

VOLUME 17 PART 1 APRIL 1998

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

*Lean Operations—The Route to
Competitiveness*

*Computer Telephony Integration
and Java*

*Future Trends in Satellite
Communications*



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 17 ■ PART 1 ■ APRIL 1998

Guest Editorial: Structured Process Improvement Keith Beales	1
<i>Structured Process Improvement</i> Lean Operations—The Route to Competitiveness Ray Hooper and Carole Jones	2
<i>Computer Telephony Integration</i> Computer Telephony Integration and Java Robert Brockbank, Stephen Peirce and Zoi Amanatidou	10
<i>Mobility and Radio Communications</i> Mobile Multimedia Applications Don Golding, Julie Harmer, Robin Mannings and Paul Randall	18
<i>Telecommunications in the 21st Century</i> Knowledge Discovery and Delivery John Davies and Ralph Cochrane	25
Intelligent On-Line Purchasing Simon Steward and Ian Videlo	36
<i>Internet and Multimedia</i> The BT Intranet Complete Service Tim Titheridge, Rob Collingridge and Graham Shorrock	43
Intranet-TV—Video Streaming for the World Wide Web Andrew Grace, Richard Jacobs, Jon Cox and Angela Barrow	49
Business Use of Internet Web Sites—Could Do Better! James Callaghan and Audrey Pie	56
Getting the Message, Loud and Clear—Quantifying Call Clarity Simon Broom, Philip Coackley and Phil Sheppard	66
Future Trends in Satellite Communications Russell Silk and Martin Bath	73
SimDS—Modelling the Performance of Future Systems Paul Sellek and David Beaumont	94
<i>Regular Features</i> pcochrane@btlabs	103
Book Reviews	104
Telecom Focus	106
Information	112

Cover Picture
Ariane 5 launch vehicle
(courtesy of Arianespace)

Board of Editors

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

Managing Editor Paul Nichols, B.SC.

Secretary Les Brand, M.ENG., M.B.A., D.I.C.

Treasurer Denis Taylor, F.C.M.A.

Assistant Editor Richard Kemp, B.SC., C.ENG., M.I.E.E. Roger Blake, C.ENG., M.I.MECH.E. Jason Chilton Martin Evans, B.SC., PH.D.

David Greenop, B.SC., M.SC. Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E, M.I.O.M. Peter McKenzie, B.SC., C.ENG., M.I.E.E., M.B.I.M.

Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT. Keith Ward, C.ENG., F.I.E.E., M.I.MGT. David White, M.I.P.D., M.B.I.M., L.I.A.A.

Theme Editors

Structured Process Improvement Keith Beales **Computer Telephony Integration** Doug Chesterman

Telecommunications in the 21st Century Steve Sim **Mobility and Radio Communications** Nigel Wall

Internet and Multimedia Bob Foster

Keith Beales

Structured Process Improvement



*structured process
improvement
works and,
indeed, is
necessary for
survival*

In common with other former monopoly telecommunications providers, BT finds itself subjected to continuing and growing competition. It is just one of an increasingly large number of players, albeit a dominant one in the UK, in a competitive market. Moreover, while BT has a lot of operational experience, it has monopoly legacy not shared by its competitors. Such a situation is not so uncommon in the world of business where fleet-of-foot new ventures can often seriously undermine established and plodding giants.

How then is BT to avoid being undermined in this way? The key is to become more like the new ventures while retaining the advantages of a corporate major player—in other words, become more agile while retaining dependability. Achieving this means, however, a big transformation in both culture, technology and processes. Such transformations have been achieved in other industries and companies when driven by intense competition. In BT this transformation has been underway and gathering pace since privatisation.

In a new themed series of articles for *British Telecommunications Engineering*, entitled 'Structured Process Improvement', theories, principles, tools, methods, techniques along with practical applications are described to give illustration of how elements of this transformation are taking place. This themed series builds upon a previous series, 'Process Management'. However, the emphasis is now upon guiding principles for operations and demonstrating how technology and process development together can bring about complementary improvements in operating costs, improved customer service and increased customer satisfaction.

The first of these articles, 'Lean Operation—The Route to Competitiveness', on p. 2, introduces some of

the principles, methods and techniques adopted by businesses worldwide over the last two decades in response to intense competition. Lean methods, whose genesis owes much to Japan and specifically its automobile industry, have been the subject of much interest worldwide and form a body of recognised best practice. While the original application was to traditional manufacturing industries, increasingly service-based organisations are joining the lean revolution.

Future articles will deal with such topics as customer service measures, access network planning methods and tools, customer service teams, process and resource modelling and agile production. These articles will not only illustrate how the particular projects are contributing to BT's operational process improvement, but also how barriers are being broken down between developers, end users, customers, process developers and operational people. In this way, process improvements are integrated with technology development and can thus be truly said to be structured. Indeed this is a feature of lean methodology when applied to a business as a whole.

Finally there is much evidence from a variety of industries that structured process improvement works and, indeed, is necessary for survival. Automobile manufacturers, aero-engine producers, automotive spare-parts suppliers, supermarkets and many more enterprises have benefited from the application of these ideas. In the increasingly strong competitive environment of the telecommunications market, BT too will need to continue to employ structured process improvement.

Keith Beales

**Engineering Manager,
Network Transport Engineering Centre
BT Systems Engineering**

Lean Operations—The Route to Competitiveness

This article introduces the concept of lean operations. It describes a set of tools and management practices which have been developed within the manufacturing industry. The authors believe that early adoption of such techniques by BT is essential for survival in an increasingly competitive environment. The article describes how teamworking is a prerequisite for quality assurance, how new product development can be improved through concurrent engineering, and how continuous improvement can flourish only in a learning organisation.

Introduction

'The lessons learned by the US automotive industry in the 1970s and 80s ...are relevant to BT today.'

For fifty years the US automotive industry enjoyed rapid growth, dominated by Ford, General Motors and Chrysler, with little competitive threat from foreign manufacturers. In the 1970s, however, the market was undermined by imports from Japanese companies, such as Toyota and Nissan. These competitors, despite their recent entry into the market, were able to produce higher quality cars at lower cost. Moreover, they were able to bring new models to market much more rapidly than their American counterparts. Indeed, it was the Japanese manufacturers' rapid response to the 1973 oil crisis, in producing smaller, more fuel-efficient cars, that enabled them to gain a foothold in the American market. By 1980, imports, mainly from Japan, accounted for over 28 per cent of the US market¹.

The US government's response was to impose import restrictions. The Japanese manufacturers, in turn, responded by setting up manufacturing plants in the US and forming joint ventures with American manufacturers. This provided an opportunity for the US manufacturers to examine the working practices used in Japan which had produced these extraordinary results. Throughout the 1980s, the US manufacturers gradually adopted most of these principles and techniques, and thus ensured their own survival.

The situation of the US automotive manufacturers in the 1970s has many parallels with the competitive

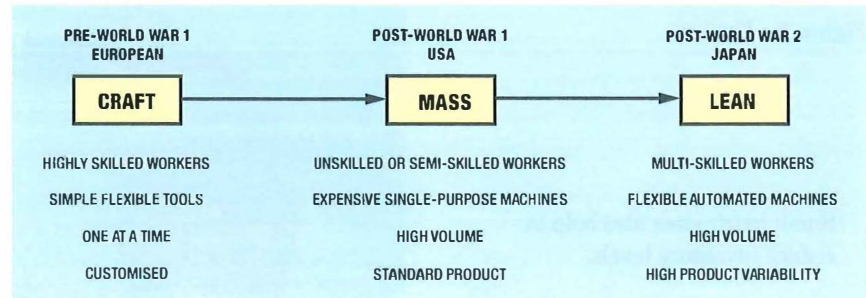
environment in which BT now operates. After many years of near-monopoly, BT is now a player in a fiercely competitive market. For example, OFTEL expects that BT's share of residential exchange lines will fall from 91 per cent at the end of 1996 to 70 per cent by the year 2000². In their book *The Lean Communications Provider*³, Adams and Willetts suggest that communications providers must adapt in a similar manner to automotive manufacturers if they are to survive. The lessons learned by the US automotive industry in the 1970s and 1980s are therefore relevant to BT today.

This article describes some of the lean operations and production management techniques, developed largely by Japanese automotive manufacturers, and more recently adopted by a wide range of other industries throughout the world. It is based on research drawn from published case studies of both manufacturing and service industries. This research work was being carried out under the guidance of the Production Design Team in *NetworkBT* in a joint programme with Systems Engineering to bring world's best practice in production engineering and management to BT. The article indicates what the adoption of these practices might mean for a future lean BT.

Evolution of Manufacturing Production Systems

Lean practice or production is a stage in the evolution of manufacturing techniques. Figure 1 charts, historically, the progress of production methods for the automotive industry. This evolution is mirrored in other

Figure 1 – The evolution of manufacturing techniques



industries, both manufacturing and service. The concept of lean production, as applied in the automotive industry, is described in a landmark study *The Machine That Changed The World*⁴.

Prior to World War 1, automobiles were made using craft methods, usually to individual customer requirements. This craft industry employed highly skilled workers, utilised simple flexible tools and made products one at a time. There was little standardisation. Much of this industry was based in Europe.

After World War 1, Henry Ford, in the USA, consolidated the moves towards mass production that he had begun with the introduction of the Model T in 1908 and the moving assembly line in 1913. The *scientific management* theories of Frederick Taylor brought about the distinction between thinkers (management) and doers (workers) during this time. Mass production methods are characterised by unskilled or semi-skilled workers tending expensive single-purpose machines, producing highly standardised product in huge volumes.

After World War 2, the Japanese began their recovery and rise as an economic power. Japanese industry could not afford the huge capital investment of mass production and, moreover, Japanese engineers regarded it as wasteful of materials and labour. Lean production methods were therefore developed as an alternative. Lean production is characterised by multi-skilled workers, tending flexible automated machines, producing products in high volume with a high degree of variety.

The current trend is to move beyond the lean phase into an *agile* phase. Agility comprises all the characteristics of lean production combined with the following four basic principles:

- products are solutions to customers' individual problems,
- virtual organisations are formed where products are brought to

market in minimum time through internal and external cooperation,

- entrepreneurial approaches are adopted so that organisations thrive on change and uncertainty, and
- knowledge-based organisations are formed which focus on distributed authority supported by information technology.

The Characteristics and Methods of Lean Practice

'Lean operators are not satisfied with being just good enough to beat the competition. They aim for perfection....'

'Lean practice has improved quality whilst reducing costs and improving throughput time, disproving the received wisdom that quality costs more and that faster is dearer.'

Lean operators focus strongly on the elimination of waste and reduction of non-value-add activities. They consider all resources including human effort, space, investment, time and inventory. Lean operators achieve their waste elimination goals by:

- **Targeting perfection** Lean producers are not satisfied with being just good enough to beat the competition. They aim for perfection, although they do not necessarily achieve it. They seek zero defects, zero inventory, continuously declining costs and endless product variety. This latter point is often referred to as *mass customisation*, where products are made in huge quantities but tailored to individual customers or customer groups.

- **Removing rectification or re-work** Lean operators avoid rectification stations or re-work loops believing them to encourage poor quality and thus high cost. A defective product is not passed down the production line; instead it is subjected to the *five whys*[†]. There are no 'white coated' craftsmen, tuning and adjusting the product or the line, on lean production lines. The removal of safety nets, such as re-work loops, is a recurring theme of lean practice.

- **Adopting Team Working** Lean operators are strongly committed to team working. Teams have team leaders, not foremen, and perform all functions, including tool repairs, housekeeping and process improvement. Teams operate continuous improvement and use structured problem solving. Team working and empowerment are considered later in this article.


Specific lean techniques

At the heart of lean production are three interrelated techniques:

- **The small batch** The lean operators realised that there were problems with mass production because components, parts and products were made in huge quantities in anticipation of demand. If there were defectives, then it might be some time before they were discovered. Making things in small batches makes defectives visible earlier. Suppliers know that any defectives they ship will be rapidly discovered.

[†] Toyota has developed a system, known as the *five whys*, in which the question 'why?' is asked until the root cause of a problem is revealed

Figure 2—Kanban



Small batch sizes also help to reduce inventory levels.

- **Just-in time (JIT)** This is the most well-known of the techniques of lean operators in Japan. The idea is that parts get delivered to the work stations when they are needed and in the correct quantity. This sometimes gets mis-applied in the west, when the inventory that would have been held by the manufacturer gets stocked in the supplier's warehouse. In true JIT the lean approach is applied all the way up and down the supply chain. Again the application of JIT reduces safety nets and encourages everyone to make sure that parts are delivered on time with no defects.

- **Kanban** 'Kanban' is the Japanese word for card, and it is the principal mechanism by which JIT operates in a factory. The basic concept is to enact a process only when the following dependent process requests, and not to make to stock. Figure 2 illustrates this. If, for example, process 1 is printed circuit board (PCB) production and process 2 is PCB component assembly, then, when a batch of PCBs (A) is launched into the component assembly process, a card or signal (B) is sent to PCB production to produce another batch. This batch (C) would be delivered to the assembly process to await processing. This arrangement is also known as a *pull* system and is quite distinct from the conventional *push* systems of mass production with their large batches and high levels of finished product stock.

Distinguishing characteristics of lean practice

These are summarised in Table 1. The third column indicates examples of what would characterise a future fully-lean BT. Many current initiatives are, of course, already heading towards greater leanness.

The two key features that distinguish lean practice from all other methods are:

- the maximum number of tasks and responsibilities is transferred to the workers on the line, and
- a system is in place for detecting defects that quickly traces every problem, once discovered, to its ultimate cause.

Lean practice has improved quality while reducing costs and improving throughput time, disproving the received wisdom that quality costs more and that faster is dearer.

So far, the basic principles of lean practice have been described. However, so that lean practice can be effective in an organisation, supporting techniques are required. The remainder of this article describes four such techniques: quality assurance, teamworking and empower-

ment, concurrent engineering and organisational learning.

Quality Assurance

'Taguchi's quality loss function allows the proper fixing of the true cost of poor quality or the cost of failure.'

Lean operators have an overwhelming concern for quality. In his humorous and instructive book³, Frank Price identifies how it is possible to achieve 'right first time' by asking and getting the answers to four simple questions in sequence:

- Can we make it OK? This asks whether our processes are capable of producing the product to the specification the customer requires.
- Are we making it OK? This asks whether our processes are actually producing the product to the

Table 1

Classic Mass	Classic Lean	Lean BT
Assembly line aisles crammed with indirect workers	Narrow assembly line aisles. No indirect workers	Multi-functional and multi-skilled teams
Uneven line distribution	Balanced line—flexible working	Smooth work flows and no 'fire-fighting'
Discarded and defective parts in bins	Defective parts tagged and in 'quality' area	Network and other faults subjected to root cause analysis
Weeks' worth of inventory at work stations	Hours' worth of inventory	Low levels of stores, 'grey van' and network inventory
Lots of rework and rectification	Little or no rework	No repeat faults
Dispirited workforce	Empowered, supported teams	Field teams fully empowered and supported

lean producers have always seen quality as everyone's responsibility, eschewing the idea of inspecting quality into processes and products and instead designing it in from the start

specification and involves process control monitoring.

- Have we made it OK? This asks whether we have met the agreed (with the customer) acceptable quality level (AQL). Defects are accepted because we know perfection, although striven for, is unattainable. AQLs tell us how far we are allowed to miss by.
- Could we make it better? This is addressed by product research and development and process evolution.

The tools to enable these four questions to be answered may be found in statistical process control⁶ (SPC) and related methods. SPC addresses the effect on products of the intrinsic variability of materials and processes. This is distinct from assignable variations due to deliberate changes to the process and material parameters.

Taguchi methods

Genichi Taguchi⁷ extended the ideas behind SPC to develop methods that enable products to be fabricated that are robust to intrinsic process parameter variations. His experimentation methods minimise the number of tests necessary to reach a satisfactorily robust design, even with quite complex processes. Moreover, the process experimentation can be done mostly within self-directed teams. These methods constitute the ways in which the 'Can we make it better?' question may be answered.

Taguchi's quality loss function allows the proper fixing of the true cost of poor quality or the cost of failure. A speculative example is given below where the ideas are applied to the process of repairing a business line. Quality loss has units of relative cost (£s).

For business repair, consider a target to effect a repair within five hours with an 85 per cent success rate. If this target were met, the processes might be thought to be

producing a robust service. However, the customer wants continuous service and, if a repair has to be done, it is wanted quickly. The target is an internal measure with little direct relation to customer satisfaction. In reality, the customer's satisfaction level declines and the quality loss grows with the amount of waiting time, as a smooth function, regardless of the internal target. This is shown in Figure 3.

As the customer waits for repair:

- revenue is lost by the business,
- countermeasures to meet the target five hours incur costs,
- customer compensation kicks in after a certain time, and
- the customer looks at alternative suppliers, incurring potential revenue losses.

The customer incurs direct costs for:

- potential loss of revenue to his/her business,
- providing alternative arrangements, and
- looking for alternative suppliers.

Taguchi postulated that the aggregation of the above, called the *quality loss function*, is quadratic (square law dependence), although other relationships may be used. The quality loss function can be used to set the internal measure or target as a balance between the reduction of quality loss and the cost of any process improvement. The use of SPC, Taguchi methods and other quality tools is the means to bear down on this latter cost.

Finally, lean operators have always seen quality as everyone's responsibility, eschewing the idea of inspecting quality into processes and products and instead designing it in from the start. The lean operators in Japan fully grasped this idea and built up self-directed, empowered and trained teams capable of using these quality tools.

Teamwork and Empowerment

'...there is evidence from a wide range of industries that self-directed teamworkingcan increase employee commitment and job satisfaction, leading to enhanced productivity and quality.'

Businesses have traditionally controlled their operations through manage-

Figure 3 – Taguchi quality loss function

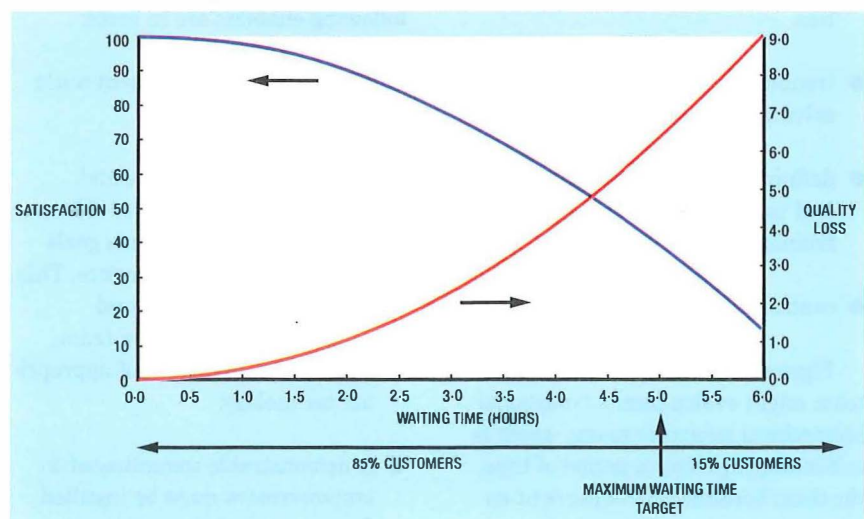
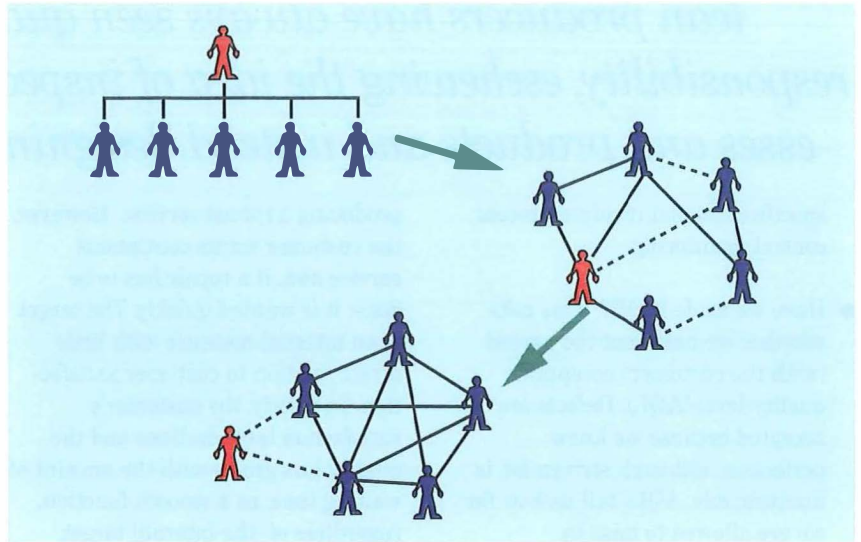


Figure 4—Typical evolution of a self-directed work team



ment hierarchies, with information and authority passed between the various levels of this hierarchy on a need-to-know basis. Organisations which have used these traditional command-and-control methods have been plagued by poor industrial relations, with the employees feeling alienated, bored and unfulfilled. Such feelings are common and understandable in organisations where the people perceive that they are not in control of their own destiny.

Self-directed teamworking has been developed as an alternative people-management system, with the aim of making working life more fulfilling, and thus improving morale. Further drivers for the adoption of new working practices include changing expectations of working life and the realisation that all the employees must play a part in improving quality. Self-directed work teams are characterised by high levels of employee involvement and empowerment in decision making, and full responsibility for achieving agreed results. Managers are fewer in number, and adopt a coaching, rather than a directing, role. The members of a self-directed work team are often multi-skilled, and rotate jobs frequently. Some of the traditional management responsibilities which might be transferred to a self-directed work team include:

- scheduling the team's work and defining individual responsibilities,
- trouble-shooting and problem solving,
- defining the team's training needs and sometimes delivering that training, and
- control over some budgets.

Figure 4 illustrates how a work team might evolve from a traditional hierarchical structure to one which is self-managing. Over a period of time, the team becomes less dependent on a manager and more close-knit.

A number of barriers must be overcome if self-directed teamworking is to be successfully adopted. These include the following:

- a lack of knowledge or expertise may render a person unable to fulfil a particular role within a team;
- operational processes, such as required authority levels, and human resource processes must reflect empowerment and teamworking;
- an outdated corporate culture can make it difficult to introduce radical new work practices; and
- managers may feel threatened, and workers exploited when authority and responsibility are moved down the management chain.

It is therefore essential that the following enablers are in place:

- Training and development must be key priorities.
- There must be accurate and timely communication of information relating to long-term goals and short-term performance. This must be coupled with good communication between team members, making use of appropriate technology.
- A demonstrable commitment to empowerment must be instilled from the top of the organisation,

leading to an environment of mutual trust.

- Clear team objectives must exist, aligned to corporate goals.

BT has begun to introduce self-directed teamworking, both among field engineers and in clerical environments, such as payment handling centres. The initial results are encouraging, and there is evidence from a wide range of industries that self-directed teamworking, when properly implemented, can increase employee commitment and job satisfaction, leading to enhanced productivity and quality. If employees genuinely feel that they are stakeholders in the business, with everyone sharing a common destiny, then they will almost certainly perform to the best of their ability.

As well as tackling issues of employee motivation, teamworking is an effective way of bringing together people from disparate functional areas to contribute to a particular project. This includes management, field engineers, clerical and other functions. The application of this concept to new product development is now considered.

Concurrent Engineering

'All stakeholders in a product or process are involved from the outset, and form a cross-functional team.ensuring that all aspects of the life-cycle have been considered'

Concurrent engineering (CE) is a team-based approach to the develop-

One important tool is quality function deployment, which translates the 'voice of the customer' into the specification of a new product.

ment of a new product and the management of its life cycle. A formal definition⁸ is:

'A systematic approach to the integrated concurrent design of products and related processes, including manufacture and support, which is intended to cause the developers to consider all elements of the product life-cycle from inception through to disposal, including quality, cost, schedule and user requirements.'

All stakeholders in a product or process are involved from the outset, and form a cross-functional team. This team is responsible for specifying the product correctly, ensuring that all aspects of the life cycle have been considered, and avoiding any problems which might occur downstream. The team might include designers, assembly workers, marketing, support staff, suppliers and even customers.

The following examples typify a concurrent engineering environment:

- When a new product is proposed, the operational processes required to deliver the product to the customer are designed concurrently with the product itself.
- Product development is led by the voice of the customer, so that the product is neither under-engineered nor over-engineered.
- Products are designed to minimise the cost and human effort required to support them throughout their life cycle.
- The product is specified correctly from the outset, removing the need for re-work.
- Data relating to the product is freely available to all stakeholders.

Several tools have been developed to support CE. Although they differ

vastly in approach, the consistent aims are to shorten the product introduction process, to give greater customer orientation and to reduce the development and whole-life costs of a product.

One important tool is quality function deployment (QFD), which translates the 'voice of the customer' into the specification of a new product. The primary tool of QFD is the *house of quality* (HoQ), which is shown schematically in Figure 5.

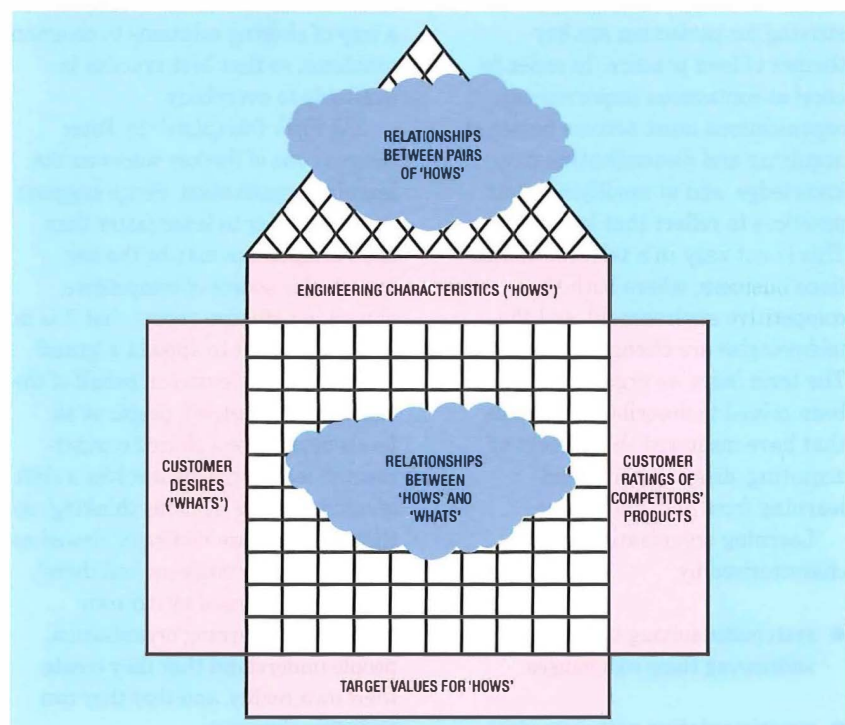
The HoQ is a series of matrices, which guides the developer through the process of capturing the attributes that the customer wishes to see in the product, then using these to generate a list of engineering characteristics. The matrices enable an exact matching between customer requirements and engineering characteristics, so that the product is neither over-specified nor under-specified. They also help to identify any conflicts between pairs of engineering characteristics, and enable the best trade-offs to be made.

Another widely used tool of CE is *product data management* (PDM),

which controls and facilitates the flow of information relating to a product across all disciplines. This ensures that all stakeholders have rapid access to accurate product information. A PDM system is basically a shared database, which comprises the only validated repository of information relating to the product. A PDM database might include the following information:

- specifications and technical drawings of the product,
- bills of materials for producing the product,
- management information relating to the product development project,
- historical data, providing an audit trail through the development project, and
- systems for monitoring, controlling and recording changes to the product.

Figure 5—Schematic representation of the house of quality



The following benefits have been reported for CE⁹:

- product development times reduced by 40–60 per cent,
- manufacturing costs reduced by 30–40 per cent,
- scrap and re-work reduced by 75 per cent, and
- engineering change orders reduced by 50 per cent.

It should be recognised that, while the tools described above, and others, can make a valuable contribution to the implementation of CE, they must be accompanied by a cultural change. This will involve embracing concepts such as teamworking, customer focus and a commitment to quality.

The Learning Organisation

'In order to excel at continuous improvement, organisations must become better at acquiring and disseminating new knowledge

Continuous improvement and striving for perfection are key themes of lean practice. In order to excel at continuous improvement, organisations must become better at acquiring and disseminating new knowledge, and at modifying their practices to reflect that knowledge. This is not easy in a telecommunications business, where both the competitive environment and the technologies are changing rapidly. The term *learning organisation* has been coined to describe businesses that have mastered the concept of acquiring, disseminating and learning from new knowledge.

Learning organisations are characterised by:

- systematic solving of problems by addressing their root causes,
- experimentation with new ideas,

- learning from both their own and others' past experiences,
- transferring knowledge quickly and efficiently throughout the organisation, and
- nurturing the capabilities of each individual within the organisation.

An earlier article¹⁰ in the *Journal* describes how one division in BT is shifting towards becoming a learning organisation. The transition involved is largely a cultural change; however, information technology systems can play an important enabling role in making knowledge more accessible. Corporate intranets are rapidly becoming repositories for information which needs to be accessed throughout an organisation. BT is developing tools to cope with the information overload that can sometimes result¹¹. These include text summarisation software¹² and the Jasper knowledge management tool¹³, which allows people with similar interests to share information over the Internet, and jointly build up a body of knowledge. Databases of problems, such as the one developed by Rover¹⁴ can provide a way of sharing solutions to common problems, so that best practice is available to everybody.

*The Fifth Discipline*¹⁵ by Peter Senge is one of the key works on the learning organisation. Senge suggests that the ability to learn faster than one's competitors may be the only sustainable source of competitive advantage. He also argues that it is no longer sufficient to appoint a 'grand strategist', who learns on behalf of the organisation; instead, people at all levels must have a shared commitment to learning. This involves a shift of mind towards 'systems thinking', so that problems are no longer viewed as being caused by someone 'out there', but as being caused by our own actions. In a learning organisation, people understand that they create their own reality, and that they can therefore change it.

Conclusion

'.....recent history tells us that success in an increasingly competitive market depends upon lean thinking throughout a business.'

This article has introduced a number of ideas that go to make up the concept of leanness, but methods and techniques on their own do not make an organisation lean. The lean organisation pursues unremittingly the elimination of waste and non-value-add activity throughout the organisation. This quest is never-ending and engages everyone through the use of continuous improvement methods. Quality and its improvement become part of everyday business, and teams are trained and empowered to use the appropriate quality tools. Quality is designed into products and services from the beginning and not inspected in later.

The design and development of products and services takes place in teams comprising all parties that have a stake in the product or service, including operational people. No longer are products and services 'thrown over the wall' to operations, but developed concurrently with them. Indeed the lean organisation has no walls; the waste elimination programme will have demolished them!

The lean organisation will learn from its mistakes and its successes. Knowledge will be disseminated in a structured and timely way. There will be almost no privileged information, and teams will share their knowledge.

This may seem idealistic, but recent history tells us that success in an increasingly competitive market depends upon lean thinking throughout a business. The US automotive industry, when faced with a strong competitive threat, discovered that the lean approach was being practised by its competitors and had to adapt very quickly. The telecommunications industry is on a similar track and those that adopt the lean

philosophy today will be tomorrow's winners.

Acknowledgements

The authors gratefully acknowledge the contributions of John Dunn and other members of BT's Production Design Team. The authors also wish to thank David Clift, Mark Gilbert and their other colleagues at BT Laboratories for their support in compiling this article.

References

- 1 FLEMING, KATHY. General Motors Corporation. Darden Graduate School of Business Administration, University of Virginia, case study. 1986, UVA-BP-0267.
- 2 GOODMARK, CHRIS. Competition hits BT harder than expected. *The Independent*, 2 Dec. 97, p. 23.
- 3 ADAMS, E. K. and WILLETTS, K. J. The Lean Communications Provider: Surviving the Shakeout Through Service Management Excellence. 1996, McGraw-Hill.
- 4 WOMACK, J. P.; JONES, D. T.; and ROOS, D. The Machine that Changed the World. 1990, Macmillan.
- 5 PRICE, FRANK. Right First Time. 1984, Gower.
- 6 MONTGOMERY, D. C. Introduction to Statistical Quality Control. 1985, Wiley.
- 7 TAGUCHI, G. System of Experimental Design; Engineering Methods to Optimise Quality and Minimise Costs. 1987, UNIPUB/Kraus International.
- 8 CREESE, R. C.; and Moore, L. T. Cost Modelling for Concurrent Engineering. *Cost Engineering*, 1990, **32**(6) p. 23.
- 9 MARKOWITZ, M. C. Concurrent Engineering Journey Starts with the First Step. *Electrical Design News*, 1991, **36**(15), p. 110.
- 10 COPE, MICK. Organisational Learning: Early Days in a Change Process. *Br. Telecommun. Eng.*, Jan. 1997, **15**, p. 356.
- 11 DAVIES, JOHN; and COCHRANE, RALPH. Knowledge Discovery and Delivery. *Br. Telecommun. Eng.*, Apr. 1998, **17** (this issue).
- 12 ProSum WWW-based text summariser: http://transend.labs.bt.com/prosum/on_line/index.html
- 13 Jasper knowledge management tool: <http://www.labs.bt.com/ourwork/jasper/>
- 14 MURRAY, KATE; and WISEHOFER, HILDEGARD. Rover Learning Business. University of Derby, European Case Clearing House case study. 1994, 494-006-1.
- 15 SENGE, PETER, M. The Fifth Discipline: the Art and Practice of the Learning Organization. 1990, Century Business.

Biographies



Ray Hooper
BT Networks and
Systems

Ray Hooper is a member of the Planning, Plant and Pilots Unit of BT Systems Engineering. He joined BT Laboratories in 1970 after graduating with a first class honours degree in Electrical Engineering from Middlesex University. He worked initially on optical-fibre communications systems, specialising in the design and development of receivers and transmitters. During this time he contributed to the original digital optical-fibre communi-

cations field trials. In 1986 he was seconded to BT&D Technologies (now part of Hewlett Packard), a joint venture company set up by BT and DuPont to manufacture optoelectronic components. He was responsible for the development and design of the early optical receiver products and managed the integrated circuit development programme. He returned to BT Laboratories at the end of 1990 to work on optical components for the access network. Subsequently he managed a team supporting new technology pilots and trials including TPON, IMS and EXTACT. He now leads a team developing network planning rules and tools, as well as contributing to the research on production management.



Carole Jones
BT Networks and
Systems

Carole Jones is a member of the Planning, Plant and Pilots Unit within BT Systems Engineering, based at BT Laboratories. She graduated from Durham University in 1984 with a first class honours degree in Applied Physics and Chemistry, and obtained a Ph.D. from the same university in 1987 for work on dielectric thin-films. After a year's postdoctoral research at Cambridge University, Carole joined BT in 1988. She initially worked on the development of non-linear optical devices, and subsequently on silicon micro-engineering. During this time she was involved with the development of several novel silicon micro-engineered devices, including a low-cost fibre-pigtailed laser package and a high-speed transceiver. In 1996 Carole began working on process management, where her main interest has been a study of production management and its application to BT.

Computer Telephony Integration and Java

Computer telephony integration is now widely used, both on the desktop to enhance office automation, and in call centres, where fast call control and associated database operations are essential. The newly emerging capabilities of the Java programming language can bring great benefits to these areas, and extend the functionality of the agent's desktop by using the Internet.

Introduction

Previous articles¹⁻⁴ in this series have introduced the theme of computer telephony integration (CTI) and described its various capabilities, both on the desk-top and integrated with networks and PBXs. So far, the articles have looked mainly at the hardware involved. This article describes the use of the Java programming language in enhancing the provision of CTI, with specific illustrations of its use in office automation on the desktop, and in centres where multiple users require flexibility in call control, with perhaps integration with the Internet.


The article starts with a brief summary of CTI, looks at the origins and nature of Java, and then shows how Java can be usefully deployed in a CTI environment, illustrated with examples.

CTI Flavours

Linking a desktop computer application to a telephone line via a local

physical interface is frequently referred to as *first-party CTI* or *desktop CTI*. Only a single local line is controlled, for incoming and/or outgoing calls; it is not possible to monitor or control any other line. First-party CTI is of particular use in an office or small business environment, or at home. Figure 1 shows various ways of configuring the hardware, using PC telephony cards, telephone cards or a telephony interface unit such as BT's Callscape⁴. Linking a customer's data to a telephone number can speed up dialling and record retrieval. For example, if a selection of customers' accounts is displayed on the screen, the caller can automatically call each customer just by selecting the account number, letting the PC find the associated telephone number and then dialling it. Incoming calls can be associated with customers' records (by using the Caller Display service), and all calls can be logged. This improves personal productivity, minimises dialling errors, and simplifies the telephony access.

Figure 1 – First-party CTI



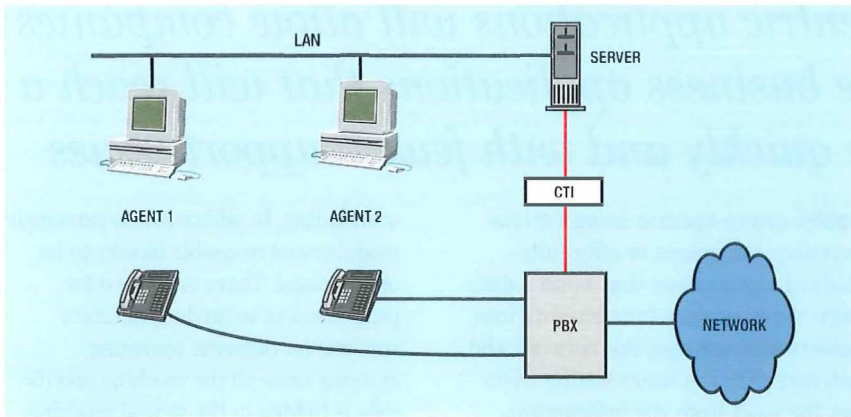


Figure 2—Third-party CTI

The second type of CTI is commonly referred to as *third-party CTI*, (or *PBX-based integration* when the controlled element is a PBX), as a desktop computer application is now able to monitor and control *any* telephone in a particular domain, via a common CTI link between a central server and the PBX or network (Figure 2). This method is used in environments catering for multiple users, such as call centres. Here inbound calls are directed to agents, depending on their skills and availability, and outbound calls can be placed automatically, and, if required, only connected to agents when the calls are answered.

Origins of Java

Java is a software technology, a computer language developed by Sun Microsystems⁵. It began life as a programming language for consumer electronic devices such as microwave ovens and toasters. This resulted in a set of unique design requirements, among which was the ability to run on different hardware and to be extremely reliable. The existing programming languages were inefficient for that purpose, so James Gosling started designing a new one, initially known as Oak, in 1990. When the World Wide Web appeared on the Internet in 1993, the design team thought that the new language would be ideal for the Internet, given the fact that it is platform independent and the Internet consists of heterogeneous systems. Java and applets (see definition later) were born. A browser developed by Sun, called *HotJava*, was the first browser to support applets and it provided an impressive demonstration of Java's capabilities.

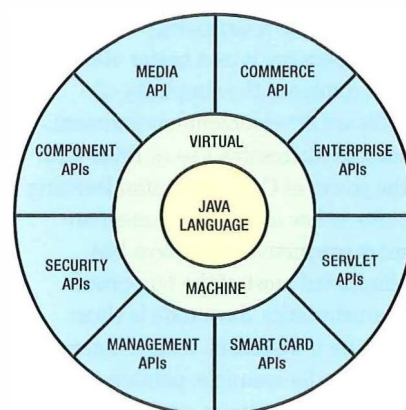
Java's Architecture and Main Characteristics

Java combines a set of technologies into a three-layer architecture (Figure 3):

- an interpreted, object-oriented language which is translated before execution into platform-independent *byte-codes*;
- a virtual machine/run-time interpreter which translates byte-codes into machine code on the fly; and
- class libraries and application programming interfaces (APIs) enabling developers to interface third-party systems with the virtual machine and vice versa.

Sun's summary of Java is a collection of buzzwords such as 'simple, object-oriented, distributed, interpreted, robust, secure, architecture neutral, portable, high-performance, multithreaded, and dynamic'. First of all the language is *simple* and *elegant*. Java's creator, James Gosling, describes it as 'C++ without

Figure 3—Java's structure



guns, knives or clubs'. Its designers removed a number of 'dangerous' features, or ones leading to bad programming practices, allowing developers to create more robust applications.

Even more importantly, Java was designed from the ground up to allow for secure execution of code across a network, even when the source of that code was untrusted and possibly malicious. Java byte-codes are downloaded from the server machine and run on the client. The security risks associated with this are dealt with by the Java Runtime System which looks at the incoming byte-codes and verifies that the code is safe. This is called the *verification process*. Another layer of security is the *sandbox model*. Untrusted code is placed in a sandbox, a secure area, where it can run with the potential for damage minimised, under the control of a security manager.

Furthermore, Java was designed for cross-platform use in compiled binary form. Since this is frankly impossible across processor architectures, Java is compiled to an intermediate byte-code which is interpreted on the fly by the Java interpreter. Thus to port Java programs to a new platform all that is needed is a port of the interpreter and a few native code libraries which form the Java Virtual Machine (JVM).

Applets and Applications

Two main types of programs can be written using Java: applications and applets.

- *Applications* are standalone programs which require the Java interpreter to run. For example, a Java version of a word processing package is an application. Java applications are analogous to C++ applications; they run independently of any Web browser. The only difference is that a Java application is portable and runs the same on PCs, Macs and Unix

Java-based network-centric applications will allow companies to build and customise business applications that will reach a wider audience more quickly and with fewer support issues

stations, whereas a C++ application needs to be recompiled for each specific platform.

- *Applets* are small applications embedded on Web pages. An applet requires a runtime environment; for example, a browser or the Sun Java Development Kit's appletviewer. In order to view or use applets, the browser must be Java-enabled; these browsers have embedded within their programs *Java runtime interpreters*, which allow the applet to communicate with the operating system (OS). Applets are more security conscious than applications. The runtime Security Manager is concerned that code downloaded from a network can inadvertently or maliciously harm or steal data, or interfere with other processes running on the client. This is why it imposes many restrictions on an applet's behaviour; for example, no access to the local file system. *Servlets* are similar to applets. Servlets are to servers what applets are to browsers.

What is Java Good For?

One thing fuelling the incredible publicity around Java is its ability to make the Internet come alive. Java's real power is its built-in control of Web pages, which are not static hypertext mark-up language (HTML) text any more but are transformed to dynamic multimedia presentations. The key point here is that applets are good not only for creating animated pictures and fancy scrolling effects, but they are also written in a real programming language allowing developers to incorporate logic and functionality alongside graphical operations. Now, with Java, Web-servers can transmit much more than static data to users. For example, if a client requests spreadsheet data from a Web-server, the requested information is transmitted with a Java wrapper applet that will run the spreadsheet. More than one

applet can co-operate using various scripting languages to offer fully-fledged applications that bundle data with the necessary functionality into objects sent through the network and interact with the users letting them get the most from the information. Some common uses are on-the-fly charting of raw data, interactive queries, or filters on large data sets. This means that Java makes network-centric (as opposed to desktop-centric) computing a reality.

Java's unique capability to work with other Internet-based technologies makes it possible to create distributed object-oriented applications that exist and function independently of any particular desktop architecture. Consequently it offers the base for a more flexible and dynamic computing model. Up to now, most client/server applications are developed on the desktop-centric model. This makes them tightly coupled to specific client architectures, and installing and maintaining them is quite complex. These desktop-centric applications have to be pre-installed or 'hardwired' onto client machines. If you want to make any changes, the new software must be installed again on a client-by-client basis. Java applets don't have to be pre-installed. They install themselves just in time, on the fly, and uninstall themselves when they are no longer needed. This makes desktop application upgrade an obsolete concept.

Why Use Java?

There are many reasons justifying the use of the Java technology. First of all, developers see it as a better alternative that combines the simplicity of software development environments such as Microsoft's Visual Basic with the power of C++. The initial learning curve is low and programmers are more productive since Java has eliminated most of the bug-prone characteristics. Java code is clean because it dispenses with 'dangerous' features; for example, pointers, multiple inheritance, and operator

overloading. In addition, it is extremely modular and re-usable thanks to its object model. There is no need for programmers to produce different versions for different operating systems since all the machine specific code is hidden in the virtual machine, allowing a Java application or applet to run on any computer where the JVM has been implemented. This is the origin of Sun's 'Write Once, Run Anywhere' mantra. Java provides the means for writing sophisticated, reliable, robust, and platform-independent applications.

In addition to that, Java reinvents the way applications are distributed to clients and executed. As explained in the previous section, Java can drastically reduce the effort needed to install, support, and upgrade applications. Thus, Java-based network-centric applications will allow companies to build and customise business applications that will reach a wider audience more quickly and with fewer support issues.

After client/server computing that revolutionised the business applications development and use, and the Internet that revolutionised the distribution of information, Java came along combining both of them in a way that had a great impact on the distributed computing model. Java is evolving into a component architecture (Java Beans) comparable with Microsoft's ActiveX/DCOM. Furthermore, CORBA—which stands for *Common Object Request Broker Architecture*—and Java's marriage is a very promising step towards fully-flexible component-based distributed systems.

Java and CTI

CTI applications use a signalling protocol to control telephony from a computer. Current computer telephony protocols can be classed as being: proprietary in nature; operating system specific (for example, Microsoft's Telephony Application Programming Interface (TAPI)); or relatively complex (for example,

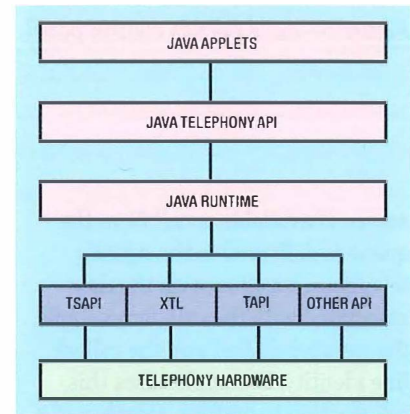
Figure 4—Sun JTAPI architecture

Computer Supported Telephony Applications (CSTA) and Versit). An opportunity exists for a simple, flexible and platform neutral application programming interface (API) to take advantage of existing computer telephony protocols. A recent offering in the commercial market for such an API is JTAPI⁶, the Java Telephony Application Programming Interface. The key elements of the JTAPI architecture (Figure 4) are a telephony application conforming to the JTAPI specification, a Java runtime environment, and a JTAPI-compliant telephony sub-system.

JTAPI (which is in fact a set of APIs created by Sun who are hoping to turn it into an international standard) provides a core set of features (for example, making, answering and clearing calls) and a number of

Java's principle of 'Write Once, Run Anywhere' lends itself to new types of communications; that is, two-way real-time multimedia. Microsoft already provides a development kit containing components of its NetMeeting and Whiteboard. This kit can form the basis of a Windows 95 conferencing application, useful for corporate intranets using Windows 95. For the wider Internet community there can be no guarantee of operating-system conformity. Java-based cross-platform customer-communications applications could extend CTI beyond just the conventional telephony (audio) by incorporating data sharing and even video components. A Java-based data-sharing application is described below.

Java's cross-platform capability will enable the user to have more



simplify back-end database integration and to provide appropriate telephony call control. Ultimately Java-based crossware, capable of coordinating business information with telephony statistics, will help to shape the call centre. Relating the type of business and success level of every call, from the perspective of the overall commercial enterprise, the organisation can appraise pricing, marketing and (call centre agent) training.

Examples

Several different examples of the use of Java in the CTI world are briefly described here, ranging from large call centres to an individual office worker's desk.

Call centres

With reduced maintenance costs and ease of distribution, Java is suited to the multi-seat call centre environment. It has the flexibility to cater for rapidly changing roles, allowing the call centre dynamically to shape its resources to meet the current load. Applications can be downloaded quickly and reliably to meet the need of the call centre at any instant in time.

A demonstration version of a downloadable Java applet controlling a call centre agent position has been developed. The applet controls the agent's telephone and provides access to a database supplying customer information and current sales promotions. The telephony API of the Java applet is based on the CSTA standard. The applet allows the agent to logon to an automatic call distribution (ACD) queue of incoming calls. When the agent is ready to accept a call, he/she can indicate this to the telephony

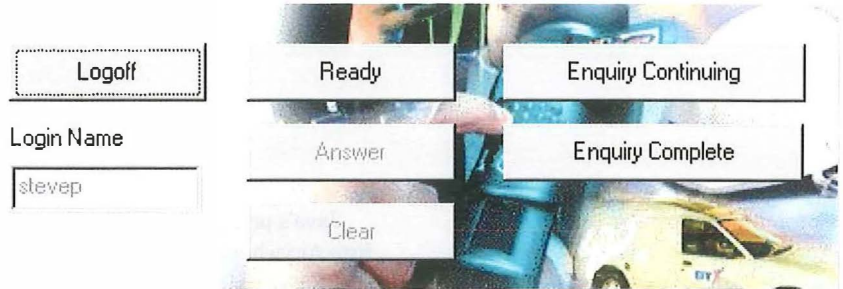
extension APIs that cover more complex features such as call centres and media-stream access. JTAPI supports both first and third party CTI. Being object-oriented, it uses a variety of objects to describe the various aspects of the telephony layer. For example, a *Provider* object manages the control of a PBX, a *Call* object represents the telephone call, and an *Address* object represents the telephone number (note that this is the logical endpoint to a call, which is not the same as the physical endpoint (terminal)). Other objects represent connections and terminals. Implementations of JTAPI can be created on many existing CTI platforms such as Sun Microsystems's SunXTL, Microsoft's TAPI, Novell and Lucent's Telephony Services API (TSAPI), and IBM's CallPath, as well as on any other independent vendor's proprietary hardware. JTAPI is a complement to, as opposed to a replacement for, existing call control protocols, hence its API is dependent on these protocols.

control over CTI office productivity applications. The ability to instruct the telephony servers delivering calls to the user, independent of current location, will assist both the users and their office or work colleagues. A Java applet controlling the users' telephony configurations has been developed and is described below.

CTI incorporating Java components adds personal productivity features to the corporate intranet. Java can extend the centrally administered company telephone directory from a static database lookup to an intelligent call initiator service. An enhanced dialling application is described below.

Potentially, any existing application could be Java enabled; for example, an agent's desk in a call centre. Within the call centre environment, the prospect of Java's cross-platform uniformity offers an escape route from proprietary messaging and transaction applications. The goal of the Java call-centre application is to

Figure 5—Java agent's control panel



server. If available, a call from the queue is delivered to the agent's telephone, together with the data associated with the call; for example, the number dialled and the caller's line identity. The applet uses this information to perform searches of a central customer database. The results of the search populate the agent's screen with all the customers associated with the caller's number, as several customers may use the same telephone line to call in; for example, different members of the same family or friends.

The call centre can stream customer enquiries by advertising a unique telephone number for each individual service or request handled. This allows the call centre quickly to match the skills of the agents to the customer enquiry. The applet uses the number-dialled information to perform a search of a central product database. The results of the search populate the agent's screen with product data, and perhaps the latest promotion or offer. The agent can then answer the call and speak to the customer. The agent will need to confirm the caller's identity and the nature of the enquiry. When the transaction is complete the agent disconnects the telephony call by using the 'Clear Call' button on the screen, and then completes any forms or adds notes before taking a new call. The applet allows the databases to be updated or the current data discarded. Figure 5 shows a typical screenshot of the demonstrator's control panel.

The BT call centre at Warrington ran a trial last year handling telemarketing calls using Java workstations in place of PCs. Each agent application was run entirely in Java, talking back to a central server. Noticeable improvements in speed were found; for example, the start-up time for getting brand new machines running, and the overall time taken from power on to being able to take the first call. The machines themselves have high reliability—the only moving parts are a fan and the on/off

switch, and later versions even remove the fan! Software security was enhanced as the risk of virus infection was virtually eliminated. The price was also seen to be a big differentiator, with the Java environment cutting the costs by 30–60%.

Data sharing

In the above example, the communication between the agent and the customer is only via basic telephony. This can be considerably enhanced by allowing the customer and agent to see and edit the same information at either end of the link. One way to do this is to use the Internet to transport the data. Figure 6 shows how an agent and customer might be connected for telephony and data. A data link is set up using the Internet, and a telephony link using either a standard public switched telephone network (PSTN) telephone or an Internet telephone. At the moment, residential customers using the PSTN are generally unable to connect to the Internet and make telephone calls at the same time, but businesses are readily capable of doing this. However, using BT's new Home Highway digital communications service, customers' existing home telephone lines can be upgraded to carry data (via the

Internet) and voice simultaneously, making this sort of transaction readily available to the mass market. Multimedia kiosks may in future also offer consumers this type of information and communications terminals.

Imagine a customer browsing the Web and coming across a company of interest. The customer then decides to find out more through personal contact. Using Java, first of all a 'Call-Me' button would be added to the Web page which would instruct the company's call centre telephone system to ring the customer (using normal PSTN telephony or an Internet-based telephone) and then make a connection to an appropriate agent. When the call is established, Java would then be used to control the pages visible to the two parties, using for example the hot keys shown on the left in Figure 7.

This customer might need some help navigating round the company's Web site, in which case the agent would select pages to show, and these would be visible at both ends of the call. Either end can choose a new page (perhaps showing specific product data), and both the customer and the agent would see the same information. Having been guided to the right product, he/she might then want to fill in an application form. This too would

Figure 6—Schematic layout of agent/customer connections

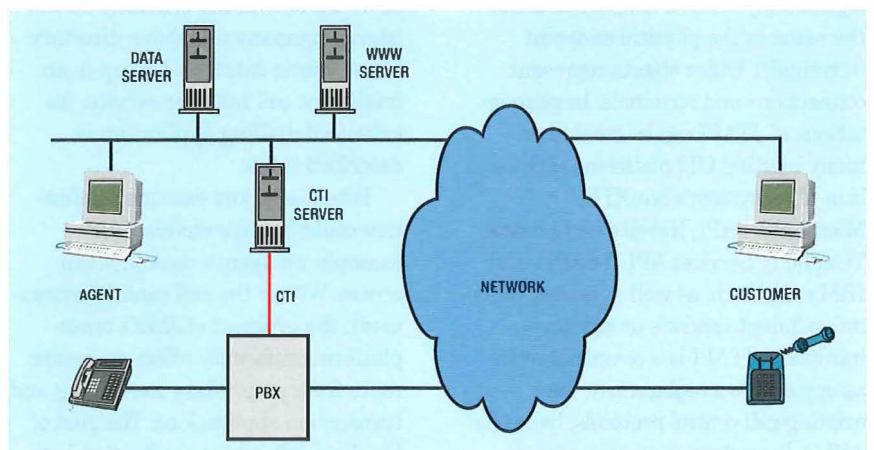


Figure 7—Sample view of elementary agent/customer data sharing

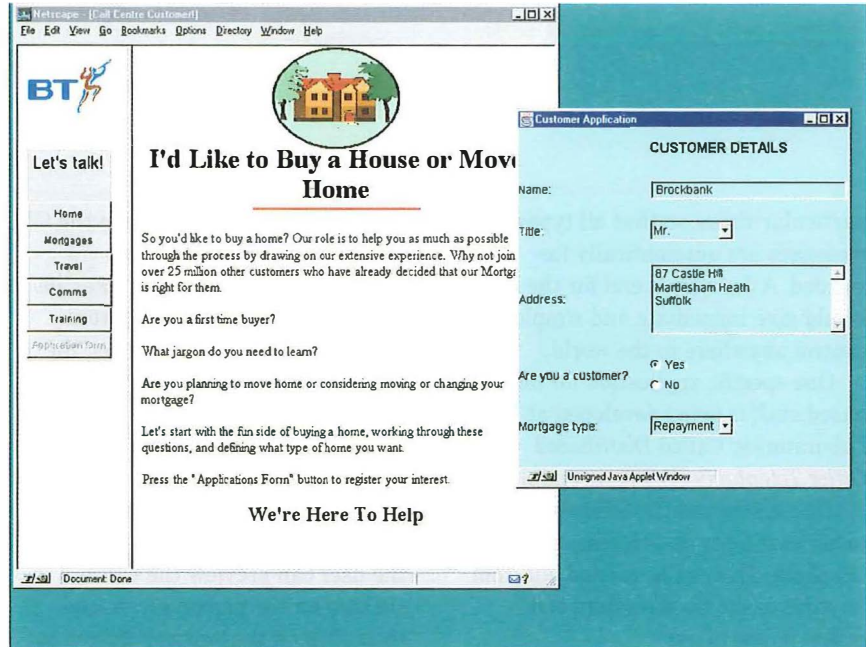
be done via Java, so that the text and controls are visible and editable from both ends. This closely simulates real life, with the customer and agent sitting across the 'table' from one another, talking to each other and sharing documents. Figure 7 shows a simple screen shot of what a customer and agent might see, with a partly completed customer details form. By enhancing the telephony channel to carry video telephony, a very realistic approach to customer/agent interaction is obtainable. Because all the transactions are carried out via the central data server, each page of information can be customised as it is delivered, and a full history of agent/customer interaction retained.

Dialling enhancements

The task of finding a telephone number and dialling it is automated in the Web browser based Directory Dialler application. It searches a remote corporate telephone directory database and initiates a telephone call using the search results. The user can customise the call initiation phase, for example, the prefix digits and the dialler application to use, in order of preference. Dialler applications include a CTI-enabled PBX, BT Callscape 100 and BT VC8000 ISDN videophone.

When the PBX is selected as the preferred dialler application, a mapping is made between the IP address of the computer searching the database to the telephone associated with that computer. Third-party CTI is used to initiate the call from this telephone to the party identified from the database search.

The other dialler application options reside on the local computer hardware, and are controlled in first-party CTI mode. Java and Netscape Navigator plug-ins (local helper software used by the browser) are used to achieve inter-application communication; that is, to bridge between the Web browser and the dialler application software. Normally the Java applet's security manager limits its operation to the 'sandbox', as described above. It is



possible to enhance the features offered by Java by compromising the 'sandbox', allowing the Java code to run on the hardware without the need for the virtual machine. Netscape Navigator plug-ins are examples of this usage, providing greater features at the expense of security.

In the Directory Dialler, user configuration options are managed using a combination of Java and JavaScript. Netscape's LiveConnect allows communication between these elements, with Java as part of the plugin component. The plug-in initiates the call using dialler applications residing on the local hardware.

Figure 8—Directory dialler

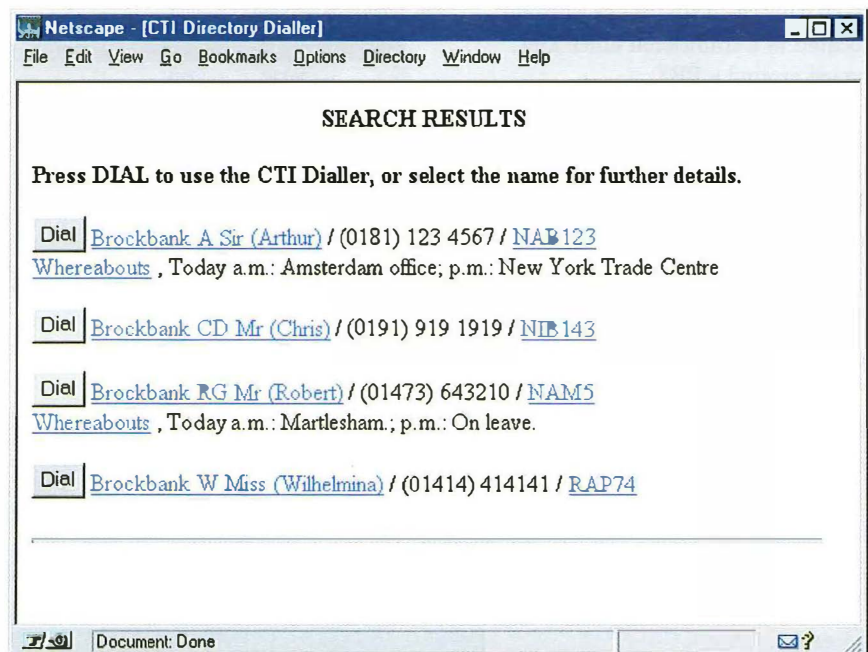


Figure 8 shows the results of a directory search—clicking the appropriate dial button sets up the call.

Call control

A major problem facing users who are mobile is how to handle incoming calls, both one's own and one's colleagues. To be able to control call routing is very desirable. For example, BT Magic⁷ is a relatively new service that gives access to fax, e-mail and voice messages, regardless of one's location. This can be done either by retrieving messages remotely (using an account number and PIN) or by telling the software package where one can be reached at

particular times, so that all types of messages are automatically forwarded. A Java front end for the user would give immediate and simple call control anywhere in the world.

One specific application for office-based staff is being developed at BT Laboratories. Called *Distributed Office Telephony* (DOT), it enables staff to define virtual workgroups and to have visibility of each other's telephony. This can be carried out from anywhere, and the Java front end makes it easy to use.

DOT addresses modern working trends such as teleworking or working when travelling, and connects colleagues together by creating a virtual office environment. When launching the application the user sees the typical screen shown in Figure 9. This is the main DOT window with the buttons allowing users to control the application's features and a collection of pictures corresponding to a user's workgroup.

The telephony status of the group members—Free, Busy or Ringing—is indicated by a label beneath each picture. The DOT user is included in his/her virtual group by default, as the first picture. The user now has control of his/her own telephony and the telephony of the work-group via the client's graphical user interface (GUI), just as if that team were co-located in a traditional office and based around a PBX.

When the user has an incoming call, the telephony status under her icon will change from 'Free' to 'Ringing' and an information window will pop up containing the CLI and the caller's details (Figure 10).

A database is used to map the CLI to the identity of the caller.

If the user decides to answer, he/she clicks on the 'Answer' button. When the conversation is over, the user ends the call by clicking on the 'End Call' button.

Users can answer telephone calls for other group members in a similar way. If a colleague has an incoming call which he/she does not answer, the user can preview the CLI by clicking on the person's icon, and then click on the 'Answer' button to answer it, so no important call is lost. If the user does not answer his/her call, a memo message window is shown on the screen when he/she returns, containing the caller's name and number, with a 'Return Call' button to allow quick call return.

A very powerful feature of DOT is location transparency. Contacting other team members when you don't know where they are is time consuming. DOT simplifies team communication—just click on a colleague's picture and make a telephone call regardless of where the called person currently is. Defining the members of your personal workgroup is easy and can be done on the fly. You simply select them from a list of user names

Figure 10—Incoming call window, ready to answer





Figure 9—DOT's appearance on the screen



that are preconfigured by the system administrator.

DOT offers a variety of additional features besides those previously mentioned. It is currently used and being evaluated by BT's CTI group on a daily basis and is a continuously evolving application.

Where Now?

As has been shown, Java has already made its presence felt in the CTI arena, and will continue to do so. It caters well for certain operations, and taking into account the predicted cost savings and the continued maturing of NCs, the role of Java on the desktop will flourish.

Java is not restricted to voice communications alone. Video and data are equally controllable by Java, and as Internet access is made more widely available by improvements to the local loop, Java and CTI will continue to go hand in hand. However, it must be borne in mind that Java is not the only way of providing this functionality, and each case must be evaluated on its merits.

Telephony via the Internet is a growing area of interest, as this allows both data and voice to be carried over the same network, with obvious benefits. Java is one clear way of controlling it.

Large call centres can expect to become significantly cheaper, and more future-proofed against changes in hardware technology.

For certain types of business transactions and enquiries, the call centre can use Java to develop cross platform applications facilitating the by-pass of its own call centre. Customers will be able to use these applications to access and/or control 'their' data without the need to contact the call centre first. Early moves into call centre by-pass are already apparent; for example, high street banks' personal PC banking facilities and parcel companies' package tracking applications. With further enhancements, a button on the Web page will allow the cus-

tomers to contact the call centre. The advantage to the call centre is the potential reduction in the number of enquiries handled by staff and the possible concentration on specific or detailed enquiries. Staff may then be able to review customer data in order to advise the customer of additional services that may be of benefit to them. Java would add support over a greater number of customer operating systems. It would also allow for multimedia content and enhance the call centre's image.

Acknowledgements

All trademarks acknowledged.


References

- 1 CATCHPOLE, ANDREW; CROOK, GARY; and CHESTERMAN, DOUG. Introduction to Computer Telephony Integration. *Br. Telecommun. Eng.*, July 1995, **14**, p. 98.
- 2 CATCHPOLE, ANDY. Computer Telephony Integration—The Meridian Norstar. *Br. Telecommun. Eng.*, Oct. 1995, **14**, p. 209.
- 3 JOHNSON, PHILIP; CATCHPOLE, ANDY; and BOOTON, LAURIE. Computer Telephony Integration—The Meridian 1 PBX. *Br. Telecommun. Eng.*, July 1996, **15**, p. 150.
- 4 HILLSON, GRAHAM; HARDCASTLE, CHRIS; and ALLINGTON, MARC. Callscape—Computer Telephony Integration for the Small Business. *Br. Telecommun. Eng.*, Jan. 1997, **15**, p. 293.
- 5 Sun Microsystems Java Web site <http://java.sun.com/>
- 6 Sun Microsystems JTAPI Web site <http://java.sun.com/products/jtapi/>
- 7 HARRIS, STEVE; MUSSON, CHRIS; and MARTIN, PAUL. Advanced Personal Mobility Services. *Br. Telecommun. Eng.*, July 1997, **16**, p. 100.

Glossary

ACD Automatic call distribution
API Application programming interface
BTL BT Laboratories
CLI Calling line identity
CTI Computer telephony integration
DCOM Distributed Common Object Model
DOT Distributed Office Telephony
GUI Graphic user interface
HTML Hypertext mark-up language
JavaOS Java operating system
JavaScript Netscape's cross-platform, object-based scripting language for client and server applications
JTAPI Java Telephony API
JVM Java Virtual Machine
LAN Local area network
LiveConnect Netscape's integration of JavaScript, Java, and plug-ins
NC Network computer
PBX Private branch exchange
PC Personal computer
PSTN Public switched telephone network
TAPI Telephony application programming interface
TSAPI Telephony services application programming interface

Biographies



Robert Brockbank
BT Networks and Systems

Robert Brockbank joined BT in 1975, after graduating from Oxford University with a First in Engineering Science, to work on analogue and then digital submarine cable systems. He was involved in laying and testing the world's first optical-fibre submarine cable system and was a member of the team that was awarded the Queen's Award for Technology in 1990 for designing and building the ultra-high reliability optical receivers for international submarine cable links. From 1990 to 1992 he was seconded to BT&D, the

joint venture set up by BT and DuPont for the manufacture of optical communication components, to develop production line test equipment. He moved back to BT Laboratories and joined his present team, the CTI Group, where he specialises in the applications, architectures and capabilities of computer telephony integration. He has filed patents on CTI, is a Chartered Engineer and a member of the IEE.



Stephen Peirce
BT Networks and Systems

Stephen Peirce graduated from the University of Strathclyde in 1985, gaining a

B.Sc. (Honours) in Electronic and Electrical Engineering. After a year at Ferranti Defence Systems in Edinburgh working on radar, he joined the Digital Networks Division at BT Laboratories. He was initially involved with the study and development of narrowband ISDN signalling protocols, then working on broadband ISDN and control networks. In 1993 he joined the CTI Group where he has been involved in several call centre projects including the development of a load tester and in the control of telephony, utilising a variety of programming languages and network topologies.



Zoi Amanatidou
BT Networks and Systems

Zoi Amanatidou holds a degree in Physics and joined the CTI Group in BT Laboratories in November 1996, after completing her M.Sc. in IT (Software and Systems) with Distinction from Glasgow University. She first worked on developing a TAPI driver in C++, and is currently involved in the Java version of the DOT application. She is working towards IEE membership.

Mobile Multimedia Applications

BT intends to exploit the current high demand for mobility by offering a new range of personalised mobile multimedia products and services. This article outlines how BT has collaborated with other organisations to develop the technical expertise that will be required if these products and services are to be commercialised.

Introduction

Mobile multimedia is an emerging technology that has become possible as mobile and fixed multimedia technologies converge. The 'vision' for mobile multimedia is that information and resulting knowledge should be available, regardless of physical location, in a form appropriate to the user's requirements and terminal capabilities. Mobile multimedia applications are commercially viable today but future mobile network enhancements will enable operators to offer a broader, more profitable, range of applications and services.

In this article both early and medium-term applications that are exploitable within the next five years are described. Middleware requirements are discussed, a mobile value-added information service is detailed and a typical customer solution involving positioning technology is presented.

Technical Developments Supporting Future Applications

Several technical developments will enable more advanced mobile multimedia applications to be offered to users in the future:

- *Network media control* A key feature of future mobile multimedia applications will be the intelligence required to support them in the difficult mobile environment. Factors such as battery life, lighting, security, driving, movement all conspire to hinder the easy use of mobile applications. Intelligent functions will help in many ways; for

example, by minimising bandwidth requirements and thus saving power and spectrum. Another example would be to convert one form of media into a more usable one; for example, while driving, e-mail could be heard via a text-to-speech converter. Information could be routed based on geographical position if the supporting intelligence knew the location of the terminal.

- *Mobile multimedia terminals* In the last couple of years several truly mobile multimedia terminals have appeared in the marketplace, and some of these devices, such as the Nokia 9000 Communicator, show how basic cellphones can be enhanced to allow a limited set of other media to be used. *In-vehicle* multimedia units are common in Japan and The MIT Media Laboratory predict that body-worn devices will become widely utilised.
- *Network approaches* Two network approaches are being adopted to support mobile multimedia applications. Firstly, the universal service approach adopted by the third-generation cellular systems initiatives (UMTS /IMT2000 etc.) and, secondly, a hybrid approach whereby enhanced GSM (with wider band logical channels and packet working) will be augmented by other networks such as Hiperlan, DECT, DAB, TETRA, Paging, etc. For this latter case, it is possible to support seamless media services but network embedded services, such as calling line identity, cannot be supported in a transparent manner.

A pragmatic approach to mobile multimedia research is to focus on key applications and demonstrate them with the objective of collecting opinions from potential customers and partners. In a similar way to BT's fixed multimedia initiatives, mobile services are likely to involve a chain of business partners who must cooperate. Collaborative trials are a good way to facilitate the formation of these chains and should enable BT to downstream its research activities.

Mobile Middleware to Support Multimedia Applications (ACTS¹ OnTheMove² Project)

In order to support mobile multimedia applications successfully, an architecture which supports the needs of mobile users is necessary: resilience to disconnection, roaming between different networks and operators, movement between geographical locations, variable bandwidth and support of a wide range of terminals.

Mobile application support environment (MASE)

The *OnTheMove* project provides mobility support for mobile multimedia applications by developing mobile middleware—a mobile application support environment (MASE) that hides the complexities of underlying wireless communication networks from applications. The MASE makes

its services available to applications through a mobile applications programming interface (API) (Figure 1). A prototype mobile multimedia business information system has been developed both to evaluate the performance of the MASE, and to gain user feedback on personalised mobile-aware applications.

The project is strongly user focused. Early in the project, focus groups and brainstorming sessions provided preliminary requirements for the MASE and mobile API. A user-based experiment also helped to validate and improve the MASE and mobile aware application design. MASE versions and mobile business information systems are tightly coupled and have been developed in two phases. The mobile business information system is designed to test and demonstrate, via field trials, the functionality of each MASE and mobile API release.

The role of *OnTheMove's* MASE middleware is to enable personalised access to multimedia applications irrespective of mobile hardware, user location and wireless network characteristics. The complexity of underlying networks is hidden from applications, and differences in these networks appear to users as changes in quality of service (QoS). Profiling is used to store and retrieve information about user preferences, and terminal and network characteristics.

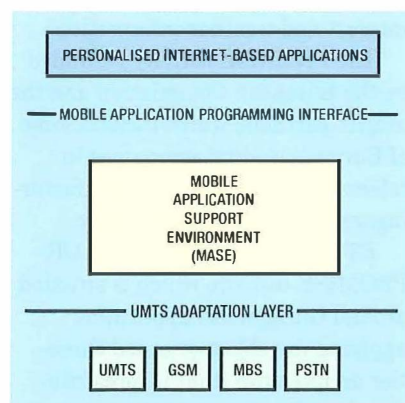
User profiles allow the user to influence the behaviour of the MASE by specifying their communication preferences which include cost, quality, time, and security. *Terminal profiles* provide a means of determining the characteristics of the mobile device which has an important impact on the way that information is presented to the application. The *network profile* stores information about network characteristics. Functionality provided by the MASE includes: session management, QoS trading, multimedia conversion, transaction support, caching, location determination, directory services, alerting, user management, and

agent support. The MASE has been developed as a distributed system, with functionality residing on both the mobile device and mobility gateways. Selected MASE functionality is also exported to information servers in order to provide mobility support to applications on the server side. The mobile API allows multimedia applications to communicate QoS and mobility specific requirements to the underlying MASE subsystems and receive information about currently available services. Maximum portability is offered by describing the mobile API in Object Management Group Interface Definition Language (IDL)[†]. The mobile API is a superset of a number of recognised APIs, in order to provide backwards compatibility for legacy applications. To maximise platform independence, the Mobile-API interfaces are made available in Java*. The Mobile-API will be a key factor in enabling the exploitation of MASE functionality.

Application field trials to test the MASE

The project includes two major field trials at five locations in Europe, including BT Laboratories, and one in Singapore. The first of these trials has already taken place, testing the first version of the MASE, its mobile API and the mobile business information system applications. The applications were a news telegram system providing reliable messaging and information filtering, a stock portfolio application providing alerting, a mobile e-mail client, and a map-based city guide application. The city guide

Figure 1—MASE and Mobile-API



[†] A complete description of the Object Management Group's Interface Definition Language can be found in chapter 3 of the *Common Object Request Broker: Architecture and Specification* document, revision 2.0, July 1995.

* Java is a trademark of Sun Microsystems Inc.

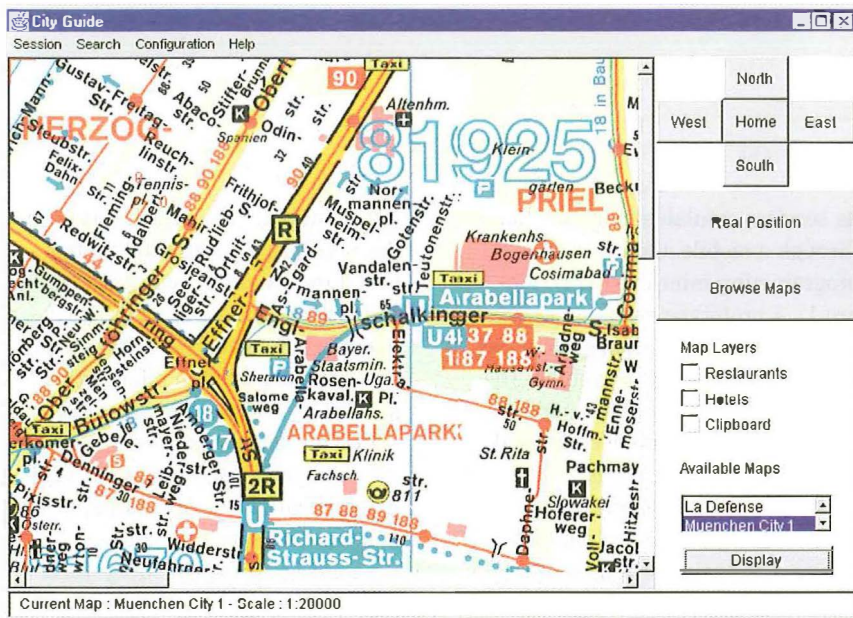


Figure 2 – City guide application

application (Figure 2) has been developed in Java and uses the MASE location manager to determine the location of the mobile device using WaveLAN cellid when indoors or global positioning system (GPS) outdoors. Maps are downloaded according to the user location, and local information can be selectively displayed on the maps. The city guide is currently being enhanced to incorporate delivery of maps which are appropriate to the network bandwidth available, and improved user preference and information search facilities.

The field trial feedback, from 154 trialists, was very favourable and is being used to refine and develop the second version of the MASE, mobile API and the applications.

The OnTheMove project has developed prototype mobile middleware to support mobile multimedia applications which will be widely used when third-generation mobile networks emerge. By specifying a mobile API, an open interface can be provided and the approach validated.

Mobile Value-Added Services

Increasingly customers are approaching telecommunications and mobile operators for more than just pure voice and data connectivity—many customers now want managed solutions that focus on their business environment. Consequently, many telecommunications operators are now moving up the


value chain and, in addition to performing the role of 'bit carrier', are acting as a value-added service provider (VASP) (see Figure 3).

In the context of this article the VASP is the organisation responsible for providing customers with personalised services based on information derived from a content provider. The VASP also performs administrative functions such as billing and user authentication. To be successful in an increasingly competitive marketplace means that VASPs must be able to deploy rapidly personalised information services that can be easily updated in the future.

Three-tier architecture to support value-added services

A development methodology that enables the VASP's business domain to be modelled is extremely attractive to organisations. This has led BT Laboratories to adopt object-oriented analysis and design techniques and embrace the business object concepts defined by the Object Management Group; that is, a business object is a representation of an entity in the business domain, including at least its business name and definition, attributes, behaviour, relationship

Figure 3 – Value chain for offering value-added services



and constraints. To provide the required flexibility the application has effectively been separated into three tiers: presentation, business logic and data access. This approach offers a number of advantages:

- the VASP focuses on the business logic (that is, the representation of the enterprise);
- business logic is reusable and is not dependent on specific user interfaces or databases (it is insulated from the presentation and data access components of the application);
- services can easily and quickly be updated; and
- a varied range of user interfaces can be supported.

Implementation (EU Telematics for Transport Project—PROMISE)

The PROMISE project partners have deployed trial systems, in six European cities, that offer reliable personalised travel and traffic information, that is position dependent, to travellers. Travellers are able to access the PROMISE services via cellular-based terminals that incorporate a Web browser, including a portable communicator and an in-car device. Based on market research and user-group initiatives the services to be offered at the test sites include trip planning, on-trip guidance, traffic and public transport information, yellow pages, points of interest and weather information³.

The PROMISE activity is funded by the European Commission and the project partners, which include some of Europe's leading companies in telecommunications, car manufacturing and information technology.

BT acts as the VASP at the UK PROMISE test site which is situated around Glasgow. An application applying the object-oriented three-tier architecture (that is, presentation, business logic and data access),

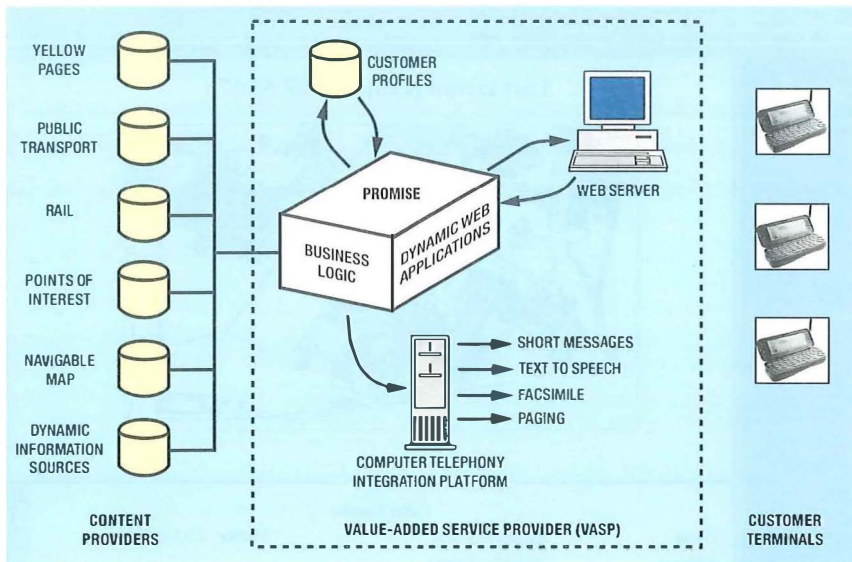


Figure 4 – UK PROMISE system architecture

environment to implement the required VASP functionality. WebObjects is extremely powerful and provides the foundation for the rapid development of Web applications. Furthermore, it allows developers to concentrate on the creation of the relevant business logic.

The adopted technical approach has proved well suited to this application and has enabled a very flexible VASP architecture to be deployed in Scotland.

Mobile Social Alarm Project (MoSA)

The pressing need for a personal safety device linked to a mobile telephone was first identified at a Care and Communications in the Community seminar organised by the BT Age and Disability Unit. Various support agencies for people with special needs claimed that the potential market for such a safety device was enormous, especially if positional information could be provided when an alarm call was made. Subsequently, development work was carried out at BT Laboratories to produce a concept demonstrator of a mobile social alarm (MoSA) service, incorporating a mobile terminal that interfaces to a service centre. The feasibility of a limited field trial using the demonstrator terminal design was accepted and a field trial is planned for the summer/autumn of 1998.

Development approach

The MoSA software developers adopted a flexible approach to system implementation—effectively *component modules*, which can be reused in other systems, were integrated together to provide the required functionality. The components used include:

- mobile voice channel (analogue or digital),
- mobile data telemetry (dual-tone multifrequency (DTMF)),

that focuses on the VASP business logic, has been deployed (see Figure 4). This system configuration is a hybrid intranet solution that allows the content of legacy databases to be personalised and made available dynamically to people who are using mobile Web browsers.

Personalisation is offered to users at the Glasgow test site in a number of ways:

- Each customer is provided with a unique home page that is linked to their user profile; this allows people easily to access relevant travel information.
- Users can specify routes that are of interest and define parameters associated with the route, such as the time of day for which information on the route would be of interest.
- The UK PROMISE architecture encompasses an alert service which allows the VASP to offer real time information to customers; that is, users are alerted of events that will affect their travel plans. The communications bearer used to send alert messages (that is, short message system (SMS), facsimile, e-mail, etc.) is defined by the user and is stored as an entry in the VASP's customer profile database.

Approximately 100 people are participating in the UK PROMISE trial which commenced in March 1998.

Adopted development environment

The object-oriented development methodology adopted by the BT Laboratories developers (the Fusion Methodology developed at Hewlett-Packard's Laboratories) has proved to be extremely flexible and well suited to the development of the VASP architecture.

Fusion is a step-by-step process that leads a development team from an initial requirements document through to the implementation of an object-oriented software system. During the process, various models are constructed, each providing a view of the software system that highlights certain issues and encodes a well defined set of decisions.

The method distinguishes three stages⁴:

- *analysis*—which produces a declarative specification of what the system does;
- *design*—which produces an abstract object-oriented model of how the system realises the required behaviour; and
- *implementation*—which encodes the design in a programming language.

The PROMISE system utilises intranet technology and browsers as the basic user interface. As a consequence of this implementation decision, BT utilised Apple's WebObjects software development

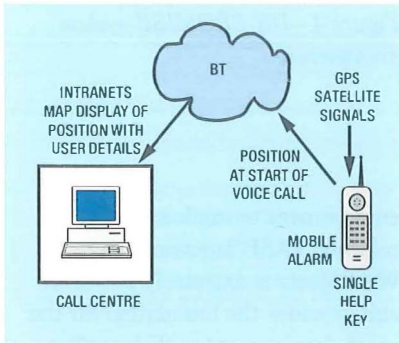


Figure 5—Schematic of the MoSA system

- second-part computer telephony integration (CTI),
- text-to-speech announcements,
- geographical information system (map information with BT's Map-Web interface),
- intranet with World Wide Web (WWW) using server push,
- mobile terminal with GPS satellite receiver.

MoSA functionality

The system consists of two distinct parts, a mobile terminal and a call centre (see Figure 5). When in use, the mobile terminal monitors the position of the user, via a GPS receiver, and buffers this in memory. If help is required, the user presses the ALERT button which automatically initiates a mobile telephone call to the service centre. At the start of the speech call, the last known position of the user is sent as a string of tones, which are decoded at the centre, and the caller's position is used to centre a map displayed on the service operator's screen. Further relevant information can also be displayed so that the operator can give the appropriate verbal assistance to the caller (see Figure 6).

MoSA Mobile Terminal

Terminals have been developed for both the first-generation (analogue TACS) and second generation (digital GSM) cellular networks.

In the initial concept demonstrator an analogue mobile telephone was integrated with a GPS satellite navigation system and a single-chip microprocessor (see Figure 7). Positions calculated by the GPS

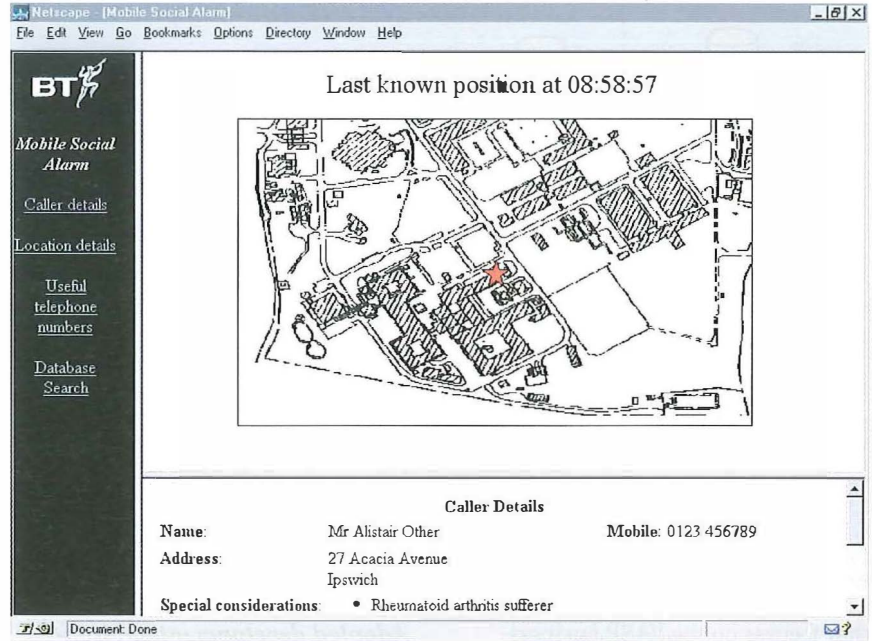


Figure 6—Example of Web page format for the service centre

receiver, which has a base accuracy of 25–100 m (dependent on the GPS system configuration), are transferred to the processor as a latitude/longitude position every second. Only when the ALERT button is pressed and the call initiated is the last known valid position with the mobile identification converted into DTMF tones for transmission over the speech channel. The speech channel is then left open for the caller to seek assistance as required.

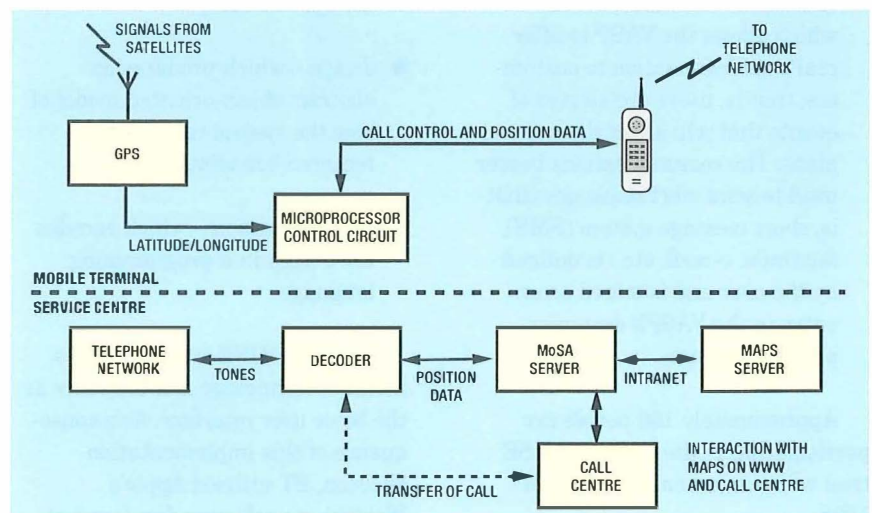
Analogue cellular telephony was chosen for use in the initial demonstrator for the following reasons:

- network independence is possible and, if required, units can be made

that will interface to the public switched telephone network (PSTN),

- a single speech call was preferred to a data call followed by a speech call,
- reliable transmission of audio coupled DTMF tones over a GSM speech path was not found to be guaranteed,
- using GSM, positions sent as a SMS message cannot be guaranteed to be delivered to the call centre at the same time as the speech call,
- minor modifications to the mobile were relatively straightforward.

Figure 7—MoSA system components



However, use of a simple interface between 'standard' GPS receivers and current-generation GSM telephones which enable full remote control facilities via a serial data port (RS232) or proprietary telephone bus provides the following advantages:

- no radio type approval is required,
- the flexibility of a distributed body worn approach is possible,
- the interface is easier to specify and downstreaming is therefore more straightforward, and
- GSM cellular technology will be in use well after the closure of the analogue service.

MoSA service centre

The MoSA service centre will be a dedicated specialist call centre that interfaces to a MoSA server. For the demonstration system, the display of the required map on the service operator's screen will be facilitated by an Internet WWW page requested from the maps server (refer to Figures 6 and 7). Decoding of the DTMF tones is carried out by a separate server and upon transfer of the relevant positional data to the MoSA server the speech call is transferred to the required service operator.

Full computer telephony integration (CTI) functionality for the linking of the display with the incoming speech call will be incorporated into a working system.

Field trial

The field trial planned for later this year will examine all issues associated with running a MoSA type service; mobile and service centre trialists alike will be asked to complete questionnaires and attend focus group meetings. An external company has agreed to be involved in the trial as the service centre operator; the MoSA system will be linked in with the company's extensive development service centre suite.

Potential mobile trialists are currently being identified.

As the purpose of the trial is to assess whether MoSA technology should be offered commercially, the trial report will form an important part of the business case to proceed to market.

Further developments

The MoSA system relies on reasonably accurate location information being available from the GPS module. However, the limitations of GPS (that is, satellite systems do not work indoors or in tunnels, positional accuracy is degraded in city centres where reflections of buildings can lead to multipath problems etc.) have been well documented. In view of the demand for position dependent services, BT is funding collaborative research into a GSM-based positioning system at the University of Technology, Sydney, and at the BT Laboratories. Early results are encouraging, giving an accuracy of 20–100 m, dependent on cell configuration, but further research and development is necessary to overcome multi-path and other problems. If this type of positioning is successful it should overcome the inherent city and building problems associated with GPS.

Conclusion

Participation in collaborative trials of mobile multimedia applications is enabling BT group to develop relationships with other complementary organisations and the technical expertise that will be required if this technology area is to be exploited commercially:

- The focus in the PROMISE project has been to use a flexible object-oriented three-tier architecture (that is, presentation, business logic and data access) that integrates easily with third-party databases and which allows BT to take the role of value-added

service provider of highly personalised dynamic information services.

- The MoSA project is an example of a solution designed to meet a specific customer need. It focuses on the role positioning will play in the delivery of mobile information. A variety of existing building blocks were integrated including mobile voice, mobile data telemetry, CTI, intranet, WWW using server push and GPS-enabled mobile terminals.
- In contrast, the OnTheMove project assumes that a third-generation network that offers a broad range of services (that is, resilience to disconnection, variable bandwidth, QoS trading, session management, multimedia conversion etc.) will exist, and then focuses on the mobile middleware and application support environment. Several typical new multimedia applications were constructed to demonstrate the type of applications that will be available in the future. These are now helping BT shape its plans for the future support of mobile multimedia applications.

Provision of seamless access to personalised mobile multimedia applications will involve a high level of network media control. Current research and development projects are resulting in a broad range of mobility-related functionality which in the future will be integrated. This work lays the foundations for the technology that will support a multimedia future on both fixed and mobile networks.

References

- 1 ACTS homepage: <http://www.infowin.org/ACTS/>
- 2 OnTheMove WWW homepage: <http://www.sics.se/~onthemove>

3 RANDALL, P.; SCHUTT, T.; GEEN, P.; and OJALA, T. PROMISE—A Personal Mobile Traveller and Traffic Information Service. 4th World Congress on Intelligent Transport Systems, 21–24 Oct. 1997, ICC Berlin, Germany.

4 MALAN, R.; LETSINGER, R.; and COLEMAN, D. Object-Oriented Development at Work, Fusion in the Real World. Hewlett-Packard Professional Books, 1996, ISBN 0-13-243148-3.

Acknowledgements

The project work described in this article could not have been completed without assistance from a large number of people at the BT Laboratories. Their technical expertise and help is greatly appreciated.

The OnTheMove project is sponsored by the European Commission in the ACTS programme. Project partners are Bonnier Information Systems AB (Sweden), British Telecommunications plc (UK), Burda New Media GmbH (Germany), Deutsche Telekom MobilNet GmbH (Germany), Ericsson Eurolab Deutschland GmbH (Germany), Ericsson Radio AB (Sweden), IBM (France), Iona Ltd. (Ireland), Royal Institute of Technology (Sweden), RWTH Aachen (Germany), Siemens AG (Germany), Sony (Germany), Swedish Institute of Computer Science (Sweden), Tecsi (France), and University of Singapore Centre for Wireless Communications (Singapore).

Biographies



Don Golding
BT Networks and
Systems

Don Golding joined the GPO as an Apprentice in 1964 and worked on exchange construction duties in North London until he transferred to the Post Office Research Department in 1971. Most of his time in research has been spent investigating different aspects of mobile communications; in particular, mobile applications for use with intelligent transportation systems (ITS). For the past five years, he has worked on position dependent systems, linking satellite navigation to mobile applications for a European Union research project to provide a navigation aid for visually impaired people (MoBIC) and, more recently, for the MoSA project.



Julie Harmer
BT Networks and
Systems

Julie Harmer joined BT in 1988 and worked at Kensington Computer Centre in London. After several years in communications software support, she moved to BT Laboratories in 1994. During the last four years she has been researching mobile communications; in particular, server technologies to support mobile applications. This has involved work in a number of EU research projects including the first phase of PROMISE, providing travel and transport information to mobile users and, recently, OnTheMove, which is researching mobile middleware.



Robin Mannings
BT Networks and
Systems

Robin Mannings graduated in Electrical and Electronic Engineering from the University of Wales in 1976 where he was sponsored by Philips. He then joined Philips and worked in the advanced systems laboratory of the mobile radio division. In 1980, he moved to the University of Bath to pursue research in mobile radio systems. In 1989, he joined BT Laboratories to work on future mobile applications and subsequently lead the team researching mobile telematics and multimedia. He currently manages the interface between BT Laboratories and Business Division's Corporate Clients Logistics Sector.



Paul Randall
BT Networks and
Systems

Paul Randall joined the Post Office as a Trainee Technician Apprentice (TTA) in 1981. In 1988 he joined the BT Laboratories after graduating with a B.Sc. (Hons.) degree in Electronic Engineering from the University of Manchester Institute of Science and Technology (UMIST). In 1991, while at BT, he was awarded an M.Sc. degree in Telecommunications and Information Systems from Essex University. Initially he helped BT's zones to deploy radio technology in the local access network. Since 1994, he has been working on future mobile applications and currently leads the Mobile Applications Group at BT Laboratories. He leads BT's involvement in the PROMISE project.

John Davies and Ralph Cochrane

Knowledge Discovery and Delivery

The World Wide Web (WWW) is seeing phenomenal growth, but the task of finding particular information is becoming increasingly difficult. Many search tools are now inadequate for users' needs. This article looks at how a combination of agent software and webcasting technologies can aid both the discovery of information and its delivery.

Introduction

Phenomenal growth has been a feature of the World Wide Web (WWW) since its inception in the early 1990s. Today, more than 60 million pages of information are available and the potential of the WWW to support electronic commerce and real-time services is starting to be exploited. Powerful search engines are available to assist people to find the information they need. However, the WWW is in danger of becoming a victim of its own success: every day millions of people trawl the network for information using any one of a dozen or more different search tools. But whether they find what they are looking for sometimes depends not only on their skill, but also on their luck. Studies indicate that many existing search engines are not now meeting the needs of users¹.

Recently, we have seen increasing interest in the use of Internet technology within organisations and the emergence of intranets². More and more information is being stored in intranet networks and this has led to the realisation that these intranets are valuable repositories of corporate knowledge. Organisations are increasingly realising the importance of using knowledge and information residing in their networks for competitive advantage. It is forecast that annual spending on corporate intranets will top \$2.1 billion in the UK in 1998; for western Europe the total is expected to be over \$7.7 billion (Inteco Corporation, 1997). Key to the usefulness and viability of both the public Internet and corporate intranets, however, will be the ability to manage the information and

provide users with the information they want, when they want it.

This article looks at the current state of information delivery and discovery before moving on to describe innovative BT solutions to the problems of information overload.

Current Technology: Some Limitations

Several problems are encountered with current information management technology. Many centre on difficulties in finding relevant information. Simple keyword queries often yield too many results of variable quality. It is usual to receive several hundred hits (documents matching the user's query) even on relatively narrow queries. Research also shows that users do not find it easy to frame the query needed to return the information they require¹.

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

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

It has long been realised that the recall of text retrieval systems can be quite low. One study³ showed that less than 20% of the relevant documents were retrieved from a database containing roughly 350 000 pages of text, using IBM's STAIRS full-text retrieval system. Interestingly, the subjects involved in the trial believed that they were retrieving a much higher fraction of the relevant documents. The study also

Figure 1—Traditional pull and push models of information distribution

claims that evaluations done on small databases do not scale linearly to large databases. The search effort required to obtain the same recall level increases faster than the size of the database. This has obvious implications for the very large information sources now becoming available on the Internet.

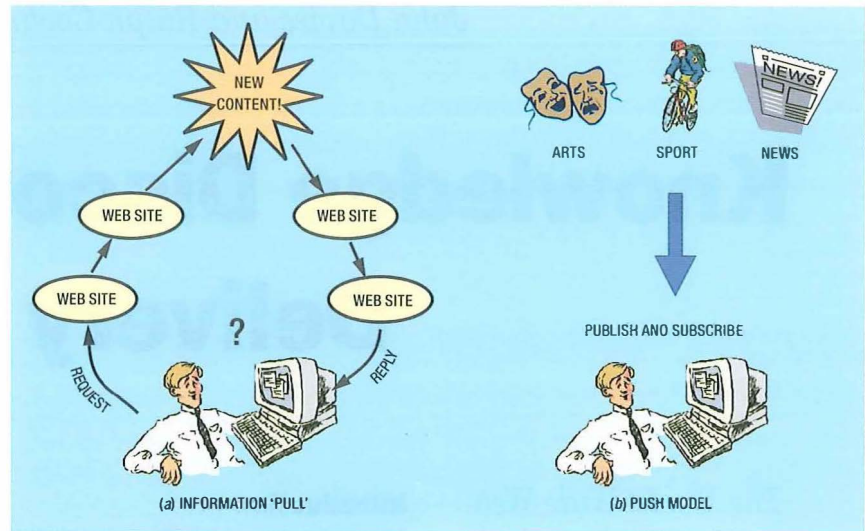
Precision (P) measures how relevant the retrieved documents are to the query:

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

The performance of current systems is disappointing even when used on a single database. When a search on multiple networked information sources is tried, several additional effects come into play, including repeated documents, multiple slight variations of essentially the same information and the problem of information granularity—current systems have no way of differentiating between detailed documents and summaries or overviews.

In addition to these problems, organising information once it has been found may not be straightforward. Most current WWW browsers (Netscape Navigator, Internet Explorer, etc.) incorporate a bookmark facility to provide some means of storing pages of interest to the user. Typically, this is done by allowing the user to create a hierarchical menu of names associated with particular universal resource locators (URLs). While this menu facility is useful, it quickly becomes unwieldy as the number of bookmarks grows.

Essentially, the representation provided is not rich enough to allow us to capture all we would like about the information stored: the user can only provide a string naming the page. As well as the fact that useful meta-information such as the date of access of the page is lost, a single phrase (the name) may not be enough to index accurately a page in all contexts.



Consider, as a simple example, information about the use of knowledge-based systems (KBS) in information retrieval of pharmacological data: in different contexts, it may be KBS, information retrieval or pharmacology which are of interest. Unless a name is carefully chosen to mention all three aspects, the information will be missed in one or more of its useful contexts. This problem is analogous to the problem of finding files containing desired information in a computer file system⁴, although in most such systems one can, for example, sort files by creation date.

Lastly, the sharing of information is not supported in most current systems. In the corporate context, when one user finds information that is important or relevant, often other users also need to know; for example, other members of a project team. Ideally, this kind of information sharing should be automated using software that knows the interests and role of each user.

Given the increasing size and availability of information sources and the limitations of current systems discussed above, novel techniques will be required for information retrieval in the future. These techniques will need to be:

- appropriate for both textual and non-textual information,
- able proactively to deliver information to users,
- able to scale up to deal with increasingly large information stores,

- capable of extracting information from documents at appropriate levels of detail, and
- able to model users' interests and able to find and share relevant information.

Push and Pull Models of Information Access

Until relatively recently, information has been obtained from the WWW via a request-and-reply mechanism, or *pull* model (see Figure 1(a)): the client initiates a request for information and the server replies. Intranets originally reduced the number of passive e-mails employees received, in turn reducing the information load on early adopters. However, the increase in the wealth of information and knowledge on corporate intranets is rendering pull as a delivery method increasingly ineffective. This is particularly true where information changes rapidly (for example, share price information) or where time is critical, as in the distribution of revisions to industrial safety documents or sales figures.

On the intranet and Internet alike, one of the most important requirements is for the relationship between content providers and end users to be mutually controlled⁵. Under the current 'request and reply' paradigm of the WWW, this relationship is very much a hit-and-miss affair. Webmasters rely on their sites providing useful information in a visually attractive format that keeps users coming back time and time again, even on the intranet. For the

The common thread is the need to reach large numbers of people with either a notification of new information and where to get it or the information itself.

user, determining where a useful piece of information is and whether it has changed is a far from satisfactory experience. *Push* aims to provide a solution by sending information to users only when it changes and only when they have subscribed to receive it. This model is known as *publish and subscribe* (see Figure 1(b)). A user requests a flow of information from a particular content provider or service with some element of basic personalisation. This is usually in the form of tick boxes. Content also follows the television set of metaphors with words like *channels*, meaning the set of content from a particular provider, and *narrowcasting*, which is the technique of broadcasting a subset of information to particular users. For example, on a major sports channel users may be able to choose from news on different sports ranging from soccer to ice hockey by ticking the relevant boxes. The semi-personalised information that they receive is typical narrowcasting which has some tailoring to users' preferences but is not fully personalised⁶.

Push technology can be thought of simply as a mechanism for delivery, a way of targeting users with pertinent information. Just as importantly it offers a way of providing feedback and hence a hysteresis loop which can be used to tailor content for subscribers. This model has obvious advantages in areas where information needs are well known and easily defined and where information may be changing rapidly. However, most push vendors today actually employ an automated pull of information from the server based on known user preferences. This approach has its roots in what was until recently the unicast-orientated Internet. However, there is a danger of excess network traffic unless such services are properly managed. Proxy servers can be placed strategically to reduce the amount of traffic duplication through the same content being retrieved by all users. Intelligent scheduling can

also be employed to conduct large file transfers and information updates at appropriate time intervals or at times when the network is quiet.

Applications that utilise push delivery are not confined to unicast networks; satellite distribution is a typical example of the dissemination of information, as it becomes available in broadcast form to a large numbers of users. There is a plethora of applications for push including delivery of software updates, news and content aggregation, event notification and file updates. It can be applied to any application where information needs to be updated. Although push is already part of basic browser software such as Communicator from Netscape and Internet Explorer 4.0 from Microsoft, it will become an even more integral part of the operating system, particularly with the arrival of Windows '98, for data delivery. Push is an important component of a knowledge-management solution. However, there is room for 'push' and 'pull' approaches in a managed enterprise environment. With all webcasting services, including push technology, video streaming and data distribution, the quality of the content is crucial. Just as successful intranets require a predefined information strategy, so the same is true for push.

True push

As the audience increases and new types of time-dependent push applications are introduced, such as event notification, it becomes apparent that the network-inefficient polling-based mechanisms for pushed information do not scale sufficiently to address future requirements. Nor can they handle the need for immediate dissemination of information. In a world billed as 'the information society', not only is the thirst for the latest information ever increasing, but the lifetime of a piece of information also is becoming shorter. Whereas in days gone by, managerial decisions could take days or even weeks to filter through an organisa-

tion, intranets have changed the way of thinking forever. It is now possible for employees to bypass the internal corporate hierarchy and gain access to another employee far removed from themselves either by rank, geographical location or work area. At the same time, employees have access to the latest information at the same instant as the decision-makers within a company. It is this thirst which is driving new ways of communicating. More information has been published in the past 30 years because of the computer revolution than was published in the 300 years before. But more importantly, published information needs to reach its audience quicker than ever.

To reach large audiences, two things are needed, a timely method of broadcast for event notification and a method of personalisation. This personalisation needs to filter the information to that needed by the individual while maintaining an efficient system for targeting large numbers of users. IP multicast is a version of the Internet protocol (IP) which addresses the problem of broadcasting to a group of users who are actually interested in the broadcast rather than simply giving a blanket coverage to everyone. Event-driven applications range from file transfer updates such as database replication to share price dealing. They can also include sporadic one-off messaging such as important corporate news or a message from the chairman. The common thread is the need to reach large numbers of people with either a notification of new information and where to get it or the information itself. Some applications such as stock quotes also require a level of quality of service (QoS) from the application. If subscribers have paid to receive the latest information, they expect to receive the information with a minimum latency. In the finance area, information must arrive as closely as possible with each dealer at the same instant, otherwise there

is an unfair competitive advantage. Customers may also pay for services that guarantee delivery or provide varying levels of reliability.

IP multicast was first proposed to the Internet Engineering Taskforce (IETF), the main standards body on the Internet, in 1989 by Steve Deering of Stanford University⁷. Its aim was to publish information to interested users based on the concept of information groups. Data providers publish to a group address rather than to individuals. They also do not have to be a member of the receiving group, merely aware of the group's address. Users who wish to receive particular information (for example, the news from CNN) subscribe to the multicast group typically using a protocol called the *Internet group management protocol* (IGMP). IGMP manages the groups dynamically so that users can subscribe and unsubscribe without affecting the data stream. This in turn removes the load from the publishing server as only one stream is transmitted without setting up connections to every possible subscriber. The work of assigning the information stream to each interested user is done through protocols like IGMP at either the end subnet (that is, at the local router) nearest the user or through multicast aware switches on each network segment (Figure 2). This method of moving the decision making process on which end hosts

should receive the information stream to the local part of the network and leaving the core network as a distribution channel makes multicast transmission inherently more scalable. Multicast also improves the broadcast model in that only users who are interested in a particular stream receive the information rather than all users on a particular area of the network, which would be true broadcast.

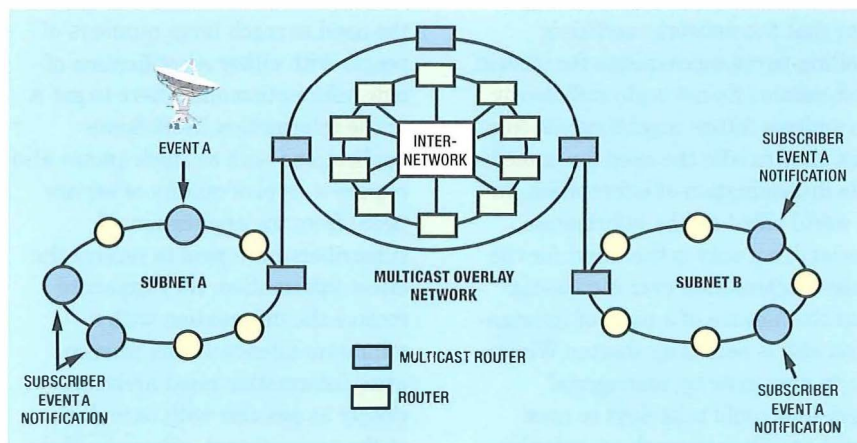
The use of multicast technology works well for information that cannot easily be tailored or that is applicable to a wide audience, such as audio or video streams. However, for news or information-based content, this method, while theoretically kinder to the network, still poses problems for the end user, who is in danger of being overwhelmed by the amount of information delivered. The break-even point for 'broadcast' multicast services is where the business model for a given service will support the mass broadcast of the same content to many users: if too much information is sent to too many people, neither the network costs nor the quality of information delivered are acceptable. Beyond this, therefore, services require some element of personalisation so that the customer gains value, the risk of information overload is lessened and the network load is reduced.

Basic multicast is unreliable, which is not a problem for livecasts

over video or audio, but for vital information that is not frequent, users could potentially miss an update, or only partially receive the information. Multicast IP is the equivalent to IP in the unicast communications stack and as such is connectionless. It disregards lost packets, which for business-critical data applications is unacceptable. Just as normal IP has transmission control protocol (TCP) to provide flow control, in these cases some form of reliability needs to be provided either within the application or within the communication stack. This reliability is crucial to event-driven applications where users may need to know that they missed some information, or request that it be re-sent. Re-transmission requests could be aggregated or scheduled for another multicast. Alternatively, specific updates may be unicast from a local proxy cache or even from the origin server itself to the end host. There needs to be an equivalent to TCP for IP multicast if reliability is to be provided for business critical applications. Early examples of reliable multicast include Pretty Good Multicast from Cisco, publish and subscribe from Tibco and the reliable multicast family of protocols from GlobalCast. Reliable protocols imply a method of feedback which impacts scalability and increases the complexity of the protocol, but which is still inherently more scalable and efficient than unicast transmission. An excellent source of further debate on the implementation of token-, tree- and cloud-based topologies for reliable multicast can be found on Globalcast's WWW site⁸.

Whether information is pushed or pulled, however, the requirement to identify and refine information relative to a particular user's interests is still a key issue. Typically, current push information services use manual (human) information selection and management techniques. The amount of information that can be processed in this way is

Figure 2—Multicast event notification



clearly limited. Later, a system is described that automates the process of content selection, summarisation and dissemination and can be applied to both server-based unicast systems and multicast services using client-side personalisation.

Personalisation—User Profiling and Information Filtering

As discussed above, one key to effective information management is user profiling—the ability to represent and reason about the interests or preferences of a user. Use of user profiling techniques will allow software automatically to filter and prioritise information from, for example, collections of Internet/intranet multimedia documents or in the future from several hundred channels of digital TV.

This section discusses techniques for representing users' interests and filtering the vast amount of information that is now available from the Internet (and other electronic sources) in a variety of media. There are three different approaches to user profiling:

- statistical keyword-based approaches;
- artificial intelligence (AI), and neural network (NN) techniques;
- social filtering (a non-explicit approach, using the preferences of other users rather than an explicit representation of a profile).

This article focuses on statistical approaches, since it is these that are used in the information management solutions described later. Experimentation with other techniques, in particular, the exploitation of neural network technology, is ongoing.

Statistical keyword approaches

Keyword (or term) matching is perhaps the best-understood approach to

determining whether some information is relevant to a user's interests. A common method for implementing a keyword matching scheme is to use a *vector space model* (VSM). Such a model constructs a vector of dimension m for each document, where m is the total number of *terms* (keywords) used to identify the content of the documents in the system. There are three main approaches to producing a document's vector:

- **Boolean** If the term appears in the document then it is given a weighting of 1, otherwise it is weighted as zero.
- **Term frequency** The weighting given to each term is based upon the frequency of its occurrence in the document. If the term is not present then it is set to zero. The resulting vector is often normalised to prevent longer documents from having a better chance of retrieval.
- **TF-IDF** The weighting given to a term under this scheme depends on the term frequency (*TF*), (which is the frequency of the occurrence of the term within the **current** document), and the inverse document frequency (*IDF*). The *IDF* is a measure of how often the term appears in **all** of the documents in the retrieval system and is often calculated as:

$$IDF = \log(N/n),$$

where N is the total number of documents in retrieval system, and n is the number of documents containing the term. Hence common terms have a low *IDF* and terms that appear in only a few documents will have a high *IDF*. The weighting for any particular term in the document is then given by

$$w = TF \times IDF.$$

As with the term frequency approach, normalisation of the

document vector is common, as is matching against phrases rather than individual terms.

A user's profile can be mapped into the vector space and the similarity between a document and the profile is then calculated as the cosine product of the two associated vectors. If the result is greater than the predefined relevance threshold for that user profile then the document is presented to the user.

The advantages of this approach are its relative simplicity and its adoption of well-understood and effective techniques from the field of information retrieval⁹. Disadvantages include the fact that the technique can be used only for text documents. Additionally, there are many ways to express a given concept (*synonymy*); the literal terms in a user's profile may not match those in a relevant document.

Conversely, most words have multiple meanings (*polysemy*), so terms in a profile may literally match those in irrelevant documents. These problems can, however, be mitigated somewhat by the use of profile enhancement (see below). Generally, this scheme (incorrectly) assumes keywords are orthogonal in vector space; that is, all terms are unrelated. The idea of synonymy shows that several keywords could describe the same concept. Such keywords are clearly related and therefore cannot be classed as orthogonal. On a less obvious scale, related concepts are also not orthogonal—see the next section for further discussion on this point.

Profile expansion

A supplementary method that has been researched for improving the performance of keyword-based filtering relies upon a technique adapted from the information retrieval technique of query expansion. The filtering system expands the profile by adding terms related to the keywords supplied by the user. These new terms can either be derived with lexical aids, such as

Figure 3 – The Jasper information agent

thesauri and controlled vocabulary lists, or statistically selected based upon the general co-occurrence of the terms in other documents (see the discussion of the ProSearch agent later).

The procedure used in Reference 10, for example, used sets of synonyms, (known as *synsets*), which formed part of a lexical system called *WordNet*. The synonyms are arranged into *is-a* and *part-of* relationships. The expansion of any given query could then proceed in several ways by adding (for example):

- only the immediate synonyms from appropriate synsets,
- all descendants in the *is-a* hierarchy, and
- all words in synsets one *is-a* or *part-of* link away from the original synset.

It was found that the expansion technique did not improve the effectiveness of queries that were relatively complete when entered by the user, for any filtering method. However, shorter or incomplete queries were significantly improved by the expansion and these still occur relatively often.

Knowledge Discovery and Delivery Solutions

Jasper

Jasper^{11,12} is a system of WWW information agents. Jasper works with a community of users to facilitate the automatic organisation, summarisation and sharing of WWW information. Figure 3 summarises the main functionality provided by the Jasper information agent.

Storing Information in Jasper

Jasper agents are used to store, retrieve, summarise and inform other agents about information found on WWW. In certain circumstances,



Jasper agents also identify an opportunity for performance improvement and seek user feedback in order to improve. Each Jasper user has a personal agent which holds a user profile based on a set of key phrases.

Given the vast amount of information available on the WWW, it is preferable to avoid the copying of information from its original location to a local server. Indeed, it could be argued that this approach is contrary to the whole ethos of the Web. Rather than copying information, therefore, Jasper agents store only relevant meta-information. This meta-information is then used to index on the actual information when a retrieval request is made.

When a user finds a WWW page of sufficient interest to be stored by Jasper, a 'store' request is sent to Jasper via a menu option on his or her WWW client. Jasper then invites the user to supply an annotation to be stored with the form. Typically, this might be the reason the page was stored and can be very useful for other users in deciding which pages retrieved from the Jasper store to visit. The user can also specify at this point one of a predefined set of interest groups to which to post the page being stored. The WWW page is added to the Jasper store. Essentially, the Jasper store is a simple term-document matrix M , wherein:

$$M(i, j) = \begin{cases} 0 & \text{if term } j \text{ does not occur} \\ & \text{in document } i \\ n & \text{if term } j \text{ appears in} \\ & \text{document } i \text{ } n \text{ times} \end{cases}$$

At storage time, Jasper agents perform four tasks:

- An abridgement of the textual information on the WWW page is created, to be held locally. This enables a user to assess the content of a page from a local store quickly before deciding whether to retrieve remote information.
- The content of the page is analysed and matched against every Jasper user's profile. If the profile and document match strongly enough, Jasper e-mails the user, informing him or her of the page stored.
- The page is also matched against the storer's own profile. If the profile does not match the content of the page being stored, Jasper suggests to the user new words or phrases they may wish to add to their profile. Thus Jasper agents can adaptively learn their users' interests by observing the users' behaviour.
- For each document, an entry in the Jasper store is made, holding keywords, an abridgement of the document, document title, user annotation, URL, storer name and date and time of storage.

In this way, a shared and enhanced information resource is built up in the Jasper store. Given that users must make a conscious decision to store a page, the quality of the information in the Jasper store is high – it is effectively pre-filtered by Jasper users.

Getting information from Jasper

- *E-mail notification* As described above, when a page is stored by a

Jasper agent, the agent checks the profiles of other agents' users in its 'local community' (the set of users who contribute to that particular Jasper store). If the page matches a user's profile with a score above a certain threshold, an e-mail message is automatically generated by the agent and sent to the user concerned, informing the user of the discovery of the page.

- **What's new** A user can ask his/her Jasper agent 'What's new?' The agent then interrogates the Jasper store and retrieves the most recently stored pages. It determines which of these pages best match the user's profile. A WWW page is then presented to the user showing a list of links to the recently stored pages, along with annotations where provided, date of storage, the storer, a link to a locally held abridgement of the page and an indication of how well the page matches the user's profile.
- **Interest groups** As mentioned above, when a user stores a page in Jasper, he/she has an opportunity to specify one of a predefined set of interest groups to which to post the page. Interest groups gather together pages of related information. Jasper users can visit interest group pages which are dynamically constructed from the pages which have been posted to them and consist of a list of links to the pages and their abridgements, along with any annotation provided by the original storer of the page. Interest groups are similar to the 'list of links' pages found in many WWW locations, with the important extensions that multiple users can contribute to the list (automatically via the storage process) and that abridgements of the information and annotations by the original storer of the link are also available.

- **Keyword retrieval** The user supplies a query in the form of a set of keywords. The Jasper agent then retrieves the most closely matching pages held in the Jasper store, using a Boolean keyword matching and scoring algorithm. The agent then dynamically constructs a hypertext mark-up language (HTML) page with a ranked list of links to the pages retrieved and their abridgements. Any annotation made by the original user is also shown, along with the scores of each retrieved page. This page is then presented on the user's WWW client.
- **Show me similar** This facility allows a Jasper user to request that his/her agent retrieves relevant pages from the Jasper store. The WWW page currently being viewed by the user is analysed and the Jasper store is searched for related documents. Given the pre-filtered nature of the Jasper store, it provides a rich resource for this type of facility.

Jasper agents are in use by several user groups within and outside BT. Jasper agents allow users to store information, along with summaries and annotations, and to share information with others with the same interests. As such, Jasper is a step along the road towards the original vision for WWW as a network which supports cooperative working and the sharing of information, rather than merely the serving of information.

ProSearch—a proactive search agent

This section describes ProSearch, an agent which automatically searches the WWW on a user's behalf, returning high-quality information based on the user's profile.

Clustering keywords

A user's profile comprises a set of words and phrases as in the Jasper system described above. This profile

may reflect an interest in more than one area: the same user could have words and phrases in his/her profile which reflect an interest in network management, ATM technology and fly-fishing, for example. ProSearch will automatically identify the topics in a user's profile and search for information about each topic separately.

Key phrases (terms) occurring in a particular document collection (such as a set of WWW pages) can be clustered into related groups. So, for example, a similarity matrix for the keywords in the Jasper store can be constructed which gives a measure of the 'similarity' of each pair of keywords in the store. For each pair of keywords, the *Dice coefficient* is calculated¹³. For two keywords K_i and K_j , the Dice coefficient is given by:

$$S_{K_i K_j} = \frac{2C}{A + B}$$

where:

$S_{K_i K_j}$ = similarity between K_i and K_j ,
 C = number of documents in which K_i and K_j co-occur,
 A = number of documents in which K_i occurs, and
 B = number of documents in which K_j occurs.

A similarity S of 0 means that two terms never co-occur, while if S was 1, the terms would occur in the same and only the same documents.

Note that this formula assumes Boolean document weighting; that is, the weight of a given document D is 1 if the keyword occurs in D , else 0. More sophisticated schemes based on keyword frequency are available.

Proactive searching

The key phrases comprising a user's Jasper profile are being used by an experimental system called *ProSearch* to search proactively for new WWW pages relevant to their interests, which can then present a list of new pages which the user may be interested in without them having to explicitly carry out a search (see Figure 3). These proactive searches

Figure 4 – The ProSearch agent

can be carried out at some given interval, perhaps weekly or monthly. Clustering is useful here because a profile may reflect more than one interest. Consider, for example, the following user profile:

Internet, WWW, html, http, football, Manchester, United, linguistics, parsing, pragmatics.

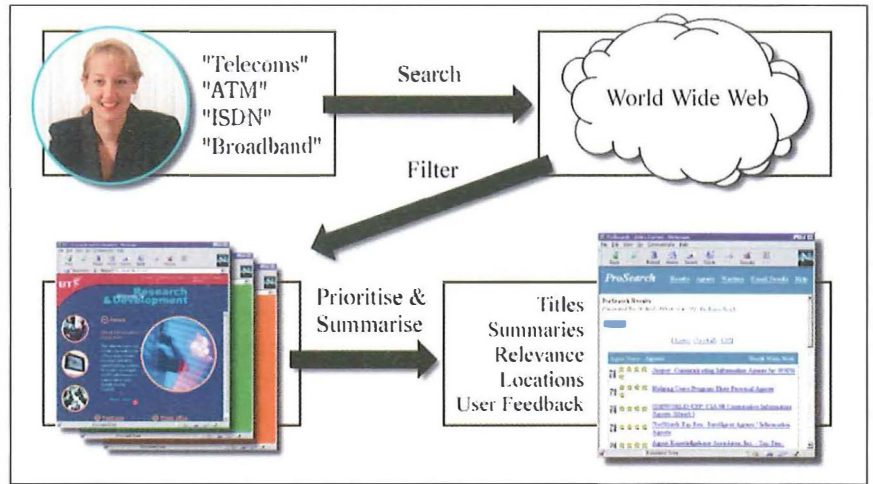
Clearly, three separate interests are represented in the above profile and searching on each separately is likely to yield far superior results than merely entering the whole profile as a query for the given user. Clustering keywords from the document collection will automate the process of query generation for proactive searching by a user's ProSearch agent.

Let us exemplify with a typical profile, the following set of keywords:

{Internet, information, retrieval, SMART, clustering, agent, intelligent, CORBA, IDL, DCE}

We can cluster these keywords as described above, using the Jasper store as the document collection. We then obtain the similarity matrix shown in Table 1.

If we use complete-link clustering¹⁴, the following clusters are obtained:



{DCE, CORBA, IDL}
 {SMART, clustering}
 {Internet, information, retrieval}
 {intelligent, agent}

The set of queries created by ProSearch from the user's profile is then presented to the user for confirmation or adjustment before the queries are submitted.

Because queries based on profile clusters are made off-line, rather than interactively by the user, there is an opportunity to analyse the results before presenting them to the user, thus providing a higher quality set of results than is possible with a traditional on-line search engine. For example, when the search results are obtained, they can be abridged and matched against the user's profile in the usual way to give a prioritised list of new URLs along with locally held abridgements. In addition, a log can be held of the pages found by Jasper for each user. By checking against this log, only new pages

(those the user has not been shown by ProSearch previously) are presented. The pages are presented to the user on a dynamically constructed WWW page which easily stores pages of particular interest in the Jasper store, thereby disseminating the information. In addition, the user can feed back whether or not a particular document was of interest. Using this information, the user's profile is improved to reflect his/her interests better. Figure 4 shows the ProSearch agent search process.

Profile enhancement

ProSearch is also being used to enhance the profiles of users of the Jasper system and other information filtering systems within BT. As discussed earlier, a user's profile can be enhanced by adding those keywords most similar to the keywords explicitly represented in the user's profile in a way reminiscent of query reformulation techniques, such as the use of spreading activation net-

Table 1 A Term Similarity Matrix

	CORBA	DCE	IDL	Internet	information	retrieval	intelligent	agent	SMART	clustering
CORBA	1									
DCE	0.62	1								
IDL	0.92	0.50	1							
Internet	0.03	0.03	0.03	1						
information	0.04	0.04	0.03	0.51	1					
retrieval	0.04	0.08	0.00	0.21	0.27	1				
intelligent	0.00	0.03	0.00	0.21	0.30	0.30	1			
agent	0.07	0.11	0.04	0.16	0.29	0.38	0.53	1		
SMART	0.00	0.00	0.00	0.07	0.07	0.19	0.10	0.14	1	
clustering	0.00	0.00	0.00	0.01	0.03	0.17	0.04	0.12	0.43	1

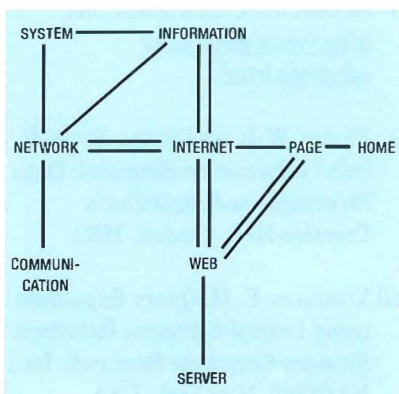
AgentCast is an experimental set of components to provide discovery of information, profiling, summarisation, delivery, and feedback and population monitoring.

works^{15,16} and other techniques based on the use of semantic nets^{17,18}.

In Reference 17, for example, the use of semantic knowledge from WordNet¹⁹ in an experimental program called *NetSerf* for finding information archives on the Internet is reported. Based on a set of 75 queries, *NetSerf* was reported to have performed significantly better than the *Smart*²⁰ information retrieval system. Interestingly, BT has performed some simple experiments with WordNet for profile enhancement but found that WordNet did not perform well with respect to the very technical vocabulary that constitutes the terms in the Jasper store. It was this technical bias that led to BT researchers looking at the store itself for a network of related terms: in effect, the clustered terms constitute a simple spreading activation network. In Reference 15, some results from cognitive psychology research are analysed to explain why it proves difficult for people to think of related terms when trying to frame a query to an information retrieval system. This difficulty is ameliorated to a degree by the automatic (hidden) enhancement of a user's profile by the system.

Figure 5 shows an example network of keywords which have been built from the keyword similarity matrix extracted from the Jasper store. The algorithm is straightforward: given an initial starting keyword, find the four words most

Figure 5—A keyword network from Jasper



similar to it from the similarity matrix. Link these four to the original word and repeat the process for each of the four new words. This can be repeated a number of times (in Figure 5, three times). Double lines between two words indicate that both words occur in the other's four most similar keywords. One could of course attach the particular similarity coefficients to each link for finer-grained information concerning the degree of similarity between words.

Thus, for example, if the words 'virtual', 'reality' and 'Internet' are part of a user's profile but 'VRML' (virtual reality modelling language) is not, an enhanced profile might add VRML to the original profile (assuming VRML is clustered close to virtual, reality and Internet). In this way, documents containing VRML but not virtual, reality and Internet may be retrieved whereas they would not have been with the unenhanced profile.

AgentCast

AgentCast builds upon the work already described in the previous two case studies to provide a basis for personalised webcasting services. It is still in the research phase and is an experimental set of components to provide:

- discovery of information,
- profiling,
- summarisation,
- delivery, and
- feedback and population monitoring.

For push technologies particularly, the ability to push a summary of interesting information is more efficient on both the user and the network. The method of summarising content is based on the BT ProSum information summarisation software²¹. The basic concept is to

provide a range of services that can be offered easily to customers using the various agent technologies as building blocks and incorporating future research.

Knowledge discovery

Knowledge discovery is no easy task when applied to the whole corporate intranet or the Internet. There are no standards for publishing information in a structure form other than physical HTML commands and addresses. AgentCast builds upon emerging Web standards such as the extensible mark-up language (XML) and parasitically through current Web search engines to find information of general interest to the population. A pool of information, some of which is delivered to each user according to his/her profile, is created. ProSearch can be applied to hunt specifically for information pertinent to each user, but at the same time general content collection is applied to popular news sites to provide the latest information in one repository.

XML was issued as a recommendation by the World Wide Web Consortium (W3C), the leading application standards body on the Web, during February of 1998²². It is a framework for describing data or information so that it can be identified and used by applications for specific tasks. XML can represent different name spaces so that the same piece of information can have different meanings depending on where it resides. For example, the <title> tag could be defined within a people directory to mean the person's title, whereas, within an electronic library, it would mean the title of a book or catalogue. XML applications, systems which use XML to characterise content, include the channel definition format (CDF) scripting language used within Internet Explorer 4.0 for the Active Desktop and Active Channels as well as the resource description framework (RDF) proposed to describe specific resource available on a network.

The benefit of XML to knowledge discovery is that it provides a standard way of defining attributes for a piece of information. This is a similar concept to recent search engines, which have attempted to use meta-tags within HTML. Unfortunately the use of meta-tags has not significantly improved search engine performance due to patchy and inconsistent usage.

Profiling and summarisation

Profiling is based on the extraction of keywords from textual documents and can be done either at any level (for example, for the population as a whole) or for the end user. It can also be used on the population to introduce a basic level of serendipity. A more complex version of this is to use clustering techniques as in ProSearch to group topics and then to apply social filtering²³ to work out from other users with similar interests what extra information should be shown to a user.

Summarisation, based on the ProSum component can be tailored specifically to a user's profile. It is based on the frequency of words and phrases within a document, based on a technique called *lexical cohesion analysis*²¹.

Delivery

Increasingly the Internet is changing from being passive to enabling mass broadcasts of the latest information. For polling-based push, server-based content creation and personalisation still works as well as for the Web itself. It is relatively simple to create dynamic Web sites that produce news tailored for an individual because the user has to request the information. True push is more difficult to tailor. Tibco's publish-and-subscribe model uses subject names to categorise content and utilise the same multicast group for many types of content. In this model, users subscribe to particular subject names such as 'uk.stock.*' which would give all stock prices in the United Kingdom.

Feedback

Feedback in a unicast environment is provided by scoring the user's reaction to a piece of information which itself was tailored to meet his/her profile. Usually users are asked to select a few tick boxes relating to their preferences when they join the service in order to 'seed' their profiles. Alternatively, the population profile can be used initially or the system can be left to 'learn' a user's interests. Such feedback can be provided by using the application specific information such as the feedback statistics in Backweb Technologies' Channel Server or through dynamic Web sites that log user requests for news stories and click-throughs. In addition to this automatic profile evolution, users can be provided with a facility proactively to view and edit their own profiles. Multicast distribution does provide better scalability, but it also removes the feedback loop that is so critical, especially for social filtering where the decisions taken by the population, or a subset of the population, are critical to the algorithm.

Conclusion

This article has discussed the issue of information overload, which is growing in importance, as ever more information becomes available online through the public Internet and corporate intranets. The importance of new proactive push models for information delivery were also explained.

The article described the Jasper agent, which facilitates the organisation, summarisation and dissemination of information through an organisation, and ProSearch, which automatically finds high-quality, relevant information for a user from WWW and other information sources.

Finally, AgentCast was described—a powerful combination of agent software and webcasting technology, which provides a novel and enhanced method of knowledge discovery and delivery.

References

- 1 POLLOCK, A.; and HOCKLEY, A. What's Wrong with Internet Searching? Designing the Web: Empirical Studies Conference, Microsoft Usability Group, Oct. 1996, Microsoft, Redmond, US.
- 2 COCHRANE, R. Unleashing the Intranet. *BT Technol. J.*, Apr. 1997, **15**(2) pp. 107–113.
- 3 BLAIR, D. C.; and MARON, M. E. An Evaluation of Retrieval Effectiveness for a Full-Text Document Retrieval System. *Communications of the ACM*, **28**(3), Mar. 1985.
- 4 JONES, W. P. On the Applied Use of Human Memory Models: The Memory Extender Personal Filing System. *Int. J. Man-Machine Studies*, 1986, **25**, pp. 191–228.
- 5 COCHRANE, R.; and CALLAGHAN, J. Building Intranets. Telecommunications Development Asia Pacific, Sept. 1997.
- 6 COCHRANE, R.; and DEVIN, P. Push Technology—Evolving the Corporate Intranet. IQPC Push Conference, Sept. 1997, Le Meridien, London.
- 7 DEERING, S. RFC 1112, Host Extensions for IP Multicasting. Stanford University. Stanford University, Aug. 1989, submission to the IETF Network Working Group.
- 8 GlobalCast Communications, <http://www.gcast.com/solutions.html>
- 9 FRAKES, W. B.; and BAEZA-YATES, R. (eds.) *Information Retrieval: Data Structures and Algorithms*. Prentice-Hall, London, 1992.
- 10 VOORHEES, E. M. Query Expansion using Lexical-Semantic Relations. Siemens Corporate Research, Inc., NJ 08540, New York, USA.

- 11 DAVIES, N. J.; REVETT, M. C.; and WEEKS, R. Information Agents for the WorldWideWeb. *BT Technol. J.*, 14(4).
- 12 DAVIES, N. J.; REVETT, M. C.; and WEEKS, R. Jasper: Communicating Information Agents, Proc. 4th Intl. World Wide Web Conference, Boston, USA, Dec. 1995. Also available at: <http://www.w3.org/pub/Conferences/WWW4/Papers/180/>
- 13 RASMUSSEN, E. Clustering Algorithms. In Reference 8.
- 14 GRIFFITHS, A.; ROBINSON, L. A.; and WILLETT, P. Hierarchic Agglomerative Clustering Methods for Automatic Document Classification. *Journal of Documentation*, 40(3), Sept. 1984, pp. 175–205.
- 15 RUGE, G. Human Memory Models and Term Association. Proc. 18th Annual International ACM SIGIR Conference. Washington, USA, 1995.
- 16 WETTLER, M.; and RAPP, R. A Connectionist System to Simulate Lexical Decisions in Information Retrieval. In PFEIFER, R.; SCHRETER, Z.; FOGELMAN, F.; and STEELS, L. (eds.) *Connectionism in Perspective*. Elsevier, Amsterdam, 1989.
- 17 CHAKRAVARTHY, A. S.; and HAASE, K. B. NetSerf: Using Semantic Knowledge to Find Internet Information Archives. Proc. 18th Annual International ACM SIGIR Conference. Washington, USA, 1995.
- 18 RADA, R.; and BICKNELL, E. Ranking Documents Based on a Thesaurus. *Journal of the American Society for Information Science*, 40(5), 1989, pp. 304–310.
- 19 MILLER, G. A. WordNet: An Online Lexical Database. *International Journal of Lexicography*, 3(4), 1990.
- 20 SALTON, G. The SMART Retrieval System. Englewood Cliffs, N.J., USA, Prentice-Hall, 1971.
- 21 DAVIES, N. J.; and WEEKS, R. Text Summarisation for Knowledge Management. *BT Technol. J.*, 16, 1998.
- 22 Extensible Markup Language, W3C Architecture Domain. <http://www.w3.org/XML/>
- 23 DAVIES, N. J.; and STEWART, R. S. User Profiling Techniques: A Critical Review. 19th BCS IR Colloquium, Aberdeen, UK, Apr. 1997.

Biographies



John Davies
BT Networks and
Systems

John Davies graduated in Computer Science and Physics at the University of London in 1981. After two years in industry, he obtained an M.Sc. in Computer Science from Essex University. After four years with GEC, he returned to Essex where he received his Ph.D. in the area of artificial intelligence. He joined BT in 1990 and currently leads BT's knowledge management research group. Technical interests include on-line information management, agent technology, intranets and their use in knowledge management applications. He is an author of, and contributor to, many books and papers in the areas of information retrieval, agents, AI and knowledge management. He is a Chartered Engineer and a member of the British Computer Society, where he sits on the Information Retrieval Specialist committee.
<http://www.labs.bt.com/people/daviesn2/>



Ralph Cochrane
BT Networks and
Systems

Ralph Cochrane currently leads the push (Webcasting) research team at BT Laboratories. His main research interests centre on information delivery and management, particularly for multimedia. An engineering graduate, he worked for Nuclear Electric in the field of wireless data communications and computing before joining BT in 1996. He can be contacted via his Web site at <http://www.labs.bt.com/people/cochrar/>

Intelligent On-Line Purchasing

New applications and services hint at the power of the Internet to bring buyers and sellers together. This article examines consumer buying behaviour. It shows how intelligent applications can help at each stage, guiding a buyer from the point of realising a need, through a complex purchasing decision, to the point where the buyer can buy a product with confidence. The roles of advertising, third party product review and independent advice are examined within this new environment.

Introduction

The Internet is introducing new ways to buy products. We can find information, suppliers and advice about purchases on-line. This is causing a fundamental shift in the relationship between buyers and sellers, and retailing will adapt to serve this new market. The nascent home-shopping market is expanding rapidly. The creation of new shopping environments that provide novel retailing mechanisms and tools for shoppers is helping to drive change.

This is just the beginning, and currently many problems face on-line buyers. It is not easy to find products or suppliers using keyword searches or by browsing directories such as Yahoo. There can be an overwhelming number of suppliers returned from a search or, owing to a peculiar way a product is described, the user may miss it entirely. Internet classified directory services help users find suppliers in predefined categories, but contain no information about precisely what products suppliers sell and what deals they are offering. A buyer must contact suppliers individually to find this out. Given these limitations, many buyers find suppliers in a haphazard fashion. There are no unifying directories or common ways to describe products.

The problems with searching are due in part to the nature of hypertext mark-up language (HTML) which controls the way a Web page is presented. What HTML lacks is any structured way to encode information about the meaning of a Web page. It lacks data about the data on the page, so called *meta-data*. Meta-data contained in a Web page would allow a search for 'flowers' to distinguish between sites selling flowers and those providing gardening tips about flowers. XML[†] (extensible meta

language) promises the widespread use of meta-data. In turn this will allow greatly improved Web directories and Web searching as search engines will be able to 'understand' the meaning of a page and its relevance to a user.

Buyers are able to find powerful suppliers with well-known brands (for example, Dell Computer Corp., for PCs) but this excludes the majority of smaller suppliers. It is a commonly held view that the Internet will cause massive disintermediation (removal of middle-men and established sales channels) and buyers will deal directly with suppliers. This will only be the case in some sectors for major brands that are capable of direct selling. The majority of suppliers and the majority of transactions will need new kinds of intermediaries to help buyers find the most appropriate deals.

A further problem facing buyers when making a purchase is understanding a complex product or product domain. A buyer may not have enough knowledge to make an informed buying decision. It is not necessary to be an expert computer user to buy and use a computer, but without a certain level of confidence about the suitability of the purchase, buyers can be reluctant to spend their money. Often the most important part of the buying process can be deciding the appropriate type of product and its most important features.

With the problems of deciding what product is needed, navigating the Net in search of those products

[†] XML can be thought of as the next generation of HTML. It allows for new mark-up tags with new meanings to be created arbitrarily. For a full description of XML, see <http://www.w3c.org/XML/>

Context and user profiles help to target adverts, and this will get more sophisticated as on-line advertisers learn more about users and their profiles.

and the difficulty of finding the best suppliers, there arises the need to create new services and tools for buyers. Services are emerging that are dedicated to satisfying buyers' purchasing needs. Indeed, these services also stimulate the purchasing need in the first place. The remainder of this article examines new tools and services being created to streamline on-line purchasing.

The Buying Process

It is very helpful to divide the buying process into stages so that technologies can be applied to aid problems encountered by buyers (see Figure 1). Well-developed models exist of consumer buying behaviour¹. It is not being suggested that all buyers follow a series of distinct steps. Rather, the stages can happen concurrently and buyers can revisit stages many times, refining their requirements.

Identify a need

This is the stage where the buyer identifies a need. The way in which the buyer is made aware of a need can vary. The realisation can either be spontaneous (for example, a person realises he/she needs food) or stimulated (for example, by advertising or by recommendation from a friend).

Product analysis

To satisfy the identified need, the buyer must gather information to decide what type of product to buy and its important characteristics. This stage may be trivial if the buyer already knows what he/she wants. This stage is rarely carried out in total isolation from the seller analysis stage.

Seller analysis

With some product preferences from the previous stage, the buyer now determines whom to buy from. This decision depends on many criteria such as price, delivery, warranty and overall perception of seller quality. The product and seller analysis stages are difficult to separate entirely; in fact, they are highly interrelated in most cases. These two stages often happen simultaneously with sellers helping to inform buyers' product choices.

Negotiation and decision

Depending on the type of purchase, there may be many parameters of the deal that must be negotiated before a purchasing decision is made. Often the price of an item is fixed and then there is less scope for negotiation. The outcome of the negotiation stage may change decisions made in the product and seller analysis stages.

Payment and delivery

In this stage, the payment (or invoicing) is carried out and the product is delivered (or collected). The type of payment mechanism affects the negotiation stage.

Purchase review

The overall buying process is reviewed in this stage. The buyer may simply learn from the process or there may be some formal feedback or rating of parts of the process. This can be used to build buyer-seller relationships, to help improve the process or to aid subsequent buyers. In an on-line environment, feedback and review are easier to gather and will play an

increasingly important role in the buying process.

Aiding Buying Behaviour

This section looks at assisting the buying process and various techniques that can lead buyers to a buying decision.

Although a buyer can identify a need spontaneously, the vast majority of needs are created by some form of stimulation. The most apparent is advertising, although this is not always the most powerful.

We are all familiar with the buying stimuli of traditional media. We are exposed to a huge array of advertisements in newspapers and magazines, and on TV and radio. Advertisements are appearing in new places too, such as train tickets. Advertisers are trying to target audiences precisely to increase advertising efficiency. On-line advertising follows this trend too. Banner advertisements are displayed dynamically depending on who is accessing a particular page. Context and user profiles help to target advertisements, and this will get more sophisticated as on-line advertisers learn more about users and their profiles.

Perhaps the most powerful buying stimulus is personal recommendation. If we are told by a friend that 'you should buy such-and-such-a-product, it's exactly what you need' then, assuming we trust the recommender's opinions, we are highly likely to follow this advice.

This type of personal recommendation can be recreated on-line to some extent. The best example of this is FireFly² which uses a technique called *electronic word of mouth* (EWoM) to find similar users and make recommendations between them. The technique can be used in many product domains. The book recommendations at the Barnes and Noble³ Web site are powered by the FireFly engine.

To demonstrate the use of EWoM for music recommendations, the myLAUNCH⁴ music community site is

Figure 1—The buying process




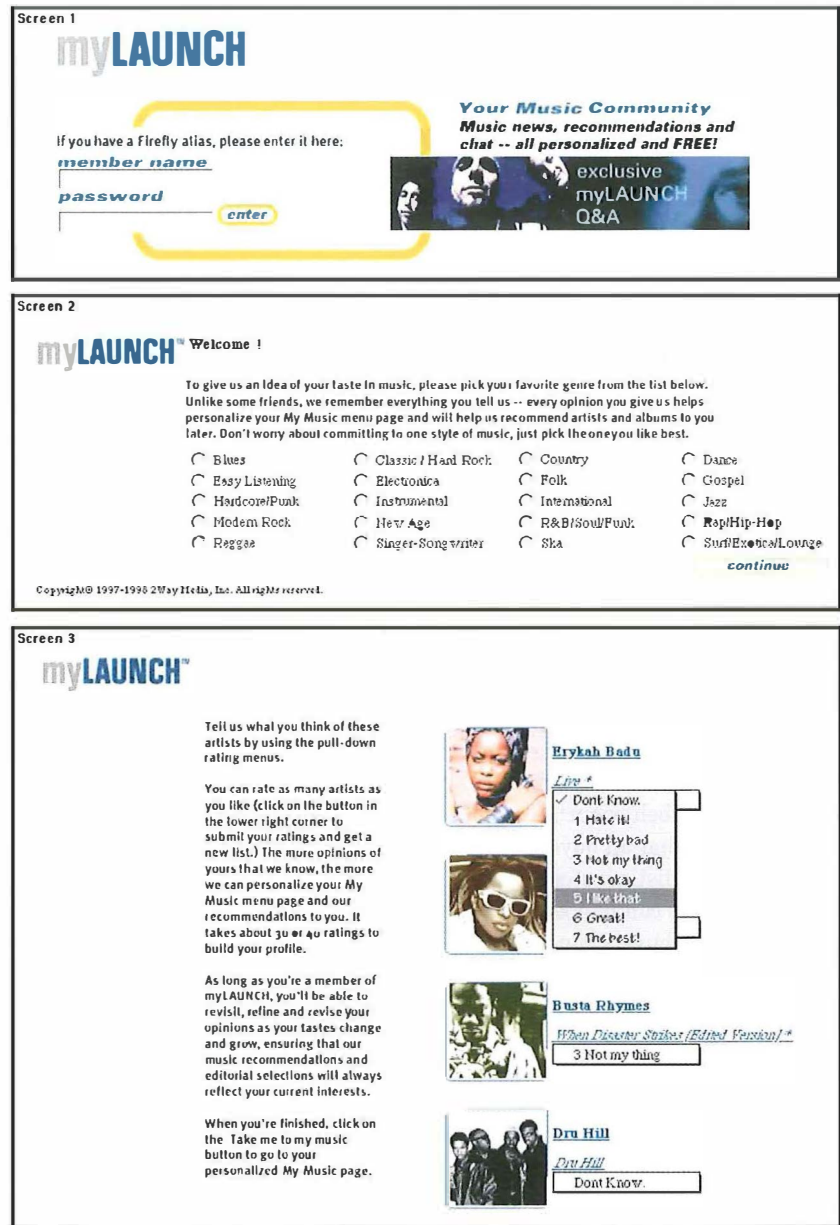
Figure 2—myLAUNCH music recommendations

used as an example (see Figure 2, which shows three screen shots). Imagine a user has a FireFly account containing personal profile information that he/she has used elsewhere. The user going to the myLAUNCH site for the first time enters his/her FireFly name and password. A rough guide to the user's musical taste is required before the preference profile can begin to be built up. Screen 2 of Figure 2 asks the user to pick one genre that best describes his/her musical tastes. This choice is not critical; it just increases the likelihood that the user will have some opinions (good or bad) about the first albums that are presented in the next screen. In screen 3, myLAUNCH presents the user with the six most popular albums in the chosen genre. The user then rates these from: 1 = 'I Hate it!' to 7 = 'The Best!'. When these ratings are submitted, the FireFly engine has a very narrow view of the user's tastes. The matching algorithms then finds a group of other users who have the 'closest' musical tastes. The system returns six albums that they, collectively, rate most highly for our user to rate. Once these are rated the matching engine has 12 ratings to match with other users. The closest users are found and their favourite albums are returned and so on. With each step the system learns more and is able to match people of ever closer musical tastes.

One further point about this system worth highlighting is that it is self-maintaining. Users are encouraged to enter new albums as they are released, together with ratings and comments. Once a new album has been submitted, users who would probably like it are alerted. Although it appears that the system 'understands' one's musical tastes, in fact it is a set of people like oneself that are making the suggestions. The system simply brokers, stores and reuses the opinions appropriately.

Product Analysis: Gathering Advice

Regardless of how a buyer realises a need, the next stage is to analyse



products by gathering information and advice. The objective is for buyers to define their requirements well enough to begin seeking suppliers.

Some buyers (particularly those with little expertise in the field) would welcome advice on what kind of a product would be most appropriate for their needs. This will be referred to as *generic* advice. Buyers would also value independent advice (say from product reviews) about the quality of individual products. This will be referred to as *specific* advice.

Generic advice

Generic advice covers advising non-expert buyers on the parameters that should be specified in their requirements, depending on the intended application of the item. For example, a

computer simply for word-processing would need less memory than one used for multimedia applications.

One technique that can be used for this task is constraint satisfaction⁵, and an example application of this is PersonaLogic⁶. PersonaLogic combines constraint satisfaction with expert domain knowledge to guide a buyer towards the ideal product based on the relative importance of certain product features. Figure 3 shows a montage of four screen details from the advisor that helps buyers choose video cameras. The user is taken through a series of screens (in this case, seven) where they are asked to enter their ratings. Each feature is described in basic terms and some screens have panels to help buyers decide the importance of features

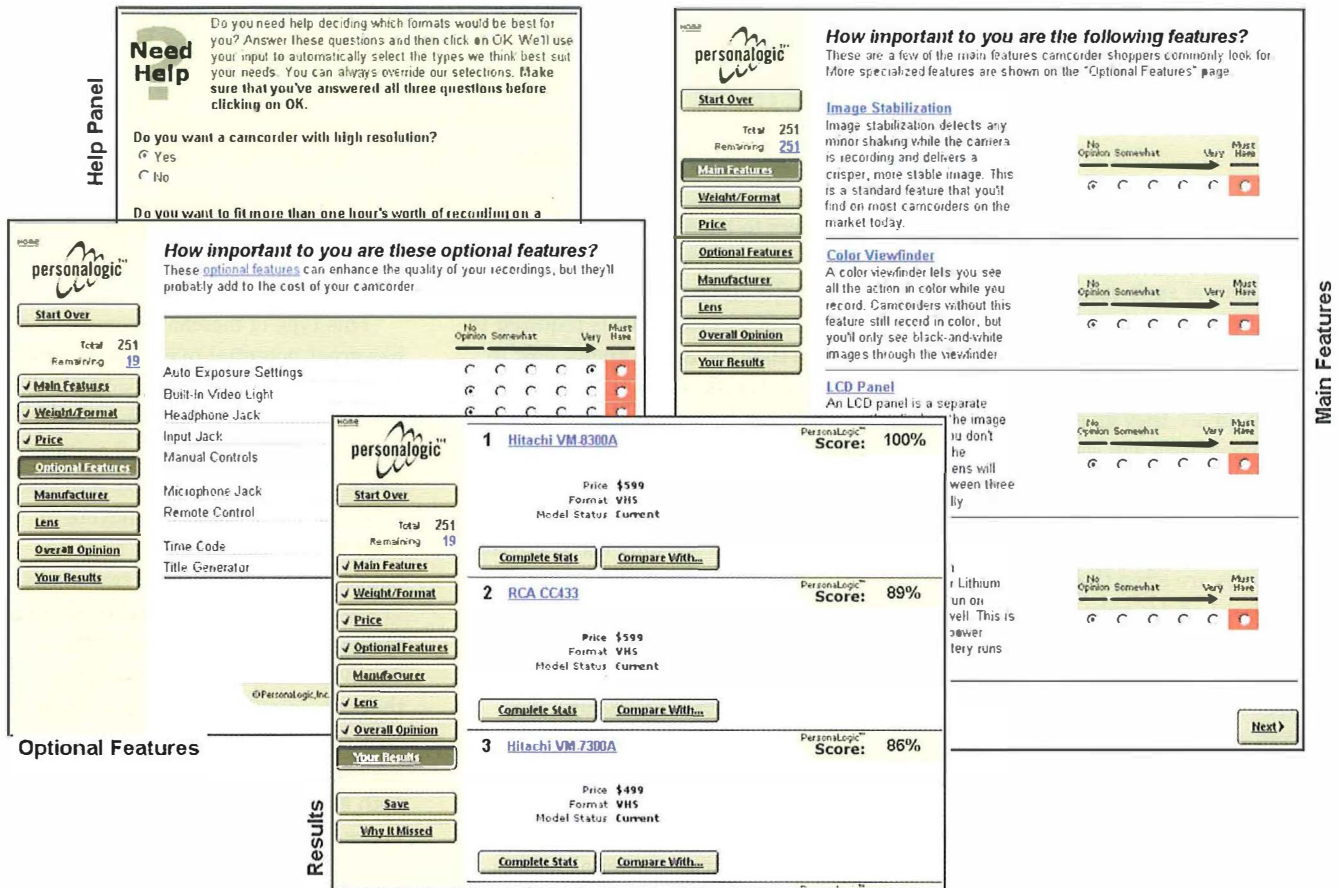


Figure 3—The PersonalLogic specific advisor

Accountable advertising

Advertising in the on-line world will undergo many changes. It will adapt to, and define the new media in much the same way as it did with print and broadcast media. We already see banner advertisements that are 'intelligently targeted', but this is just the beginning of a change which is taking advertising into the realms of one-to-one marketing†. Request for quotation services will, in effect, reverse the advertising process. Instead of suppliers pushing out messages to buyers to advertise what they sell, buyers will push messages to suppliers to advertise what they want to buy. On-line advertising will become ever more intelligent as advertisers learn about individual customers in a richly interactive relationship. However, perhaps the most significant change facing advertising is that it will become more accountable.

Advertising is a very imprecise activity. It is difficult to quantify accurately the extra sales that advertising spending stimulates. There are survey techniques to estimate a campaign's effectiveness but there is no way for the client to measure the exact return on investment of any single promotional effort, or what exactly stimulated a buyer to purchase. By contrast the Net affords many ways to measure advertising's effectiveness. When a buyer arrives at a transaction, the advertisement, advertorial, special-offer, or any other advertising oriented on-line media that brought them there, will be known. The power of any advert to induce a sale will be measurable. Advertisers will spend their promotional resources only in the places where the returns are good. This hints at the end of broadcast advertising and the increasing sophistication of on-line marketing.

† This is a principle whereby the retailer builds a mutual relationship with each customer, learning to satisfy each customer's needs on an individual basis. Current on-line retailing practices are moving towards the ideals of one-to-one marketing with the advent of many techniques to group customers into ever more finely targeted niches—the ultimate aim being a niche of one.

based on their needs. As the buyer progresses through the screens, the choices narrow. The final results screen lists the remaining suitable products in order. One very important feature of this advisor is that the user can see why products were rejected. The system requires a detailed knowledge-base about all the current products in the product domain, which calls for a carefully devised knowledge representation scheme and continual maintenance.

Specific advice

Specific advice about a product can be found from expert opinion, opinion from other users or from laboratory tests, consumer reports or ratings from respected third parties, such as consumer organisations or magazines. The key problem to solve when creating a specific advisor service will be integrating the many sources of advice. This must then be presented to the buyer in a meaningful format to help the buyer to select specific manufacturers and specific models.

Experienced Internet users may use Usenet news groups to solicit opinion about particular products

Where appropriate the buyers should be asked why they made a decision. This can then be used to advise similar buyers faced with similar decisions.

from other users. Advice is often forthcoming, but it is difficult to tell the quality of such opinion.

Magazines, particularly ones dealing with high-tech products are valued by their readers for the specific product advice they give. Many computer buyers will consult a laboratory test before buying a new machine or a major peripheral. The most popular sections of on-line magazines are often their reviews. For this type of advice to be credible, the organisation giving the advice must be independent from suppliers. There are search services designed to help buyers search for relevant magazine reviews⁷.

Supplier Analysis: Shopping Around

Once a buyer has formulated a specification of the product that will satisfy his/her need, it is time to find a supplier. As mentioned earlier, browsing and searching the Net have limitations. The ideal way to proceed with this stage is automatically to 'comparison shop' using an application (or agent) that queries many suppliers, gathers details of suitable products and presents all the options to the buyer in some easy-to-compare format. A buying decision will be based on several factors, not just price. The buyer must consider the total value add of the supplier. This will include, among other things, warranties, customer service, promotions and special offers, reputation, payment terms and availability.

There are many sites offering simple price comparison on the Net. The earliest example is the BargainFinder agent from Anderson Consulting⁸. This agent helps users find the best price for music CDs. The buyer enters the artist and the title of the CD he/she wishes to buy and the agent visits many on-line CD shops to gather prices. From the returned list, the buyer can choose where to buy most cheaply.

BargainFinder has many limitations. It is easy for the individual

suppliers to block the agent if they do not wish to be compared directly on price with other suppliers. A large amount of maintenance is required to ensure that the agent knows how to query each different CD site in the correct way. If the sites change their layout then the agent must be reprogrammed to submit the search and extract the price correctly. Finally, buyers see only the returned price and may miss much of the value add that a supplier is offering, such as special offers or promotions.

Jango is a more advanced buyer agent that is offered by Excite⁷. It is based on a similar principle to BargainFinder. This price-comparison agent can offer comparisons in many product domains as well as an index of independent review material. Jango helps buyers to choose products and merchants.

A variation of comparison shopping agents are request-for-quotation (RFQ) services. These sites take a buyer's product specification and then send it to many suppliers for them then to return a quotation. This is exactly the same process that is used in big business, but the Net makes it possible for individuals to use this efficient way to buy. The efficiencies come from only having to describe one's requirement once, and then the onus is on the supplier to match one's specification and respond with a competitive offer. An example of this type of service is Auto-by-Tel⁹ which helps buyers find the best supplier for a new car. The buyer uses detailed forms to enter precisely what type of car is required and any optional extras. The service must get enough detail from the buyer to allow the supplier to quote accurately. The specification is then sent to all the appropriate car dealers who are members of the Auto-by-Tel system. The dealers are free to reply by telephone, fax or e-mail. This gives them ample opportunity to communicate all the important aspects of the deal being offered. Buyers report large savings on cars purchased via Auto-by-Tel. This is as a result of the lower

costs to dealers of finding customers—they do not need to advertise, the customers come to them.

This type of merchant comparison has great potential to change the way many classes of goods are sold. Cars, holidays, computers, personal finance and home electronics are all suitable goods to procure in this way. Markets will be created that allow buyers to issue requests like: 'I want to go on holiday on the 2nd July for two weeks to a tropical island to do scuba-diving and water-skiing'. Suppliers will watch the market for requests that they can satisfy and will reply to the buyer with a suitable offer. In this way conventional advertising is reversed with suppliers monitoring the market to see if they can satisfy buyers' needs.

Review and Feedback

Many of the techniques described for aiding the buying process rely on gathering opinion from users and directing this opinion to subsequent users who can use it to inform their buying decisions.

At every possible point in the buying process opinion, feedback and review should be gathered from buyers. Buyers' actions and decisions must be captured for future use. Where appropriate the buyers should be asked why they made a decision. This can then be used to advise similar buyers faced with similar decisions. Of equal importance are reasons why they changed their mind or rejected advice. Buyers that indicate a requirement and then do not buy are more likely to buy later. This data can be used for targeted advertising.

Where possible, buyers should be contacted some time after their purchase to complete the purchase review stage. Buyers should be asked what they thought of the product and the quality of service from the companies involved. This information builds into a valuable 'opinion database' that is used by later buyers. Of course, not all buyers will be willing to spend time commenting

on the products. To ensure that there is plenty of feedback, the feedback process must be incentivised with, for example, future discounts for users who contribute feedback. This kind of opinion gathering is only possible when the advisor functions and trading environment are truly independent from the suppliers. If the advisors act on behalf of the suppliers then their impartiality is compromised†.

A New Buying Environment

This article has examined the stages of buying and some technologies that aid on-line buyers. How will these be integrated to form a new buying environment? Clearly, a user wants convenience and to be able to carry out the whole purchase in one place. This hints at a new kind of site. One that is part magazine and editorial to interest and stimulate the buyer, part shop assistant to help and advise on the best products, part user-group to provide opinion and advice from like-minded people, part shopping mall to allow quick comparisons and part independent expert to rate products.

Sites with some of these characteristics are emerging and are termed *virtual communities*¹⁰. They are a new kind of intermediary that acts on behalf of buyers to serve particular niches. Importantly, these sites are independent and are not run by manufacturers. Virtual communities have a strong focus (books, music, computers etc.), they provide their buyers with valuable resources (product information, reviews, EWoM recommendations, news, chat rooms, access to expert opinion, and RFQ services). The aim is to help their buyers get the best deals at the best

† Organisations such as magazines are well placed to run buyer advisor services. Although magazines get significant revenue from advertising, they manage to balance their loyalties well and appear to the readers to be acting of their behalf.

A scenario

In this scenario our buyer is a home user who needs to buy a modem. Our buyer is not technical but likes to have all the relevant information before making a decision and is wary of exaggerated claims from manufacturers. Of prime importance is ease of set-up and reliable operation, and of course, a good price. Imagine that our buyer has a computer and wants the modem to surf the Internet at home. The requirement might be: 'I need a fast modem to surf the Net using my PC'.

The generic advisor is used to 'unravel' this requirement. The advisor would take the buyer through a series of questions that turned the qualitative requirement into a quantitative one. Today a modem of 56·6 kbit/s is considered to be fast but a short while ago 33·6 kbit/s was the fastest public switched telephone network (PSTN) modem available. Expert knowledge of the modem domain is needed to allow the advisor to take a high-level requirement such as 'fast' and map that to a precise specification such as 56·6 kbit/s. (This knowledge will need to be maintained and updated by experts. For example, it could be supplied by a computer buyers' magazine or generated by periodically surveying a group of expert users.) The generic advisor would also probe the buyer about the types of uses to which the modem will be put and the machine that is it intended for. At the end of this stage, the generic advisor will be able to suggest an actual specification for the most suitable modem.

Our buyer can now request quotations (using an RFQ system) from suppliers based on the suggested specification or ask a specific advisor for reviews of actual products. Imagine that our buyer asks the specific advisor to suggest a number of suitable products.

The specific advisor would need access to on-line catalogues, reviews and laboratory tests, user feedback, discussion groups and databases of products. Each product that is suggested would have a collection of associated review material. It may even be possible to generate dynamic comparison tables. From this the buyer will be able to narrow down to, say, two or three products.

At this point the buyer could use EWoM techniques to find opinion from buyers with matching circumstances. For instance, a short review by a user with the same computer and who required the modem for similar reasons:

Review by A. N. Other; Computer: Any Co. 266; Modem use: home surfing
Modem: XXX.

'The XXX was easy to set up with true plug and play. It has performed excellently and the telephony and fax software is great – a real bonus. I do wish I'd gone for the internal version instead as it is easy to fit, a little cheaper and much neater. I also considered the YYY but I've heard that their customer help desk is not very good. Overall 9/10.'

After reading a range of reviews like this, our buyer is able to benefit from others' first-hand experience. The buyer would now be much more confident about their choice of modem and might even switch to an internal modem based on the advice. At this point there would also be an opportunity to send a question to a chat forum to clear up any unanswered questions.

Our buyer has now made an informed decision about which modem to buy and feels comfortable enough with that decision to make a purchase. The advisor has taken the buyer from an imprecise high-level requirement to the point of spending money.

prices. The book seller Amazon, and the car finder Auto-by-Tel, are best thought of as virtual communities.

Perhaps the most important characteristic of virtual communities is that members help other members. This need not be in a direct way; for example, information about one user's purchases can be stored by the community and used later. Deal histories, when shared, allow other buyers to determine if a price is reasonable. Suppliers cannot engage in price discrimination if buyers know what other buyers paid. This kind of information sharing increases buyers' power and drives prices down.

One might imagine that sellers would avoid highly organised groups of buyers but there are in fact many advantages for some types of seller in participating in these markets. Virtual communities favour smaller companies who are able to sell products through word of mouth despite larger companies' marketing efforts to shut them out. This type of channel also provides a cheaper way to reach customers. Expensive media (TV, magazines etc) promotion is not required and this can more than offset lower profit margins. Other incentives include: gathering information about customers, global reach and the risk of not participating and thereby losing out to competitors.

Conclusions

This article has shown how a range of new technologies and on-line applications are changing the way buyers choose products and make buying decisions. On-line services will allow buyers an unprecedented level of information to guide them towards the most suitable purchases. These trends indicate huge changes ahead for retailing, marketing and advertising.

A new and important type of intermediary, the virtual community, is emerging. Virtual communities help buyers to increase their buying power by providing independent information, buying environments and tools, and forums for buyers to interact. These new buying environments will provide

such clear advantages that they may dominate retail in certain sectors. Some brands will therefore lose primary customer relationships as buyers look elsewhere for information and advice.

The shift in the retail landscape caused by the Net has begun. Sellers who respond to change and exploit the unique features of the Net will perform better than those who try to resist this inevitable change.

References

- 1 ENGEL, J. F. *et al.* Consumer Behavior. Dryden Press, 1994.
- 2 FireFly: <http://www.firefly.com/>
- 3 BARNES and NOBLE: <http://www.barnesandnoble.com/>
- 4 myLAUNCH. <http://www.mylaunch.com/>
- 5 TSANG, E. P. K. Foundations of Constraint Satisfaction. Academic Press, London and San Diego, 1993.
- 6 Personalogic. <http://www.personalogic.com/>
- 7 Jango shopping agent. <http://www.excite.com/>
- 8 The BargainFinder agent. <http://bf.cstar.ac.com/bf/>
- 9 Auto-by-Tel. <http://www.auto-by-tel.com/>
- 10 HAGEL, J., and ARMSTRONG, A. G. Net.Gain: Expanding Markets Through Virtual Communities. HBS Press, 1997.

Acknowledgements

The authors would like to thank PersonaLogic Inc. and LAUNCH Media Inc. for kindly allowing their screen shots to be used in this article. For further information about these products contact:

PersonaLogic: info@personalogic.com, +1 (619) 220 5800.

LAUNCH Media: fionap@2launch.co, +1 (310) 656 4326

Biography



Simon Steward
BT Networks and
Systems

Simon Steward joined BT in 1984 as an apprentice at BT Laboratories. After his apprenticeship he joined the Submarine Optics division where he worked on supervisory and control techniques for laser amplifier repeater systems. He graduated from the University of Warwick in 1991 with a B.Eng. in Electronic Engineering. Since returning to BT Laboratories, he has worked in the cybernetics group where he studied self-organising systems, and recently in the electronic commerce research group investigating emerging new media markets and trading environments.



Ian Videlo
BT Networks and
Systems

Ian Videlo graduated from the University of Southampton with a B.Sc. in Physics. He joined BT in 1981 where he spent several years working on the reliability of electronic and opto-electronic components. More recently he has been working within the on-line applications arena and in particular focussing on electronic commerce. His main focus is in facilitating the emerging generation of on-line marketplaces and trading applications.

Tim Titheridge, Rob Collingridge and Graham Shorrocks

The BT Intranet Complete Service

BT's Intranet Complete service uses a specially-extended version of the Microsoft Commercial Internet Service product set to provide a range of Internet services and applications within a secure virtual private data network. The service's many features include administration tools that enable customers to carry out their own on-line intranet administration, configuration and ordering.

Introduction

The growth in the Internet over the last five years has been phenomenal (now over 10% per month with currently in excess of 30 million users across the world). However, there are still major questions regarding quality of service and security. Intranets address these key issues by providing a private area of the Internet, generally within the corporate domain, offering improved security, closed user group facilities and guaranteed bandwidth.

The business benefits of adopting Internet applications within a secure virtual private network are widely understood. Intranets offer greater operational efficiency with lower costs compared to traditional communications models. Intranets also provide secure access to the Internet and hence a company's employees (or commercial partners) will have access to a range of Web applications via a common browser and will be able to move between the intranet and Internet seamlessly.

BT and Microsoft partnership

Industry analysts estimate that the intranet market will be worth \$28 billion by 1999 and its growth is currently outpacing that of the Internet by 25 per cent.

In response BT and Microsoft have formed an alliance with the aim of becoming the dominant force and largest vendor of intranet services worldwide. The partnership will use the networking capabilities of BT and Concert and the software products of Microsoft.

The BT *Intranet Complete* service is based on the Microsoft Commercial Internet Service (MCIS) product set,

previously called *Normandy*¹. MCIS is supported on Windows NT servers which can be operated across multiple server farms on a global basis. MCIS is a modular, highly scalable, comprehensive suite of applications including Web, news and mail which are all accessed via a common authentication and access control mechanism.

MCIS was originally intended for single Internet service providers and additional development was required to support multiple customers on a common platform. The development was carried out by BT and Microsoft development engineers collocated in Redmond (Washington State), who rapidly developed the intranet-specific code to meet the overall service requirements.

Business benefit

One of the obvious benefits of creating an intranet is in reducing document distribution costs. However, this is only the tip of the iceberg. Many companies are now realising that there are huge cost savings in using Internet technology to build new business applications and providing simple and consistent interfaces to existing ones. Intranet applications typically have lower development costs, lower support costs and lower deployment costs. It is also possible to integrate many of the existing legacy systems with Web servers and provide access to these back-end systems through a standard Internet browser. Couple this with a secure high-performance communications network, and the cost benefits are substantial. It is worth noting that recent polls suggest that over 40 per cent of companies are seriously investigating using the Web

Figure 1—Applications suite

within the corporate environment and many have appointed directors to devise an intranet strategy.

Intranet Complete Service Features and Applications

Intranet Complete service features

BT's Intranet Complete service is made up of several industry-standard Internet applications that can be accessed via a secure high-performance network offering a range of access options and access speeds. Much of the applications development has been focused on the administration model which offers customers the tools and interface to manage and administer their own intranet. A summary of the service features is as follows:

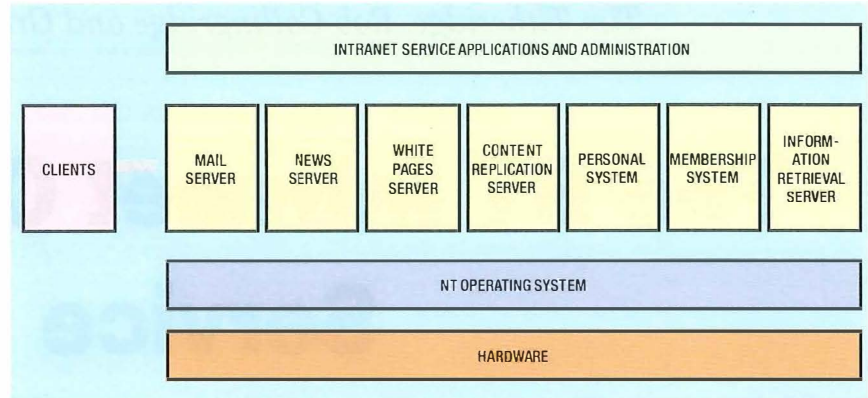
- single secure authenticated logon to all services;
- remote access via the Internet;
- administration interface based on the World Wide Web (WWW);
- closed user groups facilities available, administered by the customer;
- Internet Explorer and Netscape client browsers;
- firewall controlled access to the Internet with optional information caching;
- SMTP/POP3 e-mail facilities and mail routing for SMTP mail;
- news services with public and private news groups;
- content creation and upload tools;
- document storage;
- 500 Mbytes of space for hypertext markup language (HTML) and native files.

- search facilities (HTML content, Microsoft Office files, and newsgroups);
- software download (authenticated HTTP);
- partitioned directory service at a company level or down to closed user groups with WWW search interface;
- service helpdesk available 24 hours a day, seven days a week; and
- management reports.

Applications suite

At the heart of the Intranet Complete applications suite is MCIS, which was initially developed as an internal product to support the Microsoft Network (MSN). Some limitations of the current NT operating system meant that additional software development was required to provide a high-performance, resilient, and scalable solution within the Internet service provider (ISP) market. MCIS was designed to integrate with the NT operating system and to provide a solution to address the ISP market. This was achieved by providing a set of components or building blocks based on the common Internet technologies and support services.

The Intranet Complete applications suite is shown in Figure 1. MCIS comprises several individual servers but only seven are included in the service at this time. An additional feature is the common system administration architecture which sits above MCIS. This provides the customer with a consistent and unified way of managing users and user profiles and administering



individual services within the intranet domain.

The seven MCIS components within the Intranet Complete design are described below:

- The *Internet mail server* provides simple mail transfer protocol (SMTP) and post office protocol, version 3 (POP3) mail and is scalable to millions of users. The next version will support the Internet message access protocol, version 4 (IMAP4).
- The *Internet news server* provides a commercial grade implementation of network news transport protocol (NNTP) and supports public, read-only, moderated, and authenticated news groups.
- The *White Pages server* is a directory service, accessed through a Web browser, which allow administrators to quickly create an easy-to-use user interface. Users can create, modify and locate entries within the White Pages service.
- The *content replication system* provides an efficient way to move content across servers. It replicates any type of content from one or more remote content servers to multiple destination content servers, across the Internet or intranet.
- The *personalisation system* provides personalised content to users using a database to store users' preferences and historical information about user access.
- The *membership system* stores user information in a database to

Figure 2—The Intranet Complete service

provide a unified distributed authentication and access control system across all of the MCIS components. The system includes facilities to authenticate users, authorise access to MCIS resources, and generate billing events for processing by an external billing engine.

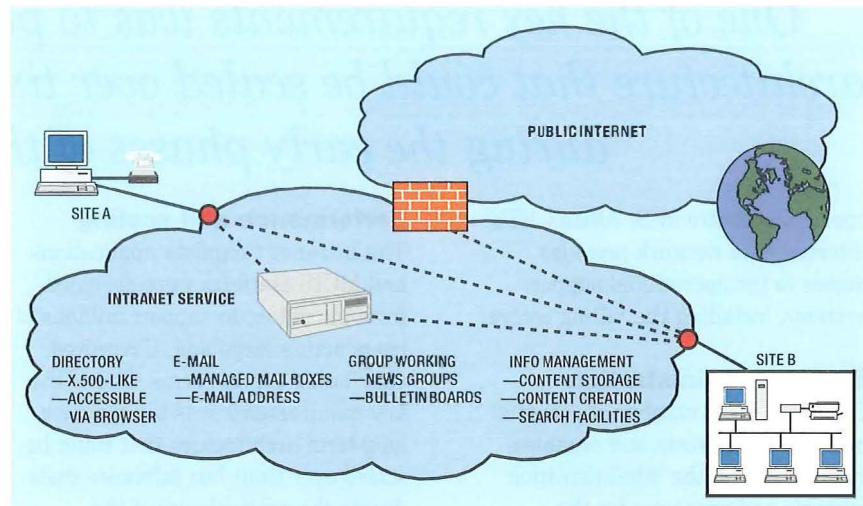
- The *information retrieval server* is a distributed search engine solution that provides full-text searches across multiple Web sites, news groups and Microsoft Office files.

BT and Microsoft are aware of the commercial opportunities in the intranet market and the customer demand for a wider range of services over the next five years. The Intranet Complete service will continue to evolve and new features will be added in response to customer requirements.

The Intranet Architecture

Figure 2 shows an illustration of the Intranet Complete service in which a virtual private data network (VPDN) offers access between a customer's sites, the Internet and the intranet-embedded applications that have been described in the previous section.

The realisation of this on a global basis represents a major engineering task. In the longer term it is expected that multiple application server farms will be established close to the major customer concentrations but this can be done only once the customer demand has been fully assessed. Deployment of multiple server farms on a speculative basis would be extremely expensive and for this reason BT has initially deployed one server farm in the UK. The aim is to establish a presence in BT's home market and then extend the service internationally. As customer demand grows, additional server farms will be established in appropriate locations such that the overall cost of providing the service is minimised. The key to this will be the



trade-off between the cost of international bandwidth against the cost of establishing and operating new server farms.

Network architecture

The network is a key feature of the service and provides customers with secure access to the Intranet Complete application servers within the data centres.

The three principal access mechanisms are as follows:

- Access to the Intranet Complete data centres or communication between communities of interest with the customer domain is provided via secure logical connections across a VPDN as shown in Figure 3.
- Authenticated dial access is provided to the data centre at which point the public switched

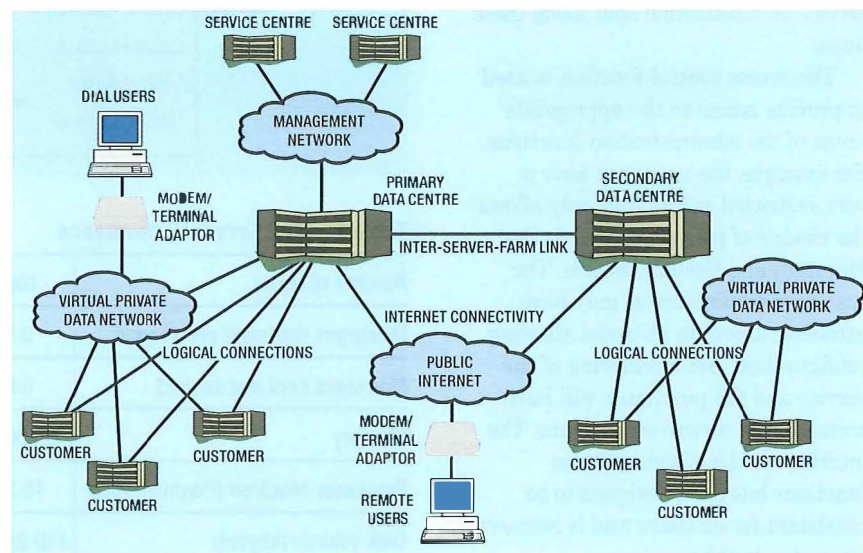
network (PSTN) or integrated services digital network (ISDN), connection is terminated by a 'digital' modem and routed to the appropriate server.

- Remote users may access the service via an encrypted tunnel across the public Internet. This access mechanism will be used in the situation where local dial access is not available.

Additional server farms will offer increased capacity and resilience, and links will be established between them to provide synchronisation of the server farms for both management data and customer content.

The Intranet Complete server farm is located at Telehouse in London's Docklands and is managed remotely by Digital at its operations centre in Newbury. The overall service is supported by BT's network

Figure 3—Top-level network design



One of the key requirements was to produce a long-term architecture that could be scaled over time but minimise costs during the early phases of the service.

operations centre at St Albans. BT's internal data network provides access to the operational support systems, including the billing centre.

Service administration

Given the large number of potential users, user interests and organisational changes, the administration systems and processes for the Intranet Complete service have been a key consideration. From the outset the aim was to provide customers with the ability to administer their own intranets by offering on-line administration, configuration and ordering capabilities.

The approach was based upon the model used in the Campus range of products, where a customer administrator is established who can configure major aspects of the service. The administrator can create new users, establish collaborative working groups, generate and store content, order more resources, all with no involvement from BT's support team. The operational process model for this reuses elements from the standard BT business model, recognising the fact that many of the processes are applicable to both the customer domain as well as BT. Thus the Configure Service process has two elements: those applicable to BT in establishing the initial intranet and those applicable to the customer for the ongoing administration of the intranet. Figure 4 shows a logical view of the service administration split along these lines.

The access control function is used to provide access to the appropriate areas of the administration functions. For example, the user may have a very restricted access that only allows the change of passwords within the Maintain and Restore process. The customer administrator may have extensive access to all areas allowing configuration and monitoring of the service and the purchaser will have access to the service order forms. The interface to the administration functions has been designed to be consistent for all users and is accessed via a standard browser.

Performance and scaling

The Intranet Complete applications and MCIS platform were designed from the outset to support millions of users across large and, if required, distributed server farms. One of the key requirements was to produce a long-term architecture that could be scaled over time, but minimise costs during the early phases of the service. A performance model has been developed to ensure that the servers deployed in the data centre (see next section) are scaled to meet the initial market predictions. The model is based upon laboratory measurements of the MCIS components, carried out by Microsoft. User

profiles are entered into the model and the results define the corresponding processor and storage requirements. Alternatively the model can accept processing power and storage parameters etc. and determine for that particular configuration how many users can be supported. In addition, the model allows the cost of supporting an individual user to be determined, which is a key input to the production of the service business case.

The results of running the model for the mail application are shown in Table 1. In the particular case illustrated, once the number of users exceeds about 15 000, multiple

Figure 4—Service administration—logical design

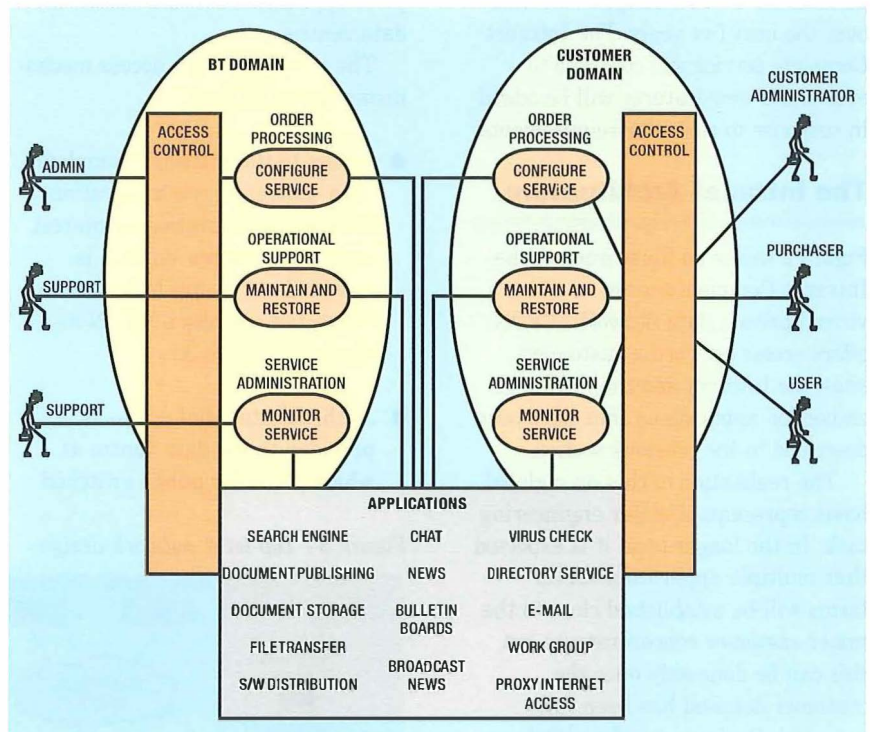


Table 1 Mail Server Performance

Number of users	1000	10 000	20 000	50 000	100 000
Messages retrieved per second	0.56	5.56	11.11	27.78	55.56
Messages sent per second	0.67	6.67	13.33	33.33	66.67
Memory	60 304	63 044	66 089	75 222	90 444
Processor required (Pentium MHz)	15.53	155	311	777	1553
Disk volume (kbytes)	120 350	303 500	507 000	1 117 500	2 135 000

Figure 5—Data centre design

processor configurations are required.

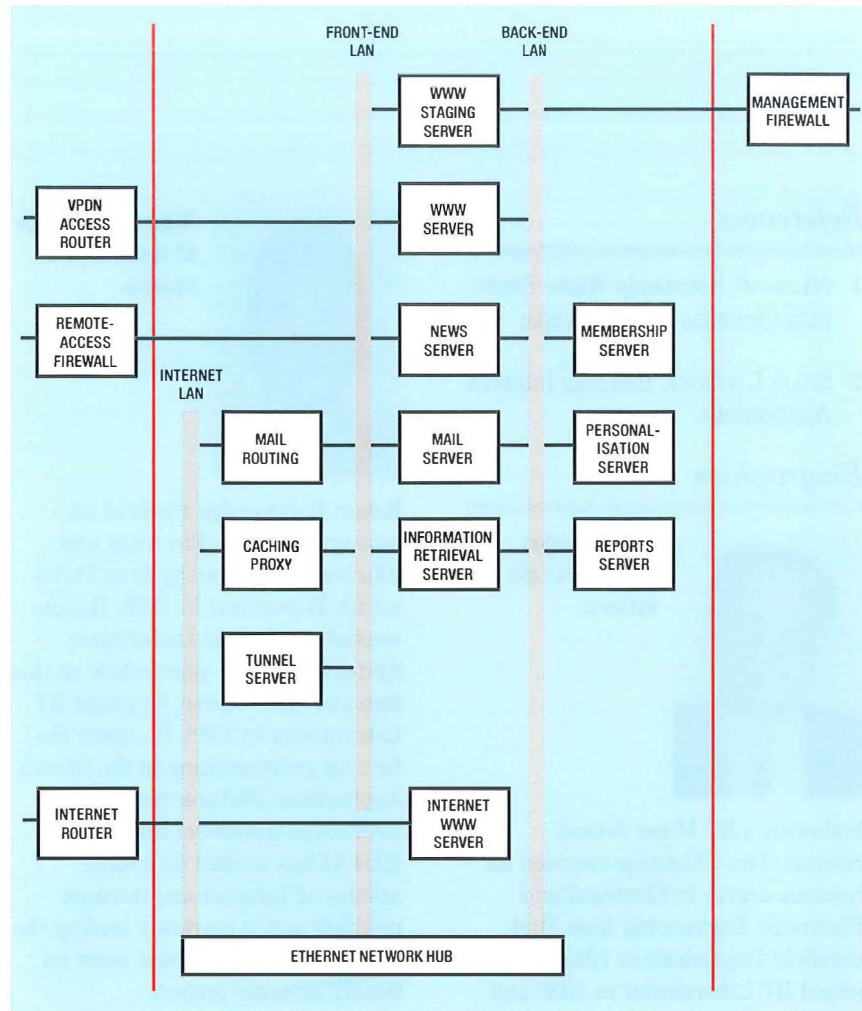
Data Centre Design

The user model within an ISP service is very different from that within an intranet. MCIS was initially designed to support the ISP user model and, as such, additional development was required to implement the intranet user model and the service and management features required. A team of BT and Microsoft engineers carried out this development.

The specific features include Web-based administration tools supporting a range of features including closed user groups. The management interfaces allow creation and deletion of these groups, control of users within groups, and allocation of resources to these groups; for example, access-controlled news groups. The administration tools also provide security partitioning between intranets. It is imperative that companies residing on shared server hardware are not visible to one another, and so secure partitioning is provided within closed user groups, on-line address books, etc.

With the Intranet Complete platform in place it is possible to produce advanced services using the basic Intranet functionality. There is a wide range of added-value services that can be targeted at specific market segments. These include electronic commerce, push technology and the integration of legacy databases or hardware interfaces. Services such as CampusWorld and HealthNet have previously required dedicated hardware to provide service, but in future could reside on the intranet service infrastructure.

Within the data centre, resilience and performance are achieved by using a set of high-performance NT servers with the various intranet applications distributed across them. The configuration is flexible and is designed to allow additional servers to be added to provide a more distributed loading on any particular



intranet component, by the use of a round-robin domain name service (DNS). If a server fails, it is removed from the round-robin DNS sequence dynamically. To avoid loss of access to data or loss of the data, content is replicated using the content replication system (CRS). Content can be replicated between servers and potentially across server farms in the same or even in different countries.

An example data centre configuration is shown in Figure 5. For simplicity the diagram does not show the exact configuration of the server farm where some MCIS components are duplicated for resilience.

Within the data centre there are three internal networks each physically isolated from the other by dual homed servers (using two network interface cards). A front-end local area network (LAN) provides access to main service features. The Internet LAN provides users with access to the Internet via the caching proxy. It also allows news feeds and mail to enter the server farm. The back-end LAN is a management LAN and can be accessed remotely by BT engineers.

The Digital support team based in Newbury can also access this LAN. All authentication and database requests also pass over this LAN.

Remote access to the server farm is gained via any ISP or from the Internet. An encrypted tunnel session is established through the remote-access firewall, through to the tunnel server.

Summary

The Intranet Complete service supports a range of commercial standard Internet services and applications within a VPDN. An intranet improves communications within an organisation and allows users to gain easy access to a range of information through a common interface. Customer management and administration tools are provided and the user interface is via a standard Web browser.

The Intranet Complete service provides a platform on which to develop corporate business plans and processes and leads to improved productivity and reduced costs.

References

- 1 Microsoft Normandy White Paper.
<http://www.ms-normandy.com>
- 2 SHAFÉ LAWRENCE. Building Intranet Applications.

Biographies



Tim Titheridge
BT Networks and
Systems

Following a BT Minor Award Scheme, Tim Titheridge received an honours degree in Electrical and Electronic Engineering from Huddersfield Polytechnic in 1982. He joined BT Laboratories in 1983 and spent the first three years designing high-speed LANs. This was followed by four years working on advanced distributed systems followed by three years in CSO as the Network Strategy and Design Manager. Since that time he has worked on several of BT's data services and is currently leading the BT Intranet project team.



Robert Collingridge
BT Networks and
Systems

Robert Collingridge received an honours degree in Electrical and Electronic Engineering from Portsmouth Polytechnic in 1989. Having worked at Marconi Underwater Systems for three years while on this thin sandwich course, he joined BT Laboratories in 1989. He spent the first six years working in the Speech Applications Division designing interactive speech services. Since 1994 he has worked on a large number of Internet and intranet products and is currently leading the applications development team on the BT Intranet project.



Graham Shorrock
BT Networks and
Systems

Graham Shorrock is a solution designer within the Internet and Multimedia Applications directorate at BT Laboratories. He graduated from Loughborough University with a degree in Electronics and Electrical Engineering. On commencing work at BT Laboratories in 1979 he worked on some of the first ISDN transmission systems and developed Modem 4192, a baseband echo cancelling modem using analogue signal processing technology. There then followed a period working on high-speed optical transmission systems and was involved in the design of the first synchronous transmission system to be deployed in the UK. More recently, he has been involved in the solution design of a number of BT multimedia and Internet-based products. Being design authority for CampusWorld CampusConnect, IntranetComplete and recently Concert's global VPDN service and the new IMS business platform.

Intranet-TV— Video Streaming for the World Wide Web

Intranet-TV is an experimental video-enabled Web site that has been running since March 1997 on the intranet at BT Laboratories. Audio and video material is delivered via the intranet using standard Internet protocols to desktop PCs. Live and on-demand access to video is provided, creating an entirely new communication channel within the organisation. This article looks at the technology behind Intranet-TV and at how the communication channel it supports is being used.

Introduction

Despite the fact that most office desks have a cathode ray tube on them, whenever there is any video-based material to watch at work, people generally stop what they are doing and gather around a traditional TV set to view it. Intranet-TV can literally change the way people work by bringing networked visual business information direct to the PC monitor on the desk.

The trial Intranet-TV World Wide Web site has demonstrated the use of Internet protocol (IP) video streaming technology within a corporate environment, providing an experimental showcase for many different Intranet applications such as on-line team briefings or business TV. More importantly, the site has given many people a glimpse of the video-enabled future of the WWW, such that they can consider how it should be exploited to deliver real business benefit in areas such as marketing, training or corporate communications.

Video-Enabled WWW

Video enabling the WWW represents a natural evolutionary step forward for the Internet. Initially the WWW supported only text, and tended to be the preserve of academics and research scientists. Next, support for graphics was introduced into WWW browsers, allowing images to be embedded within WWW pages. BT's intranet provides a good example of the

informative mixture of text and graphics on its numerous WWW sites, such that most users would find the intranet a dull place if it did not contain any graphics. By the year 2000, watching video within WWW pages will be as natural as seeing graphics on a WWW page is today, and a new generation of users will probably find it difficult to imagine the WWW without video.

Streaming video

Video streaming on the WWW involves the real-time interactive transfer of data from a remote video server to a client browser over an IP network. While it has previously been possible to download non-streamed video clips from a hyperlink reference in a Web page, this approach results in 'download delay' as the file must be downloaded in its entirety to a temporary file stored on local hard disk before it can be viewed. With streamed video, the client browser renders video and plays audio soon after data packets are received from the network. As a result, with streamed video, access latencies to visual material embedded within a WWW page are reduced to the order of a few seconds. Interactive operations such as rewind, fast forward, pause, and stop can be performed on streamed media through a control channel back to the video server. Additionally, live *webcast* video can be streamed in real-time across the network to many viewers.

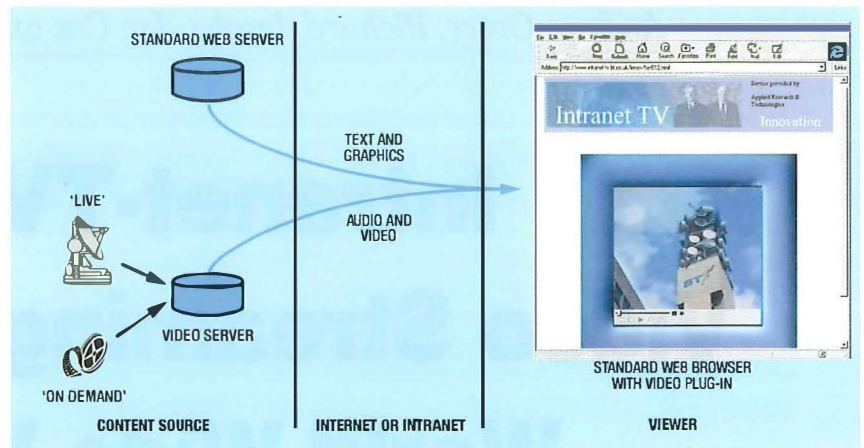
Figure 1—Schematic representation of video streaming

Figure 1 shows a sample video-enabled WWW environment. A standard WWW browser is enabled for viewing video by installing a browser *plug-in*. This video-enabled Web browser is then used to access a Web site hosting the video content. The user is presented with a video window embedded within a Web page and individual video frames are displayed as and when they are received.

Image quality

There is an intrinsic link between network bandwidth and the quality of visual image that can be delivered into a WWW browser. As video and audio are streamed in real-time, network throughput must be equal to or greater than the streaming rate, otherwise quality becomes degraded, and in the worst case becomes unacceptable. Table 1 lists typical video image window sizes and frame rates that are available at a range of different access rates over an IP network connection.

Different applications have different requirements for image quality, and image quality itself is dependent upon a number of factors, including motion within a video clip, video window size, screen resolution and viewing distance. At home, entertainment services are delivered to a TV set, with a typical viewing distance of several metres. For acceptable image quality, this would require full screen IP video at 2 Mbit/s. At work, business information would be delivered to a PC monitor, with a viewing distance of



perhaps 600 mm†. In this environment, a message can be effectively communicated using a 160×120 pixel window at narrowband rates of 28.8 or 56 kbit/s.

Convergence

Video over IP to the desktop exploits the convergence between narrowband video encoding, real-time Internet protocols, and the multimedia capabilities of client PCs.

Video encoding

There has been rapid technical innovation in the area of IP video streaming over the past few years, driven partly by advances in narrowband video-compression technology. Several streaming products are available from different vendors, such as RealPlayer from RealNetworks¹, WebTheater from Vxtreme (now owned by Microsoft), and Netshow from Microsoft².

Whereas video encoding schemes traditionally used by telcos for video-on-demand trials were based upon

† The BT display screen equipment user guide recommends a viewing distance of 500–700 mm for desktop PC monitors

standards ratified by international standards organisations, each of the Internet streaming systems mentioned above uses its own proprietary format for encoding/decoding video and audio. Video that has been encoded using one system tends to be incompatible with another system, although it is very likely that market forces will, in time, establish a de facto standard for IP video streaming. In the past, standards tended to be agreed over a period of several years. Now, the rate of innovation and change has increased tremendously, and the above vendors are enhancing their products almost every six months.

Real-time IP protocols

For real-time networked multimedia such as audio and video, timely delivery of data is more important than reliable transmission of data. If an IP data packet containing a sample of video is discarded or delayed by a network router, then waiting for reliable transmission of that packet introduces unacceptable end-to-end delay to the overall flow of media.

The IP protocol stack, shown in Figure 2, supports unreliable data transmission on a best-effort basis using unreliable datagram protocol (UDP) over IP. The real-time protocol (RTP) provides framing and timing structure to audio and video streams. In certain environments, however, network firewalls do not permit the transit of UDP traffic. In such circumstances, video and audio can

Figure 2—IP protocol stack




Table 1 Streamed video qualities

Network bit rate	Image size	Frame rate
Modem, 28.8 kbit/s	160×120 pixels	4 frames per sec
Modem, 56 kbit/s	160×120 pixels	10 frames per sec
ISDN, 128 kbit/s	320×240 pixels	7 frames per sec
LAN, 500 kbit/s	320×240 pixels	15 frames per sec
LAN, 2 Mbit/s	640×480 pixels	25 frames per sec

any architecture for Web-based video streaming must take into consideration the existing legacy of network infrastructure, network utilisation and employee demographics

be streamed successfully using hypertext transfer protocol (HTTP) over transmission control protocol/Internet protocol (TCP/IP), albeit at the expense of end-to-end delay during periods of network congestion. Typically, the client initially connects on a TCP port to negotiate client configuration and bandwidth. If UDP connections can be established then this is used. If not, then an attempt is made using HTTP and separate ports for each media type. If this fails, then all media types are multiplexed onto port 80 using HTTP. To the network this last mechanism is indistinguishable from a standard Web browser connection, although some routers may in future be able to distinguish the different flows based on traffic characteristics.

Client PC capability

As already mentioned, a standard WWW browser is video-enabled by downloading and installing plug-in software from the vendor's WWW download site. One of the main functions of the plug-in is to decode compressed frames of video. As the decode process is running in software, the client PC must have sufficient processing power to complete the decoding task.

As shown in Table 2, at 56 kbit/s a 120 MHz desktop PC can comfortably handle narrowband video streaming. As higher-rate video presents a greater challenge for the desktop PC today, narrowband video streaming is matched to the capabilities of the existing office infrastructure.

Table 2 Processing power at different streaming rates

Media streaming rate	Percentage CPU required (120 MHz processor)
28 kbit/s	30%
56 kbit/s	65%
128 kbit/s	90%
500 kbit/s	>100%

To hear an audio track, the client PC must have a sound card. In the corporate desktop market, not all PCs are supplied with a sound card and this fact may create a barrier to the use of Intranet-TV applications. However, most portable PCs are currently supplied with a sound capability, and the emergence of streamed audio/video on the corporate intranet creates the requirement for desktop machines to be supplied with sound cards pre-installed.

The client PC must also be connected to an IP network, with either a dedicated or dynamically allocated IP address. As a result, streamed video cannot be provided to PCs connected to non-IP networks, where access to the Internet/intranet is provided through network gateways.

System Design and Deployment

Very few organisations enjoy the luxury of a greenfield site development, so any architecture for Web-based video streaming must take into consideration the existing legacy of network infrastructure, network utilisation and employee demographics.

On the BT intranet, the Intranet-TV architecture has evolved on an organic basis. Streaming has been introduced at a small number of sites without any changes to network

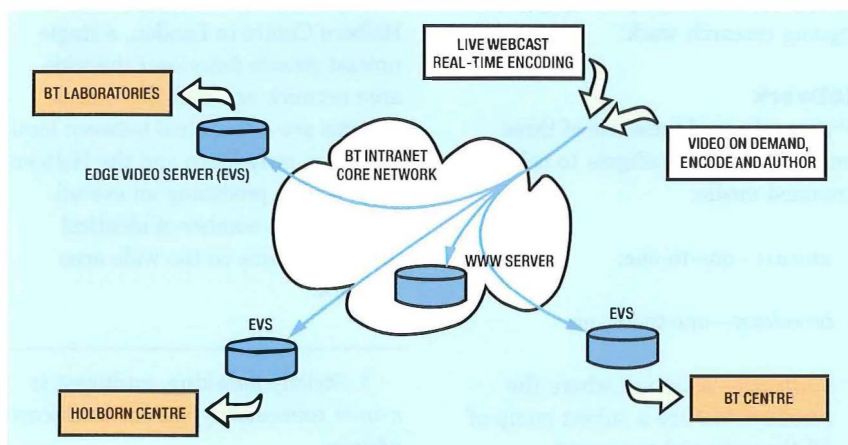
infrastructure. There is also a pragmatic reason for taking this approach: use of video streaming is likely to alter significantly some aspects of the way in which an organisation communicates, so it makes sense to follow an adaptive approach—to allow the organisation to change, and to adopt video streaming for those users for which it is most beneficial.

Figure 3 shows the architecture used for the deployment of the experimental Intranet-TV system. With this architecture, the emphasis has been to reduce network traffic over the wide-area core network by aligning the deployment of edge video servers (EVS) with employee demographics. As a result, the majority of video and audio traffic flows over the local building or site networks, not over the wide-area network. End-users first access the Intranet-TV WWW server which then dynamically generates a video-enabled WWW page that connects the video client embedded in the page to the most appropriate video server.

Content

The Intranet-TV architecture supports two basic modes of streaming operation; namely, live webcasting and on-demand access. A third hybrid mode of operation is also supported: live webcast with subsequent on-demand playback.

Figure 3—Intranet-TV architecture



For live webcasting, an event is streamed as it occurs, with the user watching the transmission at the scheduled time. Live webcasting requires a high-performance capture PC to perform real-time encoding of the video input as the event proceeds.

The production of on-demand video content is a two-stage process: it is initially captured/digitised to a file, then it is encoded into a format suitable for streaming. On-demand material can therefore be encoded to a higher quality as the duration of the coding process is not time critical.

For the hybrid live webcast with on-demand playback, the media stream is transmitted live and simultaneously stored on a server for subsequent on-demand access.

Content management

On-demand content provides a significant challenge in terms of content management. With the Intranet-TV architecture, content is replicated onto each of the edge video servers. Replication can be a heavy-handed solution and this approach can quickly become unmanageable, given the volume of content and a large number of edge servers. What is needed is the equivalent of the proxy cache used for Web documents. This provides an adaptive mechanism that automatically locates content closer to the point of use. Unfortunately, proxy caches do not currently support streamed media, and so distributed content management remains a high priority goal for ongoing research work.

Network

IP networks could use one of three communication paradigms to deliver streamed media:

- *unicast*—one-to-one;
- *broadcast*—one-to-all; or
- *multicast*—one-to- m where the receivers (m) are a subset group of all the potential receivers[†].

On-demand streaming uses IP unicast. To allow interactive control of a video clip, each receiver needs its own separate stream from the server, and so the network must be dimensioned to cope with the maximum number of concurrent accesses allowed to the video server.

For live webcasting, either broadcast or multicast could be used to allow multiple receivers to watch a single stream. In practice broadcast is not used for streaming as data packets are sent across all networks to all receivers, irrespective of whether or not anybody wishes to receive a stream. This approach would be wasteful of resources in the network and in the receiver. Multicast[†] provides the mechanism by which multiple receivers can each choose to receive a live webcast stream.

Network-layer IP multicast operates at the router level, where the routers replicate IP data packets to multiple recipients. It makes efficient and scaleable use of both server and network resources. In reality, however, the majority of intranet routers are not multicast enabled. For this reason, Intranet-TV currently uses *application-layer* multicast, whereby the video server application software, rather than the router, replicates IP data packets to multiple recipients, thus resulting in multiple unicast streams accessing a single live feed. In the experimental Intranet-TV architecture the edge video servers form a hierarchy of application-layer multicast servers. When a live event is webcast from BT Laboratories to BT's Holborn Centre in London, a single unicast stream flows over the wide area network and multiple unicast streams are established between local end-users in Holborn and the Holborn video server, producing an overall saving in the number of identical unicast streams on the wide area network.

[†] Strictly speaking, multicast is n -to- m connectivity via one multicast address

Quality of service

Live events naturally generate a peak load on the infrastructure whereas with on-demand content, demand is averaged out over time. Even using narrowband streaming, the available network bandwidth may become insufficient to support a popular live webcast. For most IP networks it is not yet possible to partition network utilisation by application or application class, although this is an area of considerable research effort. As a result, over-utilisation may well affect other intranet services, as well as degrading video streaming quality. In fact, standard TCP-based Intranet services are likely to bear a greater part of the burden as they are well behaved and back-off during congestion, whereas real-time applications such as video streaming are more likely to use UDP, a necessarily more lightweight protocol that has no mechanism for avoiding congestion.

Security and access control

Even within an intranet environment there may still be areas of access control and network security that need careful consideration; for example, router-configuration policy may vary across different areas of the intranet to separate networks with differing security requirements and ratings. Access control to video-enabled Web pages may be required if they contain business sensitive information. At a basic level, Web server access control may be used to limit the availability of the video-enabled pages although this is easily circumvented given knowledge of the path to the video streaming content. A better solution is available from some video streaming server vendors who have started to incorporate access control directly into their streaming products. As well as basic access control, these also provide support for other access models such as calendar-based, time-based, user-based and client-based authentication.

Video-enabled communication channel

Choice is probably the most important factor to consider when discussing the communication channel created by Intranet-TV. From the beginning it has been seen as an additional communications channel, not a replacement for more traditional media. The prime objective is to assess where it fits into the range of tools available to the communications community. It differs from traditional media in that it can be broadcast live or made available on demand from any user, yet it is still passive—passive, in the sense that the end user has to seek the information offered. It gives the choice and viewing decision to the user. Watch now or later? Alone or with colleagues? It also puts increased responsibility upon the users to seek out personally relevant information and not rely upon more traditional media to send it to them.

Managing the change from active to passive information requires a degree of thought, planning and communication. Making information available to users is not the same as targetting specific audiences. It also changes the balance from information as power to information as

empowering and long term can speed the decision-making process and flatten structures.

Publishers require educating in the use of this channel, to consider whether material normally produced for other channels can be adapted to make it more appropriate for Intranet-TV and whether Intranet-TV could/should become the main information source for their particular material. To minimise costs this has to be borne in mind at planning and production stages. This is particularly relevant to corporate TV programmes and any use of the channel for more personalised videos.

Intranet-TV will eventually move from a passive to an active channel, with significant corporate communications or business critical material being broadcast. This change will need to be carefully managed as too much or inappropriate material pushed to the desktop will extend the current information overload scenario with possible impact on operational efficiency. To combat this, information management standards will need to cover use of this and other intranet tools to ensure quality and appropriateness of material pushed.

Intranet-TV Applications

Since March 1997, the Intranet-TV site has hosted several different applications that utilise streamed video.

Business TV

The initial focus for the Intranet-TV Web site was the fortnightly BT business TV programme: *vision*. This programme is traditionally distributed via satellite requiring a satellite receiver, set-top box decoder and conventional TV set at all viewing locations. Intranet TV can ensure that employees are kept up-to-date with the latest corporate news accessible live and on-demand from their desktop.

Figure 4 shows the contents of a WWW browser accessing business TV on-demand, with a 160x120 pixel video window in the top left-hand corner of the screen. Underneath the video window is a table of contents for the programme, providing hyperlinks for non-linear access to different stories within the programme. On the right-hand side of the browser is a frame that is synchronised to the audio and video track. The synchronised frame

provides detailed graphical or textual information relating to the current news story. In the example, the presenter is describing a new joint venture in Italy and the graphic reinforces the story with details of the partners in the proposed venture.

Training

Streaming technology provides an efficient way to disseminate training information throughout the company; for example, safety advice. Alternatively, a sales team could be informed of a new product or sales drive.

Figure 4—Business TV, audio and video with synchronised Web pages

Company	Percentage
BT	26%
Mediaset	25%
Telenor	20%
INA	10%
Italgas	9%
BNL	10%

Figure 5—Safety training video.
Audio and video only

Figure 5 shows Intranet-TV delivering safety training material on-demand into a browser. Space on the Web page surrounding the video window can be used to carry supporting text or graphics, or in this case to carry the Intranet-TV 'brand'.

Team briefings

A team leader can use streamed video to brief his or her team 'live' without team members having to be physically present. This facility can deliver clear benefits if the team is geographically dispersed by reducing the need to travel to a common location. Additionally, if the live briefing is captured to disk for subsequent on-demand playback, team members can watch the briefing at a later time that suits them.

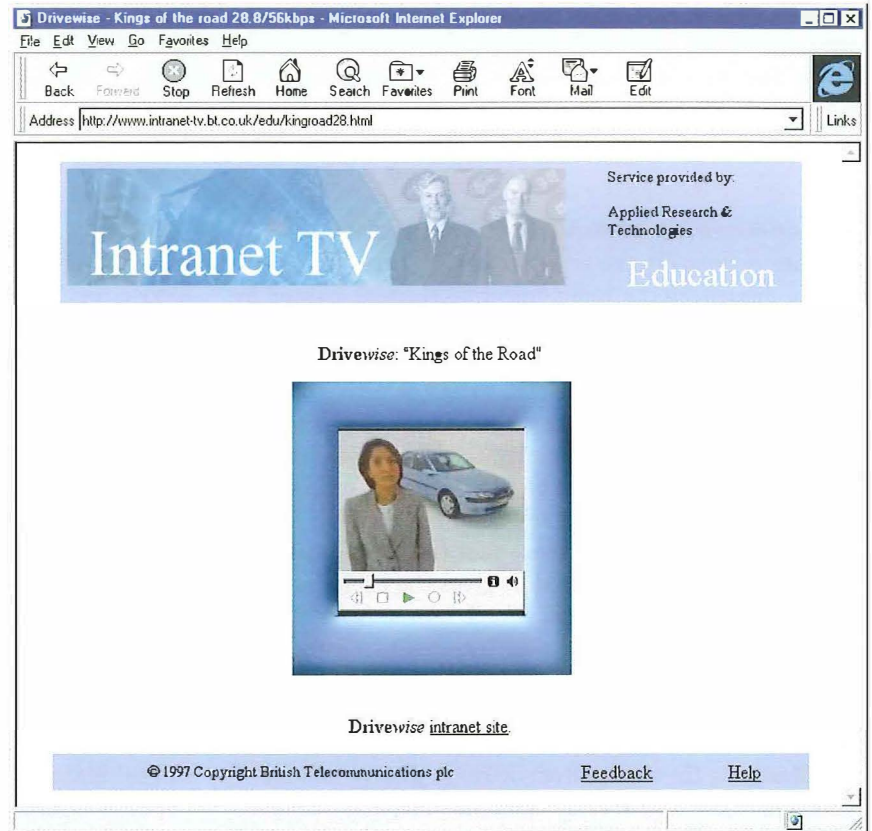
Within Applied Research and Technologies Department at BT Laboratories, the Intranet-TV site has been used to host on-line departmental briefings, for both a live audience and on-demand playback.

Personal presentations

Personal publishing of video-rich material can be linked to an individual's homepage. For example, an analyst could record a presentation, and make it available from a Web page, rather than spending considerable time and effort giving the same presentation to different audiences at multiple locations. Alternatively, senior managers could communicate their vision to employees and develop a more personal, approachable relationship with their people.

Site Usage

The Intranet-TV site receives, on average, 300 accesses per day. Over 7000 accesses of the software download page for the video plug-in have been made. Usage patterns are emerging; for example, there are more accesses on a day on which there is a live webcast such as a departmental briefing and there are



peaks of on-demand interest once a new Business TV programme has been published on the site.

The BT intranet extends globally, and Intranet TV has had the occasional user accessing the site from international locations including the USA, Australia, Germany and Japan, although the majority of site users are currently located at BT Laboratories.

User feedback

Feedback has been very positive. Several users have commented that they find it useful to play the videos such as Business TV 'in the background' while doing other tasks, for example dealing with their e-mail. They have found that they can quickly switch their attention to the video if they hear something that interests them.

Contrary to initial expectations, few people have commented about the size of the video windows. People can place the window to the side of their PC monitors and use the rest of the screen to do other things. People have said that it is difficult to read text or slides that appear in the video window during a live webcast. For this reason, slides for a live presentation are obtained in advance and are made available in a frame on the WWW page.

Conclusion

This article has described emerging technology for IP-based video streaming, and has outlined how this technology has been used to create an experimental platform for video-enhanced WWW applications on the BT intranet. Early indications of the potential for the video-enabled WWW are encouraging and have confirmed the team's strong belief in the power of this communication channel.

Concerns have been raised about the impact of streamed video with regard to increasing levels of network traffic. The Intranet-TV architecture seeks to reduce traffic over the wide area by replicating content close to its intended audience. While in the future the corporate network may support broadband multimedia communications, by streaming at narrowband rates over the existing internal network the Intranet-TV site has been able to demonstrate and explore the benefits of a video-enabled intranet today.

Acknowledgements

The authors would like to thank John Verdon and Ross Chestney in Corporate Relations Department for discussions on delivering the 'Vision'

programme via the BT intranet, and John Smith who captures and encodes this programme. Within Applied Research and Technologies, Roger Merton has provided much of the video material and Richard Jeffery has been responsible for encoding and loading it onto the Intranet-TV site. Intranet-TV relies upon network connectivity provided by the Futures Testbed at BT Laboratories, and by other internal networks. Our colleagues who run these networks are much appreciated.

References

- 1 RealNetworks.
<http://www.real.com/>
- 2 Microsoft Netshow.
<http://www.microsoft.com/netshow/>
- 3 O'NEILL, A. Internetwork Futures. *BT Technol. J.*, Apr. 1997, **15**(2), p. 226.

Glossary

HTTP Hypertext transfer protocol
IP Internet protocol
RTP Real-time protocol
TCP Transmission control protocol
UDP Unreliable datagram protocol
WWW World Wide Web

Biographies



Andrew Grace
BT Networks and Systems

Andrew Grace joined BT Laboratories in 1988 having graduated from Plymouth Polytechnic with a B.Eng. in communication engineering. He spent the first six years of his career applying novel data mining and visualisation techniques to large quantities of performance data from BT's core transmission network. In 1994 he moved into the distributed systems group to lead research work on quality of service in open distributed systems. He is currently team leader for IP video streaming research. He is a Chartered Engineer, and a Member of the Institution of Electrical Engineers (IEE).



Richard Jacobs
BT Networks and Systems

Richard Jacobs graduated from Southampton University in 1978 with a B.Sc. in Acoustical Engineering. He joined BT Laboratories in 1986 from the aerospace industry where he managed a team developing flight control software simulations. After a period working on speech-recognition applications, he joined the Internet applications team and worked on the development of CampusWorld and Healthnet. He joined the distributed systems group in 1997 where he is responsible for research into video-streaming architectures. He is a Chartered Engineer and a member of the British Computer Society (BCS).



Jon Cox
BT Networks and Systems

Jon Cox joined BT Laboratories in 1987 having graduated from the University of Liverpool with a B.Eng. degree in Electronic Engineering. He spent five years working on supervisory and fault location techniques for submarine optical systems before embarking on a three-year investigation into the design and evolution of corporate networks. After a period spent modelling the theoretical roll-out of broadband services across the UK, he began working on real-time Internet applications and joined the IP video-streaming team in 1997. His current responsibilities include the design and support of the Intranet-TV Web site.



Angela Barrow
BT Corporate Relations

Angela Barrow joined BT direct from school in 1978. The next seven years were spent in sales, moving from there to Network Marketing in 1985. After a period of three years Angela moved to Corporate Relations, responsible for community and public affairs in Yorkshire. Following a number of differing roles within Corporate Relations, Angela studied for a Postgraduate Diploma in Management at Nottingham University in 1994 and moved to her present role as Intranet Manager in 1996. She is a member of the Institute of Management and the Institute of Public Relations.

Business Use of Internet Web Sites—Could Do Better!

This article looks at how businesses are making use of the World Wide Web (WWW). Despite the fact that many companies now have a Web presence, their use of the new medium leaves a lot to be desired. Also, as the pace of change in the Internet industry is so fast, this article provides an interesting snapshot of how users, in this case companies, are keeping up to speed with this fast-moving industry.

Introduction and Background

The growth of the World Wide Web (WWW)—commonly referred to as the Web—and related developments in electronic commerce and the Internet in general present many business opportunities. At a very basic level, a new medium has been created which allows companies to change how they interact with existing customers and suppliers. This medium also provides a new route to market and can be used to attract new customers, literally on a global basis. The first step that many companies take on the Information Superhighway is in the form of a Web site which establishes them with an Internet presence. However, as the number of companies who have their own Web sites continues to grow, there is a need to look objectively at how these Web sites are actually being used.

This article presents some key results of a study that looked at how companies were actually using their Web sites. Issues such as site design ('look and feel'), perceived purpose of the site, and general usefulness of the site were examined. Eighty-five company Web sites were investigated, across a spread of industries. In addition to comparing sites within the same sector, the analysis also included cross-sector comparisons.

Despite the fact that many companies now have a Web presence, this work shows that their use of the new medium leaves a lot to be desired. Also, as the rate of change in the Internet industry is so fast—observers often comment that one 'real' year is equivalent to seven Internet years—this article provides an interesting snapshot of how users,

in this case companies, are keeping up to speed with the rate of change in the industry.

The article presents a framework for analysing how businesses are making use of the Web. The framework can also serve as a checklist for firms and their Internet service providers to discuss how they could actually avail themselves of the opportunities offered by the relatively new medium of the Internet. Typical examples of the findings under each section of the framework are presented.

The next section of the article highlights some of the characteristics of the Web and identifies some of the associated opportunities for companies. The framework used for analysis is then presented and some of the key findings are highlighted. A brief discussion of how the use of Web sites differs across industry sectors then follows. Finally some conclusions are presented.

The Opportunity

Before discussing how companies are making use of the Web, it is useful to highlight some characteristics of the Internet and identify some associated opportunities. These are illustrated in Table 1.

As can be seen from the table, there is a large range of opportunities available to companies. It should be noted that the opportunities listed in the table are mainly focused on activities related to sales and marketing. This is the area where most companies initially concentrate their efforts when they decide to establish their Web presence. There are, of course, many other opportunities open to companies who wish to

Table 1 Characteristics and Opportunities of the Internet

Characteristics	Opportunities
<ul style="list-style-type: none"> ● Worldwide network ● Millions of users—individuals and companies 	<ul style="list-style-type: none"> ● Instant international market exposure at relatively little cost ● Increase volume and value of sales
<ul style="list-style-type: none"> ● Computer network which allows customers to download information 	<ul style="list-style-type: none"> ● Better customer support ● Customers have 24 hour global access
<ul style="list-style-type: none"> ● No national borders 	<ul style="list-style-type: none"> ● Compete in new markets ● Exploit tax differences
<ul style="list-style-type: none"> ● Multimedia capability—text, data, graphics, images, photographs, animation, audio, video, virtual reality 	<ul style="list-style-type: none"> ● Display extensive and detailed information on products and services—save costs of printing, postage, etc. as the customer essentially meets the costs of downloading the information ● Put product catalogues on-line ● Create attractive Web sites to ‘seduce’ customers
<ul style="list-style-type: none"> ● Fast exchange of information 	<ul style="list-style-type: none"> ● Customer receives information much quicker ● Information can be updated very quickly ● Test market new products and services ● Quicker time to market—just-in-time marketing ● Save time—streamline information exchange
<ul style="list-style-type: none"> ● Relatively cheap communications medium ● Low cost of entry 	<ul style="list-style-type: none"> ● Eliminate paper from business processes ● Reduce dependency on advertising, sales and marketing people by servicing, via the Web site, routine requests for information ● Reduce costs ● New promotional activities—on-line games, contests, discussion
<ul style="list-style-type: none"> ● Supports two-way information exchange 	<ul style="list-style-type: none"> ● On-line shopping—take orders on-line ● Full electronic trading—including electronic data exchange (EDI) over the Internet ● Improving customer support
<ul style="list-style-type: none"> ● Can easily gather information on who accesses the site through the use of access logs, registration, cookies 	<ul style="list-style-type: none"> ● New levels of relationship marketing—personalisation of messages and content ● Improve customer satisfaction and loyalty ● Easier and quicker to gather customer data ● Learn about potential customers ● Supports more targeted marketing
<ul style="list-style-type: none"> ● Visibility 	<ul style="list-style-type: none"> ● Easy to gather marketing intelligence from Web sites of customers and competitors

exploit Internet technology. These other opportunities include cost savings, increased productivity, and new business processes through the use of Internet technology internal to the company in the form of intranets.

Many opportunities also exist for exploiting Internet technology to support collaboration and business networking among firms in the form of extranets. However, for many firms still struggling to understand the


technology and what the Internet means for their company and their industry in general, consideration of intranets and extranets is usually their second step on the Information Superhighway, Web sites being their first. Hence the focus of this article is on the use of Web sites and on how companies have so far made use of them.

Framework for Analysis

Despite the level of interest that the use of the Internet has attracted, there is little in the way of guidance for companies in the form of checklists of what their site could, or indeed should, contain. There are some useful early attempts at trying to understand how firms market themselves via the Web but the rate of change in the Internet industry renders much of this work out of date. A year or two ago company Web sites could be simply classified only in terms of no presence, basic presence, etc. However, because many sites are now quite big and cover many types of information, a more complex and detailed framework was developed. The framework used is illustrated in Figure 1.

Within each of the ellipses shown in Figure 1, a checklist was used to assess what type of information, characteristics and features were available at the various Web sites. A scoring system, where points were allocated each time an item in the checklist appeared, was used in order to provide data for analysis. Some of the key findings under each of the headings in Figure 1 are presented below, together with the detailed checklist used within each area.

Figure 1—Framework used for analysis of Web sites



General information

The checklist used under this heading is shown in Figure 2, and some of the key findings under this category are presented below.

Initially it is useful to publish some company information under all the headings identified above, since the marginal cost of adding extra information is rather small and it is very likely to be useful to customers and visitors to the site. As the company gains experience with the Web and begins to gather data on which information is being accessed and which information is not being accessed, the content of the site can then be tailored as appropriate. However, the information content should be concise, with links to more detailed information or the ability to request more information. A good example of comprehensive but concise company information is found at the Post Office's site (<http://www.postoffice.co.uk/>).

Company information

Overall, 92 per cent of the Web sites investigated have some company information, but for 23 per cent it is very limited and includes information in only one or two categories.

- *Activity/values/history* The amount of information that can be found under this heading is very varied, ranging from just one sentence on the home page to a full history and descriptions of the

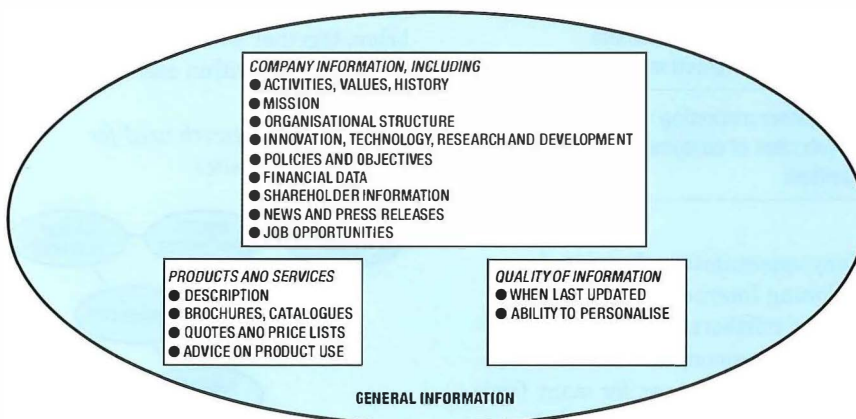
various activities of business units. However, it is rather strange to find that some Web sites do not even contain a brief introduction explaining what the company's activity is, and this piece of information has to be discovered by reading through general product information if available.

- *Mission statement* This is still uncommon, and it could be argued that such statements are not very useful since they tend to be very general objectives regarding the performance of the company and the satisfaction of the customers, etc. However, a mission statement is part of the corporate identity of a company and it seems strange that this be omitted from Web sites, part of the purpose of which is to establish a corporate presence on the Internet.
- *Organisational structure* Few companies put their organisation chart on the Web. Maybe they consider it to be confidential. This is probably correct as such charts are more likely to be useful to competitors than to anybody else.
- *Innovation and technologies or research and development* Information under this category is not very common either, and is usually very succinct. Here again, the roadblock to putting such information on the Web might be

the fear of competitors accessing it.

- *Financial results and commercial performance* This category of information is very common, and companies often provide the complete report as a pdf file (Adobe Portable Document Format) to download, instead of publishing it directly on the Web site. However, complete annual reports tend to be rather large documents and can take a long time to download. Shorter versions of annual reports are conspicuous by their absence.
- *News and press releases* This is the most common type of corporate information. Some sections are so important that a search engine is provided to find articles by keyword or date.
- *Job opportunities* Most company sites have an employment section, but not all of these have information regarding specific positions. Twenty-nine per cent have either descriptions of general positions or descriptions of positions currently available. In terms of on-line applications, 20 per cent of companies overall, or two thirds of these with job opportunities, provide an on-line application form. This seems a good result. It shows that companies are open-minded regarding the nationality of employees and that they are taking advantage of the diversity and quality of Web users to recruit the best people.

Figure 2—Checklist for categories of general information



Products and services

- *General description* Some form of general information on products/services is found in 76 per cent of the sites investigated. This leads one to question just why the other 24 per cent actually have a Web site.
- *Product catalogues and brochures* Less than 50 per cent of the sites

contain detailed product information in the form of product catalogues and/or brochures.

- **Quotes and price lists** Usually, sites with product details do not all include price information, which means that customers cannot really compare company offerings on the Web. Maybe this is just an oversight or perhaps it is a deliberate policy of the companies involved. Either way it is certainly an impediment to on-line trading.
- **Advice on product use** Information that falls under this heading can provide very useful advice and it is usually information that visitors to the site are typically expecting. Over a quarter of sites have some advice. Examples include British Gas with safety instructions, Scottish Power with suggestions for how to save water, the Co-op with diet and nutritional information, Boots with health and beauty information, and Pirelli with advice about tyres. With less than 75 per cent of sites containing information in this category, it is certainly an area where many companies could do much better.

Quality of Information

- **When updated** Only 9 per cent of sites have an indication of when a certain page was last updated, and, generally, this information only concerns a few pages of the site and there is nothing for the rest. Some sites were very out of date—one site in particular had information that was clearly at least two years old. Again, this begs the question of what such a company hoped to gain by having a Web site.
- **Personalisation** Few sites contained some form of personalisation of information; that is, when visitors can choose (or even have to choose) among various

sections. An example of this would be where sites tailor the information according to different categories of users, such as customer, journalist, potential investor, etc. However, some sites used *cookies* (visitors have to allow a cookie to be set to view the site) that are supposed to inform the visitor during his/her next visit about the sections already visited and to tell him about updates in these sections. Unfortunately, the use of cookies—where some information is stored on the user's computer—can often be viewed by visitors as being tantamount to 'spying' on their use of the site.

Sales

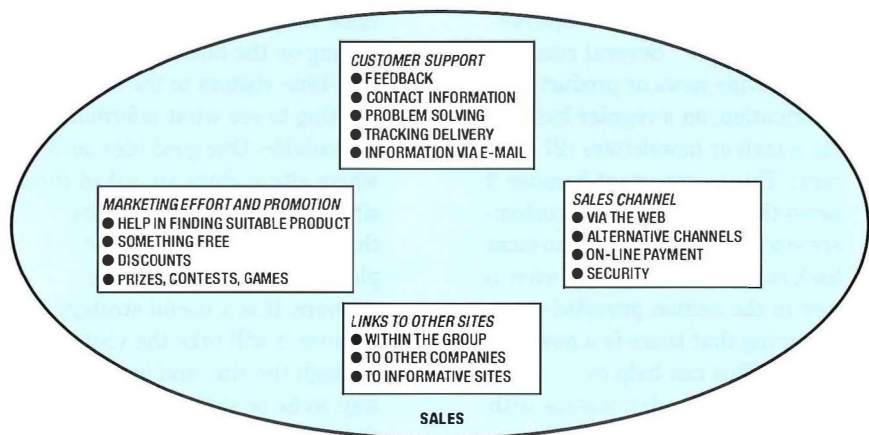
The checklist used under this heading is shown in Figure 3, and some of the key findings under this category are presented below.

The typical Web site does not provide customers with an order form. Partly this is due to the nature of the products being sold and partly due to constraints on the delivery channel. Surprisingly, in general there is no great marketing effort from Web 'enabled' firms. Most of them simply give information without trying to 'seduce' prospects. Given the fact that on some sites it is not even clear what the firm offers, there is often little scope to convert site visitors into customers.

Customer support

- **Feedback** Of course, companies use their sites to collect customer queries and comments (69 per cent), but too often it is only about the Web site itself! It is as if marketers consider the Web site as an independent variable of their marketing strategy, rather than as an integral part of their marketing activities. Using the Web site for any feedback people may have can be very valuable, since the medium itself encourages communication (people will communicate more, because it is fast and convenient compared to normal mail). Moreover, the comments obtained can be analysed more easily and accurately (it is typed so it is easy to read, clearer, and when the visitor has to click in boxes to tell the company what the feedback is about, the resulting data is already sorted!). Such data is easily stored for future reference. In terms of rapid answer to queries, *very few* companies say that they will reply promptly (though some say it is the fastest way), which is regrettable given the immediacy of the medium.
- **Contact information** Some contact information is common, but it should be 100 per cent and not 66 per cent! Again this

Figure 3—Checklist for categories of sales information and functionality



indicates that the marketing activities are not integrated. Contact information (including location) is useful to prospective customers as well as existing customers. Lists of nearest supplier or office locations are quite common and are contained on 45 per cent of the sites.

- *Problem-solving / frequently asked questions (FAQs) / technical information* Problem-solving information is also a useful type of information to have on a Web site, since it can save much time, trouble and money to both the customers and the firm. And the result is happier customers who are able to answer queries themselves and in the process learn more about the product. However, only 18 per cent of companies provide technical information, probably because the concept of a Web site as a support tool is not yet widespread.
- *Tracking delivery* Providing customers with a means to track the delivery of products is an important opportunity offered by the Web, but since it is mainly applicable to logistics companies, it is understandable that very few of the sites studied actually offered this facility. A good example of using the web to track deliveries is CargoConnect (<http://www.ccx.com/cx/vafsrall>).
- *Provide news or product information via e-mail* Several companies provide news or product information, on a regular basis, via e-mail or newsletters (27 per cent). This is important because it saves time and money for customers since they do not have to come back to the site to find out what is new in the section provided—assuming that there is a news section. This can help to strengthen the relationships with the customers.

Marketing effort and promotion

- *Help in finding the most suitable product* Only 6 per cent of sites have some form of help in choosing products. A good example is Intel's site which has a buyer's guide that is quite useful (<http://www.intel.com>). The Web is traditionally all about content, but simple information is not enough to sell, and there is great opportunity to reproduce on the Web what companies do using other media.
- *Something free* 17 per cent of the sites offered 'something for free'. Typical examples include screen-savers (BBC, Woolwich plc), publications (AA, London Regional Transport), simple software (a budget planner by Nationwide Building Society), and on-line banking software (Bank of Scotland).
- *Discount* Where the product is not suitable for selling via the Web, or where the company does not wish to sell via the Web, giving discounts via the Web site may be very appropriate to push prospects to the nearest store or office. Only 7 per cent of the sites investigated offer discounts to Web site visitors.
- *Prizes, contests, and games* Sixteen per cent of sites offer such a feature, but there is not always a prize to be won! Sometimes it is just a simple case of answering some simple questions on-line, relying on the novelty factor of first-time visitors to the site wanting to see what information is available. One good idea seen is where site visitors are asked three simple questions that require them to go to three different places in the site to find the answers. It is a useful strategy because it will take the visitor through the site, and in an active way, so he or she can learn about the company in a positive way. Of

course, there must be an interesting reward. Unfortunately, most prizes tend to be things that visitors to the site will already have; for example, a mouse mat.

Sales channel

- *Via the Web* Twenty-seven per cent of companies sell something on-line, but for some of them, it is only by-products such as promotional items like baseball caps and golf umbrellas emblazoned with the company logo. So less than 27 per cent of companies sell their core products on-line.
- *Alternative channels* Only 18 per cent of companies provide an alternative sales channel for the products sold on-line. Of course, they provide contact information, but there should be a distinct fax, telephone number or address to assure customers that their orders are processed as rapidly as via the Web. It may seem unnecessary to provide this to people since if they are on the Web, they have no reason to use another channel, but if there is a temporary problem on the server, or if they want to re-order items later on without coming back to the site, it could be very useful. It may also suit people who are not convinced of the level of security associated with ordering goods via the Web.
- *On-line payment* Only half the companies selling on-line provide an on-line payment form (13 per cent of total), which is rather small.
- *Security* Security of data relating to on-line ordering and payment seems to be something that few sites support. Some sites permit ordering via a secured server but they typically provide other means of payment too. As many users of the Web still have a deep mistrust of the level of security available, a low score in this area was to be expected.

Links

- *Within the group* Thirty-one per cent of the sites visited contained hyperlinks to subsidiary companies or sister companies. This can be useful where complementary products are produced by subsidiaries as it can help to keep the customer/prospect within the company group
- *To other companies* It is fairly common to see links to the sites of software companies that provide visitors with free plug-ins and software that enables them to read audio files (RealAudio, Shockwave), video files (RealVideo, QuickTime, ClearFusion) or even files in different formats (Adobe Acrobat Reader). However, some sites require far too many plug-ins to be installed before useful information can be viewed on their site—a sure way to deter visitors!
- *To informative sites* Links to general information sites, such as government agencies and professional institutions were found on 27 per cent of the sites visited. Such links can be very useful to Web site visitors as they provide links to related sources of information and thus can help nurture the process of strengthening the relationship with customers.

Market research

The checklist used under this heading is shown in Figure 4, and some of the key findings under this category are presented below.

Surprisingly, there were very few registration forms, survey forms and questionnaires used by the Web sites that were investigated.

- *Registration forms* Although 30 per cent of the sites had some kind of registration form, this includes forms that are not strictly registration forms; that is, are not compulsory and visitors do not have to fill them in before they can



Figure 4—Checklist for market research

access information on the site. Such forms, in which visitors leave their name and e-mail address, provide firms with data about prospects and customers that can be asked for feedback and the completion of a questionnaire later on. It is also important to note that most questionnaires focus only on the Web site itself. Typical questions include: How did you find out about our Web site? Can you find what you are looking for? Tell us how to improve our site etc. Again this reinforces the view expressed earlier that firms consider their site as a separate element of their marketing activity.

- *Survey forms/questionnaires* Only 13 per cent of the sites studied used a questionnaire, and generally there is no reward for completing it, so it is rather unlikely that many answers will be forthcoming.


Navigability

The checklist used under this heading is shown in Figure 5, and some of the key findings under this category are presented below.

Navigability relates to how easy, or not as the case may be, it is to find one's way around the Web site and to find the information that is required.

- *Assistance in navigation* Only 26 per cent of sites provide a help section.

Figure 5—Checklist for navigability




- *Use of several languages* Very few sites use several languages, and those that do are fairly big multinational firms. Since the English language is widely spoken it does not seem vital to publish versions of the Web site in other languages, but this should be reconsidered when dealing with a foreign customer base in countries with large populations such as Germany and France. Likewise, some subsidiaries in foreign countries have their site in the language of that country (for example, Chinese) and it might be worth opting for an international version of these pages.
- *Search button* Thirty-one per cent of sites have either a general or a specific search button.
- *Site map* Twenty-one per cent of sites are equipped with a map or an index.
- *General navigation* Navigation in general is quite good, but rarely very good. The most common navigational tools are search buttons (although these are not always for the whole Web site but rather for sections of it, such as catalogues and press releases) and a help or guide section. Site maps and indexes are less in use since they are necessary only for big sites. Some common problems seen on the Web sites include broken links, or links leading to a content different from the one expected.

Aesthetics

The checklist used under this heading is shown in Figure 6, and some of the key findings under this category are presented below.

Figure 6—Checklist for aesthetics



Companies should exploit fully the richness of the medium and make sites more attractive and pleasurable to visit.

Aesthetics refers to the general look and feel of the site.

- **Layout** Most sites have a rather good layout (76 per cent) but it is common to find a page in which graphics or fonts are very small and almost impossible to read. There is almost a typical layout that uses a Times New Roman font, size 10 typeface, with small photos or graphics dispersed, on a white background. It is unfortunate that so many companies have 'adopted' such a standard because there are many varieties of fonts, sizes, and colours available to give some personality to Web sites.
- **Background and general design** Just over half of the sites actually make good use of background images and colours and combine this with a good overall design.
- **Graphics and photographs** The use of graphics and photos is now very common (88 per cent). They are usually of good quality. The remaining 12 per cent of sites without graphics or photographs are at a disadvantage because they lack that typical and attractive characteristic of the Web. There is always some useful or interesting graphic or photograph that can be used to enliven text. Unfortunately 12 per cent of the sites visited have yet to discover this.
- **Animation and video** The use of animation or video is growing (57 per cent of sites). Good examples of sites that use animation are Boots (<http://www.boots.co.uk/>) and BT (<http://www.bt.com/>).
- **Sound** Use of sound on Web sites is fairly uncommon. However, as the use of technologies such as RealAudio become more widespread this is likely to change.

General Comments

The checklist used under this heading is shown in Figure 7, and




Figure 7—Checklist for general comments

some of the key findings under this category are presented below.

This final category is a general 'catch-all'. It is used to capture general impressions and information that does not fit into the earlier categories

- **Attractiveness** Less than a third of sites can be considered as attractive. Although this is subjective, it is possible to say that the combination of aesthetic factors is usually not as good as it could be, and the fact that most sites are standard and not very original, results in the visitor being bored by such similarity. Companies should exploit fully the richness of the medium and make sites more attractive and pleasurable to visit.
- **Personality** Clearly, some companies have reproduced on their Web pages the 'personality' they use off-line, but often a visitor can forget just which company site is being visited—this is usually due to the similarity of design and layout used. There is too much similarity on the Web and it is rapidly boring.
- **Clarity** Seventy-seven per cent of sites are quite clear and easy to understand. The ones which are not clear are typically those with the layout found in all magazines' and newspapers' sites; that is, with articles, headings and sections everywhere. It is unpleasant to see, and unclear because the structure is not usually apparent.

Moreover, progressing in the site can be difficult because one always encounters yet more sections containing yet more articles and even more sections! In these sites, one can typically search for the information required but then find that the search engine returns huge quantities of data—lack of structure may then make it impossible to find one's way through the data presented.

- **Usefulness** Forty-five per cent of sites have been judged as useful. Again this is a subjective assessment but the criteria used were those of quantity and quality of information concerning the company, the products and any educational information such as customer advice as well as the provision of a real benefit to visitors to the Web site. Thus, it means that more than half the sites visited did not seem to have any of these. Or to put it another way, more than half the sites visited did not prove to be useful overall.
- **Interactivity** Very few sites showed any form of interactivity. The potential to use interactivity combined with on-line catalogues to allow customers to examine and view products in three dimensions is a potential that currently is vastly untapped.

Summary of key findings

Figure 8 summarises the scores obtained for the study sites on key features. The typical Web site contains general product information (usually just an overview), news and press releases, a statement of activity or a history, financial results, a feedback space, some contact information, a few links, a few navigational buttons and an acceptable layout and design. The information provided is not comprehensive and the quantity of information under a particular heading is




Figure 8—Percentage of Web sites that score on key features

usually inversely proportional to its importance to a typical visitor to the Web site; for example, it is not unusual to find vast amounts of very detailed financial information, yet very little on the products and services.

Cross-Sector Analysis and Comparison

Having looked at the general use of Web sites across a range of industries, it is useful to re-examine the data and see if there are any

differences between the use of Web sites in different industry sectors. This section presents some results of the cross-sector analysis that was conducted. For the purposes of this work, it was decided to use the following sectors: Energy, Finance, General Industry, Logistics, Retail, Services, and Technology.

The results do vary significantly across the seven sectors. Looking at the highest scores for each variable in the framework, it was found that the sectors can be grouped into four groups as shown in Table 2.

Table 2

Scoring	Findings	General Comments
Least scoring	The industry sector has no highest score at all	<i>Industry</i> This sector is far below average for everything; there is little company information, even less detailed product information and there is no on-line ordering available. The sites are not attractive at all, the design is not very good and there is not much use of photos
Low scoring	The services sector has only eight high scores	<i>Services</i> Average or slightly below average performance, although highest score for personalisation of information; yet no company publications or white papers. Little use is made of opportunities to differentiate this sector from the 'product' based sectors
Average scoring	Energy, finance, logistics and retail obtain between 20 and 26 high scores	<i>Energy</i> Sites in this sector tend to have lots of company information. They also make a distinction between business and individual customers. The sector's highest scores are for industry information and other company-related information, as well as for links to informative Web sites <i>Finance</i> Company and product information is about right, but overall the sites are fairly average. Sites in this sector tend to provide help in finding the most suitable product for one's needs <i>Logistics</i> Information on sites in this sector is average, and there tends to be good on-line transaction possibilities. This sector also scores highest for links to other companies' sites, good in aesthetics and top for usefulness <i>Retail</i> Good company and product information, above average on-line ordering and payment, as well as marketing effort (discounts, competitions), feedback spaces, useful information (especially problem-solving information). Good use of design, graphics and animation, highest score for attractiveness and interactivity. Weaknesses: not very good navigation, worst score for problems such as broken links
High scoring	The technology sector has 49 high scores	<i>Technology</i> Almost all the features on these sites have a higher than average score, and key elements such as product catalogues, on-line ordering, on-line payment, and provision of news via e-mail have the highest incidence. However, the sector is not rated top performer for general navigation, general design and usefulness, because the sites can be rather complex. Navigation is good but not perfect despite the use of many navigational tools, so improvements can still be made


Figure 9—Variations of key elements across industry sectors

The charts in Figure 9 illustrate how some of the key elements vary across the sectors.

It is not surprising that the firms in the technology sector have more 'complete' Web sites and thus score higher than the other sectors, since they tend to be at the centre of new developments and thus are in a position to adapt more quickly. Surprisingly the Industry sector seems to make very poor use of the Web, especially as this sector probably has the greatest potential for making use of Web sites to build relationships with other companies. This sector also has a very high potential for making use of extranets and electronic trading platforms. But before companies in this sector can exploit such technologies they need to improve their understanding of the potential and capabilities of their Web sites.

Conclusions on Current Use of Web Sites

The key conclusion of this study is that, in general, Web sites are not being used effectively. For example, less than 50 per cent of the sites investigated contain useful information on the companies' products, less than 9 per cent let their audience know when information on the pages is updated and only just over a quarter do any selling on-line. Simple information is not enough to sell products, and there is a great opportunity to reproduce on the Web what companies do using other media; that is, seduce, convince and achieve sales. The idea that a Web site can be an integral component of a company's marketing strategy seems to have eluded most organisations and it appears that a Web site is treated as an independent variable. Reasons for this could include lack of understanding of the capabilities of the Web or lack of the skills necessary to design an effective Web site. Either way the opportunity is there for the companies, Web designers and Internet service providers to exploit this new medium to their advantage.



Basically, sites are not being used commercially enough—though when they are, as in the case of a company like Dell (<http://www.dell.com/>), which is selling over \$5 million worth of computers online every day, the returns can be impressive. Furthermore, the effective use of Web sites is but the first, and in many cases tentative, step on the path towards full electronic trading, where companies reap the benefits of and embrace the full range of opportunities offered by the use of Internet technology—internally in the form of intranets; collaboratively in the form of extranets; and externally in the form of Web sites that can be supporting customers anywhere on the planet, 24 hours a day. In many cases it seems that companies have a Web site because they feel that they should have one. It should be remembered that a Web site is but an enabler—it is not an end in itself.

Bibliography

The full report that served as the basis for this paper is available on BT's intranet at <http://webster.info.bt.co.uk/marketing/wwwuse.htm> (only visible within BT).

CALLAGHAN, JAMES; and FLAVIN, PHIL. Intranets: Corporate Nirvana—The End of the Traditional Organisation. *Br. Telecommun. Eng.*, Oct. 1996, 15, pp. 224–229.

SIM, S. P.; and RUDKIN, S. The Internet—Past, Present and Future. *BT Technol. J.*, Apr. 1997, 15(2), pp. 11–23.

COCKBURN, C.; and WILSON, T. D. Business Use of the World-Wide Web. Oct. 1995. (<http://www.shef.ac.uk/uni/academic/I-M/is/lecturer/paper6.html>)

HO, J. Evaluating the World Wide Web: A Global Study of Commercial Sites. June 1997. (<http://www.usc.edu/dept/annenberg/vol3/issue1/ho.html>)

KING, D. Fortune 500 on the Web: The Road to Second-Level Effects. Proceedings of the 29th Annual Hawaii International Conference on System Sciences, IEEE 1996, pp. 463–470.

Biographies



James Callaghan
BT Networks and
Systems

James Callaghan joined BT in 1986, after graduating from Trinity College Dublin with a first-class honours degree in Electrical Engineering. He subsequently gained an M.A. from the same university, an M.Sc. from the University of Essex, a Graduate Diploma in Marketing from the Chartered Institute of Marketing and an MBA from Henley Management College. His early work in BT was in network management research and he spent several years working on and leading research projects in the European RACE programme. Recent work has focused on researching the likely impact of new technologies on the way companies conduct their business. He is currently on secondment from BT Laboratories to the University of Cambridge (Judge Institute of Management Studies) where he is investigating Inter-Organisational Networks and is pursuing a Ph.D. on the role and impact of intranets. He may be contacted at jgc20@cam.ac.uk



Audrey Pie
BT Networks and
Systems

Audrey Pie graduated in 1995 from the I.U.T. of Le Havre in France (Technological University Institute) with a D.U.T. (Technological University Diploma) in Marketing Techniques. She subsequently gained an M.S.G. (Maitrise de Sciences de Gestion) with distinction at Caen University in France and an M.B.A. from Coventry University Business School in the UK. She has a wide range of experience in various areas of marketing including carrying out market studies at Piaggio Concession in Le Havre, researching business activities for the Regional Chamber of Commerce and Industry in Caen, and has also worked at BNP (National Bank of Paris) in Dreux. She has also worked as a placement student at BT Laboratories where she conducted research on the use of the WWW for gathering marketing intelligence.

Getting the Message, Loud and Clear— Quantifying Call Clarity

When we pick up the telephone today and make a call, the immediate thought in our minds is ‘are they there?’ or ‘I hope their machine doesn’t answer!’ We seldom contemplate whether we will be able to hear the person at the other end of the line clearly, or if they will be able to hear us. However, the clarity of our telephone calls is a major issue in the telecommunications world today.

Introduction

With telecommunication networks stretching around the globe and beyond the atmosphere, and with pressure on network operators to reduce costs and carry more traffic, the ability to communicate effectively is paramount. Call clarity is one of many factors that affect the quality of service (QoS) perceived by customers of a telecommunications company, but it is of vital importance.

This article discusses the issues and technologies involved in measuring call clarity in today’s modern telephone networks, and highlights a new tool that has been developed to predict call clarity—the *call clarity index* (CCI).

How call clarity monitoring has evolved, including both intrusive and non-intrusive measurement technologies, is described. The benefits and limitations of such schemes are discussed along with how complex models, such as the CCI, can be used to predict customers’ perceptions of call clarity. Finally the future development of speech quality monitoring is considered.

What is Call Clarity?

Telecommunications engineers often refer to call clarity as *transmission performance*. It can be considered as the quality perceived by the customer during the conversation phase of a telephone call, and is influenced by factors such as noise level and echo. It is a subjective measure which means it is peculiar to each individual and

modified by personal views, experiences and backgrounds. Because of this, each individual’s perception of call clarity changes with time. Averaging subjective results taken from a sample of the population is normally used to eliminate this bias and produce stable results.

However, to measure network performance reliably and repeatably requires an objective measurement method.

Call Clarity Monitoring

Historically, network performance assessment techniques were relatively simple and based on objective measurements from test calls in the network¹. This was ideal for metrics such as network availability and call failures, but assessing call clarity was limited to measuring basic parameters such as transmission loss.

Traditional methods of call clarity assessment used customer surveys or trained operators. The latter method involved a trained operator monitoring a short period of conversation at an exchange and, based on their knowledge and experience, producing a score to represent the call clarity. This method has two major limitations. The score for each call is a subjective measure, and different operators would give slightly different scores. Secondly, the operator monitors from an intermediate point in the network and does not hear exactly what the customers hear.

Customer surveys are the most common alternative method of collecting information about the call

clarity of today's networks. In this type of survey people are asked directly to answer a few questions about the telephone service that they have been receiving; in particular, the call clarity. Their major advantage is that the results show what customers really think of the services provided, including the call clarity of the network. However, they are expensive to perform and cover only a small sample of the population. Such surveys only give a general perception of quality for the service and are often influenced by factors unrelated to call clarity, such as the time since the last telephone bill. Customers tend to remember a single poor quality call rather than the numerous acceptable calls which are taken for granted. It is also difficult for network planners and maintenance staff to use customer survey results to help to improve the network quality since there is no record of the call routing or equipment used.

Intrusive measurement techniques, where a known test signal is applied, have advanced since the early days of performance monitoring. Speech-like test stimuli are starting to be used to cope with non-linear networks, and in the last few years research has produced methods of call clarity prediction using models of the human auditory system².

In recent years non-intrusive methods, which analyse live network traffic, have emerged using in-service non-intrusive measurement devices (INMDs). These INMDs are situated in the centre of the network and continuously measure various parameters related to call clarity (and network performance). These techniques cannot measure call clarity directly, only parameters related to call clarity. The CCI has been developed to fill this gap and allow call clarity to be predicted, objectively, non-intrusively and constantly.

The following sections look in more detail at the differences between intrusive and non-intrusive monitoring and where the CCI fits in.

Intrusive Monitoring

The goal for intrusive assessment is to predict speech quality objectively and remove the need for routine subjective judgements. Historically, intrusive measurement has typically been based around a transponder and responder structure. Transponders are usually located at main switching points, and responders at remote switching points within the telephone network.

The conventional method of intrusive testing is to make a web of test calls between responders and transponders and measure various network performance parameters. However, to measure completely call clarity as experienced by a particular customer the intrusive test should include the effects of the customer's equipment.

Intrusive measurement has historically involved injecting tones into the network and measuring the received signal at the other end, and more recently pink noise has been used in addition to tones. The impetus for moving away from tonal testing has come from the proliferation of non-linear elements, such as data reduction equipment and adaptive echo cancellers, in modern networks. The accurate testing of networks that include non-linear elements requires the use of a speech-like test stimulus, hence reproducing the properties of the in-service signal. This is necessary because non-linear equipment can treat non speech-like signals very differently from speech, making the

measurement path different from that experienced by customers.

Such speech material has recently been created for use as intrusive test stimuli³. This material reproduces the sounds of the human voice in speech, preserving the statistical properties of a language and representing a large population of talkers.

Modern intrusive technology, using speech-like stimuli, currently provides the most accurate objective prediction of call clarity, but requires the application of a known test signal². This type of testing is ideal for commissioning or planning purposes and is cheaper than a subjective test—the alternative method. Also, owing to the fact that a line is used solely for testing purposes, a much wider range of impairments can be tested and a greater accuracy of measurement made. However, the requirement for a known test signal limits the practical application of intrusive techniques.

Non-Intrusive Monitoring

INMDs have access to the voice channels as well as the signalling information. They perform objective parametric measurements on live call traffic, without interfering with the call in any way.

An INMD is connected at the digital 4-wire point of a telephone network as shown in Figure 1, and monitors traffic in both directions of transmission simultaneously. Objective measurements such as speech

Figure 1—Situation of an INMD in relation to a typical international telephone connection

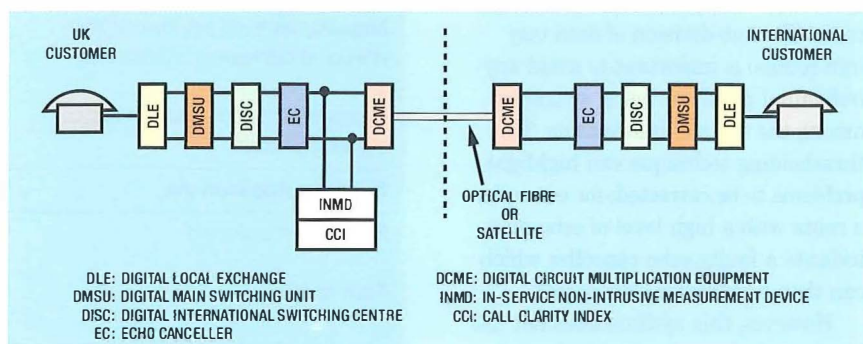


Table 1 Objective Call Clarity Parameters Measured by an INMD

Echo path loss
Echo path delay
Speech level
Noise level
Activity factor
Pulse-code modulation (PCM) threshold violations
Double talk
One-way transmissions
Front-end clipping

level, noise level, echo loss and delay (common INMD call clarity parameters are shown in Table 1) are made by the device. Access to the signalling system allows the extraction of information such as the called party number which can be associated with the objective measurements. All the information pertaining to a single call is stored in a database.

The advantage of INMDs is that they can continuously monitor real traffic within the network and over time can, theoretically, cover the entire network. The disadvantage of being located in the centre of the network is that the measured values do not directly represent what the customer hears. The signal reaching the customer's ear depends on the local network from the local exchange out to the customer, in addition to the customer's equipment. This means that inspection of the call clarity measures individually does not fully characterise the connection.

A common method of analysing data produced by INMDs is to set acceptable threshold levels for each parameter and record the number of exceptions (measurements outside the threshold) for each parameter on each route. The sub-division of data (say into routes) is important to avoid any individual problems being swamped among the mass of information. This thresholding technique can highlight problems to be corrected; for example, a route with a high level of echo may indicate a faulty echo canceller which can then be replaced or repaired.

However, this method does not use all the information collectively and

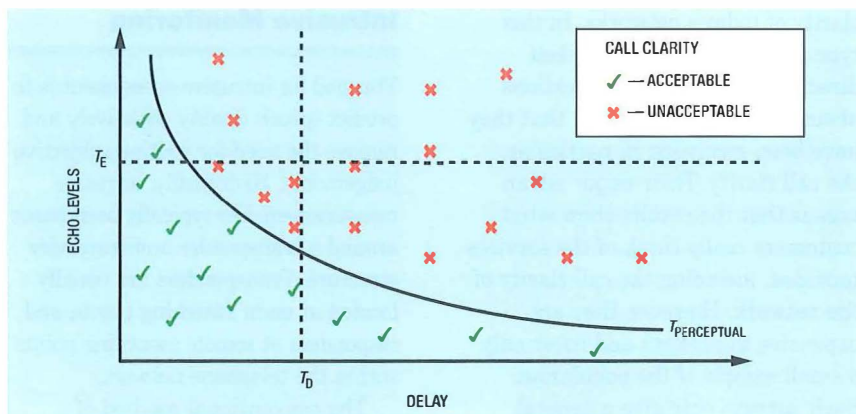


Figure 2—The limitations of conventional thresholding methods

can give misleading indications of call clarity, as shown in Figure 2. This example (for just two parameters) shows how echo and delay combine to affect call clarity. Measurements from acceptable and unacceptable calls are represented by ticks and crosses respectively. It can be seen that some calls with high delay are still acceptable provided there is low echo. Using a single

threshold for delay (T_D) misclassifies these calls as exceptions. Likewise calls with reasonably high delay and echo that fall just within the echo and delay thresholds (T_D and T_E) may be unacceptable to the customer, and are also misclassified. The threshold of acceptability perceived by the customer ($T_{PERCEPTUAL}$) is shown on the graph by the curve and it is this that the CCI emulates.

Intrusive and Non-Intrusive Applications

The goals for intrusive and non-intrusive monitoring are the same, but their principal applications are different. The table below highlights the appropriate applications of each scheme.

Intrusive	Non-Intrusive
Pros	Pros
Able to test for a wider range of network impairments	Able to measure a large number of live calls
Greater measurement accuracy per call	Measurements include the effects of customers' equipment
Cons	Cons
Test calls cost money and use network capacity	Measurement algorithms have limitations; for example, cannot always measure delay
Measurements do not include the effects of customers' equipment	Needs live traffic
Measurements not made on actual customers' calls	Network and speaker assumptions required
Principal Applications	Principal Applications
Route commissioning	Constant network health check
Fault finding	Identify problems as soon as they affect the network

Figure 3—The functional blocks that form the CCI

To determine the perceived quality of calls it is necessary to account for the combination of effects. This can be achieved by modelling human auditory perception. The adaptation of modelling techniques to INMD results to produce a 'synthetic' customer perception measure has been performed at BT Laboratories and is now deployed in the network under the CCI banner.

Call Clarity Index

Research at BT Laboratories into human perception, most notably audio perception, has led to the creation of various models of the human auditory system. These models have been applied to INMD measurements to predict the call clarity of a connection that would be given by an average customer. It is this measure that has been termed the *call clarity index (CCI)*.

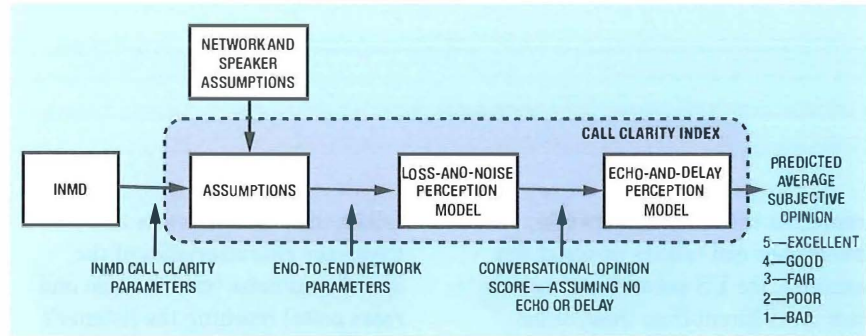
The CCI uses information produced by an INMD along with knowledge about the network, customers and results of research into audio perception to predict a call clarity measure. The CCI consists of three functional blocks depicted in Figure 3. Measurements from an INMD feed into the first of these blocks.

INMD measurements

Eight INMD measurements are fed to the CCI, four for each direction of transmission. These are the speech level, the noise level, the echo-path loss and the echo-path delay. The delay is usually found by looking for echoed speech, but if echo is well controlled there may be none. In this case the echo path loss is assumed to be high, and the CCI uses a default value for echo-path delay.

If any one of these measurements is not reported, the CCI is unable to calculate the call clarity. The reason for this is that not enough information about the connection is available to make a reliable prediction.

These measurements are passed to the 'Assumptions' block of the CCI.



Assumptions

The INMD measurements are used to fill in the gaps in the knowledge of the network connection. As there are unknown quantities, various assumptions must be made about the network and the customers.

The generic network model in Figure 4 shows how the INMD measurements relate to the signals, both spoken and heard, at each end. The network connection can be simply described in terms of the following two parameters^{5, 6}.

- *Sending sensitivity* This describes how the sound pressure generated by the talker at the mouthpiece of the telephone is transformed into an electrical signal at the digital 4-wire point of the network. This encompasses the frequency response of the telephone and the line.
- *Receiving sensitivity* This describes how the electrical signal at the digital 4-wire point is transformed to sound pressure generated by the telephone earpiece. This encompasses the frequency response of the line and telephone.

In addition to the transmitted signal reaching the listener's ear are three other signals. Two of these

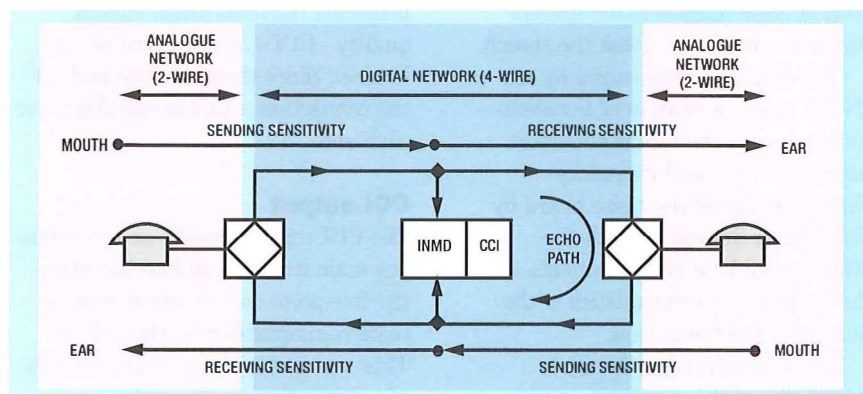
signals are caused by acoustic feedback provided to the user by the telephone handset. The first of these is a proportion of the user's speech, known as *sidetone*. This has the effect of controlling the user's speech level due to psychoacoustics. The second signal is a proportion of the room noise fed back to the user's ear via the sidetone path. The third signal is due to room noise leaking through the imperfect seal between the earpiece and the ear.

The frequency responses of the sending and receiving sensitivities assumed by the CCI are representative of a typical UK telephone connected to an average length line.

In addition to this, the sending sensitivity is adjusted depending on the INMD measured speech level. This is achieved using knowledge of the distribution of talker levels at the handset. For instance, if the INMD measured speech level is lower (quieter) than expected, it is assumed that it is partly due to a greater sending sensitivity loss in the network and partly due to a lower talking level.

These assumptions are based on BT's network and standard BT telephones. This is a realistic assumption for the UK as it stands, but in the future and for international connections it is less valid. However, owing to the nature of the CCI these assumptions can be modified to

Figure 4—Simple network model showing the positions of the INMD and CCI



represent the differing networks, telephones and talkers involved. For example, the US network characteristics are different from those of the UK, as are the telephones and possibly the talkers as well. Data describing these characteristics can be used by the CCI to account for the two different halves of the connection. This is not only useful for international connections but also for mobile and other distinct networks.

These assumptions combine to form a complete model of the network, including the customers at either end. The parameters describing this model are then passed on to the next block—the loss-and-noise perception model.

Loss-and-noise model

When someone talks at one end of a connection, his or her speech, along with room noise picked up by the microphone, is transmitted via the send path to the digital measurement point in the network. This speech and room noise also has circuit noise added, generated within the local network (telephone, line and interfaces). The INMD monitors at this point, after which the signal is then further modified by the receive path and has more circuit noise added, due to the local network, before reaching the listener's ear. The listener hears this speech and noise combination plus room noise at that end of the connection. The call clarity opinion of the listener depends on the absolute level of speech reaching his or her ear: too low means the listener will not be able to hear well, too high can cause discomfort. It also depends on the level of noise relative to the speech: high noise levels can mask the speech.

The noise level measured by an INMD is a combination of transmitted room noise and induced circuit noise. The level and frequency characteristics of the noise heard by the listener depend on both the INMD noise level measurements and the frequency characteristics of the noise and the connection.

The CCI calculates, using its knowledge of the connection and

talker, the absolute levels and frequency characteristics of the speech and noise (circuit noise and room noise) reaching the listener's ear. This 'sensation' level at the ear is transformed to a single figure value—a listening opinion index—which is converted to a conversational opinion score⁷. Using the INMD measured speech level directly within the model reduces the uncertainty in calculating the signal at the listener's ear by eliminating the need to assume sending loss.

Echo-and-delay model

In addition to the speech and noise signals, echoed speech could be present. These echoes are usually caused by imperfections in the 2-wire-to-4-wire interfaces (hybrids) and reflect a proportion of the transmitted speech back to the talker. Echo signals combined with propagation delay have a dramatic impact on the quality of telephony. In the extreme, echo and delay can cause talkers to stutter and terminate sentences early. Even long delay on its own, with no echo, can cause people to start talking at the same time making the transfer of information more difficult.

The echo-and-delay model accounts for the effects of echo signals on the call clarity perceived by the listener. This is a complex mathematical model that combines the network parameters affecting the echoed signal reaching the listener, to modify the opinion score generated by the loss-and-noise model.

The output from this block is the final score and represents the predicted conversational speech quality—CCI—as perceived by the listener. Since there are two ends of the connection a CCI is calculated for each end.

CCI output

The CCI reports results on a continuous scale from one to five, based on the five-point conversation opinion scale recommended by the ITU-T⁸. This represents the subject's opinion of a connection on the scale:

- excellent,
- good,
- fair,
- poor,
- bad.


Values are allocated to these categories from one to five, where Bad = 1 and Excellent = 5.

The telephone tends to be taken for granted, and the complexity of hearing is often underestimated. Modelling this auditory process for telephony can only be an estimate, but even as an estimate the ability to reproduce a 'synthetic' customer opinion of call quality is invaluable. This allows objectivity in what was previously a subjective domain.

People versus the CCI

To verify the robustness and accuracy of the CCI in predicting subjective call clarity, several subjective tests have been performed. The most recent of these involved emulating different telephone connections and asking subjects to converse over each one and rate the quality of the connection on the ITU-T five-point scale described above. These ratings were averaged over all subjects to produce a conversational mean opinion score (MOS_c). The CCI was then used to predict a call clarity score for each condition. Figure 5 shows the performance of the CCI against subjects' perception of call quality. The CCI score is shown on the horizontal axis and the MOS_c on the vertical axis⁹.

Figure 5—Scatter plot showing conversational MOS_c against CCI predicted score



the CCI is an effective monitor, ensuring the quality of the network is discernible to those who have the power to change it.

A strong correlation is evident between the CCI and the subjective results, demonstrating the accuracy of the CCI across a range of network conditions. An exact mapping from CCI to subjective opinion is impossible. This is due in part to the variability of subjective results; as a general rule, subjective results are expected to be within half of an opinion score if two identical tests are performed at different times.

Knowledge through information

The combination of the CCI and INMD technology is very powerful. INMDs have the potential to measure a substantial proportion of calls in a network. The CCI can use the INMD measurements to predict the call clarity on a call-by-call basis. Furthermore, this objective measure has been proven to correspond closely to average subjective call clarity metrics.

An INMD system can produce millions of call clarity related measurements every day. Considering each item of data individually loses value by not being seen as part of the whole picture. The CCI uses all the pieces of information given to it along with prior knowledge about the network and users to produce a single figure of merit for each end of the connection. These values can then be averaged for each route to provide a meaningful and useful indicator of call clarity. This reduction in information is enhanced by an increase in knowledge; by using the CCI the combined effects of those degradations present on a call can be accounted for. By analysing trends in this data it is easy to determine if the quality of a route is changing and affect the necessary remedial action.

Telecommunications operators must have their finger on the pulse of the network; the CCI is an effective monitor, ensuring the quality of the network is discernible to those who have the power to change it.

Towards the Future

This final section looks at the possible effects of new developments in call clarity monitoring focusing on how the CCI, in collaboration with INMDs, could be used in the future.

Speech quality monitoring is still evolving along with INMD technology. The three most important areas of development that will shape its future are the technology, deployment and data management.

The technology will determine the limits to which non-intrusive monitoring can be taken. Knowing how many devices are needed and where to put them in the network is a deployment issue, and managing the data generated by these devices is critical to realising the benefits.

Current non-intrusive measurement technology (that is, the hardware and software) limits the number of separate calls that can be monitored simultaneously and the range of parameters that can be measured. To enable a tool, like the CCI, to be effective across a broad range of network degradations it must be fed with as much relevant information as possible. Current INMDs cannot effectively measure parameters such as impulsive noise, non-linear codec distortion, front and back end clipping of words, or high levels of noise in speech. Increasing use of speech compression and the growth of mobile networks means that these types of degradation have the potential to become more common in communication networks; ways of measuring them and incorporating them into the CCI must therefore be found.

In the meantime, while such measurement algorithms are being developed, better use of other data collected by INMDs could be made. Often this data can give valuable information about the connection. For example, the origin and destination of calls influence network parameters such as sending or receiving sensitivity. No two coun-

tries have the same network infrastructure or equipment and so have a different distribution of such parameters. This difference can lead to errors in call clarity prediction. To avoid these errors the network (and talker) parameters used by the CCI could be tailored for each country to ensure greater reliability and accuracy.

The management of data produced by INMDs is critical for realising the benefits. To be useful the data needs to be interpreted and acted upon. Knowledge management tools such as the CCI are starting to forge a path towards truly intelligent systems that can detect problems and issue work packages to investigate and implement corrective action. Future generations could detect and treat problems as soon as they affect the network and impact on the customer.

The CCI could introduce the concept of call clarity as an objective measure into specifications for telephone operators. In the future, companies or network operators who rent network capacity may specify a certain quality that they require. Call clarity could well be one of the major factors influencing their quality specification and loyalty to a given service provider. Looking further ahead, the selection of service provider could be done dynamically using the CCI as part of a quality monitoring tool. The network operator that can provide the best price for the required quality could then be selected automatically and updated on a month-by-month, day-by-day or even hour-by-hour basis. In this sort of world, being able to monitor, and maintain, quality objectively is of utmost importance.

In summary, quality is an important market place differentiator. When all competitors charge the same, customers will select on the basis of the best service available. Call clarity is then a key factor in the service offered by the telecommunications provider.

References

- 1 ROGERS, DERYCK; and HAND, DAVID. Network Measurement and Performance. *Br. Telecommun. Eng.*, Apr. 1995, **14**, p. 5.
- 2 HOLLIER, M. P.; and SHEPPARD, P. J. Objective Speech Quality Assessment: Towards an Engineering Metric. Presented at the 100th AES Convention in Copenhagen. Preprint No. 4242, May 1996.
- 3 HOLLIER, M. P.; HAWKSFORD, M. O.; and GUARD, D. R. Characterisation of Communications Systems Using a Speech-Like Test Stimulus. *J. Audio Eng. Soc.*, Dec. 1993, **41**(12).
- 4 ITU-T Recommendation P.561: In-Service Non-Intrusive Measurement Device—Voice Service Measurements. Feb. 1996.
- 5 ITU-T Recommendation P.76: Determination of Loudness Ratings; Fundamental Principles. 1988.
- 6 ITU-T Recommendation P.79: Calculation of Loudness Ratings for Telephone Sets. Mar. 1993.
- 7 RICHARDS, D. L. Calculation of Opinion Scores for Telephone Connections. *Proc. IEE*, May 1974, **121**(5), p. 313.
- 8 ITU-T Recommendation P.800: Methods for Subjective Determination of Transmission Quality. Aug. 1996.
- 9 ITU-T Study Group 12 Question 15 Delayed Contribution D.058: Call Clarity Index Verification. Feb. 1998.

Bibliography

- MULLEE, TONY; and OODAN, ANTONY. Telecommunications Quality of Service: Principles and Management. *Br. Telecommun. Eng.*, Apr. 1994, **13**, p. 53.
- REX, A. W.; HOLLIER, M. P.; and GRAY, P. Predicting Speech Quality of Telecommunications Systems in a Quality Differentiated Market. 6th IEE Conference on Telecommunications, Mar. 1998.

Biographies



Simon Broom
BT Networks and
Systems

Simon Broom graduated from the University of York in 1995 where he gained an M.Eng. in Electronic Systems Engineering. He joined BT within the Applied Research and Technology Department where his work has centred on voice transmission performance. He is currently researching new algorithms for the quantification of transmission impairments and the prediction of call clarity using subjective models. Through this work he has become involved in non-intrusive network quality assessment and INMDs. He is an associate member of the IEE.



Philip Coackley
BT Networks and
Systems

Philip Coackley graduated from the Electrical and Electronic Engineering Department of Strathclyde University in 1988. He worked for BP

Exploration in Telecommunications before moving to BT Laboratories in Suffolk in 1989. During this time he has been involved with all aspects of voice transmission performance and has continued his professional advancement by obtaining an M.Sc. in Telecommunication Engineering at University College London. He is a member of the Institution of Electrical Engineers and is a Chartered Engineer. Over recent years he has been heavily involved with INMD developments within BT and the deployment of call quality metrics in the international network. As part of this work he represents BT at the International Telecommunications Union.



Phil Sheppard
BT Networks and
Systems

Phil Sheppard graduated from Birmingham University in 1986 with B.Sc. (Hons) in Computer Science and Electronic Engineering. He joined BT Laboratories on graduation where he has worked on ISDN transmission over copper, speech coding and recognition, network emulation, algorithm development for customer opinion modelling and non-intrusive measurement of voice networks, for which he is an ITU-T

Future Trends in Satellite Communications

Satellite communications is about to enter a new golden age. Following steady growth over the last 35 years, the field is confidently expected to undergo a period of rapid expansion as the race gets underway to provide the global information infrastructure of the twenty-first century. This article traces the development of satellite communication from its inception, discusses why it is changing and provides a glimpse of the future that is unfolding.

Introduction

Can you imagine the world today without satellite communications? It is difficult because we have become so accustomed to tuning in to television broadcasts of world events as they happen, which may indeed reach our homes via our own satellite dish. In a little over 35 years, satellites have fundamentally changed both the way we communicate and the way we learn about the world.

This article steps back and reviews how we have reached the current state, how technology has evolved and how organisations have grown to provide the services we have come to rely on. This leads us on to look at their immediate plans and how in the medium term new systems will emerge.

Many factors will determine the future shape of the industry. These include the state of technology, the changing telecommunications environment, the commercial imperatives on service providers and manufacturers, the regulators and politicians. This article discusses these factors, how they interact, the plans that are being made and the design considerations that drive specific satellite system architectures.

In the future, satellites will no longer provide simply a transmission link. As costs reduce they will be incorporated increasingly into networks in a seamless way. In some cases they may become network switching nodes themselves and so will need to be dimensioned and managed like any other network element. In addition, it is expected that new satellite-based services will evolve offering enhanced capabilities and applications.

This is an exciting time for the industry. Within the next couple of years new low earth orbiting satellite constellations will provide mobile users with voice and data to handheld terminals. Meanwhile, the race to provide a global information infrastructure is only just beginning. Today it is clear that satellite systems will play an important role in providing global services, particularly where the terrestrial networks are underdeveloped or in places which can not be readily served by terrestrial means.

Brief History

The Space Age was born some 40 years ago with the launch, by the Soviet Union, of the first man-made satellite of the Earth, SPUTNIK 1, in October 1957. This opened up the possibility of using satellites to communicate effectively over large distances and ushered in an experimental period where technology developed rapidly.

ECHO 1 was launched in August 1960 to test out radio propagation through the atmosphere and ionosphere. It comprised a 100 foot diameter metallised plastic balloon, which acted as a passive isotropic reflector. Operating in an inclined orbit at an altitude of some 1600 km, it successfully demonstrated the first live two-way voice communications by satellite.

In 1961 President Kennedy invited 'all nations to participate in a communications satellite system in the interest of world peace and closer brotherhood among peoples of the world'. The following year the United States Congress passed the Communications Satellite Act, which led to the formation of INTELSAT in 1964.




Figure 1—TELSTAR satellite

The most famous experimental satellite TELSTAR (Figure 1) was launched on 10 July 1962. By demonstrating the first live transoceanic television broadcasts it gave the general public its first glimpse of the power of satellites to carry information around the world. Developed at AT&T's Bell Laboratories it consisted of an 88 cm diameter sphere covered with solar cells and weighing 80 kg. Unlike ECHO it acted as an active repeater and was the first satellite to carry a travelling wave tube (TWT) power amplifier producing an output power of a mere 3 W. Placed into a highly inclined orbit with a period of 158 minutes it received and transmitted signals between fast tracking antennas at Andover, Maine (US), Goonhilly Downs (England) and Pleumer-Bodou (France).

During the early days there was a lively debate on whether to use such low earth orbiting satellites or to place them into the geostationary orbit (GEO) some 36 000 km above the equator. In 1945 Arthur C. Clarke had noted that three satellites located above the equator in this higher orbit could provide a global communications network. At this altitude each satellite can see more than a third of the earth's surface while its orbital period matches that of the earth's rotation, meaning that from the ground the satellite appears to hover at the same location in the sky. To test out whether it was practical to use this high orbit, Hughes Aircraft Company was asked

to develop the SYNCOM series of experimental satellites. In August 1964 the third in the series became the first satellite to operate in the GEO, demonstrating its feasibility. In view of the fact that a single satellite could provide such a wide continuous coverage and there was no need to build a series of large and complex tracking antennas on the ground, the GEO became the natural choice at the time.

The United States, Canada, United Kingdom and a number of European countries set to work on setting up a global satellite system, and this led to the establishment of the International Satellite Organisation, INTELSAT, in August 1964. Starting life with 11 members it has grown to its present membership of 142 states. The first satellite it built and operated was EARLYBIRD (INTELSAT I) which was launched on 6 April 1965. Located above the Atlantic Ocean it had a mass in orbit of 38.5 kg with solar cells planned to provide 400 W at the end of its 18 month lifetime. It used a mere 50 MHz of bandwidth split into two 25 MHz channels or *transponders*. Since the only earth stations at the time were in North America and Europe it employed an antenna with a toroidal radiation pattern that covered the Northern Hemisphere. While its total capacity was only 240 voice circuits or one television channel it clearly showed that the quality of voice was fully acceptable, although the high orbit introduced a transmission delay of some 260 ms.

The 1960s were exciting times for the space industry which, with United States public support, was engaged in a hectic space race with the Soviet Union. United States industry continued to develop larger and more powerful satellites and the National Aeronautics and Space Administration (NASA) developed launchers to deliver them. By the time man finally landed on the moon in 1969, INTELSAT had completed Arthur C. Clarke's vision of a global network of three satellites with the

addition of INTELSAT II over the Pacific Ocean and INTELSAT III over the Indian Ocean.

Like the Americans, the Europeans saw the huge potential for the commercial exploitation of space and were keen to develop their own satellites and launch capability. The United Kingdom played a major role in the satellite construction industry but, to the disappointment of many, allowed the lead on launchers to pass to other Europeans. The French in particular were keen that Europe had its own means of access to space. After a shaky start, the two earlier European organisations ELDO and ESRO were combined to form the European Space Agency (ESA) in 1975. Its major triumph was the development of the highly successful Ariane launch vehicle, which is currently the most used launcher of commercial communications satellites.

The potential for the use of satellites to provide maritime services and direct-to-home broadcasting was recognised early on by governments around the world. As a result, many intergovernmental meetings were held to establish ways in which these could be enabled while retaining the legitimate rights of individual countries. Among the most prominent issues were the allocation of limited resources of radio spectrum and access to the GEO. These particular matters were progressed through world radio conferences held under the auspices of the ITU (International Telecommunications Union).

During the 1970s the US began using satellites to transmit signals across the continental United States, while INTELSAT continued to grow and extend the capabilities of its global system. Soon regional systems were established and the Europeans set up their own intergovernmental organisation modelled on INTELSAT called *EUTELSAT*. The interim organisation was established in 1974 and headquartered in Paris.

Discussions on the formation of an intergovernmental organisation to

provide maritime satellite service between the International Maritime Organisation and telecommunications administrations started around 1969. It was not until 1979 that Inmarsat was officially founded and, to reflect the leading role played by the United Kingdom, it was based in London.

The 1980s saw rapid growth in the number of satellites built and deployed. Apart from continued growth in the American domestic market, a number of other countries developed their own satellites. These included the Australian AUSSAT, Indonesian Palapa system and satellites serving other countries such as India, China and Mexico. In Europe a number of countries looked to their own satellites to provide national services, in particular, France, Germany and Italy.

The 1980s also saw major moves in the deregulation and liberalisation of telecommunications around the world. In the early 1980s the calls for 'open sky' policies grew ever louder, especially in the USA and the UK, in the growing international satellite field. To promote greater competition and to encourage its own industry in this period, the US granted licences to several companies to develop their own satellite systems. This led to the formation of a number of privately owned systems such as PanAmSat, Orion and Columbia. In the early days the most vociferous player was PanAmSat. It was particularly critical of the unique role vested in the US Signatory to INTELSAT, the COMSAT Corporation. This established COMSAT as a carrier's carrier and so any domestic telecommunications carrier had to come to them for access to INTELSAT and subsequently Inmarsat services. While PanAmSat represented the key challenger to INTELSAT in terms of global communications, the Luxemburg-based Société Européenne des Satellites (SES) challenged the position of EUTELSAT in Europe.

Another important development during this period was the emergence of a growing market in launcher

services. From the early days launchers were only available from the US and Soviet Union. Unfortunately for the industry as a whole during the Cold War, western satellites did not have access to the highly successful Russian Proton vehicle. By the early 1980s the key competitors were the expendable launch vehicles from Atlas and Delta and the new Ariane launcher marketed by Arianespace. However, increasingly NASA looked to its reusable Space Transportation System, or Space Shuttle, to launch commercial satellites, and allowed the production of its successful expendable launchers to be scaled down significantly. This trend came to an abrupt halt with the tragic loss of the Space Shuttle Challenger in early 1986. The Reagan administration responded by banning the use of the Shuttle for commercial satellite flights and encouraging US industry to develop and market their own expendable launchers. US industry responded positively and the Atlas and Delta programmes were re-started. A number of other well publicised launch failures at the time meant that for a period of about a year there were no western launchers in operation and a backlog of satellites to be launched built up. Ariane recovered the most quickly and managed to establish itself as the most successful launch supplier since that time. Further developments are covered later in this article.

The 1990s saw yet further expansion and development in broadcasting and mobile satellite systems. Improvements in both the satellite and earth-station technology have enabled medium-powered satellites to address the growing mass market for direct broadcasting. As a result, the original distinction between telecommunications and broadcasting satellites became ever more blurred. Meanwhile, submarine optical-fibre networks have gradually replaced satellites on the thickest long-distance routes such as across the Atlantic. This trend is set to continue as the effect of cable-on-

cable restoration begins to become a reality.

In the mobile area, Inmarsat broadened its mission beyond the maritime community and started to address other users on the move, such as land mobiles and the aeronautical community. The growth in terrestrial cellular networks and technological improvements prompted the revival of interest in low earth orbiting systems, and a number of private ventures emerged to offer satellite services to small hand-held terminals. The Iridium system led the way and by February 1998 had completed the deployment of over half of its constellation of 66 satellites. Globalstar will start launching its constellation this year; and ICO, a separate company spun off from Inmarsat, is currently building its system. This is clearly going to be a highly competitive market, and since there is considerable uncertainty on the rate at which the market will grow some consolidation in the industry is a possibility.

Over the last few years there has been intense interest in much

Table 1 Some Key Milestones in development of Satellite Communications

1957	Launch of SPUTNIK 1
1961	Kennedy's Policy Statement on satellite communications
1962	Launch of TELSTAR 1
1962	US Congress enacts Sat Com Act
1964	INTELSAT formed
1969	INTELSAT establishes global network
1975	ESA formed
1977	EUTELSAT formed
1979	Inmarsat formed
1979	Ariane's first launch
1983	EUTELSAT's first satellite ECS-1 enters service
1986	Shuttle Challenger disaster
1988	PanAmSat first flight on maiden Ariane 4
1990	Inmarsat launches its own satellites
1995	ICO spun off from Inmarsat
1996	Iridium starts launching its 66 satellite constellation

higher-capacity satellites which could bring down significantly the costs of transmission. Many novel broadband systems have been proposed to meet the expected dramatic increase in demand for satellite capacity. They have generally chosen to use higher frequencies (the so-called Ka-band at 20-30 GHz) where spectrum can be efficiently reused and the band is little used at present. Even if a few of the proposed systems are built they would represent a major expansion of the satellite communications industry and could change its character significantly. These are discussed later in this review.

Markets for Satellite Communications

The key advantages of satellite communications over terrestrial systems lie in that they:

- are well suited to broadcast and point-to-multipoint services,
- provide wide-area connectivity to remote areas,
- present transmission costs that are insensitive to distance,
- provide rapid communications capability well ahead of terrestrial deployment, and
- are well suited to provide mobile and personal communications.

The ability to cross national borders so effortlessly has introduced a range of political and regulatory issues. These range from concerns over cultural intrusion to satellite bypass of an incumbent common carrier. Since these issues directly impact the rate at which a global infrastructure can be implemented, satellite issues are often at the forefront of discussions between administrations. Such matters are discussed later in the context of the changing telecommunication environment.

As noted already there is no real competition these days between the use of satellites versus submarine optical-fibre systems on the high-density transoceanic routes. Route planners will always favour optical systems as they offer very high capacity and do not suffer from the transmission delay inevitably incurred by geostationary satellites. However, there are many locations which are unlikely to be reached by cable for the foreseeable future. At present from the United Kingdom there are over 70 countries that are only served by satellites.

The range of satellite-based services continues to steadily grow. This is illustrated in Table 2. For each service

there are an increasing number of providers and products on offer.

Market Trends

Broadcasting

This has shown spectacular growth over recent years. There is now fierce competition in the marketplace which ensures that satellite costs are kept to a minimum. Originally this simply meant TV broadcasts between country gateways but, as technology has improved, TV is increasingly distributed to cable networks or direct to the home. This has led to the emergence of truly mass markets for satellites. In recent years there has been considerable growth in

Table 2: Typical Satellite-Based Communications Services

Broadcasting

- Planned global broadcasts of world events like the World Cup and Olympics
- Occasional-use broadcasting of major news stories
- Distribution of TV channels to the headends of cable networks
- Satellite news gathering to feed stories for onward terrestrial transmission
- Business TV to enable companies to keep in touch with widespread offices
- Direct-to-home TV distribution
- CD-quality radio broadcasts
- Multicasting of Internet information to many sites simultaneously

Telecommunications

- International public switched services to distant country gateways
- Provision of domestic telecommunications where terrestrial infrastructure does not exist
- Private leased line services
- Data distribution to enable a business to update its remote sites rapidly
- Specialised corporate data networks operating to small dishes
- Interconnection of local area networks to provide wide area networks
- International backbone links carrying Internet traffic

Mobile and Personal Communications

- Voice, fax, data services to the maritime, aeronautical and land mobile community
- Aeronautical and maritime distress and safety services
- Vehicle tracking and cargo management
- Wide area paging
- Wide area extension of terrestrial cellular radio
- Differential correction of global positioning satellite navigation transmissions

Rapid Deployment Services

- Early provision prior to deployment of terrestrial infrastructure; for example, Eastern Europe
- Restoration of transoceanic cables
- Communications for temporary events
- Emergency services to remote locations
- Communications to areas disrupted by natural disasters
- Communications to world trouble spots like Bosnia or the Persian Gulf

many regions including the United States, Europe, Asia Pacific, South America and Southern Africa.

A key development is the arrival of digital television. When video material is digitised, compressed and multiplexed with other channels, several hundred channels can be delivered simultaneously to domestic set-top boxes. This is sure to bring down further the transmission costs and is likely to open up new opportunities for the broadcasting industry. There is no sign yet of demand reaching saturation, even where the same digital technology is being implemented terrestrially.

What applications would use these extra channels? Among the most likely are near video-on-demand, narrowcasting of specialised information, and expanded use of business television by companies to communicate with their employees. Another important growth area could be interactive television, where a low data return link, say using the terrestrial network, could call for downloading of television or high-speed data. This in turn could be used for applications such as home shopping. Interest in overcoming terrestrial bottlenecks has stimulated considerable interest in Internet multicasting by satellite which can be used routinely to update remote mirror sites. Such services continue to blur the distinction between data transport and broadcasting.

Fixed telecommunications

Although there has been steady growth in the carriage of public switched network services on a global basis, this has slowed to around one per cent in the case of INTELSAT. The growth is principally in services between the developing and the developed world. INTELSAT still provides an important service for most countries and will continue to do so for many years.

There is significant potential for satellites providing rural communications in many developing countries. Here wireless local-loop technology

could be used in conjunction with small satellite antennas to provide quick and cost-effective national and even international connectivity.

Significantly, a major growth area is in business services for specialised communities of interest. Here improvements in technology have allowed the use of very small aperture terminals (VSATs) which can be as small as 1–2 m diameter at Ku-band (11–14 GHz). Increasingly service providers are leasing satellite power and bandwidth and designing their own networks. Multinational companies have found satellites particularly useful in extending their reach. Among the leaders here are the banks and companies in the petrochemical industry. Growth areas include the interconnection of local area networks and the provision of company-wide intranets and links

a major growth area is in business services for specialised communities of interest

with their partners or customers through extranets.

Digital transmission is almost universal these days and the thrust is to make more efficient use of the channel bandwidth. Digital compression and interleaving techniques using digital circuit multiplication equipment (DCME) can provide voice circuit multiplication factors of at least four. This has allowed users to accommodate growing demand while reducing costs significantly. As data requirements have grown there is increased interest in providing integrated services digital network (ISDN) and other higher bit rate services. Both asynchronous transfer mode (ATM) and synchronous digital hierarchy (SDH) have been successfully carried over satellite, and in future these may become key elements in the international network. Public Internet traffic is beginning to be specifically identified in carriers' plans, and in the future large satellite

bearer channels may be needed to carry the backbones for new multimedia services.

Mobile and personal communications

The growth in satellite mobile communications has tracked the growth of terrestrial cellular services. Although originally aimed at the maritime community, satellites are increasingly being used for all mobile users whether they be on land, sea or in the air. As in the other services, increases in satellite power and improvements in terminal design have allowed ever smaller mobile terminals. Examples of current offerings are described in the Inmarsat discussion below. Satellites can clearly provide fill-in coverage and provide seamless roaming. New mobile satellite systems are expected to come into operation in the next couple of years

and are described in further detail later in this article. These are proprietary and will offer voice and low-speed data. They all use hand-held terminals that will be able to work with both satellite and terrestrial cellular systems. The potential market size is difficult to predict since most of the time people will find it more economical to use terrestrial fixed services where available.

Established Satellite Systems

This section discusses the development and plans of the traditional satellite organisations. Although most countries rely on the established international satellite consortia to meet their needs, a range of new systems has matured to the point where they offer real competition. Later in the article, the new emerging satellite systems which have yet to become established are discussed.

INTELSAT

Thirty years of steady growth

INTELSAT was the first international satellite consortium and is still the world's largest. Its original membership of 11 countries has now grown to 142 states and it serves around 200 countries. For many of the developing countries it provides the only link to the rest of the world. With revenues in 1997 of \$960 million it continues to be a commercial success story, although it no longer enjoys the virtual monopoly it once had in transoceanic satellite communications.

Wide range of services

INTELSAT's strength lies in its global network which offers a very wide range of services with high connectivity. The range of fixed satellite services includes:

- point-to-point public voice/data switched services;
- video services, including occasional use and longer-term video leases;
- specialised business services serving the needs of corporate users;
- leased transponders which can be customised to meet a variety of uses;
- domestic services to supplement a country's own telecommunications infrastructure; and
- restoration in the event of interruption of transoceanic optical cables.

The expanded capability of the INTELSAT fleet

From its headquarters in Washington, D.C., INTELSAT controls a satellite fleet of some 27 satellites. These are located over the Atlantic, Indian and Pacific Oceans and more recently in the Asia/Pacific region.

From its operations centre at headquarters it controls and monitors

the entire fleet. Communication with the satellites is effected through six ground stations located in China, Australia, Italy, North America, Hawaii and Germany. In addition, an extensive communications monitoring system is in place to monitor traffic in all the key coverage areas.

The satellites have grown considerably in size, power and capability over the years. EarlyBird (or INTELSAT I) offered a choice of either 240 voice circuits or one TV channel. The current satellites in the fleet can now carry around 100 000 voice circuits as well as simultaneously three TV channels. The early satellites had an operational design life of only a few years; this meant it was possible to upgrade the fleet reasonably quickly to take advantage of technological improvements. Last year INTELSAT signed contracts with Space Systems/Loral for its ninth generation of spacecraft which will last at least 13 years and is the largest and most complex civilian satellite yet built.

At the beginning of 1998 the satellite fleet consisted of four generations designated by the roman

numerals V, VI, VII and VIII. It also operates a single INTELSAT K satellite which is for TV and operates exclusively in the Ku-band (11–14 GHz). At the time of writing it has under construction six further satellites—a variant of the VIII series for the Americas, a further TV satellite (K-TV) for the Asia Pacific and four INTELSAT IXs (Figure 2) to replace the ageing VI series. The V series is about to be retired after operating for more than twice their original seven-year design lifetime.

Table 3 illustrates how, with each generation, extra capabilities have been added. The trend is clearly towards larger masses, higher powers, more transponders and higher capacities.

Future plans

While INTELSAT has operated successfully over the years as a consortium of telephone companies, it is facing increasing competition from a growing number of private satellite systems. In addition, it is seeing its growth in traditional public switched telephony increasingly being abated as high-capacity optical fibres are installed.

Figure 2—INTELSAT IX (courtesy of INTELSAT)



Table 3 INTELSAT Satellites

	V	VI	VII	VIII	IX
Prime contractor	Ford Aerospace	Hughes Aircraft Company	Space Systems/Loral	Lockheed Martin	Space Systems/Loral
First in service	1980	1989	1993	1997	End 2000
Launch mass (kg)	1950	4200	3800	3400	4400
Satellite power (W)	1300	2000	4000	5400	8000
Channel power in zone beam (dBW)	29	31	33	36	37
No. of transponders	32	48	40	44	56
Capacities (voice circuits) *	15 000	120 000	112 500	112 500	140 000
Comments	All roles. About to be retired	Primary and major path roles in AOR, IOR	Replacement for ageing Vs	Additional capacity for all regions	Replacement for VIs

* The number of voice circuits increased significantly in later years with the introduction of DCME.

Traditionally, INTELSAT is seen as an intergovernmental organisation with the governance and day-to-day control delegated to a Board of Governors drawn from its owners or Signatories to its operating agreement. However, restructuring has been on the agenda for several years and recently major progress has been achieved.

At the end of March 1998, the Assembly of Parties agreed to the immediate creation of an independent spin-off company, temporarily called *New Skies Satellites, N.V.* incorporated in The Netherlands. INTELSAT will transfer six satellites to the new organisation during the summer. The new company will be free to pursue opportunities in the growing video and multimedia fields. This step is seen as the first step in ultimate and full commercialisation of INTELSAT.

INTELSAT itself has recently taken steps to renew its existing fleet with the procurement of the INTELSAT IX satellites. These will start to be launched in the year 2000 and are likely to operate through to the year 2015.

While INTELSAT's future direction is currently under discussion, it has taken steps to register interest in the higher Ka frequency band.

EUTELSAT

EUTELSAT was created in 1977 with an initial membership of 17 countries. Its structure was based on the intergovernmental model used for INTELSAT and its headquarters were set up in Paris. Since that time it has steadily grown although recently its membership expanded significantly to encompass countries in Central and Eastern Europe after the collapse of the Soviet Union. At the time of writing the membership stands at 46 countries, with a couple of applications in the pipeline. Its mission was to offer a broad range of European telecommunication services although its major growth area has been in TV distribution field. In setting up the organisation the Europeans were keen to achieve satellite economies of scale to match those enjoyed by the United States.

EUTELSAT faced real competition from its earliest days. In part this was

because most early members had well-established terrestrial networks and were members of INTELSAT. However, soon a rival service provider emerged to serve Western Europe. This was the ASTRA network of satellites owned and operated by SES. In addition a number of states developed their own national systems to offer direct broadcasting. EUTELSAT has risen to the challenge and, fuelled by the apparent insatiable appetite for television broadcasting, has become very successful. Its coverage areas have extended and it clearly has ambitions to grow outside Europe to meet the needs of neighbouring states. By the year 2000 it is expected to have a network of 14 satellites.

EUTELSAT services

EUTELSAT carries international and domestic fixed and mobile services. They can be broadly classified as

- TV and radio channels for direct-to-home, cable and community reception;
- domestic and international telephony;
- business communications (called *satellite multiservice system*);
- land mobile communications (EUTELTRACS); and
- European Broadcasting Union TV and radio distribution.

By far the largest revenue stream is from television distribution which accounted for over 80 per cent of its £250 million revenue in 1997. This service has shown remarkable growth over the last decade and EUTELSAT is shortly set to offer over 1000 channels as digital TV channels to homes in Europe.

System development

In 1979 the new organisation ordered five European Communication Satellites (ECS) from the European Space Agency of which four were

Figure 3—EUTELSAT HOT BIRD 3 (courtesy of EUTELSAT)



successfully launched. After pre-operational trials with an earlier version, called the *Orbital Test Satellite*, it began full operations in 1983 providing wide coverage of Europe. Aerospatiale built five of the next generation of satellites, called *EUTELSAT II*. These provided more transponders with higher powers and longer lifetimes. The first satellite was launched in 1990 enabling the key television and business roles to be maintained. When considering the third generation, the requirements had grown to the stage that it was decided to separate the television and telecommunications requirement onto two distinct types of spacecraft. To consolidate the growing television community a series of HOT BIRD (see Figure 3) satellites was ordered, four to date, from Matra Marconi to be collocated at 13° East orbital location. Telecommunications were to be met by spacecraft with wider band transponders and a new W-Series was ordered from Aerospatiale. To date four W-Series satellites have been ordered, the first of which is due for launch in mid-1998. In addition, a Russian satellite SESAT (Siberia Europe Satellite) (see Figure 4) was ordered to encompass the former Soviet countries as well as Western Europe, which will be located at 36° East. Table 4 summarises the development of the various satellite generations.



Figure 4—SESAT: A EUTELSAT satellite for East-West communications (courtesy of EUTELSAT)

Table 4 EUTELSAT Satellite Fleet

Satellite Types	EUTELSAT I	EUTELSAT II	HOT BIRD	W Series	Sesat
Prime contractor	British Aerospace	Aerospatiale	Matra Marconi	Aerospatiale	NPO-PM
Entry into service	1983	1990	1996	1998	1998
Number of transponders	10	16	20 or 22	24	18
Number in series	4	4	4	4	1
Remarks	First generation. Five ordered, four successfully launched. Two remain in residual roles	Replaced Series I satellites. Five ordered, one launch failure. Last one modified to create HOT BIRD I	Four ordered and collocated at 13° East to offer direct-to-home TV	Telecommunications and TV distribution	Wider coverage to the East

Future plans

Like INTELSAT, restructuring has been an issue for sometime. The establishment of a private company is envisaged which will take over all the operations and assets of EUTELSAT. A small continuing intergovernmental organisation will oversee the remaining public service obligations (for example, telephony).

EUTELSAT has been highly active in developing its range of products and services to serve mobile and multimedia users. In addition, it has extended its reach beyond Europe. The latest satellites have steerable spot beams that can link to regions such as Africa, the Middle East and even South America.

Inmarsat

The London-based Inmarsat organisation was created in 1979 with

approximately 40 member countries. It has a structure along similar lines to INTELSAT and has now grown to more than 81 member countries with annual revenues of around \$370 million. The space segment used in the early days was a combination of MARISAT and MARECS satellites plus Mobile Communications Service (MCS) packages flown on board INTELSAT V spacecraft. In 1990, Inmarsat launched its own series of four satellites (Inmarsat-2) and has recently completed the successful launch (February 1998) of its fifth spot-beam-capable Inmarsat-3 satellite whose capacity is more than 20 times that of the first-generation satellites. Before the advent of the newer satellites, users of Inmarsat services were usually large companies who could afford to invest in the bulky mobiles and relatively high

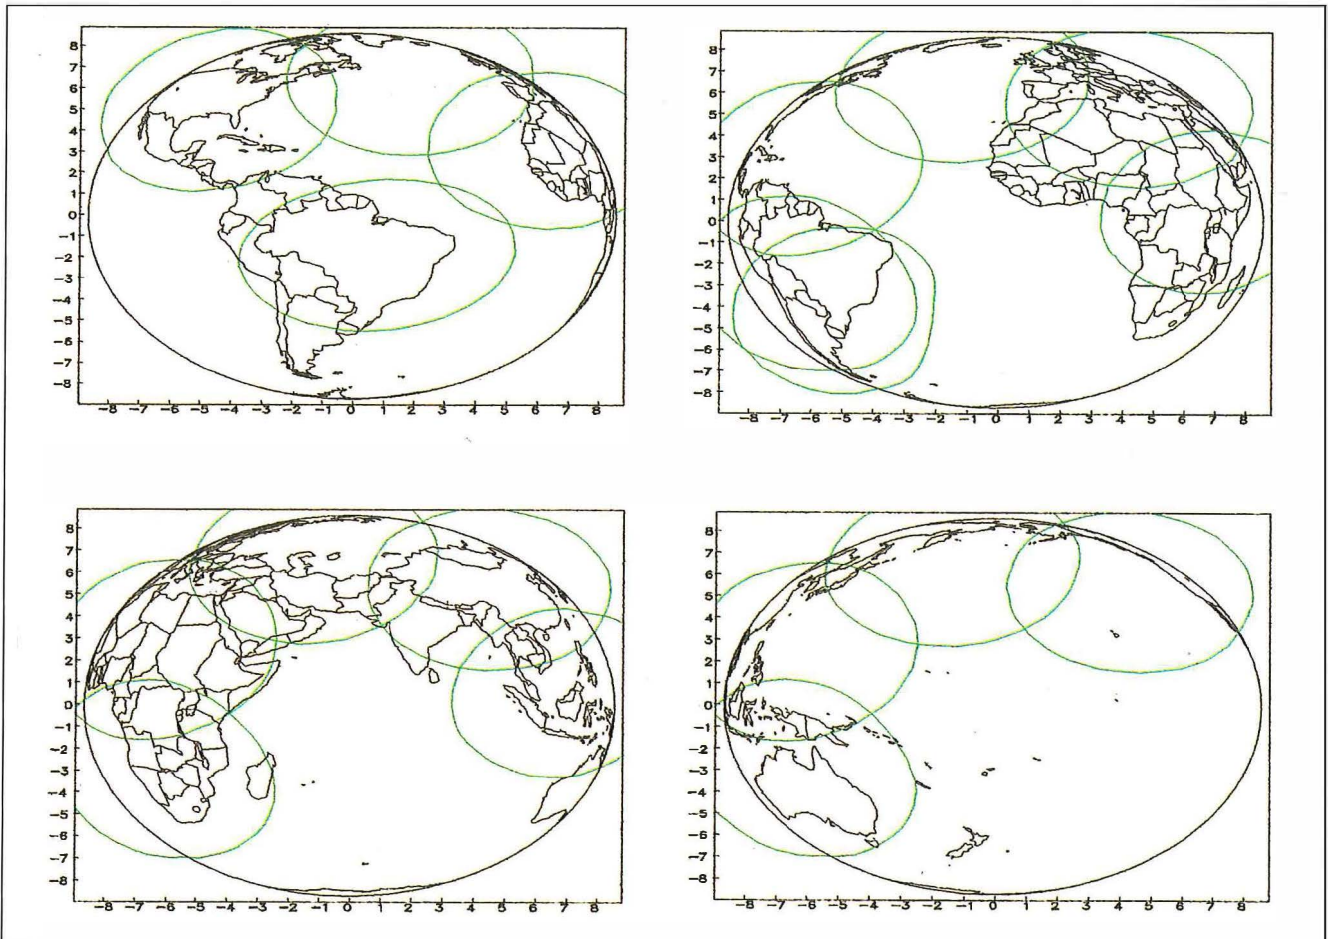
space segment charges that resulted from the low capacity of the satellites. However, with the increased power and sensitivity of the new spacecraft, the size, weight and cost of mobile terminals, along with the call costs, have fallen such that now anyone can communicate through a satellite telephone no bigger than a small laptop computer. This is especially attractive outside terrestrial mobile telephone coverage areas, which is in fact most of the geographical area of the world.

The global and spot beam coverage of the four primary Inmarsat-3 satellites is shown in Figure 5.

Services

Inmarsat is currently the premier global operator of mobile satellite systems on land, sea and in the air, and its numerous services and

Figure 5 – Coverage of Inmarsat-3 satellites. Top left: Inmarsat-3 AORW 54°W. Top right: Inmarsat-3 AORE 15°W. Bottom left: Inmarsat-3 IOR 63.9°E. Bottom right: Inmarsat-3 POR 178°E (courtesy of Inmarsat)



mobile terminal types are summarised in Table 5 along with the relevant BT brand name. The first service to be offered was Inmarsat-A which, although an analogue design, is still very much operational and currently provides communications to around 25 000 (mostly maritime) mobiles on the larger ocean-going ships. Inmarsat-B was developed as a digital replacement for Inmarsat-A and, by making more efficient use of the satellite, is significantly cheaper to use. Both Inmarsat-A and Inmarsat-B use antennae of around 0.85 m diameter that, for maritime variants, have to be stabilised against the motion of a vessel. Both types of terminal can be used for the Global Maritime Distress and Safety Service (GMDSS) and are

available in suitcase-sized land portable versions.

Inmarsat-C is a messaging system which offers access to international data services such as telex and X.25 packet data. The mobile terminals use small omnidirectional antennae requiring no pointing towards the satellite. The maritime variant of Inmarsat-C is also compliant with GMDSS. Vehicular mobile variants of Inmarsat-C are used by trucking companies for fleet tracking and load management.

Inmarsat-M is a smaller, less feature rich-version of Inmarsat-B and uses an antenna of around 0.4 m. It is also available in maritime and land portable forms.

Inmarsat-Phone is the name given to the smallest of the Inmarsat voice capable mobile services. The mobile

terminal is about the size of a portable laptop computer and weighs, with batteries, about 2 kg, and operates anywhere within the spot beams of the Inmarsat-3 satellites (which includes nearly all of the world's land masses and a large proportion of the ocean). Inmarsat-Phone is currently the fastest growing Inmarsat product with around 20 000 terminals being deployed in the first year of operation. A large dish (about 0.7 m diameter) version of Inmarsat Phone offers cheap voice, fax and data services for applications such as rural and remote communications in developing countries.

Aeronautical communications are provided by the global Aero-H and spot beam Aero-I services. An important part of the Inmarsat

Table 5 Primary Inmarsat Services

Service (BT Brand)	Mobile Terminal	Facilities
Inmarsat-A (BT Inmarsat-A)	0.85 m antenna. Maritime + suitcase versions available	Global operation of analogue voice, fax, voiceband data, telex, 64/56 kbit/s simplex data
Inmarsat-B (BT B-Sat)	0.85 m antenna. Maritime + suitcase versions available	Global digital (16 kbit/s) voice, fax, voiceband data, telex, 65/64 kbit/s duplex data
Inmarsat-C (C-sat)	Omni-directional antenna, vehicular + maritime versions available	Global messaging for e-mail, fax and telex, maritime distress and safety
Inmarsat-M (M-sat)	Briefcase portable + maritime versions available	Global digital voice (at 6.8kbit/s), 2.4kbit/s fax + data
Inmarsat-Phone (Mobiq)	Laptop sized portable. Maritime, vehicular mobile and big dish versions available	Spot-beam operation of digital voice (at 4.8 kbit/s), 2.4 kbit/s fax + data
Inmarsat-D	Small GPS receiver sized terminal	Global or spot-beam low-speed bothway messaging and paging service
Inmarsat-E	Omni-directional floating buoy mounted transmitter	Emergency position indicating reporting beacon (EPIRB) for maritime distress purposes
Aero-H (Skyphone)	Built into long haul public + executive jet aircraft	Global digital voice (at 9.6kbit/s), 2.4kbit/s fax + data, packet data, aero distress and safety
Aero-I	Smaller than Aero-H intended primarily for short-to-medium-haul aircraft	Spot-beam voice (at 4.8kbit/s), 2.4kbit/s fax + data, packet data
Point to multipoint	Various terminal sizes available depending on customers' requirements	Mostly used for conveying differential GPS correction information for oil and mineral exploration

aeronautical system is the packet data service which has facilitated the provision of aeronautical distress and safety communications, and the development of a number of value-added services to airlines including BT's own ALIS (Airline Interactive Data Services) system.

Inmarsat-D is a low-speed messaging/paging service and Inmarsat-E is an emergency position indicating and position reporting beacon (EPIRB) for maritime distress purposes.

Last but not least there are the point-to-multipoint services which are largely used to enhance the global positioning by satellite (GPS) service to provide submetre positioning accuracy.

Inmarsat ground segment

The ground segment for Inmarsat services generally consists of signatory-owned land earth stations (LESs) which access the satellites on a demand-assigned basis under the control of an Inmarsat-operated network control station; the keen competition between LES operators helps to keep the system efficient. More recently, the smaller size and wider appeal of the compact Inmarsat-Phone service has led to the development of Inmarsat service providers (ISPs) who package the mobile terminal and airtime charges in a similar manner to the service providers of a terrestrial cellular radio system.

Future plans

As with the other satellite consortia, Inmarsat has been discussing restructuring for some time. Following extensive discussions the organisation spun off a separate company called ICO Global Communications Ltd. to address the hand-held terminal market in 1995. This system is discussed later in the article.

The governing body of Inmarsat has recently agreed to convert to a public company which will be able to embark on major future investments such as the HORIZONS project,

whose service offerings will include the ability to offer affordable high-speed multimedia services to laptop-sized mobile terminals.

Other Systems

From the early days, separate domestic and regional systems emerged where the business could justify the large investment needed. Typically these were used to provide broadcasting facilities to link communities spread over large continental land masses such as the United States or Europe. The United States led the way here in the 1970s but during the 1980s this trend started to spread to other regions of the world. As deregulation in the US took hold, private systems emerged to exploit the new opportunities in the international arena. Not all have survived but the successful ones have built sizeable networks.

In Europe, ASTRA, owned by Luxemburg-based SES, has been operating since 1985. It has steadily built up a constellation of collocated satellites at the 19.2° East orbital location (Table 6). It currently operates a network of seven first-generation satellites providing direct-to-home television. This is familiar to United Kingdom households who subscribe to the Sky satellite system.

A series of much larger satellites is being planned to replace the first generation. ASTRA currently offers over 320 television channels to 68 million homes in Europe. Interactive services are to be offered on the eighth first-generation satellite. This would enable an individual user to request the satellite to downlink a broadband signal by means of a return link transmitted to the satellite by the terminal.

Two ASTRA 2 satellites, each carrying 28 transponders, are planned to be launched in 1998 to a new location at 28.2° East. ASTRA 2A is provided by Hughes and 2B by Matra Marconi.

On the wider international scene another important player is the newly structured PanAmSat. Established in the early 1980s by the late entrepreneur Rene Anselmo, it set itself global ambitions in the TV market. Starting with a single satellite in the Atlantic Ocean serving the Americas and Europe, PanAmSat went on to buy further satellites for the Atlantic, Indian and Pacific Oceans. However, after his death, the company was acquired by Hughes Electronics. Hughes is vertically integrated. It not only builds satellites but owns and operates the leading US domestic Galaxy network in the United States.

Table 6 ASTRA Constellation, 1998

Satellites at 19.2° East	Prime Contractor	Launch Date	Number of transponders
ASTRA 1A	GE Astro Space	December 1988	16
ASTRA 1B	GE Astro Space	March 1991	16
ASTRA 1C	Hughes HS601	May 1993	18
ASTRA 1D	Hughes HS601	Nov 1994	18
ASTRA 1E	Hughes HS601	Oct 1995	18
ASTRA 1F	Hughes HS601	April 1996	22
ASTRA 1G	Hughes HS601	December 1997	28
ASTRA 1H	Hughes HS601	July 1998	28

The satellites have now been integrated into a network which by the year 2000 will comprise some 20 spacecraft (Table 7). With this consolidation the founder's goal of creating a true global system has been realised.

Competition is, however, heating up as another US company, Loral Space and Communications, is beginning to construct its own global network. Loral, under the direction of Bernard Schwartz, has grown from a relatively small satellite manufacturing company into a major player in the satellite industry. It owns the mobile satellite system Globalstar

Table 7 The New PanAmSat Constellation by year 2000

Satellite	Location (degrees)
PAS-4	68.5 E
PAS-7	68.5 E
PAS-9	72.7 E
PAS-8	166 E
PAS-2	169 E
Galaxy 1R	133 W
Galaxy IX	127 W
Galaxy V	125W
SBS-5	123 W
Galaxy X	123 W
Galaxy IV	99 W
Galaxy III R	95 W
Galaxy VIII-i	95 W
Galaxy VII	91 W
SBS-4	77 W
Galaxy XI	74 W
PAS-5	58 W
PAS-1	45 W
PAS -iR	45 W
PAS-6	43 W
PAS-3	43 W

and recently acquired the US domestic Skynet satellite system from AT&T and Orion Network Systems, which currently operates one international satellite and is building another two.

On the mobile side, several satellites containing mobile communications packages have been launched including those owned by American Mobile Satellite Corporation (AMSC), which covers the North American region; OPTUS in Australasia and MTSAT in the North Pacific and Asian areas.

It is clear that both the established consortia and these new systems have plans to develop and extend their networks and services. However, even grander systems are on the drawing board. This reflects confidence that the satellite scene is set for major expansion as the industry responds to the changing telecommunications environment.

The Changing Environment

The three international satellite consortia were born and grew up in a telecommunication environment dominated by national PTT monopolies. As owners they provided the necessary capital investment and controlled the way they developed. As signatories to the operating agreements they generally provided distribution within their own countries.

This environment has radically changed and in response we now see a fundamental restructuring of the communication satellite industry. Increasingly countries are privatising their PTTs and allowing competition within their borders. With this competition, new regulatory regimes are being established with clearer separation between the established carrier and the regulatory bodies. As the new operators emerge and seek to extend their markets internationally there are growing calls for more open access to other countries' markets. Countries are naturally protective with respect to their own

industry but there has been considerable progress in this area as exemplified by the agreement to fully open European telecommunications in 1998 and the recent World Trade Organisation agreement to open up markets to some 70 countries.

The changing market has also put increasing pressure on the way the consortia operate. Fairly or not, the existing satellite consortia are seen as slow, overly bureaucratic and inflexible. The open access approach and non-discriminatory nature of these organisations are seen as outdated and inhibiting the systems from actively competing with their private rivals. Increasingly there are calls for more freedom for the organisations to find alternative sources of capital and to partner with private organisations such as satellite manufacturers. The United Kingdom administration has been one of the leaders in trying to restructure the consortia but, as discussed earlier, it is difficult to reconcile these commercial pressures with the legitimate concerns of the developing world.

The satellite supply industry is also undergoing major structural change. Since the collapse of the Soviet Union, the United States aerospace and defence industry has undergone major consolidation and refocused on civilian opportunities. This has led to the creation of very large companies such as Lockheed Martin and Boeing, which have acquired and integrated individual space-related companies. This consolidation has led to rationalisation of facilities and the opportunity to combine development efforts. This has been happening more gradually in Europe but we see a similar trend with the creation of the Anglo-French Matra Marconi in 1989 and its current plans to merge with the German aerospace company, DASA. The European space industry is still under pressure to consolidate further and we can expect further rationalisation in the years ahead. Since the margins from building satellites are

relatively small, we also see these companies seeking to move up the value chain by operating satellite systems and becoming end-to-end service providers in their own right.

As in the early days of the space age, the United States continues to lead in terms of technology. The new companies are increasingly exploiting the results of many years of large Department of Defense expenditure in advanced technology. An example of this is the Iridium mobile satellite system whose technologies were developed under the US Strategic Defense Initiative, or Star Wars programme. Many of the proposed new systems rely on such advanced technologies. As the range of satellite operators increases even the US and European military are looking to civilian satellites to meet their communications requirements.

While the appetite appears to exist to develop advanced satellite systems around the world, the pace at which they become a reality will depend critically on the resolution of a range of regulatory issues. Firstly, they need access to adequate spectrum. This complex process takes time and has been the subject of intense debate at recent world radio conferences held under the auspices of the ITU. Once spectrum has been allocated there is a need to ensure that the service provider has access to the market through appropriate national licences. Here fear of national 'bypass' of the incumbent operator has often slowed down rollout of new systems. Even when these 'landing rights' have been acquired, there is the issue of frequency sharing arrangements within the country. Often spectrum is already being used terrestrially, and existing users are reluctant to make room for new players. Once these issues have been resolved then installations have to meet both national and local planning rules. These can be difficult if, for example, the antennas are too large or raise other problems for local residents.

Although these difficulties may seem daunting, the prospects for new satellite services remain bright. The potential market is so great and the amount of investment planned is so significant that it is believed these obstacles will be overcome.

Orbits Revisited

In the early years of satellite communications, the capacity of the then available launch vehicles was insufficient to lift a useful payload into GEO and so the early communications satellites were placed into low earth orbit. As technology developed, it became possible to lift satellites into the geostationary position and until recently this has remained the most popular choice of orbit for mainstream telecommunications satellites. There are many compelling reasons for the use of the GEO orbit. These include the ease with which ground station antenna systems can be deployed and the wide area that such a satellite can see. The GEO configuration is still likely to be the best choice for broadcast, major point-to-point communications and VSAT applications. However, the choice is less clear for personal satellite communications where the antenna on the hand-held terminal is a small omnidirectional device which, therefore, does not need to be pointed at the satellite. It is also important to maintain a high elevation angle to the satellite to increase the likelihood that a satellite will be visible in towns and cities where users are likely to be in deep 'valleys' with respect to the satellite position. There is also the issue of the relatively large propagation delay of geostationary satellites which can be important for voice and some data applications where protocols may break down or become inefficient when faced with a large end-to-end delay. There is also the fact that geostationary satellites do not cover the extreme polar regions and that to produce the very small

spot beam patterns on the earth required for personal mobile operation very large spacecraft antennas are needed. This coupled with the economic need to concentrate as much traffic as possible into each spacecraft results in each geostationary satellite becoming extremely large and complex, stretching to the limit the capabilities of spacecraft and launcher technology. Of course, since the coverage of an individual satellite at the low orbit heights is much more restricted, more satellites are needed to provide global communications and, for continuous communications, some form of seamless handover is required between satellites. In addition, a much more complex ground-station system with multiple tracking antennas at each site is required. These factors tend to increase the costs of a lower earth orbit system over that of a geostationary one offering similar coverage. Nevertheless, these disadvantages have been deemed by some operators to be outweighed by the advantages and all three of the main global mobile personal communications by satellite (GMPCS) systems (that is, Iridium, ICO and Globalstar) have chosen non-geostationary orbit configurations for their systems.

Orbit height

In choosing an orbit for a communications satellite it is generally best to avoid the regions around the earth of intense radiation, the Van Allen belts, where high-energy particles from the sun are entrapped by the Earth's magnetic field. Peaks of radiation occur at altitudes of around 3000–7000 km and again at around 13 000–20 000 km where prolonged exposure can seriously shorten a satellite's lifetime. Satellites positioned below the first peak of Van Allen radiation are known as *low earth orbit* (LEO) satellites and satellites positioned in the gaps between the radiation belts known as *medium earth orbit* (MEO) satellites.

Satellites have always employed advanced technology. They are the result of very expensive research and development programmes that take many years to perfect.

Table 8 summarises the main differences between the GEO, MEO and LEO orbital configurations.

A further type of orbit is the highly inclined elliptical orbit (HEO), which can provide polar coverage and high elevation angles at high latitudes while still maintaining some of the advantages of a GEO in that the satellites have little apparent movement, as seen by mobile terminals, when they are furthest from the earth.

Trends in Technology

Satellites have always employed advanced technology. They are the result of very expensive research and development programmes that take many years to perfect. In the early days, the development was carried out under military or national programmes but as military budgets have been trimmed increasingly they are being developed by private companies. Satellite operators are generally commercial organisations and need to ensure reliable revenue from their operations. They therefore need to be satisfied that the technology deployed is fully qualified and tested before being put into service. This conservative approach has worked well in the past; however, perhaps we can see a change here. Increased urgency in getting systems into the marketplace means that manufacturers need to speed up the development cycle, and operators have to accept greater risks. Although individual early failures are more probable, the overall impact is likely to accelerate the rate of development.

In contrast to the early days when satellites were technology driven, increasingly the choice of which technology to develop is the result of the changing perception of market needs. This poses difficulties for satellite planners since a typical satellite can take two to three years to build and may need to operate for at least a decade.

Over the last 30 years major technical advances have been made at every level of integration. Satellite

Table 8 Comparison of Satellite Orbits

Orbit Type	Orbit Shape	Orbit Height (km, typical)	Typical Satellite Lifetime (years)	Typical Satellite Delay (ms)	Number of Satellites for Global Coverage
GEO	Circular	36 000	13	275	3–4
MEO	Circular	10 000	10	75	10–12
LEO	Circular	750–2000	6	10	45–70

network architectures have developed from a satellite operating to a few large national gateway earth stations to the highly varied and heterogeneous networks that are evident today. It is notoriously difficult to forecast exactly how future satellite networks will evolve. A relatively minor technical improvement in one area can have a dramatic impact on the overall system cost structure. This in turn can modify the service offering and create new commercial opportunities.

To illustrate this, consider how the concept of direct-to-home broadcasting has evolved. In the late 1970s it was assumed that this would be implemented by deploying, on each satellite, a handful of very powerful travelling-wave tube amplifiers with output powers in excess of 200 W. This was at a time when satellites were particularly constrained in terms of power and mass. In the event this approach was overtaken by modest increases in satellite powers and earth station receiver sensitivity. These allowed direct-to-home broadcasting with much more modest powers of around 50 W. This in turn enabled the satellites to carry many more channels and consequently brought down transmission unit costs significantly. This greatly stimulated the market. Subsequent improvements in platform capability have been exploited to continue to drive down unit costs.

Steady incremental improvements appear to be the rule although when combined in novel configurations they can introduce apparently revolutionary change. To see where the future is taking us we should

therefore look at the developments currently in the pipeline.

A number of key technological developments are outlined below.

Improvements in radio frequency component technology

More efficient solid-state amplifiers will reduce power requirements both on the ground and in space. Solid-state amplifiers are already displacing travelling-wave tubes at lower power levels. They offer the added benefit of offering more linear channel characteristics. Low-noise devices employing high-electron-mobility transistor (HEMT) technology have enabled more sensitive receivers allowing smaller antennas to achieve the same performance. On the satellite an increased level of component integration has enabled mass, power and space savings. Although modest changes in themselves, when multiplied throughout a system they can have a significant impact. Increased interest in the use of higher frequencies such as Ka-band (20–30 GHz) has spurred development in radio-frequency components in this band. However, there will need to be significant progress in this area to enable costs to be reduced to the point where they can be economically incorporated in the millions of multimedia terminals envisaged by the proposed new satellite systems.

Extended use of digital signal processing

Almost all services these days use digital transmission. New digital signal processing technology can be used to provide coding, filtering and

Figure 6—Ariane 5 launch vehicle
(courtesy of Arianespace)

modulation. These allow better and more reliable transmission performance without the need for adjusting analogue components. Another advantage of digital processing is the ability to reconfigure easily the function of devices in software, say to increase the degree of coding used or the transmission rate. This capability is particularly interesting in the case of mobile applications where, by using *software radio*, terminals can be reconfigured over the air. This is already apparent in some new closed-user satellite networks where system-wide network software changes can be made by downloading new code to terminals in the field.

Increased use of on-board processing

Availability of powerful space qualified processors has opened up opportunities for more practical on-board processing. Not only are they being used to monitor and control the satellite, but they can also be used in the payload to greatly enhance its capabilities. Although transparent transponders are very flexible there is growing interest in demodulating the radio signal down to baseband signals, switching them into different downlink streams and remodulating them for onward transmission. This allows more efficient bit transport through the satellite, and the provision of a fast switch would allow signals to be directed to many beams in a very flexible manner.

Development of larger more powerful satellite platforms

All the major manufacturers have developed their platforms so that they can launch larger payloads, thereby reducing the incremental cost of adding extra transponders. For example, Hughes has come a long way from SYNCOM where its latest bus design, the HS702, promises to be able to offer payloads up to 1000 kg and, by using arrays of efficient gallium arsenide photocells, can deliver up to 15 kW of solar power. Such a spacecraft could carry



up to 70 typical transponders. The new product line platforms are more modular in design so that greater use can be made of production line techniques. Simpler integration and testing are already reducing delivery schedules allowing the satellites to respond more quickly to market requirements.

Increases in launcher capability

As satellites have grown in physical size and mass, the launch providers have had to increase their lift capabilities. These have been readily exploited and led to further calls for higher lift capabilities, and so the cycle continues. All the major launch service providers have plans to upgrade their capability and so they hope to increase their market share. This is still an expensive business and it costs typically \$20 000 to lift one kilogram into geostationary orbit.

In 1997 it is estimated that the European Arianespace launched around 40 per cent of the world's

communication satellites. The competition is, however, hotting up with the Russian Proton and American Atlas and Delta offering launch services. Arianespace has built up its business on the highly successful Ariane 4 rocket but has staked its future on the more powerful Ariane 5 (Figure 6). Unfortunately the first launch was a well-publicised failure, but the second launch, illustrated here, suggests that the development programme is back on track.

During the Cold War, the West had no access to Russian launcher technology. That has now changed and Lockheed Martin's International Launch Services now markets the Russian Proton vehicle through its joint venture with the Russian companies Krunichev and Energia. The Proton has a very successful history and has already demonstrated that it can launch Western satellites. A more capable version called ProtonM/Breeze M was recently selected by INTELSAT to launch one of its next generation satellites in the year 2000.

the number of systems and satellites have shown steady growth and this shows no sign of slowing down. There is therefore considerable optimism in the industry on its long-term prospects.

Apart from the existing suppliers, an intriguing new player is about to enter the market. The SeaLaunch limited partnership has been established by Boeing Commercial Space Company, the Ukrainian KB Yuzhnoye/PO Yuzhmash, RSC-Energia of Russia and Norwegian Kvaerner Maritime a.s. This company plans to use the powerful Ukrainian Zenit rocket from a refurbished North Sea oil rig that will be floated out into the Pacific. This innovative approach is being seriously backed by Boeing and its first test launch is planned for late 1998.

Other countries are looking to provide launchers and these include the Chinese Long March and the Japanese HII. Entry into this business is often beset with difficulties and problems, and it may be several years before these new entrants will be considered technically acceptable for regular commercial use. Although the US Space Shuttle could deliver satellites into low earth orbit it has not been allowed to do so on a regular basis since the Challenger disaster in 1986.

Although competition is increasing, the demand for launch services is likely to outstrip supply over the next few years. In this case, a real decrease in launch costs is not expected in the near future.

Other technology improvements

These include:

- *Advanced antenna designs* that allow more efficient use of satellite power. The use of shaped reflectors is already widespread and very large antennas are needed to generate spot beams for mobile users. Inefficiencies in power combining have meant that extensive use of phased arrays is still slow in coming.
- *Wider use of intersatellite links* to provide interconnectivity between satellites can reduce the number of ground stations needed. This technique is being

used for new mobile satellite systems such as Iridium. Initially they will use radio frequencies but in the future optical links are envisaged.

- *Advanced means of propulsion* can either extend lifetime or allow larger payloads. Since much of a satellite's mass in orbit is in the form of fuel for keeping the satellite on station, a more efficient form of propulsion can be seen to have immediate benefits. Low thrust ion propulsion systems have been developed over the years. Unfortunately they have been dogged by technical problems, and with the larger and longer-lived platforms currently available the need for their use has been somewhat reduced.

Emerging New Satellite Systems

Conventional wisdom contends that in the years ahead we will see a rapid expansion in global communications. Typical estimates put the worth of the worldwide communications industry as \$100 billion a year by the year 2000. At present the commercial satellite industry has been estimated at \$6 billion annually and so even if satellites addressed a few percent of the growing market this would represent a major expansion. As we have seen, the number of systems and satellites have shown steady growth and this shows no sign of slowing down. There is therefore considerable optimism in the industry on its long-term prospects.

This positive prospect has encouraged both existing industrial parties and new players to contemplate radically new types of system. New ventures have already sprung up to meet the challenge. These differ markedly from the international consortia that established the industry. They are typically run as private companies that are free to invest their own resources or seek finance in the open market. Freed to

pursue purely commercial goals they can move quickly in search of the most profitable business. Interestingly, in developing these systems they have had to form partnerships with a variety of companies and telcos. As a result they are beginning to resemble consortia themselves.

Unencumbered by operating existing networks and technologies many have proposed radical if not visionary approaches. Not surprisingly United States companies have been leading the way supported by access to advanced technology and tacit US government support. To date, most of the proponents are existing aerospace or communications companies who see the potential for such systems offering long-term growth potential. Traditional telecommunications operators are clearly interested in their plans and have invested selectively in some of them. However in the more ambitious and expensive systems there has been a more cautious approach.

There is no shortage of planned systems. Since availability of spectrum is a long lead-time item, the first hurdle has been to secure a sufficient allocation. As wireless communication for terrestrial use is also set to expand in the years ahead, finding available and suitable spectrum has been a major issue. Existing satellite allocations at traditional bands such as C-band (4/6 GHz) and Ku-band (11-14 GHz) are already heavily congested. As a result there has been considerable interest in higher frequencies; that is, Ka-band (20/30 GHz) and more recently V-band (40/50 GHz). Higher frequencies not only provide considerable bandwidth but can employ smaller antennas on the ground and on the satellite. One of the principal reasons they have not been used to date is that they suffer from severe propagation impairments in the presence of water such as rain.

The new systems can be differentiated into two types. Satellites providing mobile services offer narrowband communications (up to

say 144 kbit/s) while the fixed terminal services can offer broadband, which usually is taken to start at T1 rate (1.544 Mbit/s). The following sections discuss these two types. The narrowband systems described are well into construction and the first Iridium system is expected to go into service in 1998. The broadband systems are more ambitious and are less well developed. They are mostly in the design phase. Some may never be built.

Narrowband systems

There are several new narrowband satellite communications systems scheduled to be deployed over the next few years. Firstly we have the GMPCSs of Iridium, Globalstar and ICO whose salient characteristics are given in Table 9. All of these systems will employ non-geostationary orbital configurations and all will offer voice plus low speed fax and data services to handheld mobile terminals around the same size as today's cellular handsets.

To take an example, the ICO system has 12 satellites (10 working plus two spares) arranged in two orbital planes inclined at 45°. The orbits are of the MEO type with satellites approximately 10 000 km above the surface of the earth (see Figure 7). This results in an orbital period of six hours with each satellite being visible from one point on the earth's surface for typically about 20 minutes. Clearly, to achieve continuous communications, handover between satellites is required which is analogous to a terrestrial cellular radio




Figure 7—ICO constellation (courtesy of ICO)

system handing off between cells. Each ICO satellite has a capacity of around 4500 voice channels.

Globalstar and Iridium both use the LEO configuration with the result that each satellite covers a smaller area of the earth and therefore demands many more (but of individually lower capacity) satellites than ICO. Apart from the satellites, a significant difference between the Iridium and Globalstar systems is in the ground station network. The Iridium system has a relatively small number of ground stations (around 12) and uses intersatellite links and on-board satellite switching to route calls back to a satellite that can 'see' a ground station. Globalstar's solution is quite different and avoids complex intersatellite routing by having a very large number of ground stations.

There are also a number of proposed regional geostationary satellite systems using *super GEO* satellites such as EAST (Euro-African Satellite Telecommunications) backed by Matra-Marconi Space Ltd. The satellite for EAST, shown in Figure 8, is planned as a multipurpose design providing a range of services including narrowband




Figure 8—EAST satellite (courtesy of Matra Marconi Space)

communications to hand-held user terminals and higher bandwidth services to VSAT antennas. An idea of the size of the satellite can be gained from the fact that the antenna is around 13.5 m diameter.

Other embryonic satellite systems include the so-called *little LEO* systems such as ORBCOMM, which will use up to 36 small (40 kg) satellites in low earth orbit to provide mobile tracking, remote monitoring and commercial and personal messaging services.

Broadband systems

Over the last few years there has been considerable debate over the construction of a global information infrastructure. This has been fuelled by rapid development of the public Internet and affordable computing facilities. Availability of broadband satellite systems can play an important role in making this concept a reality.

At present the population of computer users of the Internet is estimated to be 60 million. The ITU estimates that this figure will increase to 300 million by the year 2000. In addition, there is a rapid increase in computer processing power, the availability of cheap data storage and increased demands being placed on local area networks. As a consequence, communication between distant users may well become a limiting factor.

Already the existing telecommunications networks are struggling to

Table 9 Global Mobile Personal Communications Systems

System Name	Orbit Type	Number of Operational Satellites	Inter-satellite Links	Number of Ground Stations for Global Coverage	Scheduled Service Start	Rough System Cost (\$ billion) (see note)
Iridium	LEO	66	Yes	12	3Q 1998	5
Globalstar	LEO	48	No	100–150	1999	2.6
ICO	MEO	10	No	12	2000	4.5

Note: From tele.com, January 1998

The growth of the Internet has clearly suggested that a multi-media market is indeed waiting to be developed. Consequently, all emerging satellite systems claim to address this market.

provide a reliable and efficient service and so a number of solutions are being sought. A range of terrestrial wideband access technologies is being developed but all such solutions are capital intensive and time consuming to implement. This is where satellites can play their part. They can overcome terrestrial bottlenecks, provide backbone capacity to the terrestrial networks, reach remote locations, augment existing facilities relatively quickly and provide 'last mile' access.

The growth of the Internet has clearly suggested that a multimedia market is indeed waiting to be developed. Consequently, all emerging satellite systems claim to address this market. There is no shortage of ideas on what this market will look like, but using the terminology of the computer industry there is as yet no consensus on a 'killer application'. Among the many applications envisaged are videoconferencing to support distance learning or telemedicine, and on-line reference services to warehoused data. Every market sector will no doubt develop its own set of applications and this is likely to be a major growth industry for the next century.

At present the most promising satellite uses appear to be:

- provision of high-capacity international backbones to large country gateways,
- extension of corporate networks to remote locations,
- provision of broadband access to the growing number of small-to-medium companies such as the fast-growing communities of Internet service providers,
- network access to the growing population of small-office/home-office users,
- broadcasting of multimedia data to many locations simultaneously, and
- provision of broadband links to residential users.

The key satellite features needed for such services are:

- high capacity, leading to reduced cost of bit transport;
- flexible bandwidth assignment (for example, bandwidth on demand);
- small inexpensive terminals (less than \$1000);
- global or wide area coverage to capture maximum number of users;
- availability of high bit rates (at least two-way transmission at 1.544 Mbit/s);
- ability to carry the widest range of multimedia applications; and
- seamless provision of high-quality and availability of an end-to-end service.

These features have typically been translated into:

- significant reuse of limited bandwidth by means of using multiple spot beams;
- on-board fast packet switching to maximise satellite throughput;
- use of available higher frequencies, although mitigation techniques are needed to overcome propagation impairments;

- integration of satellite network management into the rest of the network; and
- new earth terminal designs to handle new satellite system requirements for tracking and use of new frequencies.

Since data is now seen to be the main traffic type, the problems encountered by satellite delay or network 'latency' need to be addressed. These are legitimate concerns since the data protocols currently used are not optimised to accommodate satellites. A mismatch here can reduce data throughput, especially if there are errors on the satellite link. However, new Internet data protocols are being developed in the relevant international fora.

Recently there has been a lively debate over whether satellites can actually carry Internet-type traffic. In one sense this is strange since satellites were involved in the early developments of the Internet to link diverse local area networks. This is being actively investigated and it is worth noting that even terrestrial systems have their own delay problems. To avoid this issue altogether some of the broadband systems have opted for constellations of low earth orbiting satellites. Others have opted for traditional geostationary orbit or even proposed hybrid geostationary and non-geostationary constellations.

A wide range of systems have been proposed. Table 10 indicates a

Table 10 Some of the Proposed Broadband Systems

System	Proponents	Satellite Constellation
Astrolink	Lockheed Martin	9 GEO satellites — Ka band
SPACEWAY	Hughes	8–20 satellites — Ka band
Celestri	Motorola	63 LEO satellites — Ka band 9 GEO Satellites
CYBERSTAR	Loral Space & Communications	3 GEO — Ka band
SkyBridge	Alcatel	64 LEO satellites — Ku-band
Teledesic	Craig McCaw/Gates/Boeing	288 LEO — Ka-band



Figure 9—Teledesic system (courtesy of Teledesic)

few of the them and shows how they have adopted widely different strategies. The estimated costs for constructing these systems run into several billion dollars. It is clearly a challenge for these proponents to raise such finance although this has been achieved in the case of the narrowband systems. In making the business case for these it should be noted that for the low earth orbiting constellations most of the satellites need to be in place before a reliable service can be offered to end-users.


Although there is wide variation in detail, a few characteristics are common. Typically the satellite serves users with dish sizes down to 600 mm. At the higher frequency Ka-band (20–30 GHz), they have to operate at higher elevation angles to provide visibility and reduce propagation losses. For terminals operating to the low earth orbiting constellations they have to use fast tracking techniques and be able to handle satellite handover seamlessly while handling connections. The uplink terminals usually offer capacity from tens of kilobits per second up to a few megabits per second. On the downlink, terminals can receive signals up to 155 Mbit/s. To make a business case, most of the systems require an immense growth in satellite user base and expect to be serving millions of end-users.

The satellites provide multiple small spot beams that on the earth corresponds to cells between 300–500 km diameter. Most of the satellites provide extensive on-board processing including ATM-like switches. Bandwidth is delivered as and when required rather than being permanently assigned to the terminals. To cut down the need for extensive ground networks a number have to use intersatellite links. They all claim to be ready to offer service in the next 3–5 years

Probably the most ambitious of these systems is Teledesic (see Figure 9) and this is now discussed.



Top left: orbits; top right: satellites; bottom: network



Teledesic's 'Internet in the Sky'TM

The vision of satellites providing cheap global access to the information society led Craig McCaw to found Teledesic in 1990 in a suburb of Seattle. Backed by Microsoft's Bill Gates and the giant Boeing Company, this private company has set itself an immense challenge. It sees itself literally providing an Internet in the sky.

The Teledesic network consists of a constellation of some 288 satellites in low earth orbit. Covering nearly 100 per cent of the world's population, it is designed to support millions of simultaneous users. At an altitude up to 1400 km the constellation consists of twelve planes each containing 24 satellites. A personal computer would access the network via a small terminal that tracks satellites while they pass overhead at an elevation of greater than 40° above the horizon. The terminal-satellite communication links operate within an internationally agreed portion of the Ka-band (28.6–29.1 GHz in the uplink and 18.8–19.3 GHz in the downlink). Each satellite acts as a switching node routing traffic via intersatellite links to neighbouring satellites and then down to a terrestrial service provider

connected to the Internet. Within any circular area of 100 km radius the network would support at least 500 Mbit/s of data to and from user terminals. Calls would be established on demand allowing capacity to be allocated only when needed.

As can be appreciated, to realise such a system requires significant advances in technology over conventional geostationary satellites in a range of areas. A key development is to bring the cost of the Ka-band terminals down to a level that can be afforded by a truly mass market. Since this is likely to involve phased arrays, many experts believe we need a technology breakthrough to realise this economically.

The overall cost of the system is large, with its own estimates put at \$9 billion. Again some experts believe with current technology even this figure is an underestimate. Funding such an investment is itself a challenge. Its business case depends critically on the emergence of a mass market that may or may not develop.

Despite a considerable degree of scepticism on the feasibility of this proposal, perhaps the real merit of the Teledesic system is that it thinks the previously unthinkable. It provides a grand goal that may indeed become

achievable sometime in the next century. Even the Teledesic engineers are realistic enough to realise that its design may need revision. In 1997 they simplified the system architecture considerably by scaling down the original constellation of 840 satellites to 288 operating in a higher orbit.

Conclusions

We have come a long way from Telstar and the need to use large 30 m ground stations to communicate. With this in mind, who can foresee where we will be in 30 years time? What seems certain is that satellites for communications are set for major growth.

It is clear that the major consortia will continue to develop and grow. Plans are in hand for all these systems, and satellites are being built and launched. During the next couple of years we will see hand-held satellite communications become a reality. In fact these networks have a limited lifetime and so follow-on networks are already being planned.

New technology, markets and the changing telecommunication environment mean that the prospect for newer broadband systems is good, although at present it is unclear what form these will take.

It is expected that the cost of providing satellite capacity will continue to drop in real terms in which case satellites will find themselves used for an increasingly large number of applications, perhaps combining with other satellite provided services such as global positioning to provide further products. They will be more ubiquitous than at present and users will have more services and providers to choose from.

There are many technical, political, regulatory and commercial challenges ahead but it is hoped this article has made clear that a new era in satellite communication is almost upon us.

Finally, it is amusing to note how satellite development appears to have come full circle. The earlier debate over low earth orbiting versus geostationary

systems has been rejoined and it looks as if the future will contain both systems and possibly hybrid systems.

Glossary

ATM	Asynchronous transfer mode
GEO	Geostationary orbit
GPS	Global positioning by satellite
DCME	Digital compression and multiplex equipment
DTH	Direct to home
ECSS	European Communication Satellite
EPIRB	Emergency position indicating reporting beacon
ESA	European Space Agency
GMDSS	Global Maritime Distress and Safety Services
GMPCS	Global mobile personal communications by satellite
HEMT	High-electron-mobility transistor
HEO	Highly-inclined elliptical orbit
ISDN	Integrated services digital network
ITU	International Telecommunications Union
LES	Land earth station
LEO	Low earth orbit
MCS	Mobile Communications Service
MEO	Medium earth orbit
NASA	National Aeronautics and Space Administration
PTT	Post telegraph and telephone
SDH	Synchronous digital hierarchy
SES	Société Européenne des Satellites
SESAT	Siberia Europe Satellite
VSAT	Very small aperture terminal

Bibliography

For an 'Introduction to Satellite Communications' by Russell Silk, see unit 8.2 of the Institution of British Telecommunications Engineers' (IBTE's) *Telecommunications Engineering: A Structured Information Programme*, Issue 17, July 1995.

Further details on the various satellite organisations and manufacturers can be found on the public Internet sites whose addresses are given below.

Current systems

INTELSAT:

<http://www.intelsat.int/>

EUTELSAT:

<http://www.eutelsat.org/home.html>

Inmarsat:

<http://www.inmarsat.org/inmarsat/>

PanAmSat:

<http://www.panamsat.com/>

ASTRA:

<http://www.astra.lu/>

Orion:

<http://www.orion-nsi.net/orionmn.htm>

Emerging mobile satellite systems

Iridium:

<http://www.iridium.com/>

Loral:

<http://www.loral.com/>

ICO:

<http://www.ico.co.uk/>

Globalstar:

<http://www.globalstar.com/>

Broadband Systems

Teledesic:

<http://www.teledesic.com/index.html>

SkyBridge:

<http://www.skybridgesatellite.com/>

Celestri:

<http://www.mot.com/GSS/SSTG/projects/celestri>

SPACEWAY:

<http://www.spaceway.com/>

Astrolink:

<http://www.astrolink.com/>

Satellite industry

Matra Marconi Space:

<http://www.matra-marconi-space.com/>

Hughes (for PanAmSat, Galaxy):

<http://www.hughes.com/>

Lockheed Martin Telecommunications:

<http://www.lmtelecom.com/>

Aerospatale:

<http://www.aerospatale.fr/groupe/indexa.htm>

Arianespace:

<http://www.arianespace.com/>

Acknowledgements

The authors would like to thank colleagues within BT for their support in the production of this article; in particular, Derek Aldous, John Wroe, Mike Seymour and Peter Bolingbroke. Thanks are also due to the companies who have provided the graphical material.

Biographies



Russell Silk
BT Global Commu-
nications

Russell Silk graduated with a degree in Mathematical Physics at the University of Sussex and obtained a Ph.D. in X-Ray Astronomy at the University of Leicester. He joined BT in 1975 where, in the Network Strategy division, he worked on advanced network services for System X. In 1981 he moved into the satellite communications field in the International division of the company. He subsequently worked on a wide range of satellite projects including the definition of EUTELSAT business services, ran the UK's first VSAT trial and played a significant role in the development of the first major satellite TV distribution system within the UK using the INTELSAT system. In 1987 he became Head of INTELSAT System Studies representing UK interests in the INTELSAT Technical Committee. Between 1994-96 he took on the Chairmanship of the Technical Committee. As Satellite Development Manager in Global Communications he is currently looking at new satellite opportunities for BT. Russell is a Chartered Engineer and a member of the IEE.



Martin Bath
BT Aeronautical and
Maritime

Martin Bath works in the Aeronautical and Maritime division of BT Networks and Systems. He gained a B.Sc. in Telecommunications Engineering from the University of Essex in 1982 and an M.Sc. in Satellite Communications Engineering from the University of Surrey in 1987. Following a spell in BT's North Sea Task Force designing satellite communications systems for the offshore oil and gas industry he worked on the implementation of broadcast and interactive VSAT satellite systems. In 1992 he joined BT Aeronautical and Maritime to work on the design of mobile satellite communications. He is a Chartered Engineer and a member of the IEE.

SimDS—Modelling the Performance of Future Systems

SimDS is a new 'end-to-end' service modelling tool which allows service and system designers to test, by computer simulation, new interactive multimedia services. The current version of the program, SimDS 3, supports connection-oriented retrieval services using a classic client/server architecture. SimDS 4 will extend the functionality to support connectionless, multicast and peer-to-peer services such as videoconferencing. The use of SimDS opens up a number of avenues for 'workflow' in service design as the same tool may be used by systems analysts, service designers, systems integrators as well as marketing teams.

Introduction

BT is constantly developing new and innovative interactive multimedia services for a very competitive global market. This rapidly changing environment poses demanding challenges for service and system designers to develop new effective services on time and to budget.

Service performance issues are of key importance to the success of any interactive service, a factor that was clearly identified in BT's interactive TV (iTV) Marketing Trial¹, which was run in Ipswich and Colchester (UK) in early 1996.

SimDS (simulation of distributed services) is an end-to-end service modelling tool which has been specifically created to address these issues. It can predict what response times a user can expect to see from a given overall system and demand profiles, as well as providing a system simulation and capacity planning toolset. The SimDS software program is the result of a co-development between BT Laboratories and K-PAR Interactive Media Systems (K-PIMS) Ltd². There is a continuing development programme to improve the functionality and appropriateness of SimDS so that it meets the needs of BT and, indirectly, its suppliers. The current version, SimDS 3, is described in detail in this article. A future release, SimDS 4, is now under development.

From the outset, SimDS 3 has been designed with a graphical user

interface (GUI) to make it easy to use and relatively intuitive to operate. The tool has the flexibility to be used for 'first-cut' analysis of a service or, provided it is supplied with detailed configuration data, a more in-depth analysis. SimDS also offers the potential to analyse the cost of implementing the service.

This article is divided into two major sections. Firstly, there is a description of what SimDS is and, secondly, a simple explanation of how SimDS works. Additional sections describe possible uses and future enhancements.

What is SimDS?

Overview

SimDS is an end-to-end modelling system. Its purpose is to inform systems designers and planners how things will actually work when all the components of the system are connected together. In the case of an interactive TV service, it could be used to identify potential network 'bottlenecks' or report the likely response time to a user's request or action. SimDS models activity in four major system elements concurrently: users, networks, servers and services.

SimDS is a service-simulation and capacity-planning tool, based on a discrete-event simulation approach. It has been designed to help in the planning and management of networked computer systems which may include multiple clients, networks, and servers.

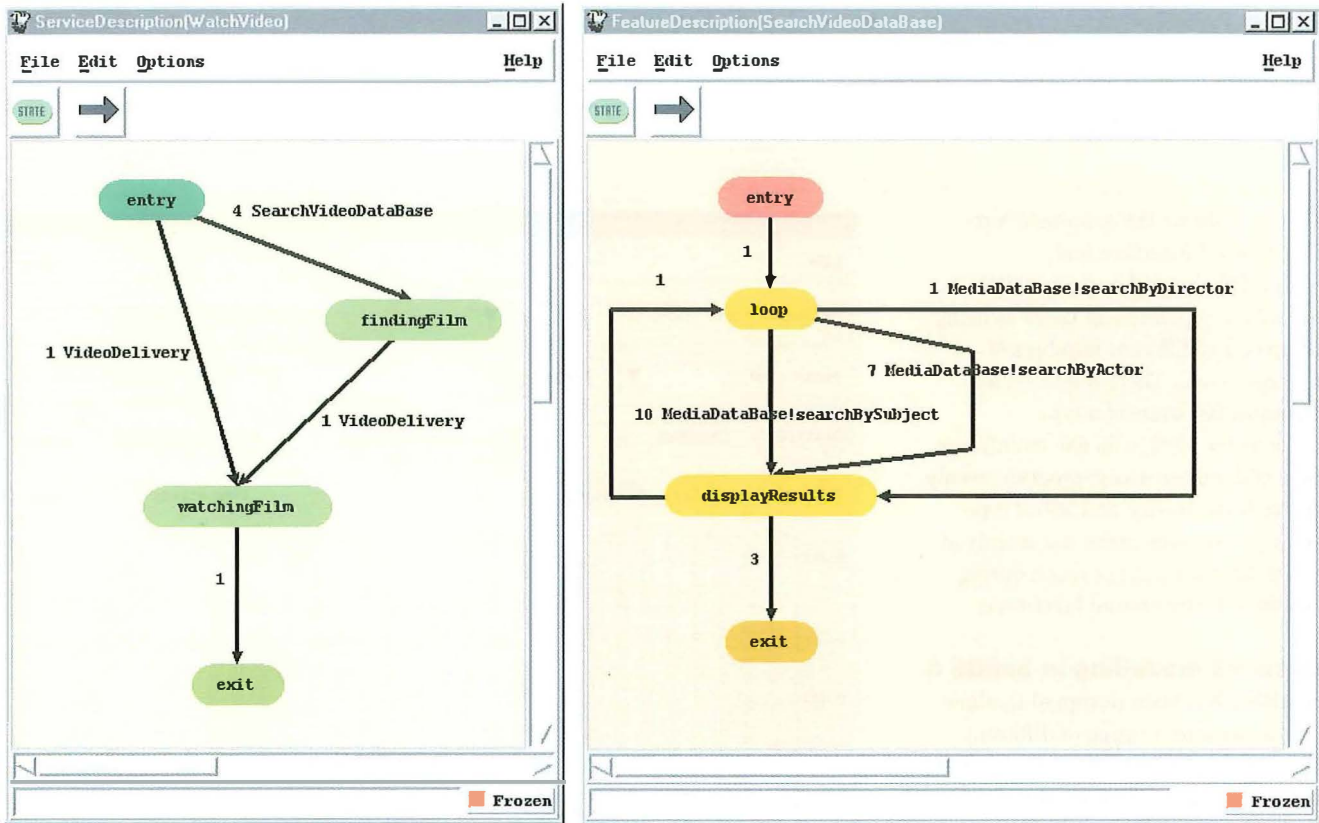


Figure 1 – State machine editors

What does SimDS offer?

SimDS 3 provides the following major features:

- It models in detail the behaviour of client/server type systems and their users. The client/server systems can, for example, be a traditional database access/update system or a large scale iTV system.
- It caters for a range of server architectures. Servers are modelled as a collection of processors, buses and disks that share the workload.
- The server model supports a range of services, such as video and audio stream delivery, World Wide Web (WWW) browsing, home shopping, banking etc.
- It permits characterisation of services in terms of server resource requirements. This is a flexible approach which allows further services to be added easily and quickly.
- It models the requirements of differing types of user. Users are characterised in terms of their

patterns of usage of the available services.

- It can support sophisticated network models which accurately simulate highly-loaded connection-oriented broadband networks.
- It can log statistical and diagnostic output information. This information may be used to analyse the performance of the simulated system.
- It can provide a breakdown of system costs. SimDS provides the capability to assign a cost to each simulated component (for example, server, link, etc.) so that a part or the total cost of the service can be calculated.

SimDS 3 Description

SimDS 3 uses four sub-models called *User*, *Network*, *Server* and *Services* that interwork concurrently to model a complete system. These sub-models are designed to model the behaviour of the users, the underlying network, the network server(s), and the actual service(s) respectively.

The following sections give further detail on each of the sub-models.

User and service modelling in SimDS 3

SimDS 3 allows a fair degree of precision in the way that user behaviour is represented, and it has been designed to make it relatively easy to represent large user populations, if required. At the heart of the SimDS user model is a state machine mechanism, that allows the SimDS operator to describe, in probabilistic terms, what a user is likely to do at a given moment while using a particular service.

The operator might specify, for example, that, while the service *WatchVideo* is being used, the viewer will, with a 80% probability, choose to search the video database and, with a 20% probability, know exactly which video is required and start watching it. A graphical editor makes it easy to draw the appropriate state machine diagram, and to attach to it the appropriate probabilities. Figure 1 shows examples of the state machine editor.

Having described how a user will behave while using a particular service, SimDS lets the operator describe how the user's behaviour will vary through the day, using a graphical bar-chart based interface.

Figure 2 shows the graphical bar-chart based interface tool.

SimDS also makes it possible to describe a population of users as being composed of different numbers of different types. There might be, for example, 200 users of a type *order_entry_clerk*, who use mainly one type of database access program evenly throughout the day, and 500 of type *sales_person*, who make use mainly of a Web browser, but not much during the three hours around lunchtime.

Network modelling in SimDS 3

SimDS 3 has been designed to allow it to represent a range of different types of network infrastructure; from shared media such as Ethernet, through to switched media such as asynchronous transfer mode (ATM) networks. It allows complex systems to be represented, composed of combinations of user populations, server machines, networks and subnetworks, links and switches. A simple-to-use GUI is used to define a top level picture. Figure 3 shows an example of a top-level configuration.

Of course, there are numerous network modelling tools already available which allow detailed analysis of network behaviour. The network modelling functionality within SimDS 3 is not intended to compete with them; its primary design aim is to make it as easy as possible to capture important aspects of network behaviour as an integrated part of an end-to-end systems model, that tells the systems designer what he or she really needs to know about how the system as a whole will work. So, for example, the SimDS network model is designed so that several different levels of detail can be used to describe a network or subnetwork. Furthermore, the network can be described hierarchically as networks and subnetworks, and detail can be switched on or off at each subnetwork level, using an interactive GUI. Figures 4 and 5 show an example of a network hierarchy with different levels of detail shown. The arrows in the diagrams detail the direction of flow across the links.

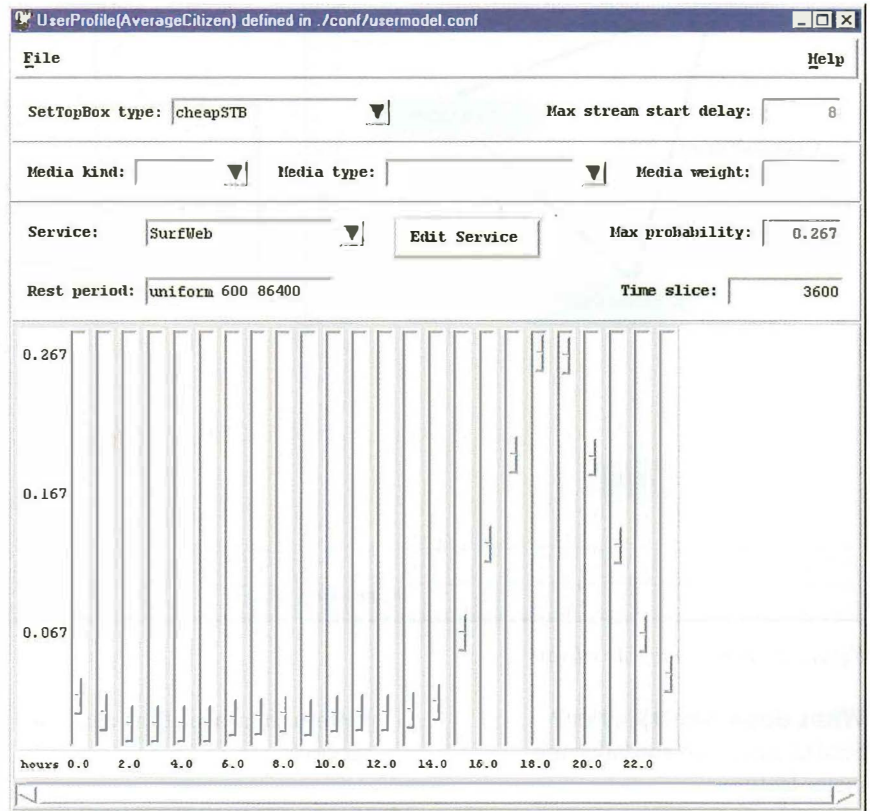
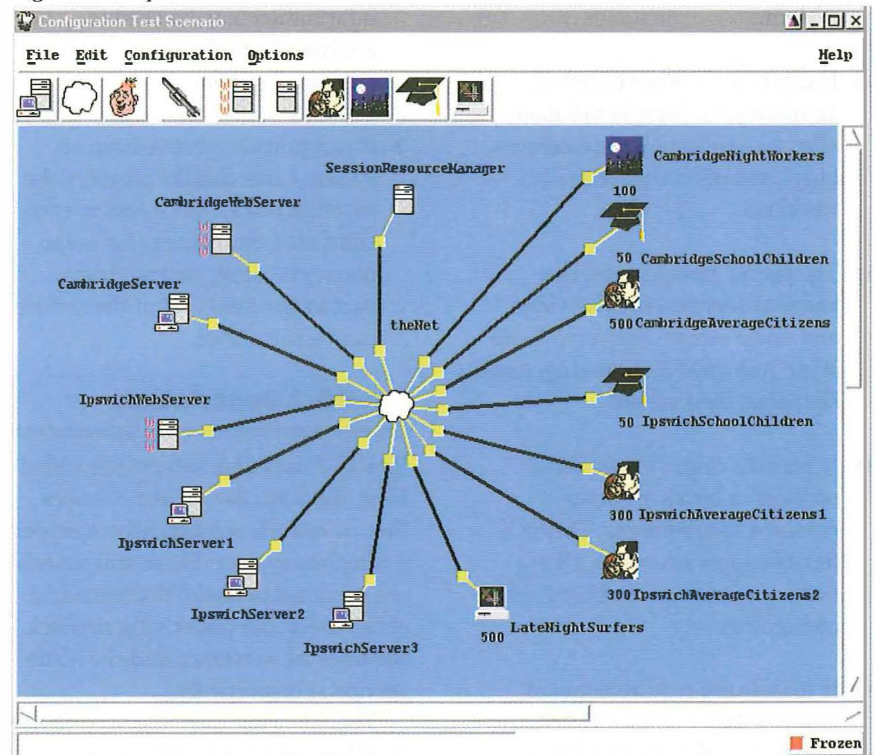


Figure 2 – User profiles

Figure 3 – Top-level view



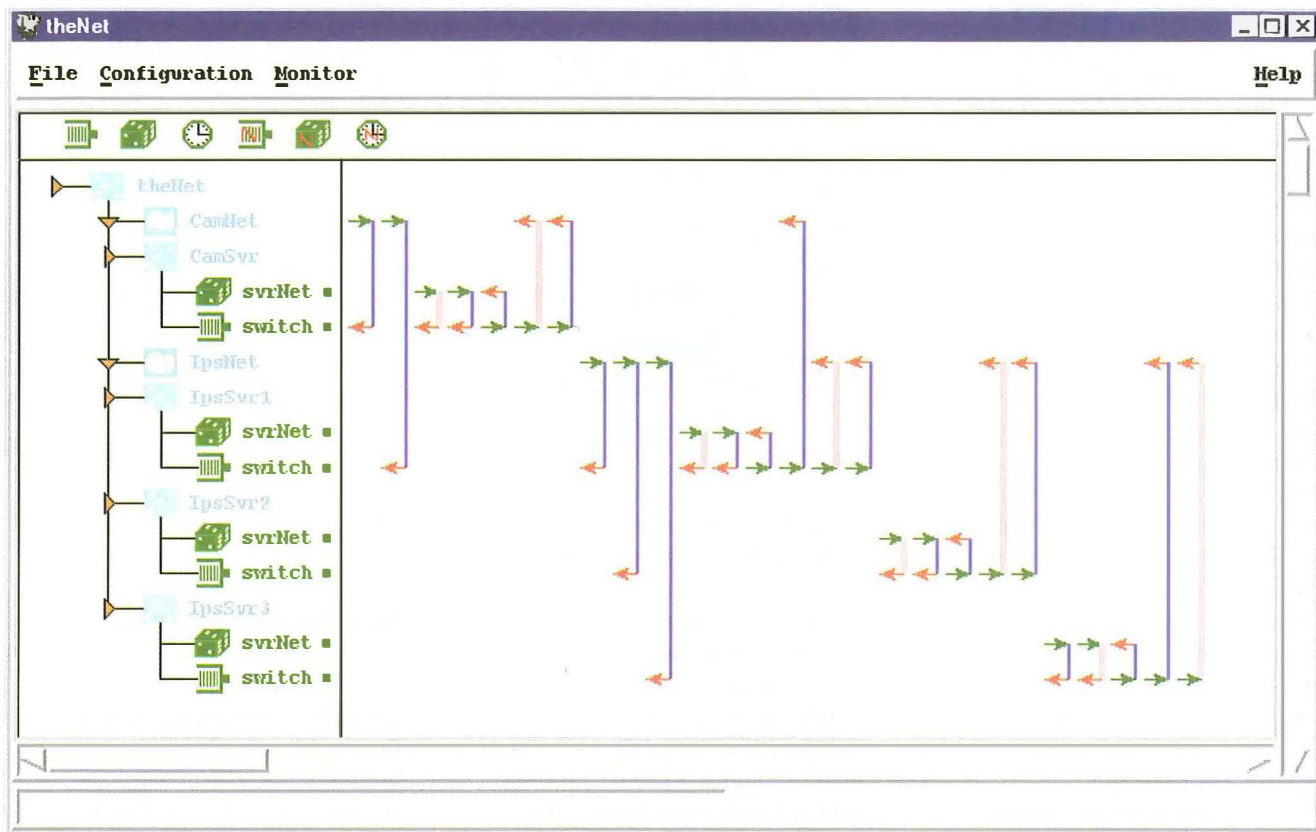


Figure 4—Network view (reduced amount of detail)

Figure 5—Complete network view

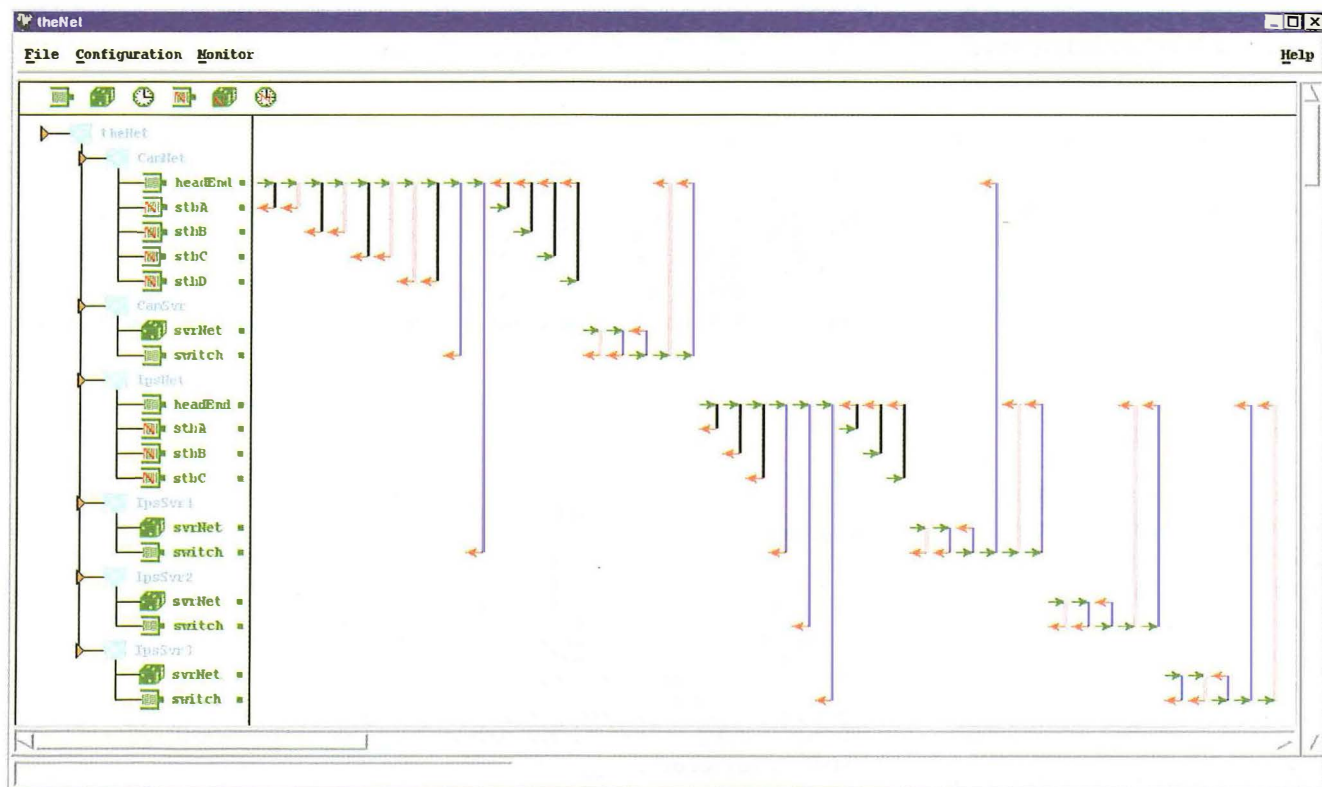


Figure 6—Server view

Server models in SimDS 3

Server machines within SimDS 3 can be fairly complex, and can feature multiple processors, disks, external connections and internal interconnections; either via shared bus technology, or via internal switched networks. Servers can be anything from a single processor PC being used as a hypertext markup language (HTML) server, through to massively parallel video-on-demand media streamers. Figure 6 shows a top-level graphical view of a server configuration. The icons represent the numbers and types of processors and media items held on the server.

SimDS 3 server models can include representations of disk caches, and levels of hardware detail up to, for example, the rotational latency of the disks being used and the response times for disk head movements.

Result reporting

The most important part of the simulation is the set of results that it produces. SimDS 3 is able to report results to the operator from two points of view.

- **User-oriented reporting** SimDS 3 can report on how well services are being delivered, as far as the users are concerned; for example, how long it takes from the mouse-click to the Web page being downloaded and displayed, or how often the flow of a video stream is interrupted.

- **System-oriented reporting** SimDS 3 can display machine loadings, cache hit rates and network loading.

Numerous mechanisms are provided within SimDS 3 through which results can be viewed, both at simulation run-time, and for later analysis. A top-level run-time monitor screen reports on certain key variables; for example, how many users are currently engaged in which kind of activity and how often excessive response times are being experienced.

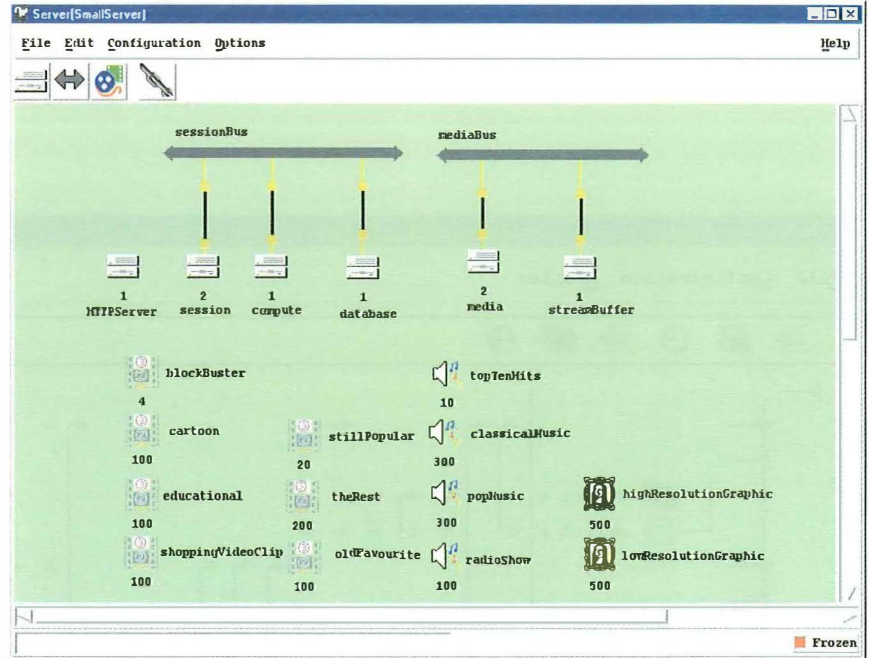


Figure 7 shows an example of this run-time monitor screen.

A number of different graphing mechanisms are also provided to allow easy visualisation of how certain variables are changing over time, and mechanisms to allow both graphical and numeric data to be stored to an external file(s), for later analysis. Figure 8 shows run-time feedback of network load on the various links, as the simulation progresses, using a thermometer type

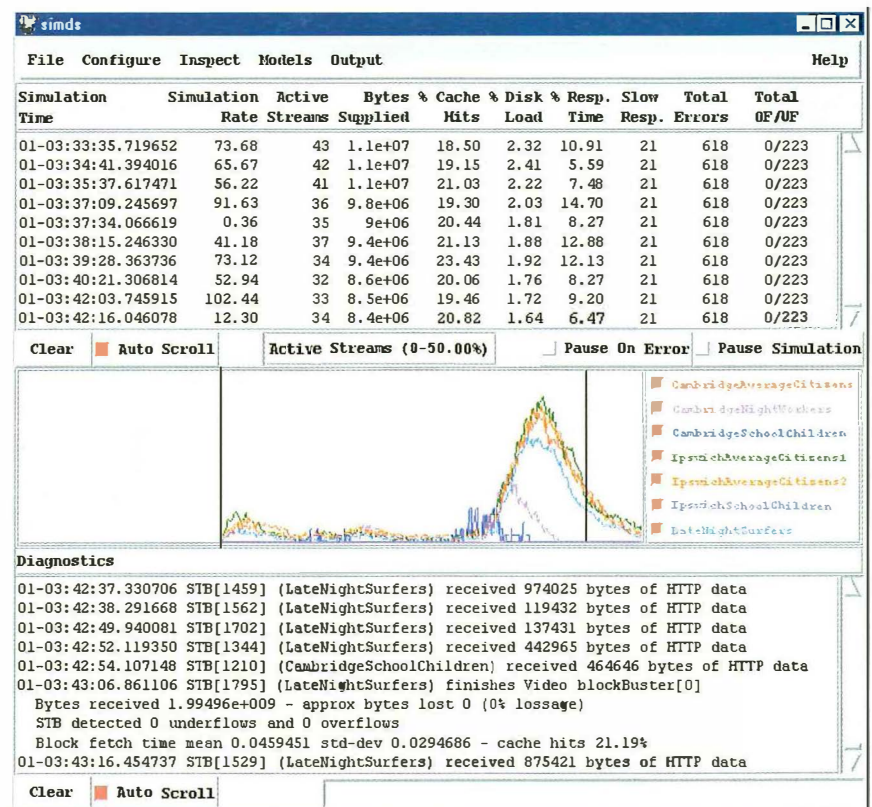
analogy. A designer can then see whether a specific link is reaching 'boiling point' or overload.

How SimDS 3 works

Object-Oriented Discrete Event Simulations

SimDS 3 is implemented as an object-oriented discrete event simulation. The various entities to be simulated are defined in C++ classes

Figure 7—Main simulation view



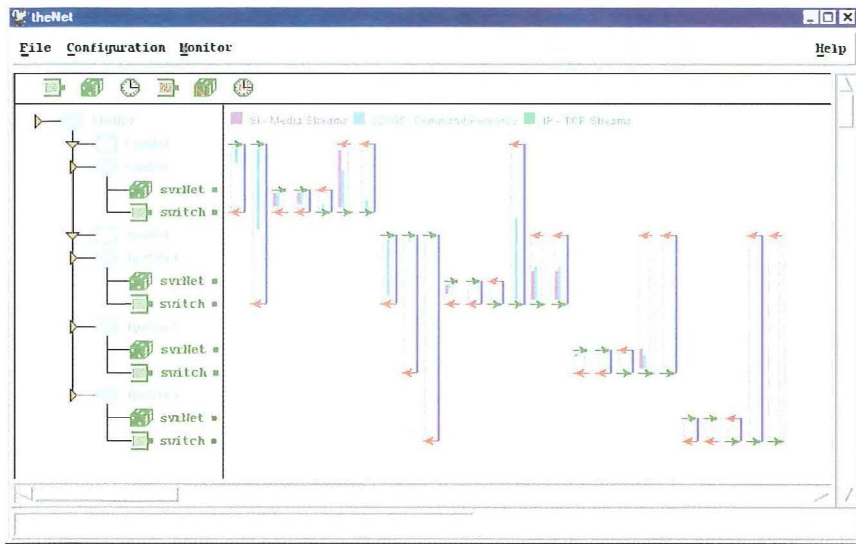


Figure 8—Monitoring the usage of the Network View

Network—scope of the model

The SimDS 3 network model is currently designed to support only connection-oriented broadband networks. These networks can carry both constant-bit-rate (CBR) traffic (for example a video stream), and variable-bit-rate (VBR) traffic (for example a transmission control protocol/Internet protocol (TCP/IP) packet). The VBR traffic can be thought of as individual messages of different duration.

A broadband network carries traffic from many types of users, both business and private. Some operations such as local area network (LAN) to LAN interconnects generate high-bandwidth messaging traffic which would have a dramatic effect on CBR media streams if multiplexed onto the same link.

Whatever the traffic sources, the technique of providing separate paths for messaging (VBR) traffic and CBR traffic is important where the quality of CBR traffic is to be preserved. It is, therefore, possible to configure links so that they are restricted to carrying CBR traffic in the network model.

and subsequently are dynamically instantiated as simulation objects.

A discrete-event simulator provides a mechanism whereby *Events* can be scheduled to occur in a particular order. These events are ordered by increasing 'time of occurrence'. As the simulation progresses the simulated time increments and the simulation runs as quickly as allowed by the amount of computation done at each event.

A SimDS 3 simulation object maintains its own role-specific state information and provides a public functional interface in a similar fashion to objects in a conventional object-oriented program. The difference here is that an operation to be carried out on a simulation object is constrained to occur at a particular point in simulated time. This coupling of an operation (or action) and a time at which it is to execute constitutes an event.

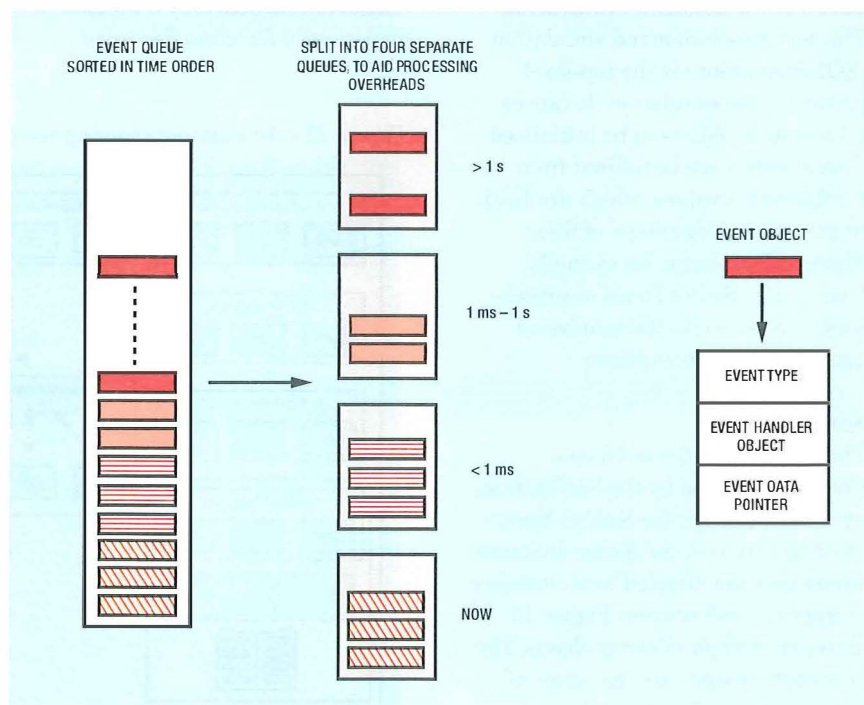
Events and queues

SimDS 3 maintains the sequence of pending events in a queue and invokes the required operation on the appropriate simulation object as the specified point in simulation time is reached. The execution of an event normally causes other events to be generated.

As the simulation runs, events need to be added to, and removed from, the queue. Every event added to the queue causes the queue to be re-ordered, to maintain the correct time sequence. This event queue manipulation takes up a large amount of computing time. In order to reduce this computing overhead,

SimDS 3 has four event queues (see Figure 9). The first queue is for events that will happen instantaneously. The second queue is for events that will happen less than 1 ms in the future. The third queue is for events between 1 ms and 1 s in the future and the fourth queue is for events scheduled greater than 1 s in the future. As events are added, only a single queue needs to be reordered. Events that are cancelled and need to be removed are marked as 'dead' and left on the queue. As the events are not removed then no reordering is required. When a 'dead' event is scheduled for execution, it is simply discarded.

Figure 9—Event queues






Figure 10—An example of event bursts

SimDS 3 uses a burst-rate model³ to simulate the network. Burst traffic is characterised by periods of constant flow rate between abrupt changes in rate. The rate change instances become the discrete events in this model. The model keeps track of channel rates over every link in the network, updating them with each new event (Figure 10).

Delays in the model manifest themselves as delays in the propagation of rate-change events through the nodes (because of queuing delays) and through the links (because of propagation delays).

Simulation objects

As previously mentioned, SimDS 3 contains several different simulation objects which interact with one another during the simulation process. This section only covers the major objects and what they are used to model, as shown in Figure 11.

Services-on-demand simulation

The services-on-demand simulation (SODSimulation) is the top-level driver for the simulation. It causes all the other objects to be initialised. These objects are initialised from configuration values, which are held in an external 'database' of files. These values range, for example, from *contextSwitchTimes* of processors, in a server, to the number of users forming a population.

Server objects

The *Server* operates as follows. Commands issued by the *SetTopBoxes* are routed through the *SubNet* hierarchy and arrive at the *Server* instance where they are directed to an instance of *ApplicationProcessor*. Figure 12 shows an example of server objects. The *sessionProcessors* are instances of the *ApplicationProcessor* class.

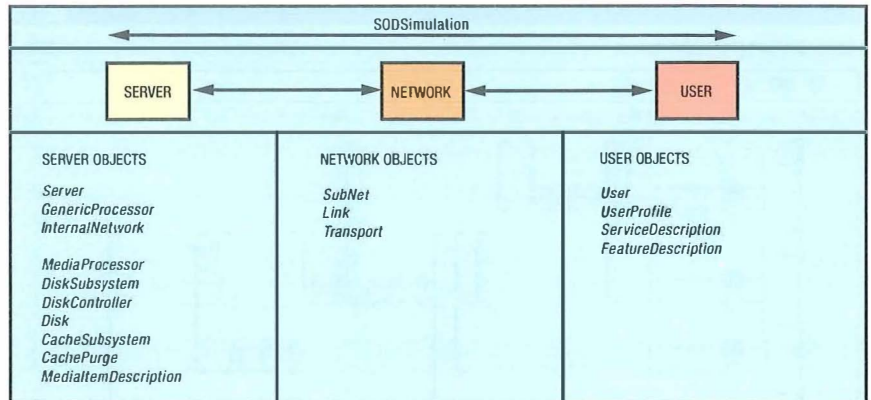


Figure 11—Relationship of major simulation objects

Each command received makes some demand on the *Server's* resources. Depending on the nature of the command (and the associated resources required), the *ApplicationProcessor* either services the requested operation itself, or passes the command to the *Server* subsystem which has indicated that it is capable of servicing that particular operation.

A configurable number of instances of *MediaProcessor* model the media stream subsystem. A *MediaProcessor* is composed of *CacheSubsystem* and *DiskSubsystem* sub-objects that model the performance-critical subsystems. The distribution of the media data across the disk system is defined by a number of *MediaItemDescriptions*. Figure 12 shows *mediaProcessors* containing both cache and disk subsystems.

The *Server's* database requirements are modelled by a number of instances of *DatabaseProcessor*.

These handle the session tracking, billing requirements and media database searches.

Network Objects

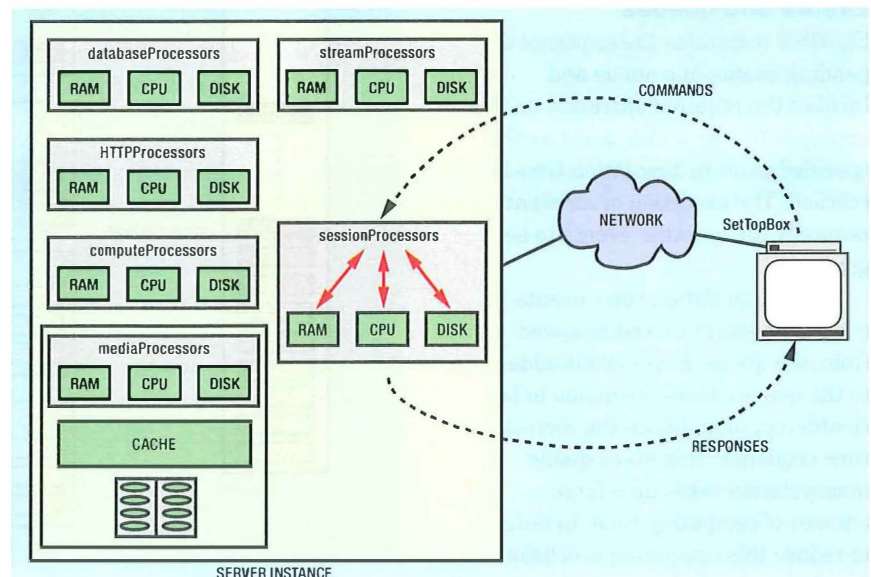
A SimDS 3 network is made up from *SubNets*, *Links* and *Transports*. Figure 13 shows the relationship of these objects to a network diagram.

SubNets

The network is modelled as a hierarchy of *SubNets* (see Figure 4). Each *SubNet* represents a network node (switch element, router, LAN segment). The *SubNet* names represent the network hierarchy (just as file pathnames reflect the directory structure). Leaves in the hierarchy represent connection points. Links are joined to these connection points.

The granularity of the simulation model can be set for each *SubNet* to either Coarse or Detailed:

Figure 12—An example showing server objects



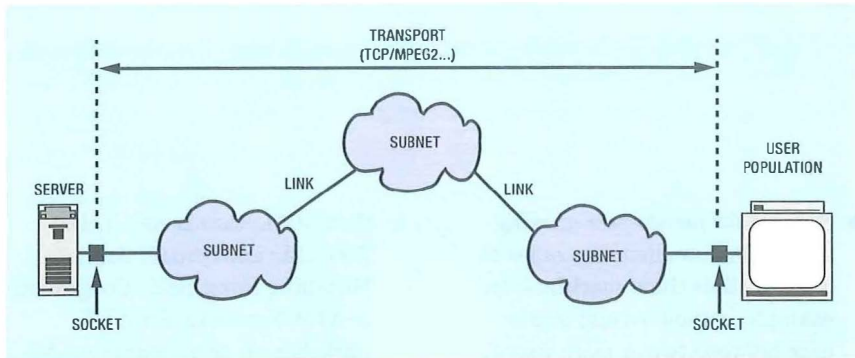


Figure 13—An example showing network objects

- *Coarse*—The *SubNet* handles all traffic, all its lower level *SubNets* are inactive.
- *Detailed*—The *SubNet* is inactive, its lower-level *SubNets* handle all traffic.

SimDS 3 supports three traffic modes: burst, synchronous and stochastic. These modes characterise the traffic delays at each *SubNet*:

- *Burst*—periods of constant flow rate punctuated by abrupt changes in the rate.
- *Synchronous*—burst with extra delay determined by the amount of *Link* bandwidth sharing.
- *Stochastic*—burst with additional random delays determined by traffic load. Lower level *SubNets* inherit the configurable delay parameters from the higher level *SubNet*.

Links

Links model the physical connections between *SubNets*. They are uni-directional connections. One of the *Links* parameters is a buffer size. This buffer is situated at the *SubNet* output, so feeding into the *Link*.

Transports

A *Transport* represents the protocol in use at one end of a point-to-point connection. SimDS 3 supports coarse and detailed TCP and coarse model MPEG-2 protocols. Coarse model TCP and MPEG-2 are modelled as passive protocols which are defined only by their data overhead. Detailed model TCP is an active protocol which models the behaviour of TCP's congestion control algorithms.

User objects

Each instance of *User* class models the behaviour of a single user of the system (as defined by an instance of *UserProfile*) and is associated with an instance of *SetTopBox* which provides the means to access a *Server*.

The information specified in the *User* object's associated *UserProfile*, *ServiceDescription* and *FeatureDescription* objects determines the manner in which the *Server*'s services are utilised. Using this information, each *User* object issues service requests through its *SetTopBox* and measures the time taken for the system to respond. The response time is logged for later analysis and compared against the appropriate maximum acceptable value. A diagnostic message is produced if the response exceeds the maximum acceptable value.

SimDS Development

The SimDS tool development is a joint undertaking by British Telecommunications plc and K-PIMS Ltd. K-PIMS Ltd. is one of the new breed of 'virtual' software houses that has a base in Bristol, England, but operates using experts from around the country who usually telework from their homes. K-PIMS was chosen to help develop the tool from a short-list drawn up by project leaders in Applied Research and Technology's (ART) Future Systems Group working within the Interactive Multimedia Services Futures Campaign.

Validation

Computer simulations rely on detail and accuracy of input information in order to obtain a meaningful and useful result. The old saying 'rubbish in, rubbish out' is very appropriate when applied to computer simulation

systems which have been supplied with limited information. Furthermore, a simulation tool has very little real worth if it has not been tested to represent real-world events. SimDS 3 was proved to be effective by validating it against real-life data. The validation of SimDS 3 was achieved by loading it with the configuration of BT's iTV trial and performing a number of simulation runs for periods of a 'week in the life of the trial'. The results were then compared to the actual 'real-life' data from the trial. The simulated results compared favourably with the actual results of the iTV trial and this success was regarded as validation of the modelling tool.

Uses of SimDS

The flexibility of SimDS to model end-to-end services provides service and system designers with an invaluable tool, not just for testing the performance of new service scenarios, but also for many other functions. SimDS allows designers to try out new service concepts and architectures, implement 'what if' analysis, investigate network and system dimensioning for an evolving service, and project when service capacity has to be increased to support an increased customer base.

SimDS also offers the potential for new ways of working. A 'workflow' process can be envisaged whereby a sales team may use SimDS for requirements capture and then pass on the SimDS configuration to their design team for in-depth design configuration. The completed configuration may then be passed on to the analysis team to test the design's performance and scalability.

SimDS may also be used by a consortium of companies to test their candidate design model prior to responding to a service provider's invitation-to-tender. The benefit to the consortium is the added confidence factor that SimDS demonstrates that their candidate model will work to the required performance criteria.

What Next?

Following testing, use and exposure of SimDS 3 with various teams across BT Laboratories, certain small enhancements have been identified which will make the toolset more widely applicable. An enhanced version of SimDS, to be called SimDS 4, has been specified, and development work is underway at the time of writing.

SimDS 4 features

The new features of SimDS 4 will allow a greater range of system scenarios to be modelled. For example, SimDS 4 will be able to model network computing scenarios, multicast services, connectionless networks such as traditional IP networks, audiovisual conferencing scenarios between client's machines and combinations of the above.

It is planned that SimDS 4 will contain all the features of SimDS 3, plus the following enhancements:

- *Client machine modelling*
The basic 'dumb' client provided in SimDS 3 will be further enhanced in SimDS 4 to support intelligent clients; for example, personal computers.
- *Generic protocol modelling*
A generic approach to protocol modelling to allow more than just TCP and MPEG-2 protocols to be modelled. This will be done by offering network service options; for example, the UDP/IP protocol consists of connectionless routing and unreliable delivery service options.
- *Enhanced Service Support*
Support for the following network service options will be included:
 - connection-orientated routing,
 - connectionless routing,
 - reliable delivery,
 - unreliable delivery,
 - constant bit rate,
 - available bit rate,
 - unrestricted bit rate, and
 - multicast group addressing.

- *Support for peer-to-peer working*
This will allow client machines to 'talk' to other client machines; for example, an audiovisual conference between two or more 'users'.
- *Support for group working*
Support for data conferencing scenarios such as those identified in T.120 standards and audiovisual services like those in the T.130 standards.
- *Enhanced data-logging facilities*
Improved and enhanced data logging of error conditions, so that detailed off-line analysis may be undertaken.

Conclusion

It is vital to be able to predict the performance of large-scale interactive multimedia systems well before architectures are fixed and the equipment installed. SimDS, an end-to-end service modelling toolset, meets this need. This software program provides service and system designers with an invaluable tool for predicting the system performance of new interactive multimedia services identifying future 'bottlenecks' in current systems as demand increases. The flexibility of SimDS's modular design means that it can be further enhanced to meet future simulation needs. SimDS has the potential to become a major simulation tool for the new multimedia service industry.

Acknowledgments

The authors wish to acknowledge the skill and dedication of their development partners, K-PIMS Ltd.

References

- 1 KERR, G. W. A Review of BT's Interactive TV Trials. IEE Colloquium on Interactive Television, London, 2 October 1995, pp. 6/1-5.
- 2 K-PIMS Ltd. Web site. <http://www.k-par.co.uk/simds>

- 3 PITTS, J. M.; CUTHERBERT, L. G.; BOCCI, M.; and SCHARF, E. M. Modelling Burst Scale Congestion in ATM Networks. First UK workshop on performance modelling and evaluation of ATM networks, June 1993, UK.

Biographies



Paul Sellek
BT Networks and Systems

Paul Sellek joined BT Exeter Area in 1978 working in the Drawing

Office. In 1987, he graduated from Leicester Polytechnic, achieving a B.Sc. in Computer Science. He moved to BT Laboratories and started working in the area of data communications, networks and network protocols. He joined the Future Systems group in 1996 to work on the design and development of the SimDS Tool.



Dave Beaumont
BT Networks and Systems

Dave Beaumont joined BT Laboratories (formerly the Post Office Research Dept.) in

1968, as a Trainee Technician Apprentice and has worked in the Microwave, Cable, Semiconductor and Visual Communications divisions. In the last of these he specialised in video compression for which he holds joint patents. He has actively participated in ITU-T, COST, ATM Forum, RACE and ETSI on standardisation of video compression and its transmission over ATM networks, leading international teams in RACE and ETSI. He became Head of Future Systems in the Applied Research and Technology department in 1995. He currently leads the development of SimDS as part of his research effort into future systems for interactive multimedia services.

Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his regular column in the Journal. In this issue he considers a world without telephones.

No Telephones

Now and again someone takes my breath away by a monumental statement lacking any understanding of our technology-dominated world. Most recently it was a pundit on the perils and dangers of the telephone, and how we could live without it. Let us consider such a world for a moment.

In many respects the invention and patenting of the telephone by Alexander Graham Bell in 1876 was a technology spin-off from his work on harmonic telegraph systems and speech therapy with the deaf. These studies had created a deep-seated desire to understand the mechanisms of sound, speech and hearing. So it was his refinement and bringing together of the carbon microphone and fixed-magnet earphone for therapeutic applications that enabled him to create the first telephone. At that time no one could have guessed that the demands of the telephone would trigger such an incredible series of follow-on developments of even more remarkable technology. These now underpin every aspect of human existence, with the telephone network an absolutely vital component of our civilisation.

To contemplate a world without the telephone we also have to contemplate a world without microphones, earphones and loudspeakers. Why? Because, as soon as these modules are realised the telephone is not only possible but inevitable, and vice versa. So, imagine a world without radio and TV, only silent movies, and no computers. Why no computers? Well, without the telephone Strowger would not have invented the automatic switching machine, and the technology that enabled the building of Colossus (the first significant United Kingdom

computer of World War II created for code breaking) and other early computers. But in a more seminal sense, without the development of the intricate electromechanics for telephone switches, and the simplest of electronic components, we would not have the ability to automate many of our production processes. And automation is a key factor in the creation of effective and efficient mass production, and most importantly, a viable electronics industry. Unless there is an ability to repeat intricate processes continually and accurately, electronics on any scale is fundamentally impossible.

Without the demands from telephone and radio broadcasting networks, it is highly unlikely we would have developed thermionic valves to the height of perfection they reached. Their use would have been focused on telegraph radio stations and copper line regenerators. Nor would we have seen the transistor, integrated circuit, and an effective computer industry. No coaxial cables, optical fibre, lasers, microwave radio, satellites or space race. We would be living in a world of early Victoriana, with masters and servants, telegraph systems and copper cables, messenger boys and telegrams. It would be a world much slower than today, and unable to sustain a population now fast heading towards six billion. Most likely the world would have stabilised at around four billion people and become a stilted place. Extinction would be staring us in the face with disease and disaster a much greater threat than today.

The music and entertainment industry would also be puny. There would be no amplifiers, radio, hi-fi, tapes or CDs, just plastic cylinders and discs scratching out their music with diamond tipped pins feeding acoustic exponential horns mounted in mahogany cabinets. No 'Pop' industry, with just the rich and privileged getting to see and hear opera, orchestras and bands. No television, video recorders, camcorders, computer games or pocket calculators, personal digital assistants, lap-top computers and



personal computers. Offices would be dominated by paper, typewriters, typing pools, and all business would be subject to the delays of the postal network.

It gets much worse. International travel would be for only the very few on aircraft that are a far cry from today's jumbo jets. Trains, ships and aircraft are organised over telegraph systems operating over copper wires and primitive wireless links. Travel schedules strictly are rationed by unbelievably limited and expensive data transport. It is likely there would be a lack of trained telegraph operators, bottlenecks of information, ignorant populations starved of information and experience, dictatorships and empires, wars and famine.

We would be living on a polluted planet consuming ever more hydrocarbons to support incredibly inefficient manufacturing and logistic systems, with paper consumption destroying huge tracts of forest. Our science, technology, and basic understanding limited by the crude technologies produced without electronics. No radio telescopes, no scanning-electron microscopes, no medical scanners, no accurate time keeping, no space race and no significant advance in materials and production techniques.

If you want to get a partial insight into this hypothetical scene, just visit the third world today. With no telephones you cannot travel easily. You cannot book a ticket for train, boat or plane without travelling to the ticket office. When your plane takes off the destination has no idea of your status until you are within 200 miles

or so. The supply and delivery of food, clothing, heat light, power, water and all the things we so take for granted is then a big deal. And quite naturally life becomes focused on survival alone. There is just no time to think and do anything else.

The only people who complain about modern life and the conven-

ience and security afforded by technology are those who have it, but do not understand it. Life in the third world is not Utopian—it is tough. Alexander Graham Bell's legacy to mankind is enormous, much bigger than most of us can imagine. Thanks to him this article was typed on a laptop computer and


sent over a GSM telephone link from a train travelling at 90 mph between London and Ipswich. Had I not had this facility, not only could I have not done this, I could not do my job either, which would not exist anyway.

No telephone, no telecommunications companies and no industry!

book reviews

The Ethical Global Information Society—Culture and Democracy Revisited

Edited by Jacques Berleur and Diana Whitehouse



Do not read this book for a definitive definition of the Global Information Society. However, if you want some challenging ideas and views on how new information technologies could effect the society we live in, then it is certainly well worth a read.

I should point out that the book contains papers presented at a conference in May 1997, titled 'Working Conference on Culture and Democracy Revisited in the Global Information Society'. The event was sponsored by the International Federation for Information Processing (IFIP), a non-governmental, non-profit umbrella organisation for national societies working in the field of information processing.

The papers in the book are grouped into themes: ethics, the global information society, cultural challenges, democratic challenges, and

the ancient Greek agora. The book concludes with an overview of the main conclusions from the conference. In this review I can only give a flavour of a few papers that caught my attention.

The paper by Julie Cameron and Karin Geiselhart 'A Charter for Citizens of the Global Information Society' gets to the heart of who will have control in this new society and what are the responsibilities and obligations of government and industry. The power of integrated information and communication technologies to collect, store, combine, manipulate and disseminate information to and from any part of the world now affects every individual. There is an implicit assumption that the ordinary citizen will need protection from excessive exploitation, particularly from the new converged media industries. In their paper they lay out a set of principles, which aim to reduce the vulnerability of individuals and societies and to promote new sustainable societies.

Each principle is only a beginning, but each has been designed to address a key issue arising from the impact of the new technologies upon our societies.

My second choice of paper is Colin Beardon's, 'What does it mean to be virtual?' The author contends that we have to examine very closely the language we use when we talk about, or describe, new technologies. When we think about the future of a new technology, we usually do so in words, through presentations, and written reports etc. But how many times have we left our audience confused? We need to realise that the words we use are a form of tool, which used incorrectly can lead to misleading conclu-

sions, so easy in companies with people of mixed technology literacy.

Take the term 'virtual' as an example; it has had a long history, being used in computing from the early 1960s as a premodifier to computing terms—virtual addressing, virtual machine, etc. When the term 'virtual reality' was coined it achieved spectacular notice. It originally meant 'morally virtuous' and before being associated with terms such as 'unreal' or 'potential'. A modern dictionary definition of virtual describes it as: almost or nearly the thing described, but not completely. However, technologists developing virtual reality systems use a definition hinging around the word 'stimulation'. But postmodern-ists have introduced another meaning to the term 'virtual reality' based upon their concept of 'virtuality'. For them it expresses a 'potential'; for example, image data stored on a computer disk is virtual, but once displayed upon the screen it is actual.

What we are seeing is language being used to enable culture and technology to come to terms with each other. From this interaction new terms and words appear, but the author warns that without care about the language we use our thoughts about the future are frequently trapped.

My final choice of paper is by Ellen Christiansen, titled 'Gardening: A metaphor for sustainability in information technology-technical support' based upon a study of how technical support people work.

Bonnie Nardi developed the gardener metaphor to illustrate the potential that may be found by facilitating end-user computing by employing 'super users' as local

book reviews

gardeners. Such people play three different roles simultaneously: as experts when working with technology; as colleagues, engaging in collaborative learning when working with users; and as parents, taking pride in seeing the users walk away with the results, even though their contribution is often invisible and may not be acknowledged.

This gardening metaphor can be extended to four distinct areas of technical support work: *cultivation* of the technical equipment, programs and infrastructure; *nursing the users*, the peer learning through which technical support of the production is accomplished; *harvesting*, presenting high-quality products of which the technical support is an important, yet transparent feature; *composting*, making sure that by-products of the production process, and parts of all kinds, are kept for potential reuse.

What attracts me to the author's gardening metaphor is the close relationship between a sustainable support environment, illustrated by the metaphor and the important issue of corporate learning. In her study, members of technical support teams worked alongside the users, providing a human-to-human interaction based on mutual understanding. The advantage is that the user gains in technical competence and understanding of the systems which results in a high level of sustainability in user work activities.

The book highlights the conclusions of the conference by indicating that the information and communication technology society is potentially extremely inclusive, but in practice it is becoming very exclusive, elitist, and sectarian. To empower people there should be universal service and equity of access that provides

meaningful, quality information via diverse channels. In the long term, a real democracy presupposes an information society where citizens are able to participate in the decision-making process.


*Published by Chapman & Hall
ISBN 0-412-8296006*

Price: £80.00 xx + pp. 379.

Reviewed by David Greenop

Distributed Applications and Interoperable Systems

Edited by Hartmut König, Kurt Geihs and Thomas Preuß



The challenges of designing and managing complex information systems effectively across heterogeneous networks are substantial, in intellectual and practical terms, but overcoming them is essential to the future success of any company seeking to play a major part in the developing global information economy.

This book is a report of the proceedings of an October 1997 conference covering a broad range of topics. It contains papers and presentation summaries, many with multiple authors. The tone is academic and international (that is to

say European and Far Eastern) with no examples of contributions from the key United States IT commercial market leaders and no strong overall editorial linkage. A casual reader, with limited time, seeking a lucid and coherent guide to the practicalities of distributed information systems is likely to be disappointed.

However, there is much of interest for the committed systems professional with previous background in the topics covered, who might benefit from exposure to a more abstract analysis of the problem space. There are sections on management of distributed applications, agents and Internet applications, CORBA, language support, and security and reliability. Occasional beams of light are shed on various current 'religious wars', such as Java scalability. Perhaps surprising is the absence of any major recognition of the importance of asynchronous messaging for practical implementation of distributed architectures, though some specific references are made, for example, to the possibilities of mobile agents in overcoming some of the issues of asynchronous processing. The overall impression left by this book is that any manager entrusted with charting a course through this minefield with a view to delivering shareholder value will wait some time for a clear consensus to emerge about the best technical solutions, while IT professionals who can go some way towards mastering the complexities of large-scale distributed systems and present credible proposals will be much in demand.

Published by Chapman & Hall

ISBN 0-412-82340-3

£75.00. vi + pp. 294.

Reviewed by Graeme Mathieson

telecom focus

Agreement Reached on 56k Modem Standard

The International Telecommunication Union has agreed on the technical specifications for pulse code modulation (PCM) modems (also known as 56k modems) and

has initiated the formal approval process.

The ITU, a United Nations agency, coordinates global communications standards. Study Group 16 of the ITU Telecommunication Standardization Sector (ITU-T) where the work on the modem standard was carried

out, drives the development of standards for multimedia systems.

The new Recommendation, designated V.90, is expected to be widely used for applications such as Internet and on-line service access. V.90 modems are designed for connections which are digital at one

end and have only one digital-to-analogue conversion. Download speeds of up to 56 000 bit/s are possible, depending on telephone line conditions, with upload speeds of up to 33 600 bit/s. Manufacturers currently producing modems based on proprietary schemes have already stated they will rapidly migrate to the new standard.

According to industry analysts, the ITU agreement will boost modem sales significantly. VisionQuest 2000, a market researcher, estimates the number of modems shipped each year is likely to rise to 75 million by the year 2000 from 50 million in 1997.

Work began on the development of V.90 (previously referred to as *V.pcm*) in the ITU-T in March 1997. 'This is the shortest period of time ever taken for an ITU-T modem Recommendation to achieve "determination" approval status, and demonstrates a commitment by the ITU-T to respond quickly to urgent market needs', said Mr. P.-A. Probst, Chairman of Study Group 16.

The V.90 modem harmonizes the two competing proposals submitted last year. Customers who have purchased 56k modems based on either technology may be able to get software to make their devices compatible with the new ones developed on the V.90 standard.

Sir Iain Vallance to Become Part-time Chairman at BT

BT has announced that Sir Iain Vallance is to change to a part-time role as Chairman of BT from 31 July. This follows the appointment of Sir Peter Bonfield as Chief Executive of BT in January 1996, a post previously held by Sir Iain.

Sir Iain said: 'It has been a great privilege to serve for over 10 years as executive Chairman of BT, while the telecommunications industry has gone through a major transformation.

'I believe that with Sir Peter now serving as Chief Executive for over two years, and the rest of the top team in place, it is appropriate for me to move on to become part-time Chairman.

'I am pleased that my new role will enable me to continue to play a part in ensuring that BT achieves its ambition to become the world's leading communications company.

'While my role will be changing significantly and consequently the number of days I will spend on BT business will decrease, my commitment to the company's success will be as strong as ever.'

Sir Peter Bonfield said: 'I cannot over-estimate the impact which Iain Vallance has had on BT and on the telecommunications industry as a whole.

'The company's reputation both at home and abroad as a symbol of quality and excellence stems directly from his vision. He has achieved in 10 years as Chairman what others could not accomplish in a lifetime.

'However, I should emphasise that while Iain is retiring from his full-time role, I will continue to benefit from his advice and judgement in the months and years ahead.'

New Europe-United States Cable System

An international consortium consisting of AT&T, MCI, Telecom Italia, Telefónica de España, Trescom, and over 30 other telecommunications carriers have signed a construction and maintenance agreement to build a new fiber-optic undersea cable system linking southern Europe and the United States.

The Columbus III Cable System will span 10 000 km and cost an estimated \$236 million. Tyco Submarine Systems International Ltd., Alcatel Submarine Network Systems, and Maristel SpA will build the system, which is expected to be completed and in service by September 1999.

The synchronous digital hierarchy system will consist of two optical fibre pairs, and will include within-system restoration. Each fibre pair will initially operate at 2.5 Gbit/s per second, per wavelength, with full path duplication in each direction, for a total of 10 gigabits per second, allowing the transmission of approximately 120 000 simultaneous calls.

The system eventually can support up to eight wavelengths per fibre pair, or 40 Gbit/s per second, the equivalent of approximately 500 000 simultaneous calls.

The system's full capacity will be more than 20 times that of the transatlantic portion of Columbus II. That 12 000 km system went into operation in December 1994, and handles up to 20 000 simultaneous transatlantic voice calls.

'Columbus III is designed to satisfy the exploding demand for faster and improved transatlantic communications, especially in the areas of Internet and high-bandwidth use, between the United States and southern Europe into the first decade of the 21st Century' said Thomas McInerney, managing director for AT&T's international cable planning group.

The cable will be configured as a ring connecting Hollywood, Florida in the United States to Lisbon, Portugal; Conil, Spain; and Mazara del Vallo, Italy. AT&T, Telefónica, Telecom Italia and Marconi will own and operate cable stations in the United States, Spain, Italy and Portugal, respectively.

The memorandum of understanding to build Columbus III was signed in April 1997 by AT&T, MCI, Trescom, Telecom Italia, Telefónica de España, and Companhia Portuguesa Radio Marconi. Construction of the system is expected to begin shortly.

Further Steps Towards The Information Superhighway

BT has indicated that it will be the first European communications company to join the Universal Asynchronous Digital Subscriber Loop (ADSL) Working Group (UAWG) driving world standards for DSL Lite. This group comprises Microsoft, Compaq, Intel and a number of United States communications companies and equipment suppliers.

DSL Lite, a lower cost development of ADSL technology, will work up to 30 times faster than most existing modems but will operate on an

existing copper telephone network. It will be compatible with the standardised ADSL line cards which BT is planning to install in exchanges.

Evaluation of DSL Lite equipment is already well advanced, and BT will begin tests in the UK as soon as production prototypes are available.

Chris Earnshaw, Managing Director of BT's Network and Systems, said: 'Our early evaluations are extremely promising, and we are determined to be the leading European company for exploiting low-cost ADSL technologies to bring a world of high-speed information services to our customers.'

BT announced last year that it will be mounting a major trial of ADSL services in West London, working closely with suppliers Fujitsu and Alcatel. The DSL Lite development has the advantage that the equipment at the customer's home can be built into a PC, just like today's modems, and plugged into an ordinary telephone socket.

BT and Independents in Service Trial

BT has announced that it is to run a technical trial of a new service which would enable independent companies to use BT's network to serve their own customers.

The trial will involve two companies—Long Distance International (LDI), a provider of indirect telephone calls, and UniqueAir, a mobile telecommunications company. It will cover the 0171 and 0181 areas of London.

Ian Morfett, BT's Director, Strategy and Business Services, said: 'BT hopes that this trial will pave the way for a full national service, where BT's telephone lines and network equipment can be used to improve the range of choice for UK customers.'

'We believe that a successful national service will help to expand the UK telecommunications market by developing imaginative new products for new and existing customers. BT would be delighted to provide those services to an emerging breed of retailers in the UK telecommunications industry.'

It is expected that a total of about 1000 customers will be involved in the trial.

BT will supply and maintain customer telephone lines in the trial area for LDI and UniqueAir, carrying out installation and maintenance, and providing records of the calls made on those lines for charging purposes. BT will charge the companies for services provided and for calls.

UniqueAir and LDI will use the lines to sell telephone calls and services to their customers, under their own branding of the service. They will set their own prices for their service, together with their own billing and customer services.

Nick Rowley, Group Managing Director of UniqueAir, said: 'Gaining direct access to all residential and small business customers through the trial will allow UniqueAir to offer a total fixed and mobile telecommunications package using a combination of networks. The 'single bill' package will allow operational savings to be passed on in our prices, enabling UniqueAir to become not just a simple reseller of cellular airtime but a complete communications service provider.'

Bill Noseworthy, Managing Director of LDI in the UK, said: 'We are delighted to have been selected for the trial. The initiative marks a major breakthrough in the UK telecommunications market. We now have the opportunity to gain proper ownership of the customer with our brands and deliver the services they want.'

'Customers will no longer have the hassle of two bills—one from BT and one from us for long distance—they will now be bundled into a single LDI bill. It is a significant development in the continued liberalisation of the UK telecommunications market.'

Fixed and Mobile Licences in Singapore go to BT Consortium

A consortium comprising BT, NTT of Japan, Singapore Technologies, and Singapore Power, has been awarded a fixed and mobile telecommunication

licence by the Singapore government. Under the terms of the licence the company will begin services on 1 April 2000.

Sir Peter Bonfield, BT's Chief Executive, said, 'BT is delighted with the news. It endorses what we have believed all along—that BT and its partners have put together the best mix of skills, experience and technology to meet Singapore's needs.'

Internet Micropayment Service

BT is to trial a revolutionary new payment service, *BT Array*, aimed at internet users who wish to make small value purchases.

The trial, scheduled to last six months, is available and can be contacted via e-mail at btarray@prospect.labs.bt.com

BT Array is aimed at internet users who wish to buy goods, services, information or software that are priced between a few pence and a few pounds. The service will provide an easy, secure and economical way of making these 'micropayments'.

Companies selling their services or content during the BT Array trial include Which? Online, Hemmington Scott Publishing, Internet Magazine from EMAP and Amco Technical Services.

Unlike wallet-based micropayment systems, BT Array works in an easy and intuitive way for buyers to pay for their purchases over the Internet.

Internet users first register and forward their credit card details and their e-mail address on-line—a once-only process over secure facilities.

Users choose their own account name and password when they register, and are then able to purchase items from any merchant displaying the BT Array micropayment logo.

BT Array will consolidate the account and charge it to the customers' VISA or Mastercard periodically.

Simon Champion, Head of BT Electronic Commerce, said: 'Internet traders until now have not had an

easy method for charging for small-value goods, services, information or software delivered over the Internet from BT.

'BT Array provides an easy and secure micropayment system which extends the use of VISA and Mastercard credit cards into a new area of the market.'

BT Array users can manage their accounts and make purchases from any World Wide Web access device. In addition, they see their own credit balance and get a listing of their purchases on-line.

Security for On-line Payments

A law on the use of 'electronic signatures' in on-line transactions was proposed by the European Commission. It would set minimum requirements on security and liability, ensuring that electronic signatures—which enable the identity of senders and the integrity of the data to be checked—are recognised throughout the European Union. The Directive would remove one of the remaining obstacles to the widespread take-up of electronic trade, explained the Commissioner Martin Bangemann.

The proposed Directive is available at <http://europa.eu.int/comm/dg15/en/index.htm>.

Iridium Satellite Constellation Complete

The launch of five Iridium communications satellites on 17 May 1998 marked the technical completion of the venture's 66-satellite constellation. This is a major milestone in the schedule aiming for the start of service in September this year. Just 12 months ago the first five satellites were launched in one of the most ambitious launch campaigns ever undertaken in the space industry.

Iridium's 66-satellite wireless personal telecommunications network is designed to offer full global communications coverage and a variety of communications services, including voice, data, fax and paging,

to customers with Iridium hand-held telephones and pagers.

BT and NHS Partnership for On-Line Information

Frank Dobson, Secretary of State for Health, supported by BT, has launched the National Health Service 50th Anniversary Internet web site, marking an important step towards the future of public health information.

Sponsored by BT, the web site (<http://www.nhs50.nhs.uk/>) will provide a fast and innovative means of communicating with patients, staff and interest groups about the NHS and health care at home and in the community.

This is the first time the full information resources of the NHS have been accessible via the Internet. It marks another stage in the Government's drive to modernise the Health Service and take advantage of information technology.

People will be able to get clear, accurate and authoritative information about their own health and about the NHS in order to stay healthy and to get the best out of the service when they need it.

Improving public health is one of the Government's top priorities. The imminent publication of a Green Paper consulting widely about a key public health initiative, 'Our Healthier Nation', will increase the importance of making good public health information available.

The web site will provide information on the NHS, health-related issues and events forming part of the NHS 50th anniversary. It will evolve during the year, including an educational initiative with schools during Science Week, health education interactive games and personal health information.

The NHS is already trialling other online information systems, including an intranet at the Pinderfields and Pontefract Hospitals NHS Trust in Yorkshire. A number of fully functional Web servers are now in place, providing real-time information to doctors and other staff on bed availability,

waiting lists and a corporate directory.

Vision for Regulation of Communications Industry

Before he left Oftel, Don Cruickshank, Director General of Telecommunications, set out how a new simplified system for regulating the UK electronic communications industry could benefit consumers and encourage investment.

He called for the industry to be regulated by two new bodies:

- An Electronic Communications Commission, responsible for competition, economic and social policy issues, including consumer protection, and
- An Electronic Communications Standards Authority, responsible for content regulation, including public service broadcasting

And he called for a new statutory Consumer Council to promote consumer interests across the whole of the sector.

He pointed out that electronic communications covered not just traditional broadcasting and telephony, but also the way in which information technology controls access to networks.

He recommended a complete overhaul of the way content is regulated, because present arrangements were rapidly becoming impractical and unenforceable.

The framework for a new regulatory structure is set out in the final part of Oftel's evidence to the Culture Media and Sport Select Committee inquiry into the regulation of broadcasting. It covers the special rules needed in the new converged world of broadcasting, telecommunications and information technology, and how they should be enforced.

Publishing Oftel's evidence, Don Cruickshank said: 'The present regulatory system works by granting a privilege—the spectrum or a licence to operate—and being able to impose a set of detailed rules in return. This system breaks down as

capacity constraints disappear and the asset becomes less valuable. And this is happening now.

'We simply can't dither over sorting out what the new system should be. The new world, in which a flood of data and images will flow over networks spanning national and regulatory boundaries is already with us. The present regulatory structure is creaking, and the overlaps are creating uncertainty. 'No change' is simply not a viable option. The European Commission has already launched its debate into these issues, and we need to be able to lead that debate.

The UK has a head start because our telecommunications and creative industries lead the world. These industries will only invest if there is clear, stable and transparent regulation. We cannot afford to throw away that head start. We need firms to invest in these new services for the sake of job creation, and to make sure UK consumers get the benefits of new services in the Information Age.

The goal is an 'Open State' in communications, where there are no material capacity constraints on businesses setting up to offer these new services, provided they stick within a simple set of rules. In this 'Open State' all consumers must be able to pass smoothly and at an affordable price through the electronic turnstile which gives them access to the new education, information and entertainment services.

To make regulation clear and simple we should scrap the Telecommunication Act and most of the Broadcasting Acts. We should replace it with new legislation with relatively few rules designed to ensure fair competition, protect consumers and make sure there are no "have nots".

We cannot move to the 'Open State' without facing four big issues: how to make sure there will be no 'have-nots' in this new world; making sure there is interoperability, for example between set-top boxes; making sure we can keep the benefits of public service broadcasting; and making sure we have arrangements for filtering out potentially offensive material. Oftel's evidence to the

Select Committee deals with all of this.

'All our citizens should be able to be connected to these new services through an access control system in an affordable way, if necessary through Universal Service Fund arrangements.

'On interoperability, whoever is designing and licencing set top boxes for example must not be allowed to user that privileged position to cut back consumer choice by freezing out someone else's programmes and services.

The things consumers like about public service broadcasting, such as balanced reporting and high-quality programming, should be protected during the transition until we see what the new market delivers. That means giving public service broadcasters the resources to meet spiralling production costs, and that will need changes to the current licence fee system.

The trickiest issue of all is how to deal with potentially harmful content. The regulator should not tilt at windmills and attempt what cannot be achieved. The days in which material arrived in our homes from relatively few channels to which the Government controlled entry is disappearing fast. In the new world, consumers will be able to browse content, including video or programmes on demand, from thousands of different sources whenever they like.

The focus of content regulation must shift to reflect reality, by giving the consumer control over what services they use. This means we need a system in which all content is classified, based on codes drawn up by industry bodies working with the statutory bodies. That code should mean consumers are only able to see the material they choose to see, with safeguards like PIN numbers so children cannot get hold of unsuitable material. There must be a regulatory 'back stop' behind this which gives the regulator effective enforcement powers enshrined in law.

The final big issue is how the industry should be regulated. The existing structure of regulation, in

which regulators have overlapping and sometimes conflicting duties, needs reform. The present doesn't provide the clear, transparent and stable regulation which is needed; neither does it provide for an effective means of consumer representation or consumer redress.

The skills needed to ensure fair competition in the market are entirely different from those needed to deal with 'taste and decency' content issues. So two new regulatory Commissions should be set up. Competition, economic and social issues would be dealt with by an Electronic Communications Commission. Cultural content and standards issues, including the supervision of public service broadcasting, could be dealt with separately by a new Electronic Communications Standards Authority.

Underpinning all of this, we need a new consumer representation body with a remit to cover the whole sector, an Electronic Communications Consumer Council, to promote consumer interests and with the power to recommend change.

There would inevitably need to be close working between the two regulators because some content issues might have economic effects. However the area of overlap would be small, and the benefits outweigh any costs. These proposals are infinitely better than what we have now, and preferable to the alternatives of a single regulator, or a split between infrastructure and services regulation.

'If we put all this in place UK consumers and businesses can reap the full benefits of the Information Age.'

Oftel Plans Better Deal For Payphone Users

Plans to give payphone users a better deal have been announced by Oftel. Launching a consultation on payphone provision in the UK, the Director General of Telecommunications, said 'Growing competition means that there are more payphones in a variety of different places and this is good news for

users. But not all payphones in public places offer the same facilities even though they may look very similar. This is confusing and frustrating. Therefore, I am consulting on whether facilities such as the ability to call freephone numbers and access other network providers should be available from more payphones.

'We must find a balance between telephone services that operators should be obliged to provide from payphones and what should be left to their discretion.

'Too much regulation might discourage competition in payphone provision, but too little could leave users vulnerable.

'Lack of price information is a particular problem, especially on payphones operated on private premises by private individuals. The consultation document also seeks views on whether any special measures are needed to deal with unauthorised advertising in public callboxes.'

New Rules On Charges For Access To Digital Set-Top Boxes

Guidelines on charges to broadcasters for access to set-top boxes in digital television have been published by Oftel.

The Director General of Telecommunications said, 'The conditional access system (the electronic turnstile in the set top box) is crucial to the development of pay TV. As pay TV begins to play an increasingly important role in the UK, it will be vital to ensure that broadcasters are able to use conditional access systems on fair, reasonable and non-discriminatory terms.

'If we are to have effective competition in this market it is essential to ensure that control of conditional access systems is not used to distort competition in television markets by favouring some broadcasters over others. The guidelines set out the framework for achieving this objective while making sure that those who invest in conditional access systems are able to get a fair reward for their investment.

'I am committed to the government's objective of ensuring that there is a single framework for all digital services. As the next step toward that objective there will be a consultative document setting out Oftel's proposals for extending these rules to cover charges for access control for interactive services.

'This will lead to a single set of guidelines which will cover charges for both conditional access for digital television and access control for interactive services. However, given that the launch of digital television is now only weeks away I took the view that publication of the guidelines for conditional access charges would be helpful by giving broadcasters and conditional access providers a clear idea of what rules will apply. I hope that this will remove a potential stumbling block in negotiations between broadcasters and conditional access providers.'

Accounting Rate System Crisis

The ITU has released a report on the crisis in the accounting-rate system. The report is the product of the Seventh Regulatory Colloquium of the ITU, which was held in Geneva last December. Participants in the Colloquium included; telecommunications regulators, policy makers and telecommunications industry executives, Neil McMillan, chairman-designate of the forthcoming World Telecommunications Policy Forum, and T. Matsudaira, Chairman of Study Group 3 of the ITU Standardization Sector and ITU top officials.

The Colloquium's main conclusion is that the accounting-rate system, and the level of accounting rates, is not the real problem. The real problem is the lack of telecommunications development in many countries, due to such causes as inadequate investment, inefficient pricing policies, monopolistic industry structures and a lack of clear and effective national regulation.

To address these underlying problems, the Colloquium suggested

five measures, for consideration by the Second Policy Forum and the Second World Development Conference in Malta. These measures relate to investment and telecommunication development, national pricing policies, the 'new modes of operation' the role of the national regulator, and transition issues. The Colloquium also proposed several specific tasks that the ITU could perform including the creation of a 'one-stop shop' for funding of network development, which would help countries apply for concessionary and commercial funding available for such development, as well as offering technical advice and training.

In the past, the Regulatory Colloquium has dealt with subjects such as regulatory structures, universal service, global mobile personal communications systems, interconnection, the impact of the WTO trade agreements, and telecommunications convergence. In each case, it has sought to develop straightforward analyses and practical recommendations on how such matters should be addressed in an era of telecommunications deregulation. Both the Chairman's Report of the Seventh Colloquium and previous reports are available on the Internet via the ITU Home Page (<http://www.itu.int/itudoc/og/colloq.html>).

In addition, a much more detailed Briefing Report, prepared by Michael Tyler of the United Kingdom as a basis for discussion, is available for sale in English, French and Spanish.

World First for BT OnePhone

The UK is set to become the first place in the world to benefit from a breakthrough in telephone technology which merges a GSM mobile handset with the domestic phone, using a single number.

BT's **OnePhone** service will go on sale to the public later this year following the successful development of a version for business.

While in the home, the BT **OnePhone** logs on to the fixed

telephone network, acting as a high quality digital cordless phone. But once outside its 300 m range, it switches to a GSM mobile network to become a fully-functional cellular telephone.

Customers can have a new single number which reaches the BT **OnePhone** regardless of whether it is in home or mobile mode.

Eric Guilloteau, General Manager of BT Mobility Solutions, said:

'**OnePhone** is a mould-breaker which will transform the way we communicate.'

The integration of fixed and mobile telecommunications is one of the most important recent industry developments and BT's announcement gives the United Kingdom a dramatic lead over the rest of the world.

Scientists from BT's research laboratories have been working with the Swedish company Ericsson to develop the technology.

BT announced business trials for BT **OnePhone** in June 1997 at its own sites in Heathrow, London, Leeds and Hemel Hempstead. Following the success of these it has supplied the product for use in a number of large corporations.

The early success of the business service has enabled BT to develop a residential version more rapidly.

Over time, the BT **OnePhone** service will be developed to offer new services such as access to e-mail, fax and the Internet.

BiB Progress on EC Approval

BT has confirmed that it has given undertakings to meet the concerns of the European Commission (EC) in approving the formation of British Interactive Broadcasting (BiB).

Rupert Gavin, managing director of BT's Consumer Division, said: 'We believe the way is now clear for BiB to go ahead, including BT's participation in this venture. BiB will be a major step forward for the UK and will help bring interactive TV services to millions of viewers of digital television, starting later this year.'

BiB, a joint venture between BT, BSkyB, Midland Bank and Matsushita Electric, was formed to

provide a digital interactive TV platform over which content providers will be able to offer a range of interactive applications. These include home shopping and banking, education, games, internet and e-mail services.

As part of the approval package proposed by BiB and its shareholders to meet the Commission's concerns, third parties would have access to the BiB-subsidised set-top boxes and the software needed to create and run interactive services.

This was anticipated, and the detail of the proposal has now been worked out and provided to the Commission.

Also within the approval package is a proposal from BT to divest itself of its broadband cable TV interests in Westminster and Milton Keynes. The EC considers that BT's control of the existing broadband delivery mechanism in these areas raises competition issues in the light of BT's participation in BiB.

In the divestment of its broadband cable interests, BT will ensure a seamless transfer of customers to new operators.

Rupert Gavin said: 'This is great news for the UK and BT. BiB is a world first which will bring the UK to the forefront of the information age and enable digital TV viewers to enjoy the benefits of the information society.'

New Multimedia Opportunities For BT

BT has welcomed the Government's intention to lift the restrictions on national telecommunications companies providing broadcast services to homes over their networks.

Bill Cockburn, Group Managing Director of BT UK, said: 'Customers of the new multimedia services will reap the same benefits of competition as telecommunications customers.'

'With terrestrial, satellite, cable and interactive TV due to go digital this year, it is increasingly important that outdated regulatory restrictions are not allowed to slow down or distort the development of the UK information society.'

Bill Cockburn said that BT particularly welcomed the clarifica-

tion that broadcast services delivered over the Internet would not be considered as breaching the current restrictions.

He added: 'These kinds of services were never envisaged when the original broadcast ban was introduced, which shows the danger of trying to regulate by artificial restrictions on services and technologies.'

'BT will now be better able to plan for the implementation of multimedia services to the home with the certainty that from 2001 any broadcast element to these services will no longer be prohibited.'

Telephone Users Shop Around

More telephone users are reaping the benefit of competition in the telecommunications market, according to new data published by OfTel.

More than 750 000 householders chose telecommunications companies other than BT for their telephone lines last year, doubling the number of households with non-BT telephone lines in the last two years.

BT's share of the market for households has dropped by 3% over 12 months from its 91% share.

David Edmonds, Director General of Telecommunications, said, 'These figures show competition is developing well in the UK. Domestic telephone users are now taking advantage of the choice of telephone companies competing for their custom.'

notes and comments

The next edition of this *Journal* will coincide with the FITCE Congress being held, this year, in London. Papers from this congress will be published in the summer edition of the *Journal* which will, as a consequence contain significantly more than the normal number of pages.

For this edition only, *Telecommunications Engineering: A Structured Information Programme* will **not** be published as a supplement to the *Journal*. It will, however, return in the Autumn.

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by The Institution of British Telecommunications Engineers, Post Point G012, 8-10 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 *British Telecommunications Engineering* or the *Structured Information Programme* unless any such statement is made specifically by the Board.

© 1998: The Institution of British Telecommunications Engineers.

British Telecommunications Engineering and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

Subscriptions and Back Numbers

Annual subscriptions (including postage)—companies, universities, libraries and other bodies £60.00 UK, £65.00 overseas; private individuals (paid by personal cheque) £36.00 UK, £41.00 overseas.

Single copies and back numbers—companies etc. £15.00 including postage UK, £16.25 including postage overseas; private individuals (paid by personal cheque) £9.00 including postage UK, £10.25 including postage overseas.

Overseas customers should pay by a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London.

Price to BT employees: call +44 171 356 8050 for details.

Orders, by post only, should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: G012, 8-10 Gresham Street, London EC2V 7AG.

Remittances for all items should be made payable to 'IBTE' and should be crossed '& Co'.

An index of recent back issues is on <http://www.ibte.org/ibtepubs.htm>

Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: G012, 8-10 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, 8-10 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 the Board of Editors' permission is necessary for reprinting extracts.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by The Institution of British Telecommunications Engineers for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5.00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/98 \$5.00 + .00



THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone/Centre	Contact	Telephone Number
London	Terry McCullough	(01707) 601151
Martlesham Heath	Howard King	(01473) 642583
Midlands	Eric Pedersen	(0121) 230 4344
North East	Geoff Jenkinson	(01226) 248141
North West	David White	(0161) 600 2468
Northern Home Counties	Kevin Rees	0802 195219
Northern Ireland	David White	(01232) 215354
Paris	IBTE Office	+44 171 356 8008
Scotland	Raymond MacMillan	0800 671991
South West	Peter Shaw (in Bristol)	(0117) 972 3216
Southern Home Counties	John Dymott	(01202) 554001
Wales	Geoff Oatway	(01492) 531370

IBTE International HelpLine: +44 171 356 8008

IBTE On-Line (Internet): <http://www.ibte.org>

(BT intranet): http://networks.intra.bt.com/ibte_nat/ibtehome.htm

BRITISH TELECOMMUNICATIONS ENGINEERING

Journal Subscription Order Form

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager
Post Point G012
8-10 Gresham Street
London EC2V 7AG
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£60.00 UK; £65.00 Overseas**
Private individuals (to be paid for by personal cheque): **£36.00 UK; £41.00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', should be enclosed with the order. Overseas customers should pay a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London. Alternatively, customers may wish to use a local subscription agent.)

Title (Mr/Mrs/Ms etc) _____ Initials _____ Surname _____

Mailing Address _____

Postcode (Zip Code) _____ Country _____

Signature _____

Date _____

Please state the issue you wish your subscription to commence: _____
(April, July, October or January)

VAT Registration Number (zero rate): 243 5917 54 *Photocopies of this form are acceptable.*



FEDERATION OF TELECOMMUNICATIONS ENGINEERS
OF THE EUROPEAN COMMUNITY (FITCE)



37th European Telecommunications Congress

Institute of Education, London, 24–28 August 1998

‘Diverging Roles in a Converging Marketplace’

Telecommunications is converging with computing and broadcasting, opening new opportunities both for traditional operators and new arrivals. By the time of the Congress, the European telecommunications market will have been open for eight months to new entrants from all industries. The Congress will include papers exploring the technical and commercial opportunities that this situation offers.

The programme includes keynote speakers, a round table discussion and papers divided into 14 sessions: The Market Place, The New Order, Working Across Boundaries, Managing Across Networks, Intelligent Network Services, Mobility, Tools for Management, IN Management, The Broadband Platform, Broadband Access, Internet Opportunities, Applications, Futures, and Into the Millennium. Technical visits, a cultural programme and partners programme are also included.

Further information, including details about how to register for the Congress, is available from national FITCE associations, the FITCE'98 Web site <http://www.btwebworld.com/fitce98/> or the Congress organising office, BT Conference Unit on +44 181 477 2842.

IBTE London Colloquium 30 October 1998

CALL FOR PAPERS

‘Emerging Telecommunications Core Network Technologies’

Papers are invited for the above IBTE Colloquium. The colloquium will focus on the core transport network (as opposed to local access technologies).

Authors are requested to submit an abstract of 200–400 words to Paul Nichols, IBTE Office, Post Point G012, 8–10 Gresham Street, London EC2V 7AG; Telephone: 0171 356 8022; Facsimile: 0171 356 7942; e-mail: paul.e.nichols@bt.com. The deadline for submissions is 1 July 1998. The submission should include the author's name, affiliation, address, telephone and facsimile numbers, and electronic mail address.

Tickets (including proceedings and light refreshments) are free to IBTE Members. For all others, tickets are £50 each (by cheque payable to 'IBTE').

For bookings and further information, contact the IBTE Office, Post Point G012, 8–10 Gresham Street, London EC2V 7AG; Tel: 0171 356 8008; Fax: 0171 356 7942; E-mail: london_colloquium@ibte.org. In all communications, IBTE Members should quote their membership numbers.

Contents

Guest Editorial: Structured Process Improvement	1
Lean Operations—The Route to Competitiveness	2
Computer Telephony Integration and Java	10
Mobile Multimedia Applications	18
Knowledge Discovery and Delivery	25
Intelligent On-Line Purchasing	36
The BT Intranet Complete Service	43
Intranet-TV—Video Streaming for the World Wide Web	49
Business Use of Internet Web Sites— <i>Could Do Better!</i>	56
Getting the Message, Loud and Clear—Quantifying Call Clarity	66
Future Trends in Satellite Communications	73
SimDS—Modelling the Performance of Future Systems	94

