of Java applets (and, potentially, other languages) to be included within VRML worlds. No

Figure 11—Visualisation of neural network channel assignment algorithm¹⁵



longer will the downloaded worlds remain static and sterile; objects will be brought to life with pre-programmed behaviours and responses, and independent applications will have their own existence within the world. The user will be able to interact with menus and control panels, just as they would in a conventional computer interface.

Potential services include:

- educational experiments and simulations,
- interactive data visualisations and modelling (Figure 11),
- remote control and navigation interfaces, and
- interactive adventure games

The introduction of interactivity and behaviours raises new security issues, since code in addition to the browser is now running within the local environment. A full discussion is beyond the scope of this article, but the potential risks are highlighted by a security loophole in the Java-enabled version of Netscape 2.0¹⁶.

Multi-user environments

A further extension to the functionality of servers and browsers would enable multi-user environments¹⁷. Such applications are particularly attractive to service and network providers, in that participants contribute to the interest and 'content' of the world, thereby reducing the demand for regular updates, revisions, and costly new material. Parallels can be drawn with

telephone conversations, in which the service provider supplies an 'empty' audio connection, and all content is generated by the participants.

A multi-user interactive space could support services such as:

- multiparticipant games,
- participatory review of product or building designs,
- a virtual office and communications interface for a distributed project team, and
- a general-purpose environment for computer supported co-operative working (CSCW) applications (Figures 12 and 13).

In addition to increased security concerns, multi-user applications must also address issues of scalability, with respect to both server performance and network traffic^{18, 19, 20}. Further discussion is again beyond the scope of this article, but there is a range of technical challenges in developing a system architecture to support the required functionality, particularly in areas such as service creation and billing which have received minimal attention in early demonstrators.

Audio

So far this article has concentrated on extending the functionality of the visual interface. However, the critical contribution of audio must also be considered, and, in particular, the role of spatial audio in enhancing user experience of virtual environments. Recent Internet developments include RealAudio which streams



Figure 12—Physical CSCW environment




Figure 13—Virtual CSCW environment¹⁹

audio with real-time decompression²¹, and Internet Phone which provides a full duplex audio link, albeit with unacceptable delay for most applications²². Future developments will support spatial audio within distributed applications. This will include sound as an aspect of object behaviour, discrete audio events, and user input in the form of speech. Such functionality will greatly enhance applications such as games or shopping, and be an essential component of CSCW environments.

Discussion

Current developments in Internet standards will promote a transformation in applications. Jacaranda and Portal provide only an early indication of the diversity of services that will be supported by future server and browser functionality. Moreover, with standards will come the critical mass required to grow major new markets out of current trials and niche applications. These developments will lead to currently unimagined services, with far-reaching implications for the Internet, at least equal to the changes resulting from the development of the WWW. Potential for a profusion of new services notwithstanding, such developments also raise important commercial and network-related issues. This section briefly discusses some of the key considerations, with particular attention to the role of the network provider.

While near-term technical developments can be predicted with a degree of certainty, commercial

Internet developments are largely unpredictable even on a timescale of months. Anticipating the successful services or 'killer applications' entails a significant element of speculation, and, although there is some scope for major companies to influence developments, the essentially open and uncontrolled nature of the Internet will ensure a continuing healthy diet of the unexpected. The market is already a complex tangle of alliances, joint ventures, established computer and telecommunications companies, and promising start-ups.

Although an unpredictable wealth of services will come and go in this unstable market, network providers will nevertheless need to provide guarantees on quality of service, facilities for security and billing, and perhaps a common infrastructure of network-based storage and processing capability. In parallel with the emerging content and application standards such as VRML and Java, new network protocols are being developed to meet the traffic requirements of selected services. Examples include RSVP (ReSerVation set-up Protocol) which allows Internet applications to obtain a specified quality of service, typically by reserving resources along the data path²³, and real-time transport protocol (RTP) which provides support for applications with real-time properties, including timing reconstruction and loss detection²⁴.

However, the distributed approach to management and ownership of the Internet will inevitably give rise to critical issues which are essentially

'everybody's problem', and yet are simultaneously 'nobody's problem' when it comes to prescribing a solution. Examples include service-impairing congestion, and blocking or filtering of pornographic material. Such considerations will create a role for Internet service providers to maintain islands of high-quality access, which would include an editorial function to filter and catalogue the wealth of available information. They will also provide scope for the same cost-effective IP network technology to be used within the managed environment of a closed user group on a private network: so called *intranets*.

One secure prediction for the future Internet services is that market instability and commercial uncertainty will endure, and the pace of change will be sustained or even accelerate. Committing to traditional type is a hazardous undertaking, as evidenced by the overwhelming majority of references to this article, which are only available in electronic form. This simple observation is an early example of the shifts in business and working practices which have been provoked by early Internet applications, but once again is only a taste of the changes yet to come.

Conclusions

A key reason for the success of the Internet is that the protocols and standards are as generic and low-level as is practicable. These sturdy building blocks provide powerful enablers to the wider community of application developers.

Returning to the stages in Internet development outlined in Figure 1, it was observed that stage two, retrieval, was built on URL addressing and the HTML document standard. This is a very basic level of functionality, and yet has spawned an abundance of applications and universal interest in Internet services.

The increase in functionality offered by developments such as VRML and Java will lead to the third stage, perceptualisation, transforming the interface and enabling a wealth of new services, far in advance of current WWW browsers and existing passive Internet databases. Early demonstrators such as Portal and Jacaranda provide a glimpse of future services, which will include interactive visualisation and multi-participant virtual environments.

Acknowledgements

The authors are grateful to the many researchers at BT Laboratories who contributed to the development of Portal and Jacaranda, and in particular to Adrian Sharples, Andrew McGrath, Dave Burraston, Rob Taylor-Hendry and Martin Crossley.

References

- 1 BT Laboratories Portal. <http://virtualbusiness.labs.bt.com/vrml/portal/home>
- 2 BT Laboratories Jacaranda. <http://www.labs.bt.com/innovate/informat/jacaranda>
- 3 BELL, G.; PARISI, A.; and PESCE, M. The Virtual Reality Modelling Language: version 1.0 specification, 1995. <http://vrml.wired.com/vrml.tech/vrml10-3.html>
- 4 RAGGETT, D. Extending WWW to support platform independent virtual reality. <http://vrml.wired.com/concepts/raggett.html>
- 5 VRML repository. <http://www.sdsc.edu/vrml>
- 6 MITRA; HONDA Y.; MATSUDA, K.; BELL, G.; and MARRIN, C. An Overview of the Moving Worlds Proposal for VRML 2.0, 1995. <http://webpage.sgi.com/moving-worlds/over3/overview.main.html>
- 7 GOSLING, J.; and MCGILTON, H. The JAVA language environment: a white paper, 1995. http://java.sun.com/whitePaper/javawhitepaper_1.html
- 8 Sun microsystems Java WWW pages. <http://java.sun.com/about.html>
- 9 *JavaWorld*, on-line Java magazine. <http://www.javaworld.com>
- 10 Microsoft VBScript (Visual Basic scripting language). <http://www.microsoft.com/vbscript/vbsmain.htm>
- 11 Microsoft ActiveX (Activate the Internet development kit). <http://www.microsoft.com/corpinfo/press/1996/mar96/activexpr.htm>
- 12 General Magic Telescript. <http://www.genmagic.com/Telescript/>
- 13 MORELAND, J. L.; and NADEAU, D. R. The Virtual Reality Behaviour System (VRBS): Using Perl as a behaviour language for VRML. Proc. 1995 VRML Symposium, San Diego, Nov. 1995.
- 14 MORPHETT, J. Presence in Absence: Asynchronous communication in virtual environments, proc. 3D and Multimedia on the Internet, WWW and Networks, Bradford, UK, April 1996.
- 15 WALKER G. R. Challenges in Information Visualisation. *Br. Telecommun. Eng.*, April 1995, 14, p. 17 and <http://www.labs.bt.com/innovate/informat/infovis>
- 16 Potential vulnerability in Netscape products. Netscape Standards Documentation. http://home.netscape.com/newsref/std/java_security.html
- 17 HONDA, Y.; MATSUDA, K.; REKIMOTO, J.; and LEA, R. Virtual Society: extending the WWW to support a multi-user interactive shared 3D environment. Proc. 1995 VRML Symposium, San Diego, Nov. 1995.
- 18 Sony Virtual Society project. <http://vs.sony.co.jp/VS-E/vstop.html>
- 19 GREENHALGH, C. MASSIVE teleconferencing system. <http://www.crg.cs.nott.ac.uk/~cmg/massive.html>
- 20 Swedish Institute of Computer Science. Distributed Interactive Virtual Environment project. <http://www.sics.se/dive>
- 21 RealAudio. <http://www.realaudio.com/>
- 22 Internet Phone. <http://www.vocaltec.com/iphone.htm>
- 23 RSVP overview. <http://www.isi.edu/div7/rsvp/rsvp.html>
- 24 RTP overview. <http://www.fokus.gmd.de/step/employees/employees/hgs/rtp/faq1.html>

Glossary

As with any technical discipline, the Internet has developed an extensive vocabulary of acronyms, familiar and convenient to those working in the area, but superficially intimidating to newcomers. It would have been impractical to avoid these terms in this article, but the following glossary (based on <http://www.cwru.edu/help/webglossary.html>) is offered in mitigation:

Applet A small Java application embedded within a WWW page, and

requiring another Java program (such as a browser) in order to run.

Browser Software that allows you to navigate information databases; examples are Netscape Navigator and NCSA Mosaic.

HTML HyperText Markup Language; Used to tag various parts of a WWW document so browser software will know how to display that document's links, text, graphics and attached media.

IP Internet protocol. A set of standards that control communications on the Internet. An IP address is the number assigned to any Internet-connected computer.

Java A programming language specifically designed by Sun microsystems for distributed computing environments.

Perl A programming language designed for manipulation of text, files and processes.

Protocol A set of standards that define how traffic and communications are handled by a computer or network router.

TCP/IP Transmission control protocol/Internet protocol. The basic protocol controlling applications on the Internet.

URL Uniform resource locator. The addressing system used in the WWW and other Internet resources. The URL contains information about the method of access, the server to be accessed and the path of any file to be accessed.

VRML Virtual Reality Modelling Language. Intended to support interactive multiparticipant simulation on the World Wide Web.

WWW World Wide Web. A hypertext-based Internet service used for browsing Internet resources.

Biographies



Graham Walker
BT Networks and
Systems

Graham Walker joined BT as a sponsored student and, after graduation from Oxford University in 1986, spent six years researching into coherent optical transmission systems. He was involved in the world's first field trial, and made advances in aspects of polarisation control and the noise performance of optical amplifiers. This work resulted in numerous publications, and the award of a Ph.D. from Cambridge University in 1992. More recently, he has been leading a group within Advanced Applications and Technologies, working on information visualisation and shared information environments. Current interests include visualisation of call record data, and the development of shared spaces for collaborative working with data and applications. He has recently completed an MBA at Cranfield School of Management and is a Member of the IEE.



Jason Morphet
BT Networks and
Systems

Jason Morphet completed his BSc in Software Engineering at Suffolk College (associate college of the University of East Anglia) following an HND in the same subject. He joined BT as a summer student in June 1995, staying on to become a Research Fellow and subsequently a full-time employee. His interests include both the network and human interface issues of distributed virtual environments, and he was closely involved in the AA&T Portal development.



Marco Fauth
BT Networks and
Systems

Marco Fauth received a B.Sc. in Computer Science from Fachhochschule (Polytechnics) Karlsruhe, Germany, in 1994, and is currently working for an M.Sc. in Telecommunications Engineering from the University of London. He spent his final B.Sc. project at BT Laboratories developing a visual interface to a complex modelling environment. He stayed on as a Research Fellow working on parallel programming using PVM, and then moved onto Java for delivering distributed applications. Jacaranda was the first graphical control and reporting application to be developed and is the subject of a patent application. He joined BT as a full-time employee in 1995, and has been working on a range of WWW-based customer interfaces, all involving Java.



Paul Rea
BT Networks and
Systems

Paul Rea joined BT in 1987 after graduating from Aston University with a degree in Computing Science. He has been involved in several areas of software engineering including the development of a graphical interface for a software engineering tool, and the study of system development methods. More recently, he has been working on visualisation and distributed virtual environments, taking a leading role in Virtuosi, a DTI-supported CSCW project. He has been closely involved in BT's VRML activities, and is a member of the AA&T Portal team.

Moving to a Paperless Quality Management System

Aardvark to Zulu

Electronic management systems offer important advantages when compared with paper-based systems. Improved accessibility and rapid updating are key factors, and lead to increased use and greater adherence to requirements. New authoring software and the World Wide Web are making it easier to realise the dream of a paperless management system.

Introduction

Management systems are normally home grown, and it is often the case that first versions are over-prescriptive bulky tomes which gather dust on sagging office shelves. The enthusiasm of authors, and there may be many, exceeds that of readers. Indeed, authors may be experts from relevant fields, but are not necessarily the greatest scribes in the world. Their contributions may be difficult to read and use, impossible to maintain, and likely to fall quickly from favour. The presentation of information is often considered unimportant, or not considered at all. This is a fundamental mistake. Any publisher knows that books are judged to some extent by their covers, and that a reader's interest is lost when pages are badly designed or inappropriately illustrated. A poorly presented system just does not get used. A management system must be written for its readers, and for a particular medium. Half a dozen pages of quality manual are of little use in an electronic hypertext system—they merely submerge the reader in motherhood statements and scroll bars. The author should think again. But is there a way of serving paper and electronic media with a single format?

The Need for Change

In a research and development environment, it is particularly important for managers to encourage innovation, flair, enthusiasm and

flexibility. Dump a few volumes of bureaucratic quality procedures on a desk and the green shoots of a new idea may get squashed flat. At BT Laboratories' Network Engineering Department it was soon apparent that innovation was important for management systems as well as technical work. For some years, management-system information had been spread between numerous large and heavy A4 ring-binders. These were updated at different times by different authors—a situation unlikely to lead to consistency in a fast-changing organisation.

Improvement began with a decision to draw together all the top-level information from a procedure-based management system, and convert it into a process-based form. It was also decided to achieve as much organisational independence as possible, and to add a comprehensive index. This would make future maintenance more practicable and enable people to find more easily the information they required. However, initial work revealed something more important. Key requirements were addressed in widely separated locations, and even repeated wherever they had significance. A single change of organisation or requirements could therefore lead to the reissue of several documents. Radical change was justified.

Organisational Independence

In recent years, BT has been restructured and reshaped to enable it to

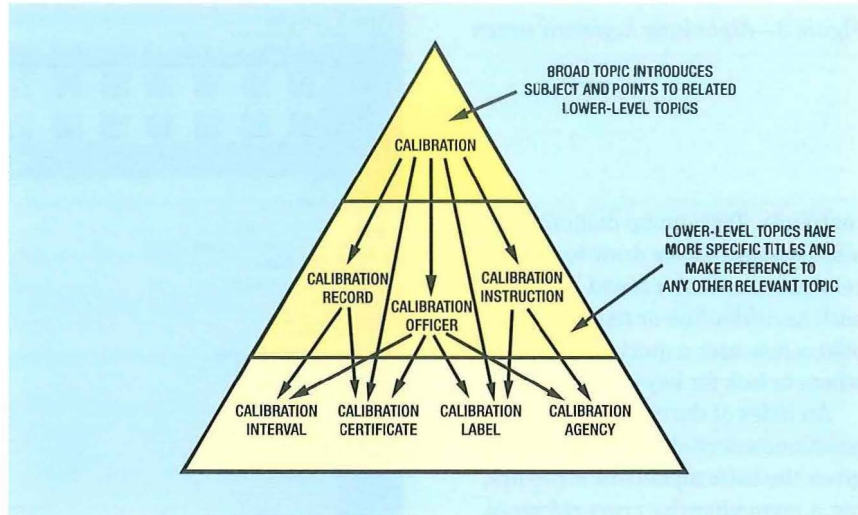
Figure 1—Pyramid topic structure

focus better on specific market sectors, and to recognise the needs of different types of customer. This has brought significant change to internal organisational interfaces. To avoid continuous updating of organisational information, a number of changes were made to the management handbook. References to the department's name were removed from the text and front cover. Those who must observe the requirements of the system are named in a distribution list held on a central database. Personal name, role and contact information was removed to a separate list. Finally, the scope of the system was widened to encompass, in decreasing detail, work carried out beyond department boundaries. This replaced a well-defined information boundary with a softer decline in the level of detail.

Going Alphabetic

In a two-month period, a single author, aided by a word-processor, tore apart several A4 volumes of procedure-based documents. The basic objective was to cut the quality manual and all the procedures into paragraphs, and collate topics into single-subject piles related to processes or sub-processes. Thus, *corrective action* of all types formed one pile, and all discussion of *issue control* formed another. The result confirmed the extent to which subject matter was repeated, sometimes in numerous inconsistent forms.

The next stage was to rewrite each pile of similar paragraphs into a single process description, using words appropriate to all applications. That done, a startling fact emerged. The number of words in the whole system had reduced by 60 per cent. Why? Everything was now addressed once only, and the requirement for broader application of process descriptions had resulted in identification of key principles rather than in fine detail. These were presented succinctly with heavy use of bullet-pointed lists. Title pages, contents



lists and history information had also been discarded.

The new topics were now given short descriptive titles likely to occur in text. Topic titles were chosen in small pyramid structures, Figure 1, a generalised top-level title referring out to related lower-level topics under more specific headings. This allowed topics to be kept short, and cross-references to the top of a pyramid of topics then led to relevant lower-level material. All topics were then assembled in alphabetical order. The result was a process-based encyclopaedia of concise single-topic sections. Having spread subject matter to the four corners of the alphabet, the challenge was then to

interlink them. This was achieved by italicising each occurrence of every topic heading within the whole encyclopaedia. Thus, the existence of a topic entitled *delivery* is advertised by italicising every occurrence of the word *delivery* throughout the handbook, Figure 2. Derivatives such as *delivering* and *delivered* may be italicised at an editor's discretion, but a degree of consistency is important. The *delivery* topic must of course read appropriately for every instance of the word *delivery* and its derivatives from which a reader may have been referred.

A hierarchy of top-level process diagrams points to lower-level diagrams or text entries to provide

Figure 2—The use of italics and cross-referencing

DESIGN DOCUMENTATION

Any standards needed for design documentation are defined in this handbook or local instructions. Documentation required as part of a final product or service, as well as other documentation which is produced but not *delivered*, must be defined in a relevant plan. All changes to design documentation must be made using appropriate change controls. The documentation produced must be consistent with:

- ensuring that everyone involved has a clear specification and understanding of objectives;
- allowing both a project team, and customer, using agreed methods, to assessments of product or service quality;
- providing sufficient information

Page 47

DELIVERY

No product should be delivered to a customer until all the activities specified in a relevant plan have been satisfactorily completed, and the associated documentation and release authority are available. The delivery itself should take place in a manner agreed with the customer and be appropriate to the product concerned.

Written records showing deliveries of product to a customer must be maintained as a deliverables list. These records should show the date of delivery, and the quantity and system testing status of product involved.

Page 43

Figure 3—Alphabetic keyboard screen

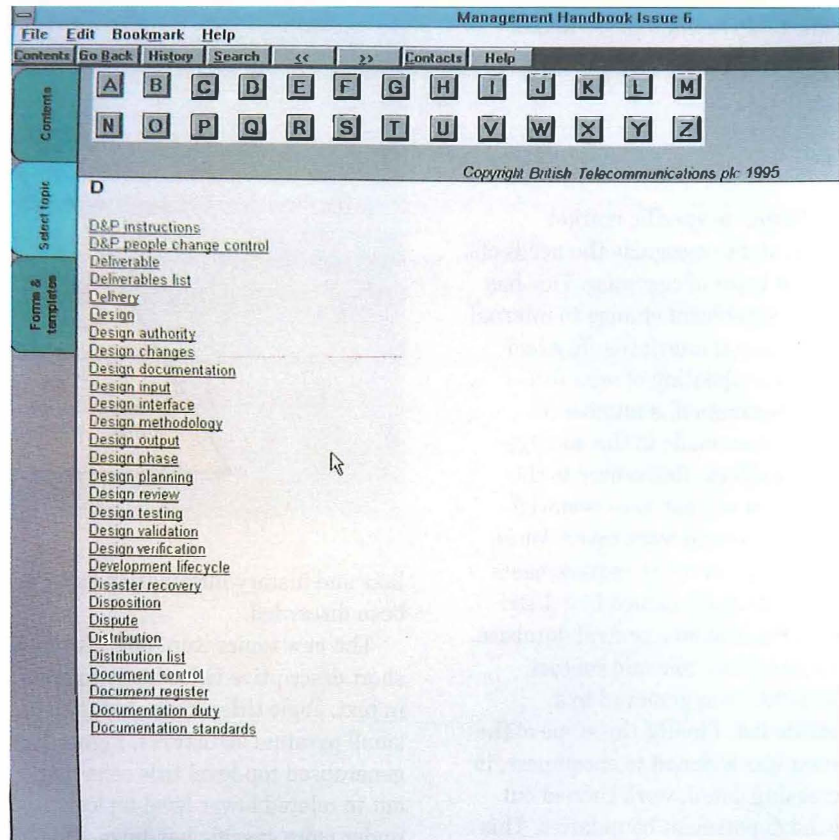
continuity. Topic group diagrams are also included. These draw together related topics under broad headings such as *calibration* or *testing*, and offer a new user a quick overview of where to look for key information.

An index of the complete encyclopaedia now served little purpose given the basic alphabetic structure, but a comprehensive cross-reference proved invaluable. This lists the location of every italicised word. Under *review*, topics such as *design*, *review record* and *corrective action* would appear as sub-entries because those topics refer to aspects of review. A reader is then directed to a single location for a description of a particular process, and is presented with a list of related topics to which the process refers. An alphabetic change log, listing all significant changes since the previous issue, is included at the back of the book. This points to text marked by change bars within topics, and to relevant change request references to give traceability back through the authorisation process to the originator of each change.

A Pocketbook

The complete document is bound into a 9 cm × 18 cm hard-cover pocketbook modelled on a business diary. It is labelled with a simple issue number, and printed on 60 g/m² paper to keep it slim enough to fit into a small pocket. The pages are tab-cut to reveal index letters, and a silk marker is provided. Two thousand copies were printed at £3.50 each—considerably less than the cost of printed pages for an A4 ring binder. When the cost of a ring binder, cover and spine inserts and divider tabs are considered, the difference is even greater.

Distribution of the pocket book was followed by a significant increase in usage and hence adherence to requirements. Many users see the attractive and manageable physical appearance of the pocketbook as reflecting lightweight requirements.



This is true, but to a lesser extent than users seem to realise. If the doors of perception were cleansed, some advantage might well evaporate. Feedback from users also indicates that the short-form bullet-point information is easier to find and quicker to absorb.

New users have some initial difficulty in knowing where to start or what to look up. The topic groups help in this respect, and it is not long before familiarity increases to the point where readers go directly to the information they require. Most important of all, people use the pocketbook because it is convenient to do so. Finally, the realisation that managers were carrying the pocketbook to meetings and on business trips in the City prompted the inclusion of the London Underground route map.

Going Electronic

The next step was to transfer the pocketbook into electronic hypertext form. A number of authoring tools are available, and Microsoft Viewer was chosen principally because it has all the necessary features, and offers a free run-time licence. The system is PC based, and uses screens of information taken directly from

topics in the paper pocketbook. Text is stored in a single rich-text-format file, and topics are separated by page breaks. This single-file structure aids maintenance, which is facilitated by good search and replace functions. The alphabetic structure lends itself well to this conversion, the italicised words being converted directly into hypertext links which guide users through the system. Embedded codes are different for each medium, but the techniques involved are similar. This is a major advantage because it makes practicable parallel paper and electronic versions without an intolerable maintenance burden. This underlines the need to plan the format of information with target media in mind.

The PC screen is divided into two panes as shown in Figure 3. The small area on the left-hand side is used for personal organiser tab-style hot-spots that provide rapid access to various categories of information. The system uses several sets of tabs, and automatically displays those most appropriate to the selected topic. The other much larger area is used for presentation of the topics themselves. The philosophy adopted calls for several alternative methods of accessing the same information. Topics are listed alphabetically via an



Figure 4—The electronic and paper management handbooks

on-screen A–Z keyboard, and may be accessed via hot-spots in top-level process diagrams and topic groups. Browse buttons allow users to browse sequentially, and in either direction, through topics listed under a particular letter of the alphabet. A word-processor-style search facility is also provided. This enables users to search the entire management system for instances of a particular word or sequence of characters, and reports back with a list of topics within which the specified characters are found. Occurrences of the required characters are automatically highlighted in the relevant screens. The search operators AND, OR, NOT and NEAR can also be used. Hence, one could search for instances of the word *authorisation* NEAR to the word *document*, in either order, where NEAR is definable as being within a specific number of words.

The top-level menu system also points to a number of supporting functions. The first of these is an ISO 9001 cross-reference which works in both directions. Users choose a paragraph of the ISO standard, and are presented with a brief introduction and a filtered alphabetic list of hypertext links to related topics. In the other direction, each topic includes a hot-spot which generates a pop-up window containing a reference back to the relevant part of ISO 9001. This is particularly useful for auditors, and during assessments or

surveillance visits. References are handled in a similar way: click on a reference number anywhere in the text to see a pop-up window containing title, author and source information. A complete list of references is also available.

A full on-line hypertext help system, structured in the normal Windows help-file format, is provided as part of the Microsoft Viewer package. By running a second instance of Viewer from within the first, help information can be pro-

vided in a separate window while the management system is in use.

The change log, which is broadly the same as the one used in the pocketbook, uses an on-screen A–Z keyboard. Click on a particular letter of the alphabet and a list of changed topics is presented. Each entry includes a brief description of a change, and a reference back to its source—usually to a problem-report database. Changes are listed in alphabetical order, and are hypertext linked to the relevant topic where change bars are used to mark changes.

Electronic form templates are embedded in the Viewer system.

Specific forms are hypertext-linked from numerous points within relevant topics. Forms are also listed in their own expandable menu. The templates are stored in Microsoft Word 6 format, and the system starts Word 6 from a declared directory, loaded with the appropriate template, when a form is requested. A number of other features, such as installation and de-installation instructions, and various figures are also included. The automatic set-up program for the system provides for networked or stand-alone PCs, and is delivered on two floppy disks in a compiled and compressed format.

The management system has been approved, in both paper and electronic form, Figure 4, by Lloyd's Register Quality Assurance.

The Future

The most recent stage of development has been to transfer the management system to the World Wide Web (WWW)—a world-wide hypermedia system and part of the Internet. The basic structure used is the same as

The most recent stage of development has been to transfer the management system to the World Wide Web

for the Viewer system. However, for reasons of speed, it is best to have each topic, or screen of information, stored as a separate file. Once again, the embedded link information takes a different form, but conversion can be achieved quite rapidly. All the features of the Viewer system are duplicated in the WWW system, including the search capability which must now be able to deal with about 1500 separate files. The number of files involved is an inconvenience from the point of view of maintenance, but the advantages in terms of cost, convenience of delivery, and the ability to make rapid changes are enormous.

Recent development has added a Tikit cross-reference similar to the earlier ISO 9001 facility, and a topic-by-topic change log. The latter allows users to see the change history of each screen at the click of a button. In similar style, a very brief summary of each topic has been added under a button labelled **IN BRIEF**. This reminds users of the key points only. Guidance and suggestions are also available on a screen-by-screen basis.

Other options are opening up all the time. Adobe Acrobat is soon to be available for WWW—this could allow a return to a single data file supporting WWW or platform-independent stand-alone applications. Such a system can offer improved graphics and a more powerful search capability.

People now spend less time searching for information in the management handbook. There is no doubt that users are more inclined to make reference to it. Usage is monitored from the network server and the statistics provide useful maintenance information. Adherence to requirements has certainly improved. And what happened to sagging office shelves? They still sag under the weight of computer manuals, but I get the impression that the book-shelf is a threatened species. Maybe just one of the chain saws in the Amazon rain forest will now fall silent.

Acknowledgement

The author wishes to acknowledge the contribution made to this work by John Brisbane of Applied Information, Ipswich.

This article was first published in *Quality World*, Volume 21, Issue 12.

Contact


Further information may be obtained by contacting the author:

Mr R. H. Macmillan
Room 46, MLB2
BT Laboratories
Martlesham Heath
Ipswich
Suffolk
IP5 7RE

Tel: (01473) 645233

Fax: (01473) 649192

Biography



Roderick Macmillan
BT Networks and
Systems

Roderick Macmillan, MISC, joined BT Laboratories at Martlesham Heath, Suffolk, in 1970, working first on data handling systems for satellite propagation measurements, and then on the design and development of signalling receivers using custom large-scale integration techniques. Since 1987, he has been responsible for the development of quality management systems using desk-top publishing, hypertext systems and the World Wide Web. In recent years, he has been particularly concerned with the reduction of complex management systems to more accessible forms.

Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his column, which commenced in the January 1996 issue of the Journal, by looking at the explosive growth in information and at our ability to deal with it.

Is Serendipity Finished? by Peter Cochrane

There you are looking for something quite specific, and usually urgent, when you come across something else you will need in the future or have been unsuccessfully looking for at some time in the past. This is often a delightful experience and gives a warm feeling of discovery and reassurance. In its raw form it is nowhere bettered than the Dickensian library full of books and manuscripts. It is difficult to pinpoint when it became a useful mechanism of discovery, major preoccupation, and even sport. Presumably its heyday arrived with the transition of the library, from a disordered and chaotic place for the storage of parchment, to the well-ordered Dickensian edifices we enjoy today.


It is interesting to pose the question of how much serendipity you can have, or indeed how much can you stand, and what is the right level of serendipity. The medieval library had to employ a librarian as the guardian and regulator of information. He was the human contents list, the index, filing system and retrieval mechanism. Not only did he perform the curatory role, but he was the agent who accessed the library and located the very manuscript you required. He alone decided who saw what, when and where! In medieval days there was not open access to information—it was strictly regulated and controlled. So what chance serendipity—not much! For the past 200 years at least, we have enjoyed increasing levels of serendipity with the introduction of order and decimalisation into paper libraries and the transition of the librarian's func-

tion from guardian to assistant and information agent. The openness of information and access to it has meant that most of us have experienced a high percentage of serendipity by merely walking through rows of shelves and spotting the odd title, that obscure journal, article, and paper. Even the daily newspaper affords us a high degree of serendipity through the attention catching mechanism of headline and picture.

Contrast all of this with the seemingly infinite information world of Internet. Here we have almost 100% serendipity—an abundance of data and access with an overriding lack of information, no order, no sign posts, no eye catching indicators. Actually finding what you want is now the great challenge. Being totally awash with serendipity—things you are not looking for—poses a new and serious problem. In this environment, information seems to come in two dominant classes: that which is of no interest, and that which is distracting, interesting, but still of no direct benefit! The problem is now to find anything that you actually require, and when you do, to have some notion of its value and fidelity.

In another corner of the electronic world, and poles apart from Internet, we have the CD-ROM. These generally embody an almost total lack of serendipity. So well organised, sterile and deep is this medium that drilling down from the opening page to the information you require can involve over five clicks of the mouse and a total lack of visibility of information either side of the mine shaft you have just dug. In this medium navigation is a major problem! After three to five clicks you can find yourself totally lost and disorientated, with no sign posts, indicators or frame of reference to help you. Totally lost, you resort to Control Quit and start again!

Somewhere between the extreme of Internet and CD-ROM lies the world we would like to live in. A world with the right degree of serendipity, the right percentage of chance that allows us to optimise our creativity and rate of work. The



question is how are we going to realise it? Creating serendipity by design has got to be a major challenge, for in our recent past we have created such worlds by accident. Perhaps we have to wait for the electronic world to evolve to a point where it spontaneously creates some serendipitous environment for us. But I suspect not! For while the world of electronic information is on a scale so colossal it defeats the human mind to contemplate its vastness, we have already experienced the delights of serendipity, we know what is, we have some measure of its value to us. Intuitively we feel we should be able to manufacture it.

Around the year 1600, the Pope had one of the largest libraries on the planet with just over 400 books. Today many of us enjoy a personal library of more than 400 books and could perhaps usefully employ a librarian! The Library of Congress in the United States has an estimated 24 million volumes and takes in over 3.5 km of new bookshelf worth of books each year. In the not-too-distant future we will be in a position to have 24 million books in our homes, but obviously not in paper form. Today we can buy 2000 of the classics on a single CD. The complete Works of William Shakespeare is now available for 2.5p! New generations of CDs or similar storage media will exceed this density 100-fold and allow us to create the Library of Congress, and even more, in a volume that will not overpower our living space. Undoubtedly, we will not choose to

do this, and will resort to accessing a vast selection of distributed libraries across the planet. Without a team of librarians, how then will we search and find what we are looking for, how then will we enjoy the serendipity of wading through the bookshelves of paper?

Perhaps we will have to look to artificial agents who will learn about our habits, interests, wishes and desires through direct observation. They could then take on the role of a very intelligent and intimate


librarian that will roam the data banks on our behalf and bring to our attention related and unrelated text, documents and materials to aid our creativity. Perhaps we will have to spend more time with other human beings discussing our problems and formulating our views to create new degrees of serendipity that have so far escaped us. Either way, we face a major challenge with mankind's knowledge base now reckoned to be doubling in a period of less than two years. This is especially so with the

creation of increasing levels of short-term information that just act as a fog and barrier to us spotting what we are really looking for. We are going to need help through increasing levels of computing power and artificial intelligence. There are no human attributes that will enable us to cope with the massively increased levels of serendipity we now have and will increasingly see. We have to hope that the machines can help us or face an increasingly sterile and less creative world.

book review

Source-Matched Mobile Communications

by W. C. Wong, R. Steele, and C-E. W. Sunberg



This book starts by introducing the principles of mobile communications, and then shows how to determine link performances for different types of speech coding, error control coding and modulation methods. Source-matched coding and modulation schemes are introduced, demonstrating how to take advantage of the source and channel characteristics to improve the quality of the signal in adverse channel conditions. The authors are established world-class leaders in the field of speech coding and modulation research with many joint publications on the subject.

The book is structured into nine chapters, each supported with references. The reader is introduced to the concept of advanced cellular architecture including macrocell,

microcell and picocells. Fortified with the cellular concept, basic speech coding techniques including pulse code modulation (PCM), differential PCM, subband coding and linear predictive coding are reviewed.

There is a study of transmission errors and digital noise power in binary modulated PCM systems for both Gaussian and Rayleigh fading channels. The authors conclude that noise performance is improved by using a code sensitive to the nature of the noise process.

In chapter 5, the authors mark a shift in the book as the focus is turned to the subject of soft demodulation in binary modulated PCM systems.

Chapter 7 evaluates the effect of transmission errors for differential PCM (DPCM) transmissions and extends the treatment to embedded DPCM. Theoretical derivation of the performance of both systems operating in Gaussian and Rayleigh fading channels are presented.

The final chapter of the book covers linear-predictive coded speech combined with 64-level quadrature amplitude modulation (QAM) and source-matched transformed binary pulse excited linear predictive coding (TBPE-LPC) speech transmission. It investigates the suitability of a coherent 64-QAM scheme over both Gaussian and Rayleigh fading channels using theory and simulations.

This book is a good reference for readers seeking a deeper understanding of the relationship between speech coding, error-control coding and modulation, especially under Gaussian and Rayleigh fading channel conditions. It is rich in mathematical derivations and requires a good background understanding of digital signal processing and probability theories to appreciate the material. While the book keeps a system perspective throughout with a mobile radio theme, it addresses limited mobile system application and implementation issues. The focus generally lies with analysing generic link performance using a combination of coding and modulation techniques. The methodical approach of the authors in addressing the subjects is particularly commendable. For this reason, graduate students and research workers will find the book a good overview of the subject and the references provide a good source for further in-depth studies. However, due to its strong theoretical bias, the immediate benefits of the book to practising cellular and mobile radio engineers may not be so easy to establish. Overall, it is a good reference book and some supplementary exercises would make it a good course book for graduate students.

Published by
IEEE Press and Pentech Press,
Typical price £38.50 viii+256 pages
ISBN 0-7273-1908-6

Reviewed by Stanley Chia

BT and ITJ to Offer Concert Virtual Network Service in Japan

BT and International Telecom Japan Inc. (ITJ), one of the leading international carriers in Japan, have announced they will jointly provide Concert Virtual Network Service (Concert VNS) to customers in Japan beginning in late July.

Alfred Mockett, Managing Director, BT Global Communications, said: 'We are very excited about offering global corporate voice to our customers in Japan in partnership with ITJ. Concert VNS will strengthen our service in Japan, where we already successfully provide managed data services—Concert Packet Services and Concert Frame Relay Services.'

Mr. J Takahashi, ITJ Executive Vice President, said: 'Large corporate users in Japan are diversifying. The requirements for global services is increasing and we are happy to be able to comply with these users' needs.'

BT will be responsible for the operation of the global backbone network, service support for customer installations and service delivery outside Japan in conjunction with its global partner MCI.

BT will also conduct joint sales and marketing activities with ITJ in Japan and with ITJ worldwide for Japanese customers.

ITJ will be responsible for joint sales and marketing activities with BT, and installing and operating customer networks in Japan through its subsidiary ITJ IT.

New Offers of Help to People who are Disabled or Elderly

New products and services to help disabled or elderly people use the telephone are in a new BT guide.

The 1996 'BT Guide for People who are Disabled or Elderly' is available free from BT shops or by calling Freefone 0800 800 150 (voice) or 0800 243 123 (textphone).

Braille, large print and audio-tape versions of the guide are also

available from BT on request. The guide is also accessible via the Internet on <http://www.bt.net/home/community/>

The guide is divided into sections listing products designed for people with impaired hearing, speech, sight, mobility and dexterity. It also details ways in which BT can make it easier to manage telephone bills.

Geoff Knight, of BT's Age and Disability Unit, said: 'The 1996 guide contains more items than ever before to make it easier for people to use the phone whatever the difficulty.'

'BT's latest Freestyle cordless phones (Freestyle 80 and Freestyle 120) automatically select one of four channels to give the clearest reception. In addition the Freestyle 120 has a back-lit keypad; every time the phone rings or a button is pressed the entire keypad is automatically illuminated.'

Expansion of BT Global Network Services

BT Global Network Services, provided by BT and MCI's joint venture, Concert, has recently expanded its managed data network service to India, bringing the number of countries where BT GNS is provided, and supported, on an end-to-end basis to 38. BT GNS will initially be available in Bombay and Madras, with service due to be extended to other major cities, such as Calcutta and Delhi. GNS provides instant connectivity to more than 1300 cities worldwide over a single network. Further countries planned to come into service this year include Taiwan, Russia and Greece.

BT and MCI Offer New Concert Call-Centre Service

BT and MCI have announced that they are introducing a new service allowing multinational companies to integrate multiple call centres in Europe and the US while offering local access to customers. The new service from Concert, the BT and MCI joint venture global services company, called *Concert Inbound*

Service, allows companies to increase efficiency, reduce the cost of their operations and give greater flexibility in their call-centre marketing strategies. By dealing with one vendor for end-to-end service, billing and regulatory issues, call-centre customers receive a high quality of service and standardised features. These include customer service, one bill convenience and call controls that are not restricted by geographic boundaries.

Alfred Mockett, BT's Managing Director Global Communications, said: 'With the globalisation of corporations and their products, the new service allows customers to give callers a local in-country appearance, while siting their call centres where it makes the most business sense. This is a further addition to our extensive Concert portfolio of global communications products meeting the needs of the world's multinational businesses.'

McQueen, a direct marketing, customer service and technical support provider, said: 'Concert Inbound Service allows us the flexibility to have PSTN or Freefone access in many different countries without the hassle of having to deal with many different operators.'

Craig Saddlington, General Manager of Marketing and Technology Services for McQueen, one of the first customers to sign up for the service, added: 'We can get consistency of service across Europe and deal with one vendor for billing and service.'

BT Responds to UK Competition Law Proposals

The Government has published proposals for reform of UK competition law. BT welcomed the Government's decision for reform of the law which needs to be updated to be relevant to today's business environment. BT thinks that it is important that regulation should have clear and transparent rules.

The proposed reform of the law on restrictive agreements is logical and sensible. It expressly envisages a right to have a decision of the

Office of Fair Trading reviewed by an independent tribunal and for points of law to be referred to the High Court on appeal. Each of these points has been proposed by BT in its representations to OFTEL on the proposed new fair trading provision. OFTEL has rejected them.

The Government's proposals on abuse of dominant positions suggest that there are no serious gaps in the current law and procedures but these are in need of updating. The proposed changes seem appropriate and proportionate to the weaknesses which are acknowledged to exist. BT, again, has proposed to OFTEL a similar set of procedural improvements. OFTEL has rejected them.

BT believes OFTEL's proposals are inconsistent with the proposals set out by the Government and will suggest to OFTEL that they should change their proposals to be the same as the new law.

Environmental Accounting In Industry

A BT report adds to the debate BT is spearheading to draw up standard rules for companies to report environmental costs and benefits. The company has brought together some of the UK's leading experts on environmental accounting in a publication which examines key questions the issue raises for companies today.

Environmental Accounting in Industry comprises two reports. The first, authored by Geoff Lane of Coopers and Lybrand, concentrates on the potential for establishing a comprehensive set of company environmental accounts, taking BT as a case study.

The second report, by Martin Bennett of the Wolverhampton Business School and Peter James of Ashridge Management Research Group, takes a broader remit and examines current trends and practice in environment-related management accounting by leading companies in North America.

Ian Ash, BT's Director of Corporate Relations, said: 'Over the past

four years BT has received national and international recognition for its approach to environmental reporting. 'In that time it has become increasingly apparent to us that there are no standard rules for reporting environmental costs and benefits, nor any consensus view on how such data would be evaluated.

'For example, how will a high environmental cost be interpreted? Will it be seen as a sign of environmental commitment, or will it symbolise an earlier failure to integrate environmental factors into planning processes?'

Globetel Partners Create New Company For Israeli Licence Bid

The partners of the Globetel consortium today announced that they will form a new company called, Newtone—the Israeli Company For International Telecommunications Ltd—to address international telecommunications in Israel. Newtone is a joint venture between BT, MCI, Tadiran Ltd (Tadiran), Idan Software Industries I.S.I. Ltd (Idan) and Darcom Ltd (Darcom). It will submit a response to the request tender from the Ministry of Communications for granting of one of two licences for international telecommunications services in Israel. The joint venture shareholding will be: BT—25 per cent, MCI—15 per cent and the three Israeli partners—20 per cent each. The consideration will be less than 1 per cent of BT's consolidated net assets.

BT Backs NCC Report on Information Society

BT welcomed calls by the National Consumer Council for swift action to ensure wide access to the information society. Rupert Gavin, BT's director of multimedia services, said: 'The speed of convergence between telecommunications, broadcasting and computing holds enormous potential to open up an

ever-widening range of services for customers.

'The NCC recognises the value and benefits of services already on offer or under development in areas such as education and services for the disabled. BT is proud of its work in this field and aims to remain at the forefront of developments designed to make the information society available and accessible.

'BT invests more than £270m a year on research and development, the major part on developing new and improved networks and products. Products such as Campus-World, with more than 2 500 British schools linked to the Internet, lead the world. Developments in videoconferencing, speech synthesis and virtual reality are already having an impact on education and medicine across the country.

'However it is essential that if the UK's communications companies are to develop such services for the future, the regulatory regime must allow adequate returns on capital to support and encourage adequate investment.'

BT's 'Mass Market' Internet Access Service

BT has launched its mass market dial-up Internet service. The company has sent out registration software to thousands of people who have enquired about the service—BT Internet—since it was announced at the end of February.

BT Internet is aimed at residential and small business customers as well as users new to the Internet. The service emphasises ease of use, competitive subscription costs with no hidden charges, comprehensive customer support and a jargon-busting approach suitable for the non-technical user.

A full range of Internet services is offered, including worldwide electronic mail, file transfer, and access to vast quantities of information through the World Wide Web and discussion groups.

BT Internet will bill customers directly for their subscription only.

Calls to service are charged at local rate throughout the UK and are billed separately as part of the customer's regular telephone bill.

The BT Internet customer enquiry line, Freefone 0800 800 001, is available for anyone interested in subscribing to the service.

BT Introduces Interactive Multimedia Kiosk

BT has announced the planned launch of a significant number of touch screen, interactive multimedia kiosks in London this summer. The kiosks, called *Touchpoint*, are aimed primarily at residents and English speaking tourists.

Information held on BT Touchpoint for customers to use will be predominately free to access and browse. The contents will range from a guide to what's on and where to go (including restaurant, cinema, theatre guides as well as a ticketing facility) to up-to-date news, sport and local news, horoscopes, as well a street guide and shopping area.

By simply touching the screen, customers will be able to access any area they choose, to see and

hear the content in a mixture of text, picture and video formats. They will also be able to talk to companies featured on the system to find out more and even buy goods and services through dedicated Freefone 0800 telephone link.

Each Touchpoint kiosk will comprise a colour touch screen, a printer from which vouchers, coupons, and maps can be printed. The Touchpoint kiosks will be located across central London in areas such as travel terminals, shopping locations and leisure and tourist centres.

Independent research conducted by NOP and commissioned by BT has shown that the Touchpoint concept will prove popular amongst customers, with more than 90 per cent of those researched saying they would definitely use it.

Rupert Gavin, Director of BT's Multimedia Services, said: 'BT's activity in the fast-moving multimedia market is intensifying. Touchpoint kiosks represent an important early step in bringing advanced multimedia applications into the reach of a mass market audience.'

New Seal Of Approval For BT Metering Systems

The British Approvals Board for Telecommunications (BABT) has renewed its approval certificate for BT's bill metering systems, confirming that they meet BABT's stringent control requirements and OFTEL's Standard for Public Telecommunications Operators' Meter Systems.

The certificate was first issued in 1994 and BT is still the only telecommunications company in the UK to have received such approval. Charles Paull, BT's Director of Financial Accounting, said: 'BT has always placed great importance on accurate billing. This further independent stamp of approval will give customers even greater confidence in our systems and control processes.'

'The standard is very challenging, but we were always confident in our ability to meet it.'

David Clarke, Director of BABT, commented: 'BABT has worked closely with BT to assess the accuracy and reliability of its metering and billing systems, and we are satisfied that the services covered by the approval comply with the standard.'

industry news

France Telecom Unveils Internet Services

France Telecom has announced a series of innovations that will offer powerful and economical access to all types of on-line services, from Mintel to the Internet.

France Telecom is rolling out solutions for connection to Internet access providers, all offering uniform rate structures throughout France. Users can connect to their Internet access provider from anywhere in France for the cost of a local call.

Additionally, using the same principle of uniform call charges from anywhere in France, a service is provided involving 'kiosk' type revenue sharing with service providers at 1.29 francs per minute.

Providers can offer Internet access, without requiring a subscription, using this service.

Monopolies and Mergers Commission's Report on Number Portability

The Director General of Telecommunications, Don Cruickshank, announced the findings of the Monopolies and Mergers Commission's (MMC) investigation into telephone number portability—the ability of a customer to keep their existing number when they change operator.

The Director General said: 'This is a good day for customers. I am pleased to announce the MMC has endorsed my view that number portability is in the public interest,

and that BT should not be able to recover from other operators all its costs in bringing about portability, as its licence provides at present.

'The MMC has concluded that modifications to BT's licence are necessary, enabling me to direct BT not to recover its costs in full from other operators. The MMC's detailed recommendation on the allocation of the costs between BT and other operators is very close to my original proposal, which would have resulted in a 75–25 per cent split of BT's total portability costs over the next five years. The MMC recommends that the split should be approximately 70–30 per cent. In contrast, BT's best offer to date would have resulted in a 15–85 per cent split in its favour.'

'Furthermore, during the course of the inquiry BT's cost estimates

have fallen considerably and it has made commitments to introduce cheaper methods of providing portability. A major barrier to choice will be removed. More customers will be able to enjoy the undoubted benefits of competition in telecommunications. However, I do regret the lengthy delay in delivering something which the great majority of customers wish to have.'

The Director General went on to comment on two important specific elements of the MMC's proposal. 'The MMC has proposed that where BT has said it will introduce a cheaper solution by a certain date, I should have the power to require them to charge only that cheaper cost from that date—regardless whether the cheaper service is in fact available. This 'efficient operator' concept is a very important form of incentive regulation, and I will be looking at whether a similar approach should be adopted elsewhere.

'The MMC have also proposed that I should have discretion to disallow recovery by BT of any cost which is unreasonable—for instance because it reflects the use of an inefficient process. Again, this is an extremely helpful and important discretion for me to have.

'This issue of discretion is one which has been raised both in the context of this inquiry and in respect of regulation of telecommunications generally. Let me be clear that in this area as in others, I have never sought discretion other than that which is proportionate—in this case, for the purpose of securing the early introduction of portability.

'I am now launching a period of statutory consultation on a licence modification based on the MMC's recommendations. Once this is completed, and the licence modified, I hope that BT and other operators will work closely with me to implement the MMC's findings and deliver portability to customers as quickly as possible.'

The MMC's proposals for cost-allocation are as follows:

- BT should bear the initial costs of modifying its network (system setup costs).
- BT should be able to pass on to the other operator concerned, the costs of enabling individual customers to port their numbers (per-line setup costs).
- After October 1997, BT should bear the 'additional conveyance costs' associated with routing a call to a ported number using the 'call dropback' method. During the period up to October 1997, the extra costs of 'tromboning' over and above the 'call dropback' costs should be shared equally between BT and the other operator.

BT's estimate of its costs fell heavily during the course of the inquiry. It made a commercial offer to other operators which OFTEL estimates could have resulted in other operators facing a bill of more than £500 million between now and the year 2000 if BT had recovered its costs in full. By the end of the inquiry, the estimated total costs had fallen to £200 million. On the MMC's proposal, other operators would now face a bill of some £60 million for this period, with BT picking up the remainder.

OFTEL's Universal Service Proposals

Don Cruickshank, the Director General of Telecommunications, published a consultation document on Universal Service.

Presenting the document, Don Cruickshank said: 'Universal access to basic telecommunications services is becoming an increasingly important national objective. It provides economic benefits and is an essential part of most people's daily lives.

'Universal service means basic telecommunications services readily available and affordable to all customers who want them. The basic principles are geographical accessibility, affordability and equal

opportunity. With operators' cooperation and imagination there should be an increase in the 91 per cent of UK households now connected to the network and less people denied its many economic and social benefits.

'The document also considers how the universal service principle could ensure that people with disabilities can benefit from the services now being offered. At this stage the focus is on the needs of deaf and speech impaired people. There is now a wide range of special facilities and equipment available for those who are unable to use standard telephony products. But these need to be more readily affordable and more widely publicised.

'The document also looks at the telecommunications needs of the education sector. If the education sector and the government agree that schools for 5—16 year olds should receive a higher level of service, and the costs are not too great, a Universal Service Fund could help provide affordable advanced services to schools. I welcome a wide-ranging debate on these important issues.

'There is, of course, a cost associated with delivering universal service. The evidence suggests that the net cost is low—less than one per cent of UK telecommunications revenues. But it is still large enough to justify an effective funding mechanism. The document explores how a Universal Service Fund should operate: who should pay into it; who should receive contributions from it; the basis of calculating the amounts of payments into and out of the Fund; and who should administer it.

'I am also keen to see universal service being delivered competitively wherever possible. "Pay and Play" and competitive franchising are two ideas which the document considers as feasible options for maintaining customer choice and quality of service while balancing the interests of BT and other operators. These could be used for service to uneconomic areas, special

services to disabled people and advanced services to schools.

'As technology develops, advanced services spread and costs drop, the level of universal service may need to be reviewed. The document therefore looks at what criteria might be used when considering an upgrading of universal service and what form any review body might take.'

Comparing the Performance of Telecommunications Companies

Don Cruickshank, Director General of Telecommunications, has published quality-of-service indicators which will help UK telecommunications customers compare the performance of different telecommunications companies. OFTEL has published two reports—one for business customers and one for residential customers—showing how eight leading telecommunications companies perform across five areas:

- installing services on time,
- providing reliable networks,
- repairing faults on time,
- handling complaints promptly, and
- issuing accurate bills.

Don Cruickshank said, 'UK telecommunications companies are the first in the world to produce accessible, customer-focused information of this sort on a comparable basis. This represents an important first step in giving customers better information on which to base their choice of telecommunications supplier. The ability of telecommunications companies to provide this information is a welcome sign of their increasing focus on customers. Our intention is to build on this initial information, in subsequent publications, and include other indications of value to customers for a

wider range of telecommunications companies.

'The publication is the result of a voluntary, cooperative effort between the telecommunications companies, consumer organisations and OFTEL. The data reported have been checked by independent auditors to give customers confidence in their accuracy.

'The results are for fixed-link companies, but mobile telecommunications customers also need to compare performance of their telecommunications supplier to make more effective choice. I would encourage the mobile industry to follow the lead given by the fixed-link telecommunications companies.'

OFTEL Consults on Duct and Pole Sharing

Don Cruickshank, Director General of Telecommunications, has issued a consultative document which looks at BT's current policy on duct and pole sharing and its implications for the telecommunications market. The main objective of the consultation is to seek views from network operators, site owners and customers on whether duct and pole sharing would enhance competition.

Don Cruickshank said: 'The consultative document reaches a number of conclusions on different aspects of sharing. The broad thrust is that in some circumstances, notably where there is limited access to customer sites, there could be economic benefits in shared access. These need to be explored further in trials which BT has said it is ready to conduct.'

The EU and Mobile and Personal Communications

The European Union (EU) market for mobile telephone networks will be opened to full competition as a result of a Directive formally adopted by the Commission under Article 90 of the Treaty. It requires Member States to abolish all exclusive and special rights in the area of mobile communications,

and to establish open and fair licensing procedures for digital services. The Directive also removes all existing restrictions on the use of facilities for mobile networks. Member States will have nine months to notify implementing measures from the date the Directive came into force.

Cellnet, Barclaycard and the Merchants Group Sell Mobile Phones in Joint Project

Mobile phones are being sold in a joint project that includes Cellnet, Barclaycard and the Merchants Group.

The project encompasses a targeted offer to Barclaycard customers to buy a mobile phone and subscribe to the Cellnet cellular phone network. The Merchants Group, who are managing customer response to the offer are responsible for the sale of the mobile phone handsets.

Barclaycard are offering selected customers the opportunity to purchase a Cellnet digital mobile GSM phone with a free connection, a competitive monthly rental and a 20 per cent discount on calls made on Cellnet's Regular Caller Plus tariff. As part of the offer, all ongoing charges are conveniently debited to cardholders' Barclaycard accounts, automatically, each month.

As an additional benefit, the Cellnet mobile handsets feature an exclusive 'Barclaycard button', which with one touch provides direct access to Barclaycard services such as Road Assist, International Rescue, customer services and bill payments.

Survey of UK Internet Usage

The number of adult 'surfers' on the Internet could more than double to three million by the end of 1996. While men outnumbered women by nearly four to one among the more active users at the end of 1995, the ratio is expected to decline quite

sharply. These are some of the findings by market research firm IRB International in a new study of UK current and prospective usage of the Internet.

The survey carried out by IRB International, on behalf of a number of leading financial services companies, provides a wide picture of how the 'Net' is currently being used and how people expect to use it in the future. The Net offers considerable business potential, finds the survey, but so far is 'grossly under-utilised' as a marketing/communication medium.

Based on a random sample of 13 599 individuals across the UK, IRB estimates that, at the end of 1995, there were 1.3 million adults (15 years-plus for the purposes of this survey) who had accessed the Internet at one time or another

(equivalent to 2.8 per cent of the adult population of 47 million). Based on 'intended' users, it is estimated that this figure will more than double to around 3 million (6 per cent) by the end of 1996.

IRB has, however, defined a 'user' as someone using the Internet regularly or having connected to it several times. The survey also concludes that the Internet is **not** the domain of the young single male: already, at the end of 1995, 46 per cent of the more frequent users were over 35 years of age, 53 per cent married or cohabiting and 41 per cent with children at home.

With 79 per cent of the more frequent users utilising the Internet for work purposes and 69 per cent for personal purposes, there is active use for **both** these reasons. Overall, 73 per cent have used the

Internet to transfer information via e-mail, 52 per cent to find information on products and services, 51 per cent to access market and industry information and 47 per cent to communicate with friends and relatives.

Expressed concerns, among users and intended users, included security, particularly of credit card payments, Internet cost, lack of personal interaction, and the difficulty of trying out or testing products before purchase. There is also a resistance to more blatant forms of advertising, more so among intended users. The survey suggests that, in addition to these concerns, companies marketing on the Internet need urgently, in cooperation with other interested parties, to address the problem of Internet site awareness, which at present can be a 'hit and miss' affair.

Could You Write an Article for British Telecommunications Engineering?

Contributions of articles to the *Journal* are always welcome. **Anyone** (BT or otherwise) who feels that he or she could contribute a telecommunications-related article (either long or short) which may embrace technological, commercial and/or management issues, is invited to contact the Managing Editor, BTE Journal, Post Point G012, 2-12 Gresham Street, LONDON, EC2V 7AG. (Tel:(0171) 356 8022; Fax (0171) 356 7942). Authors are advised to contact the Managing Editor before committing significant effort to preparing articles. Guidance notes for authors are available on request.

Field Focus

Field Focus is a feature in *Journal* comprising short articles (up to about 600 words) on specific local BT projects. For example, novel solutions to field problems have, previously, formed the basis of very interesting articles. Anyone interested in writing a Field Focus item is encouraged to contact the Managing Editor, or their local IBTE representative (see list inside back cover).

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by *British Telecommunications Engineering Journal*, 2-12 Gresham Street, London, EC2V 7AG. (Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in this *Journal*, unless any such statement is made specifically by the Board.

© 1996: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Price: Single copy—companies, universities, libraries and other bodies £9.25 (£10 including postage UK, £11.25 including postage overseas); private individuals (paid by personal cheque) £5.00 (£5.75 including postage UK, £7.00 including postage overseas). Annual subscriptions (including postage)—companies etc. £40.00 UK, £45.00 overseas; private individuals (paid by personal cheque) £23.00 UK, £28.00 overseas.

Overseas customers can pay by sterling drafts drawn on London banks or by sterling Eurocheques.

Price to BT staff: £2.25 per copy.

Back numbers can be supplied if available at the single copy rates given above.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 2-12 Gresham Street, London EC2V 7AG.

Remittances for all items (except binding) should be made payable to 'BTE Journal' and should be crossed '& Co'.

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: +44 171 356 8050. Fax: +44 171 356 7942.)

Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: G012, 2-12 Gresham Street, London EC2V 7AG. (Telephone: +44 171 356 8050. Fax: +44 171 356 7942.)

Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and special permission is necessary for reprinting long extracts, but editors are welcome to use not more than one-third of any article, provided that credit is given at the beginning or end, thus: 'From British Telecommunications Engineering'.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by the *British Telecommunications Engineering Journal* for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/96 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local-Centre Contacts

The following is a list of Local-Centre contacts to whom enquiries about the Institution may be directed. Where no Membership Secretary details are shown, membership enquiries should be directed to the Local Secretary.

Centre	Local Secretary	Membership Secretary
Aberdeen	Steve Murray (01224) 752475	
East Midlands	Dave Bostrom (01908) 656215	Ian Bethell (0115) 947 8587
Lancs & Cumbria	Trevor Hughes (01706) 40840	
London	Jim Skinner (0171) 843 8552	Peter Martin (0171) 492 6138
Manchester	David White (0161) 600 2468	Mike Ku (0161) 833 2669
Martlesham Heath	Peter Shaw (01473) 224488	Nisha Moudgill (01473) 224801
North Downs	Nick Smith (01227) 474823	
North East	Richard Jackson (0191) 496 0251	Frank Balloch (0191) 261 3226
North Wales & Marches	Norman Wallace (01691) 666377	
Northern Home Counties	Terry Birdseye (01702) 372557	Stephen Graham (01753) 564179
Northern Ireland	Allan McDougall (01232) 215301	Brenda Sharvin (01232) 320092
Paris	Delphine Lalire +33 1 46 67 27 89	
Scotland East	Graham Neilson (0131) 668 5878	Keith Doig (0131) 345 4710
Sevenside	Mike Stenning (0117) 920 5691	Dean Shephard (0117) 920 6457
South Downs	Graham Lindsay (01273) 222633	
South Wales	Peter Coleman (01222) 514524	Keith Voyce (01222) 555075
Staffordshire	Ray Piper (01785) 813483	
West Midlands	Paul Jervis (0121) 230 4181	
West Scotland	Ian McReady (0141) 220 7118	Jill Maxwell (0141) 220 5837
Westward	Rob Rand (01392) 267114	Chris Gould (01392) 267113
Yorks & Lincs	Steve King (0113) 237 8124	Paul Horncastle (0113) 237 6542

Associate Section Zone Committee Contacts

The following is a list of Associate Section Zone contact points to whom enquiries about the Associate Section should be directed.

Zone Committee	Contact	Telephone No.
London	Terry McCullough	(01707) 601151
Midlands	Mark Broughton	(0116) 242 3846
North East	Geoff Jenkinson	(0113) 237 8163
Northern Home Counties	Graham Lovell	(01473) 646056
Northern Ireland	Brian Walker	(01232) 232413
North West	Phil Holt	(0161) 339 1971
Scotland	Raymond MacMillan	(01463) 665913
Southern Home Counties	Martyn Harvey	(01273) 568388
South West	Roger Garland	(0117) 920 5485
Wales	Ferg Brown	(01952) 249886

IBTE International HelpLine — +44 171 356 8008

Contents

Guest Editorial: Making Breakthroughs	2
Process Re-engineering: Measures and Analysis in BT	4
TRIADS: Planning Changes to BT's Operational Support Systems Infrastructure	13
Cashless Services Replacement System Project A network perspective	21
The Future Isn't What It Used To Be	28
Education for Changing Times	32
Key Technologies for the Information Industry	38
Internet	39
Distributed Processing—Managing the Future	46
TINA—A Collaborative Way Forward	54
Software Agent Technology	59
Navigating On-Line Service Environments	66
From Books to Bytes—Managing Information in the Information Age	72
Security in the Information Age	78
The Information Needs of Network Communities	85
Interactive Visualisation and Virtual Environments on the Internet	91
Moving to a Paperless Quality Management System From Aardvark to Zulu	100

