

VOLUME 16 PART 1 APRIL 1997

BRITISH TELECOMMUNICATIONS ENGINEERING

Included in this Issue

Start of New Theme:

Optical Technologies

*World Trade Organisation Talks on
Basic Telecommunications*

*An Evolutionary Strategy for New
Communication Services*



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



BRITISH TELECOMMUNICATIONS ENGINEERING

Contents

VOL 16 ■ PART 1 ■ APRIL 1997

Editorial: Optical Horizons	1
David Smith	
<i>Optical Technologies</i>	
Building the Road to Optical Networks	2
Goff Hill, Lars Fernandez and Roberto Cadeddu	
<i>Perspectives</i>	
World Trade Organisation Talks on Basic Telecommunications	13
Neil McMillan	
<i>Operational Support Systems</i>	
Operations and Maintenance Centre: Delivering Network Services	18
Andrew Dawson-Maddocks, David Cooper and Colin Scobie	
<i>Telecommunications in the 21st Century</i>	
The Mirror—Reflections on Inhabited TV	29
Graham Walker	
Whither Speech? Convergence and the Future of Telephony	39
Alwyn Lewis and Fred Westall	
A Computer Vision System for Natural Communication	45
Dave Machin and Phil Sheppard	
An Evolutionary Strategy for New Communication Services	49
Heikki Sundquist and Jari Karttunen	
Authoring and Automating Content with its Associated Metadata for Multimedia On-Line Services	56
Stephen Searby, John Salmon, Janie Taylor, Paul Foster, Royston Walker, Richard Gepp and Sally Mossman	
Customer Access Signalling on Concert VNS	66
Sanjiv Angrish	
MuseNet	70
Sue Nicholas	
<i>Regular Features</i>	
pcochrane@btlabs	71
Book Reviews	72
Telecom News	75
IBTE Contacts	Inside Back Cover

Cover Picture:
Telegraph Colour Library

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.; Martin Evans, B.SC., PH.D.; Warwick Flury; David Greenop, B.SC., M.SC.; Peter Howell, B.SC., C.ENG., M.I.E.E., M.C.I.B.S.E, M.I.O.M.; Murray Trantor, B.SC., C.ENG., F.I.E.E., D.M.S., M.I.MGT.; Keith Ward, C.ENG., F.I.E.E., M.I.MGT.; David White, M.I.P.D., M.B.I.M., L.I.A.A.

Theme Editors

Optical Technologies David Smith

Perspectives Chris Earnshaw

Operational Support Systems Sinclair Stockman

Telecommunications in the 21st Century Peter Cochrane

David Smith

Optical Horizons



*The capacity of
BT's existing glass
network can be
upgraded by
between one and
three orders of
magnitude*

Cyberspace, Internet and 'Information Superhighway' are now part of the everyday language used in a world where computers and telecommunications have at last started to converge. The two industries have different backgrounds and perspectives. The cheap silicon computing technology now on every desktop spun out of the intensive research programmes of 40 years ago, driven by cold war military and space programmes; the Internet itself was originally developed as a robust defence network. However, the shared vision of a global information network also assumes a falling cost curve for bit transport. This cost reduction is only possible because of the outstanding success of the civil telecommunications optical-fibre research programmes initiated in the late-1960s and early-1970s. Vintage issues of this *Journal* describe some of the ground-breaking contributions made in those early days.

Today, glass circles the globe and in the developed world most inter-exchange traffic is carried over fibre. Technology first developed for civil telecommunications is now finding its way into cable TV systems, surveillance systems, computer local area networks (LANs) and avionics. Interestingly, the basic semiconductor laser research that originally underpinned optical communications went on to seed the compact disc player and CD-ROM; the latter, arguably, one of the most important technologies to increase the home use of personal computers.

The special series of articles on 'Optical Technologies' starting in this issue of *British Telecommunications Engineering* is not intended to be a celebration of past heritage but a pointer to the future. The capacity of BT's existing glass network can be upgraded by between one and three orders of magnitude by using combinations of wave-division multiplexing (WDM), optical amplifiers and ultra-fast time-division multiplexing. The increasing richness of the available optical technology allows optical networking to be considered as a complement to existing electronic networking techniques, such as synchronous digital hierarchy (SDH). These optical networks will not only have increased capacity but should be more resilient to major faults, such as cable breaks, because optical switches can be used to rapidly reconfigure and reroute huge blocks of capacity without complex electronic demultiplexing. The first pilot applications of WDM networking in BT are to be tested in Scotland later this year.

Future articles in the series will describe some of the advances within the laboratory and the new applications. Techniques to self-route optical packets at 100 Gbit/s, which is 1000 times faster than today's LANs but necessary to meet the projected performance of tomorrow's workstations, have been demonstrated. New concepts for optical logic processors which may ultimately out-perform electronic computers are being researched. Optical technology is being applied to revolutionary new applications such as fibre-fed radio picocells, potentially eliminating any electrical connection to the remote base station—not even for power.

Research and development is increasingly a shared activity with no single team ever holding all the pieces necessary for holistic success. In this vein, the first article of the series, on page 2, describes research into optical networking carried out across Europe within the RACE collaborative research programme.

David Smith

**Manager, Networks Research Unit,
BT Networks and Systems**

Building the Road to Optical Networks

The transmission network is not seen by customers, but their service would not be possible without it. While network capacity is increasing, the costs of transmission must be held down. Collaborative research has shown how a new kind of network layer based on optical network elements can provide a solution. The network layer is fully manageable and transparent to underlying electrical signal format and bit rate. From an operator's point of view, these features could form the bedrock for broadband networking into the next millenium.

Introduction

In January 1992, the RACE Multi-Wavelength Transport Network (MTWN) project set out to prove that a new kind of network layer could be constructed from optical network elements. The network layer was to be based on optical signal routing, optical amplification and optical cross-connect switching. It was to be fully manageable and transparent to underlying electrical signal format and bit rate. From an operator's point of view, these features could form the bedrock for broadband networking into the next millennium. At the outset this was a novel, radical and highly speculative prospect but by March 1996, the consortium (Table 1) had concluded its work, having successfully developed key technologies, demonstrated the operation of a number of versions of a flexible, managed optical network and identified possible evolution routes. The project has had a major impact on the wider development of transport networks and this article presents an overview of its achievements and future optical networking potential.

Table 1 The MWTN Consortium

BT Laboratories, UK
Ericsson Telecom AB, Sweden
Ellemtel, Sweden
Telia AB, Sweden
Ericsson Telecomunicazioni, Italy
Pirelli Cavi, Italy
Italtel, Italy
CSELT, Italy
CNET, France
University of Essex, UK
University of Paderborn, Germany

Wavelength-Division Multiplexing

Wavelength-division multiplexing (WDM) allows an optical fibre to carry many optical transmission signals together over a single fibre by setting each signal at a different wavelength (or 'colour'). The different wavelengths can be defined by very pure laser sources, each modulated by a separate information signal and the resulting signals combined (or multiplexed) onto a single fibre. At the receive end the signals can be separated using either an optical splitter/filter arrangement or dispersive element which gives a physical separation of the different wavelengths in the way a prism separates visible light into colours. This concept has been understood since at least 1968 when Kinsel and Denton¹ proposed the use of wavelength multiplexing to increase the capacity of an optical transmission system, but it was 1977 before the first practical experiments were staged using a blazed grating². However, it was not until the mid-1980s that the first practical systems trials were staged, including a series of two wavelength trials in the UK network and soon afterwards in Sweden (some of these systems are still in operational use, including one as a 140 Mbit/s protection circuit between Ipswich and Colchester!).

It was several years later that the first proposals were made to interconnect WDM systems³ and to develop flexible WDM networks⁴ incorporating optical amplification and switching. The success of the early experiments⁵, the first field demonstrations in the London Fibre Network (1989) in which individual wavelength channels were switched⁶ and related demonstrations

Figure 1—Concept of the multi-wavelength transport network


in the European RACE I programme⁷ stimulated widespread international interest and was one of the primary factors to shape the formulation of a collaborative project called *Multi-Wavelength Transport Network* (MWTN) within the European Union's RACE II programme.

Principles of MWTN

The MWTN network (Figure 1) offers a set of optical routing nodes that are interconnected by fibre^{8,9}, in which optical wavelength channels can be set up, modified or cleared down under the control of a management system. The architecture may be likened to that of an SDH network except that instead of digital channels being multiplexed and switched, optical wavelength channels are being controlled.

The fibre layout can be adapted to any physical structure (such as a mesh or ring) while the logical channel connectivity is independent of physical topology and can be set according to the traffic loading. A key property of the channels is that they are 'transparent' and can take a payload of any format or bit rate (within design limits). These properties enable a network layer based on WDM principles to be developed that is capable of supporting a variety of client layers having signals of different bit-rates and formats, such as synchronous digital hierarchy (SDH), plesiochronous digital hierarchy (PDH), asynchronous transfer mode (ATM) etc.

An MWTN node can provide wavelength channel add-drop or cross-connect functionalities and the WDM signals may be processed either as a multiplex or as individual channels. Many component and node configuration options are possible using the basic elements of optical amplifiers, wavelength multiplexers and space switches. To keep the task to manageable proportions, the MWTN project developed two types of node: one that demultiplexes a multiwavelength signal and allows channels to be cross-connected (termed an *optical cross-*



connect node) and a simpler version that allows individual wavelengths to be added to or dropped from a multiplexed path (termed an *optical add-drop node*). For each node a variety of component technologies were developed and compared, while architectural studies focused mainly on developing the optical layers and their relationship to the SDH layers.

Optical Cross-Connect Node Configuration

The optical cross-connect node configuration chosen for the MWTN demonstrations (Figure 2) contains four principle sub-systems: a line sub-system with erbium fibre optical amplifiers, a transmission sub-system with wavelength switched or tunable

Figure 2—MWTN cross-connect node architecture





Figure 3—MWTN in the Stockholm Gigabit Network

sources, an optical switching sub-system and a management sub-system.

Fibres entering an MWTN node carry multiple wavelength signals. These composite signals are first amplified, then demultiplexed and routed as separate wavelengths via a space switch array, either to outgoing fibres or dropped to local receivers. Locally generated signals can be introduced via the space switches and can be multiplexed with other 'local' or 'through' signals. The practical work in MWTN is based on 4×2 nm wavebands spaced at 4 nm intervals in the erbium fibre window, each with a bandwidth up to approximately 200 GHz. The wavelength plan adopted allows for twice this number to be operated within the 30 nm window allowed by erbium fibre amplifiers. This was a modest target in terms of numbers of wavelengths, but there were many technical risks to be faced to achieve the project objectives. A



further technological approach (not illustrated) based on *fine grain WDM* with channels spaced at 10 GHz (~ 0.1 nm) intervals was also developed which had the potential to substantially increase the capacity of the network.


Field Trials

During the course of the project, two optical cross-connect nodes and two wavelength add-drop nodes were constructed. The main trials, Figure 3, took place under the leadership of S. Johansson of

Ellemtel (now Ericsson Telecom) from September 1994 to September 1995 in the Stockholm Gigabit Network (provided by a separate Swedish consortium and part of the Telia AB network). Two cross-connect nodes were installed on opposite sides of Stockholm, one at Älvsjö and the other at Kista. They were interconnected via two physically separate fibre routes. The trial network is shown in Figure 4.

The optical equipment incorporated a variety of leading edge integrated optical technology components such as 8×8 space switches

Figure 4—MWTN demonstrations in the Stockholm Gigabit Network



based on lithium niobate¹⁰, wavelength tunable DBR lasers¹¹, 4×4 amplifier gate array switches¹² and multi-grating filters¹³ in indium phosphide and acousto-optic filters also in lithium niobate^{14, 15}. It also included high-power wavelength flattened amplifiers. A management system based on TMN principles¹⁶ was developed by K. Elander and team of Ericsson Infocom to control the devices and to integrate optical-layer and SDH management processes. The management system (Figure 5) uses a two-layer functional model as the basis for its object definitions and provides a selection of fault, configuration and performance facilities. The two nodes communicated via an X.25 packet network. The management system controlled the network configuration and enabled a detailed network status analysis to be carried out.

The trials showed both service configuration and network protection functions. Videophone connections were set up and cleared down via the management system. The optical cross-connect switches could be configured either as 1-to-1 switches or as optical broadcasting elements and the management system was also used to configure a broadcasting reference circuit that made five passes through an optical node (twice through the Kista node and three times through the Älvsjö node). A fibre failure was also simulated: the management system detected the loss of signals and automatically restored the failed channels via an alternative routing. A detailed analysis of the alarm and configuration status of the network could then be made via the management system. This demonstration was a major milestone for the project and it was a world first demonstration of a managed optical network following TMN (Telecommunication Management Network) principles, proving the practicality of a managed optical network platform. It also helped to highlight the need for architectural standards for optical networks.






Figure 5 – MWTN management system

A further demonstration showed a unidirectional, survivable wavelength add-drop ring¹⁷. This used a novel node configuration (Figure 6) designed by S Merli of Ericsson Telecomunicazioni based on acousto-optic lithium niobate wavelength switches which were developed within the project together with 2×2 electro-mechanical optical switches. The acousto-optic tunable switch

(AOTS) routes selected channels for onward transmission or for dropping. The 2×2 switch selects either the working or the protection mode. Management operations were integrated with the SDH management system. Two optical add-drop nodes were constructed and demonstrated, first in the Stockholm Gigabit Network and then at the European Exhibition of Optical

Figure 6 – Wavelength add-drop node





Communication (EEOC) in Brussels, September 1995 (Figure 7). The system was used to carry coded video traffic between end terminals. This was a world first demonstration of a TMN managed WDM ring network.

Case Studies

As well as showing that the technology was feasible, it was equally important to check that there would be benefits in realistic applications. The following case studies examined a range of design issues, embracing traffic, cost and physical dimensioning.

A metropolitan-scale network

A key measure of the effectiveness of an optical network layer was made by comparing a representative ATM/SDH network design to one in which the ATM/SDH network was supplemented by MWTN nodes. A case study of a Milan metropolitan network was carried out by M. Giorgi and R. Cadeddu of CSELT¹⁸. The basic network contained 42 nodes and was the result of an earlier RACE I study (Figure 8). The first question to be decided was: should all nodes include optical routing or just a few? If the latter, which should they be? A routing strategy was developed to exploit the high throughputs possible with MWTN nodes. The criterion that was used for the optical node selection, was based on the identification of the set of k -optical nodes which gave the minimum global network cost, and an exhaustive analysis for all node combinations from four through to 24 was carried out.

The results were striking. The lowest-cost network found used 20

optical nodes and eight wavelengths. This was not an optimised network design in terms of node size but in terms of overall network cost, yet the sizes of the digital cross-connects were reduced up to 80% (Figure 9) and on average, 35%. The cost assumptions were conservative, given the long-term high-volume traffic scenario considered. This work showed, for the first time, that modifying the routing strategy to take account of the higher-capacity optical-channel paths could

Figure 7—Survivable optical add-drop multiplexer ring demonstrated at EEOC '95

have a major impact on overall network complexity and hence on network investment costs.

A national-scale network

A major case study carried out by R. Brändström, S. Karlsson and P. Lindberg of Telia Engineering based on the Swedish national network considered both the capacity demands and the transmission limitations together. A pure SDH network and a MWTN solution were considered. Both solutions were capable of single failure protection against node failure using 1+1 sub-network connection protection (electrical or optical). The pure SDH network solution included future network elements like STM-64 and STM-256 line terminals and regenerators, making cost estimation difficult. A simple transmission model was used to test validity of the three solutions for a variety of amplifier spacings. The approximate transmission modelling

Figure 8—Milan area network: the selected MWTN network is highlighted





Figure 9—Digital cross-connect size savings in the Milan area network






Figure 10—National-scale transport network




Figure 11—Overall cost comparison of pure SDH and MWTN/SDH solutions

suggested a maximum of 15 intermediate nodes with 60 km spacing for the MTWN solution, corresponding to a maximum transmission distance of 960 km. The target networks were chosen accordingly with 16 wavelengths, each carrying 2.5 Gbit/s. Four of these channels were reserved for later upgrade. The network that was studied²³ had a capacity demand of 1401 unprotected VC-4 channels and a common cable structure according to Figure 10 with three or more cables between the 24 nodes.

The study showed that the MWTN solution was significantly lower in cost than the pure SDH solution (Figure 11). One of the key factors was the cost difference between using few optical amplifiers rather than many regenerators in the longer-distance links. Furthermore, under conditions of growth, the optical




Figure 12—Network growth with SDH requires a larger digital cross-connect and more line equipment than an MWTN solution

cross-connect solution offers savings in the size and cost of the digital cross-connect and of the line terminations (Figure 12). In this study the average saving in terms of VC-4 ports on the digital cross-connect was 39%.

These results were independent of the Milan study and they provide corroborative evidence that MWTN solutions can lead to significant cost savings owing to the use of amplifiers rather than regenerators and to optical cross-connect networks rather than digital. In higher-capacity and longer-distance networks the introduction of an optical layer will lead to even greater savings in capital expenditure.

A European-Scale Network

An illustration of the long range potential of a WDM network came

through joint studies, led by M. J. O'Mahony from University of Essex, between MWTN and the European COST 239 (Co-operation On Science and Technology) project led by E. le Coquil from CNET. This recognised the possibilities for optical networking at a European or global level and developed a hypothetical European scale network based on MWTN principles¹⁹. This network was based on a relatively small number of nodes that provided gateways to national networks and which were interconnected by very high capacity pipes that would carry all the international traffic. It assumed a futuristic scenario for traffic levels and then developed an outline network design. One of the resulting networks is reproduced in Figure 13.

Figure 13—A European-scale network based on the principles of MWTN




Figure 14—WDM platform development

The diameter of the network is more than 3000 km and this stretches the capabilities of optical networking towards its limits. Physical effects such as noise build up, crosstalk and non-linearity will limit the numbers of wavelengths that can be transmitted together across these distances without regeneration. On the other hand, too few wavelengths in the network could result in poor blocking characteristics. One solution to this problem was to partition the network¹⁹ into independent islands with high-capacity links between the islands and regenerators at the ends, allowing the same set of wavelengths to be reused in each island.

Opportunities with WDM

MWTN has shown that it is possible to carry PDH and SDH signals together on the same WDM infrastructure and the indications are that this can be achieved at a realistic cost. This opens up new ways of upgrading and restructuring transport networks. The first way this might be used is to increase the efficiency of utilisation of fibre cable resource, with WDM offering an alternative to installing more cable. This could aid network planning by

providing spare wavelength channels in the optical network layer as an alternative to reserving equivalent capacities in SDH network layers. A planning approach based on over-supplying the network with low cost WDM capacity could potentially lead to faster, lower cost service provisioning. (Current ITU Recommendations anticipate up to 32 wavelengths being used in the short-term.)

Just as a cable can be used to support PDH and SDH signals together, so too could a WDM network platform, as long as the appropriate management systems support is available. Then, with much higher capacities being carried by fibres, rapid and reliable bulk capacity protection systems will be needed to guard against cable damage by civil works etc. Optical protection switching can give fast protection against fibre failure and is able to re-route wavelength channels either individually or together. The inherent speed of optical switches gives a high potential for fast protection and a range of possible options. Just as in

SDH the switching can be either single ended or dual ended and protection can be either 1+1, 1:n or at the network layer. In essence, any protection scheme that is possible in digital network layers can also be implemented in the optical layers.

If future broadband network requirements dictate that a different client network architecture is needed, this could be grown above the WDM platform and alongside the PDH and SDH layers to allow a controlled migration to the new network (Figure 14).

The WDM platform could provide wavelength channels which are accessible directly via digital or optical interfaces or through a transponder/wavelength converter. New architecture developments can be facilitated by providing additional functionality in the optical layers such as managed optical add-drop multiplexers (fixed or flexible) and managed optical cross-connect networks. Some of the options for ring or cascaded networks are illustrated in Figure 15. These





Figure 15—Some configuration options for MWTN nodes






Figure 16—The London-East Anglian Network (LEANet) test bed

components give rise to a range of new options for network infrastructure. For example, when using a fixed add-drop multiplexer (with no switching) the protection and network configuration control remain in the SDH (or PDH) layers. This leads to a need for duplicate line cards and optical channels. By adding on a pair of 2×2 optical switches the protection function is moved to the optical layer with the effect that fewer line cards are needed.

The manufacture of larger scale optical switches provides a technological challenge to achieve the necessary performance and cost. Networks requiring optical cross-connect switches will therefore take time to develop. However, the benefit is that optical cross-connect switches can become a natural complement and provide load relief to SDH or ATM cross-connect switching and can be used for protection switching, capacity routing and even broadcasting. Networks with optical cross-connect switches require a more complex control and management system and the technology is still being developed. However, as network capacity grows, so too will the need for these components and in the long term they could provide a powerful solution for a survivable optical network.

Beyond MWTN

Although the MWTN project has ended, several strings of work are continuing. Within BT, a networks test bed known as *LEANet* (London-East Anglian Network) is being developed to support network and application initiatives. In Italy,

Telecom Italia is working with some of the project partners (CSELT, Ericsson Telecomunicazioni, Italtel, Pirelli) to trial different types of WDM ring technology in a project known as the *PROMETEO* (PROgetto METropolitano Telecomunicazioni Ottiche). Within the ACTS programme²⁰ several projects are building on the expertise developed within MWTN.

LEANet

LEANet consists of a fibre network which extends from Martlesham Heath to London using fibres in operational cables (Figure 16) and includes a WDM layer, a 2.5 Gbit/s SDH add-drop ring and element control and management systems.

The northern part of this test bed (highlighted in red) has already been used to demonstrate a Ciena 16 wavelength system, optical protection switching, PDH system restoration, support of SDH channels and optical system test facilities. The demonstrations have helped to build confidence in the technology and have emphasised the potential to reduce network cabling costs. In the future, LEANet will be linked to ATM switches and IP routers in BT's futures test bed to support new applications research to off-site facilities. It is also providing a base infrastructure for ACTS projects WOTAN (Wavelength Agile Optical Transport and Access Network) and COBNET (Corporate Optical Backbone Network) which build on the optical routing concepts developed by MWTN. WOTAN is examining how WDM transport systems can be integrated with other network platforms such as ATM and wavelength agile PON access systems while COBNET explores how to build

corporate network islands based on optical cross-connects that could be linked together by means of a transparent transmission network.

PROMETEO

The success of MWTN has resulted in two metropolitan network field trials coordinated by Telecom Italia in the Rome and Turin networks. Know-how from the project has been funnelled and integrated in the PROMETEO project. To verify the technological maturity of WDM equipment, prototype optical add-drop multiplexers (OADM) and optical protection systems are being employed in ring networks to test and validate the performance in a metropolitan network environment. The two experiments are positioned at different time instances of a possible development path towards optical networking in ring networks. Both experiments involve network planners and network management teams so that a balanced operational assessment can be reached.

The Rome unidirectional ring²⁰, shown in Figure 17, uses three nodes to perform optical channel add/drop and protection of STM-16 SDH tributary channels. It has a total capacity of 10 Gbit/s (4×2.5 Gbit/s) and represents a complete optical layer in that it is possible to establish an optical channel trail connection between nodes of the ring. In case of a line fault, all traffic is rerouted on the standby fibre opposite the fault to reach its original destination.

The Turin bidirectional ring¹⁹ is at an intermediate evolution point. The aim is to demonstrate that installed transport network capacity can be upgraded by relieving the client SDH layer of the need to assign spare capacity for protection purposes: WDM channels are only used for protection. In this implementation SDH VC-4/VC-12 path trail connections are used for communication between the nodes. Three optical protection systems have been developed and interfaced to ADM-4. These systems are bit rate and signal




Figure 17—PROMETEO optical unidirectional self-healing ring in the Rome metropolitan network

format transparent and can be used at higher transmission rates (2.5 Gbit/s or 10 Gbit/s). The network is shown in Figure 18.

Two separate wavelength channels are used in a bidirectional node-to-adjacent-node configuration on the two fibre ring. Each wavelength can be back-routed on the fibre opposite the fault to occupy the available wavelength slot. Using just two wavelengths and using a simple functionality of an optical layer, it is possible to meet capacity demands from traffic growth immediately, upgrading rather than replacing existing SDH ring networks and increasing the transport capacity available for traffic.





Figure 18—PROMETEO optical bidirectional self-healing ring in Turin metropolitan network

Figure 19—METON demonstrator outline in the Stockholm Gigabit Network



METON

The ACTS METON project²⁰ takes a wider, more futuristic view and is developing an optical WDM-based transport layer aimed at supporting an ATM/SDH client network to provide a significant number of subscribers with a multiservice network which can support heterogeneous services (Figure 19). METON will stage a demonstration in 1998 in the Stockholm Gigabit Network National Host which will consist of a service application layer based on ATM together with an optical transport layer. Five sites, locating the OADMs will be equipped with ATM and optical add-drop multiplexers to provide service access. Services will be available through a centralised ATM switch located in the central node. The optical cross-connect node determines the routing of the optical channels over the network while the optical add-drop serves dedicated optical channels to each ATM multiplexer.

The project takes off from the achievements of MWTN. In the area of network management and control, the focus is now raised from the management of the optical elements to a more comprehensive network element management, with particular attention to the issues of management information flow within an

optical network using dedicated optical control channels.

Conclusions

MWTN has provided a longer-term vision for transport networks and has generated a range of evolutionary options which are now widely influencing strategic transmission network developments. The high international profile of this work has had a significant impact on other research programmes and on the development of optical network technology around the world. Studies have shown that a multi-wavelength approach in the transmission network can lead to better utilisation of fibre and has the potential to simplify the network through reduced numbers of regenerators, through lower quantities of digital switching equipment and through the ability to enable growth and facilitate network restructuring. The technology developments, demonstrations and analyses made in this project have all shown that wavelength routing is feasible, flexible and manageable and should now be considered as a short-to-medium-term technology rather than a long-term one. Since the completion of this project, long-distance network operators have begun installing operational WDM systems with up to 16 wavelengths and some of the technology developed in the MWTN project has been incorporated into some of these earliest systems. The interest in flexible networking is now migrating from research into development programmes and the way in which BT is exploiting optical networking technology will be covered in future issues of the *Journal*.

The MWTN project has led the world in this field and its success was marked in November 1995 by receiving the RACE Award for its extraordinary contribution to the European Union RACE programme in technological progress.

Acknowledgements

The content of this article is based on contributions from all of the MWTN project partners (Table 1) and their supporting teams. It is not possible to mention all of the 50 or more contributors to the project individually, but the authors wish to thank the individual members of the project teams for their enthusiastic contributions, the participating companies, the Stockholm Gigabit Network consortium, NUTEK (the Swedish National Board for Industrial and Technological Development) and the European Union RACE programme for their strong support of this work.

References

- 1 KINSEL, T. S.; and DENTON, R. T. Terminals for High Speed Optical Pulse Code Modulation Communication System: II Optical Multiplexing and Demultiplexing. *Proc. IEEE*, Feb. 1968, **56**(2).
- 2 TOMLINSON, W. J.; and AUMILLAR, G. D. Optical Multiplexer for Multimode Fiber Transmission Systems. *Appl. Phys. Lett.*, Aug. 1977, **31**(3).
- 3 KOBRINSKI, H. Cross-Connection of Wavelength Division Multiplexed High Speed Channels. *Electron. Lett.*, 1987, **23**, pp. 974–975.
- 4 HILL, G. R. A Wavelength Routing Approach to Optical Communications Networks. *Br. Telecom Technol. J.*, July 1988, **6**(3), pp. 24–31.
- 5 CHIDGEY, P. J.; and HILL, G. R. Experimental Demonstration of Wavelength Routed Optical Networks Over 52 km of Monomode Optical Fibre. *Electron. Lett.*, 1989, **25**, 1451–1452.
- 6 WESTLAKE, H. J.; CHIDGEY, P. J.; and HILL, G. R. Reconfigurable Wavelength Routed Optical Networks: A Field Demonstration. *Proc. ECOC*, 1991.
- 7 JOHANSSON, S. et al. Optical Cross-connect System in Broadband Networks: System Concepts and Demonstrators Description. *IEEE J. Lightwave Techn.*, 1993, **11**(5/6), pp. 688–694.
- 8 HILL, G. R. et al. A Transport Network Layer Based on Optical Network Elements. *IEEE J. Lightwave Techn.*, **11**(5/6), pp. 667–679.
- 9 CHIDGEY, P. J. Multi-Wavelength Transport Networks. *IEEE Communications*, Dec. 1994, **32**(12).
- 10 GRANESTRAND, P. et al. Pigtailed, Tree Structured LiNbO₃ Switch Matrix with 112 Digital Optical Switches. *Phot. Lett.* 1994, **6**(1), pp. 71–73.
- 11 STOLTZ, B.; DASLER, M.; and SAHLEN, O. Low Threshold-Current, Wide Tuning-Range, Butt-Joint DBR Laser Grown with Four MOVPE Steps. *Electron. Lett.*, 1993, **29**(8), pp. 700–701.
- 12 LARSEN, C-P., et al. Transmission Experiments on Fully Packaged 4 × 4 Semiconductor Optical Amplifier Gate Switch Matrix. *Proc. Photonics in Switching '95*, 1995, paper PThD2, pp. 95–97.
- 13 WEBER, J-P., et al. Four Channel Tunable Notch Filter Using InGaAsP/InP Reflection Grating. *Phot. Lett.*, July 1994, **6**(1), pp. 77–79.
- 14 TIAN, F. et al. Polarisation Independent Integrated Optical, Acoustically Tunable Double Stage Wavelength Filter in LiNbO₃. *IEEE J. Lightwave Techn.*, July 1994, **12**(7), pp. 1192–1197.
- 15 WEHRMANN, F. et al. Fully Packaged, Integrated Optical Acoustically Tunable Add-Drop Multiplexers in LiNbO₃. *Proc. 7th European Conference on Integrated Optics*, Apr. 1995, Delft, pp. 487–490.
- 16 HUBINETTE, C.; ALMSTROM, E.; and JOHANSSON, S. Results from the Stockholm Gigabit Network—WDM Networking. Broadband Superhighway, NOC '96, IOS Press, pp. 80–86.
- 17 MERLI, S. A Transparent Self Healing Ring Architecture: Analysis and Dimensioning. *Proc. IOOC '95* (Session Title ThC2-2).

- 18 CADEDDU, R. et al. Multi-Wavelength Optical Transport Networks CSELT Technical Reports special issue 'Towards all-optical networks', June 1996, **XXIV**(3), pp. 367–382, ISSN 0393-2648.
- 19 O'MAHONY, M. J. et al. The Design of a European Optical Network. *IEEE Journal of Lightwave Technology*, 1995, **13**(5), pp. 817–829, ISSN 0733-8724.
- 20 ACTS Project AC073. METON 'METropolitan Optical Network'. Project summary, <http://intec.rug.ac.be:8080/u/horizon/projects/meton/>
- 21 MARICONDA, A.; MERLI, S. An Optical Add-Drop Multiplexer (OADM) Node Architecture in a Fully Transparent Self Healing Ring Network. ECOC '96, Oslo, 15–19 September 1996, 4, pp. 139–142.
- 22 CADEDDU, R. et al. An Optical Bidirectional Self-Healing Ring with Increased Capacity Using WDM. ECOC '96, Oslo, 15–19 September 1996, 3, pp. 257–260. (Or see http://www.csel.it/Csel/prom/index_e.htm)
- 23 BRÄNDSTRÖM, R.; and LINDBERG, P. Using Successive Smooth Approximations in Optimisation of SDH Networks. Networks 1994, paper T9.3.

A complete list of MWTN publications can be found at: <http://intec.rug.ac.be:8080/u/horizon/>

Biographies



Goff Hill
BT Networks and
Systems

Goff Hill joined BT in 1961 (then the General Post Office) as a Youth-in-Training and went on to gain an honours degree in Electrical Engineering at the University of

Newcastle-Upon-Tyne in 1969. He then moved into research, developing digital access networks and later video performance assessment techniques. His early interest in optical transmission began in 1983 when he moved to the Optical Communications Systems Division where he took a special interest in developing the concept of an optical network layer. This later became the subject of the RACE II Multi-Wavelength Transport Network project which Goff led. The project played a pioneering role in the development of multi-wavelength networks and in 1995 won the RACE Award for Technological Progress. Goff leads the Future Networks Infrastructure Group in the BT Laboratories at Martlesham Heath where his interests now include the management of WDM networks and practical evolution paths towards multi-wavelength network platforms.



Roberto Cadeddu
CSELT

Roberto Cadeddu received the B.Eng. degree in engineering Physics from McMaster University of Hamilton, Ontario, in 1989. He worked as a research assistant at the McMaster Campus Nuclear Reactor on neutron activation analysis and flux mapping of the reactor core. In 1990, he joined CSELT in Italy, where he was initially engaged in studies and simulations of coherent optical transmission systems, also as part of the European RACE I programme. He has taken part in RACE II programme, project 2028 MWTN (Multi-Wavelength Transport Network), engaging in transmission simulations of HDWDN optical communication systems and in

dimensioning and cost analysis of optical networks employing wavelength routing techniques. He is author and co-author of technical papers in the field of wavelength-division multiplexing and holds one international patent. Currently, as a senior researcher in CSELT, he is part of the All Optical Network Architectures group. His interests include the evaluation, applicability and introduction of photonics transmission techniques in the telecommunication transport network and their impact on the network architectures, dimensioning and management.



Lars Fernandez
Telia AB
Sweden

Lars Fernandez was born in Stockholm, Sweden. After receiving his M.Sc. degree in Electrical Engineering in 1981 from the Royal Institute of Technology, KTH, Stockholm, he worked on analogue optical transmission at Matsushita Electric Central Research Laboratories in Osaka, Japan, from 1981–82. During 1982–1989, while at Ericsson Telecom, he was responsible for technology studies of optical detectors and for development of digital optical receivers. Since 1991 he has been responsible for technology and systems studies of optical switching and optical networking at Telia AB. He has been responsible for the work undertaken by Telia in the RACE II project R2028 MWTN and is currently responsible for Telia contribution in the ACTS project AC073 METON. He is a member of the Swedish society of electrical and computer engineers, SER.

Neil McMillan

World Trade Organisation Talks on Basic Telecommunications

Telecommunications is already one of the fastest growing sectors of the international economy, spurred on by technology and liberalisation. The author of this article, who has recently been chairing the Group on Basic Telecommunications in the World Trade Organisation, charts the progress of international negotiations to achieve a new agreement which promises further to extend liberalisation of the world's telecommunications markets and benefit the economy as a whole.

Introduction

At around 5.15 p.m. on Saturday 15 February this year, I was able to bring down the gavel on one of the most important international negotiations in recent years on basic telecommunication services. Since May 1994 the Group on Basic Telecommunications in the World Trade Organisation (WTO) had been discussing the liberalisation of telecommunications worldwide, along with some of the main regulatory issues which flow from liberalisation.

The telecommunications industry at present accounts for some US \$660 billion turnover worldwide, and is growing faster than almost any other major sector of the economy; namely, some 10% last year. It is therefore one of the most important industries not hitherto covered by multilateral trading rules, and, without unnecessary hyperbole, I think I can safely say that the deal we did that Saturday evening will make a major impact on the global economy, not just in telecommunications but in almost every other sector. In this article I would like to give a short overview of some of the issues we faced and how we were finally able to resolve them.

At the end of the Uruguay Round in 1994, ministers from some 130 countries had agreed to set up the World Trade Organisation, covering both trade and goods, and, for the first time, services, which in the most developed economies make up about 70% of GDP. The new WTO also set up a framework for the



Neil McMillan, Director, International Communications Policy, Department of Trade and Industry

legally binding settlement of disputes between trading partners, making it possible to enforce market opening commitments made by WTO members both quickly and, more importantly, with a degree of legal certainty which was not present in earlier agreements.

When the conclusion of the GATT Uruguay Round was agreed in Marrakech in 1994, there was a number of services sectors where ministers considered that the new WTO should continue negotiation to endeavour to make further progress. Four negotiating groups were set up, with the target of completing negotiations by mid-1996. These covered financial services, free movement of people, maritime transport, and basic telecommunications. The aim was to



Delegates at the World Trade Organisation Talks on Basic Communications

bring these services fully under the General Agreement on Trade in Services (GATS), part of the new WTO's range of agreements.

First Phase of Negotiations up to April 1996

When ministers set up the Negotiating Group on Basic Telecommunications (NGBT) they set it a target of completing its work by 30 April 1996. This group had voluntary membership, and covered some 61 countries. It started work in May 1994, and it achieved a considerable amount in the 22 months over which it met. At the end of the group's work, there were 48 offers on the table, covering around 90% of the world telecommunications market in terms of turnover. It also worked up a reference paper with a set of regulatory principles covering the main elements which any new market entrant would be anxious to see handled fairly as markets are opened, including:

- an independent regulator;
- interconnection on transparent and fair terms;
- ability to obtain a licence within a reasonable time, and against a transparent set of criteria;
- some assurance that incumbent operators will be regulated to prevent abuse of any market power they enjoy; and
- where needed, fair access to scarce resources such as frequencies and numbering.

Unfortunately we were unable to reach agreement in the last days of the negotiations in April due to concerns, particularly from the US, at the level of commitments on offer. It was feared that, with the opening up of the market, operators in some countries who had made no offer or offered only partial liberalisation could exploit the situation unfairly and cause commercial damage. These concerns were particularly focused on satellite services and international telephony.

We were, however, able partially to rescue the situation by reaching agreement that we would adopt the protocol necessary to add basic telecommunications to the services covered by the GATS, and lock in each country's best offer for a further six months. We aimed during this period to try to negotiate a better overall deal, and sort out some of the outstanding problems. If we failed after that time, the protocol, and offers scheduled to it, would fall.

Negotiations between July 1996 and February 1997

We set up a new group, which I was delighted to be asked again to chair, with the task of carrying forward these negotiations against a new deadline of 15 February 1997.

Satellite and international services

The first thing we did was to focus on these two outstanding issues, and there was a series of informal meetings with interested delegations to see how far some of these difficulties could be overcome. At the same time, all delegations were urged to

improve their offers, and to encourage other countries to put forward theirs.

In October 1996 we organised an informal exchange of views between members of the group and the main consortia offering fixed and mobile satellite services to discuss their concerns on access to markets for new systems to be launched over the next few years. This meeting was an unqualified success. Both sides went away from it, I think, with a very much better understanding of each others' concerns, and we heard shortly after the meeting from the industry itself that they considered a deal covering satellite-based telecommunications services as helpful in establishing their new ventures, and that they would fully support a successful outcome in February. This positive outcome was reinforced by the ITU Policy Forum on Global Mobile Satellite services that same week.

On international services, where discussion had always been dogged by the potentially distorting effect of the existing accounting rate system, we were able to make similarly encouraging progress. Part of this was, I believe, a realisation among some of the larger industry players in the main negotiating countries that their global ambitions would be harmed considerably by another failure to agree on international services. It was also helped by a number of other parallel developments, both the very welcome increase in the number of countries making market-opening commitments in basic telecommunications, and the decision of the US Federal Communications Commission (FCC) to launch its initiative on benchmarking of accounting rates.

Although the latter did cause some nervousness among a number of the negotiators, it was clear by the end of 1996 that we had a very much better possibility of including international services in the deal than had been the case in the previous April. Now that the negotiations are over, I

can perhaps admit that, when we started in 1994, I was very uncertain as to whether any countries would agree to include international services in the agreement. This was based on the latent problems arising out of the very high profitability of these services, the distorting effects of accounting rates and the regulatory framework they must operate in (proportionate return, parallel accounting). The rapidly accelerating trend towards liberalising basic services throughout the world, which has undoubtedly become more evident over the last 2½ years, made it much easier for countries to feel relaxed about these services. I am certainly delighted that ultimately so many countries included international services in their offers with no major limitations attached.

More and better offers

With these two obstacles cleared away, the group was able to concentrate on trying to encourage as many new and improved offers as possible. Just before the WTO Ministerial Conference in Singapore in December last year, the US and EU put forward improved offers, and, during the course of the Ministerial Conference, a large number of countries, many of whom had not participated actively until then in the work of the group, came forward with promises of new offers.

Spectrum management

As is the way of such things, it was not all plain sailing after that. There were concerns on some countries' part about the treatment of radio frequencies. A large number of countries had made their offers of market opening for mobile or other services requiring radio frequencies conditional upon the availability of those frequencies. This was to meet the legitimate concern that, as radio frequencies are a limited resource, it would not be possible to license an unlimited number of operators in a particular service. As the negotiations went on, however, people

became more and more nervous that this could still leave them open to challenge by an operator wanting to enter their market where, for perfectly normal spectrum-management reasons, they had held back frequencies which otherwise might be made available for that service. There was, equally, concern that the broad wording of this limitation could be used as a hidden trade barrier by countries claiming that they had no spectrum available when this was not really the case. After a number of weeks of discussion, we finally came to a useful result. I circulated a note setting out an understanding of how frequency management could be continued under the WTO disciplines and recommending that the limitation in countries' offers should be removed. This advice was followed in almost all cases in the final schedules

the WTO to resolve issues on accounting rates, and suggesting that an MFN exemption might thus not be necessary in this area.

At the end of our negotiations we had an outcome with which I think everyone who took part in them can be pleased and proud of their contribution. Instead of the 48 countries which made offers in April, we had 69, including a large number of countries from the developing world, among them some of the fastest-growing markets in the world. Overall, the agreement now covers some 93% of world turnover in telecommunications. No less than 57 of those countries included in their schedules of commitments the full text of the reference paper on regulatory principles which, as I mentioned above, the group put together before the end of the April

*Overall, the agreement now covers
some 93% of world turnover in
telecommunications.*

of commitments which we now have on the table.

Accounting rates

A similar last minute problem arose on the question of the accounting rate system and its compatibility with the most favoured nation (MFN) requirements under the GATS. We did not seek to resolve the question of whether accounting rates came under the aegis of WTO or not. They are clearly reciprocal arrangements, but there was discussion earlier in our negotiations as to whether they were purely commercial arrangements (which are not covered by the GATS agreement) or Government measures (which are). In the end, some countries entered MFN exemptions in order to avoid any doubt on this issue, but the majority of delegations accepted a further Chairman's note I produced proposing that delegations should not use the dispute settlement procedures under

round of negotiations. Another seven countries adopted some of these principles, particularly those covering interconnection, which are central to the proper functioning of competing services and network.

Implications for the Global Telecommunications Market

The importance of telecommunications services

Telecommunications has a larger turnover and employs more people than most major manufacturing sectors in the developed world. In the UK, for instance, it employs nearly twice the number of people as our car industry. Companies worldwide spend more on telecommunications services than on oil products. It is also often forgotten that this is a technology that all industries and

commercial activities, regardless of their business, rely on to do that business quickly and efficiently. That dependence will not go away: in fact it is growing, as information systems are networked. Financial transactions, the transfer of sales figures or basic manufacturing-design information all depend more and more on telecommunications links. People are also using the Internet for day-to-day communications where once they would have sent letters or telexes. Telecommunications services are therefore not only an important economic force in their own right, but a vital input into the economic well-being of each of our countries.

Globally-traded service and liberalisation

This has in turn changed the nature of telecommunications services from a national-based utility, commonly run by state-owned monopolies, to a highly mobile, globally trading service, in which network technology makes it possible, unless prevented by regulation, to offer a service to anyone anywhere, *from* anywhere.

I think, therefore, that in reaching an agreement on 15 February, we hit exactly the right time for a global trade deal covering a globally-traded service. Regulatory barriers are already coming down all over the world. In the US, the division between local and long distance services has already begun to be removed following the passage of the 1995 Telecommunications Act. In Europe, the full liberalisation of all services at the end of this year will create a major shake-up in the way telecommunication services are offered both to business and domestic users. In Japan, foreign ownership rules have been relaxed and more competition allowed, and in many other Asian-Pacific countries new impulses towards liberalisation are emerging.

This pattern of change adds up to major opportunities for all operators. At the same time, it will also change

the way international telecommunications services are offered across the world, including in those countries which at present have no plans to open up their markets to competition. Liberalisation on the major traffic routes will provide opportunities for operators to set up their own service and network businesses in countries which they have hitherto served through the traditional correspondent agreement. These agreements at present involve accounting rates being paid when the traffic is handed over from one operator to another. Everyone accepts that these accounting rates, for better or worse, no longer bear any clear relationship to the cost of supplying this service. Equally, they

costs if it is not possible, bilaterally, to agree lower accounting rates.

This trend, and the agreement which we reached in February will, I believe, thus have an effect on the process of change in all countries. If we accept that telecommunications has a growing role in the economic competitiveness of economies overall, it will not be long before those countries most anxious to see their economies develop will take advantage of the opportunities provided for under the WTO. This correlation between telecommunications and economic competitiveness was brought home to me recently by a study undertaken with a representative cross-section of companies in a wide range of manufacturing and

*we hit exactly the right time for a
global trade deal covering a
globally-traded service*

represent an important source of revenue for many developing countries to build out their networks.

The agreement and its impact on change

Our agreement and the general trend towards market opening will speed up the process of change. On routes where liberalisation is introduced at both ends carriers will have a choice of terminating traffic on their facilities or negotiating commercial arrangements based simply on local interconnection and the opportunity cost of building facilities for themselves. Competition will inevitably also drive down the traditionally high margins enjoyed on international call tariffs, creating new pressures on operators in liberalised markets to seek lower accounting rates even on those routes where these are retained. I believe that this, indeed, will become a commercial imperative for operators in those markets, and force them in some cases to seek unilateral action, to bring down their

commercial activities. After political stability and a trained workforce, the third most important factor cited by the majority of these companies as determining where to locate a new foreign investment was not roads, or electricity supply or airports, but the availability of a high-quality and reliable telecommunications network. A number of countries, not least in Africa, took this to heart, and made offers in the last few weeks of our negotiations with the explicit rationale of seeking to encourage foreign investment in their telecommunications network.

Competition for investment

This made a lot of sense. Unlike many investments, telecommunications requires a large upfront investment before any cash flow comes on stream. It involves sunk costs which are really sunk: if you change your mind as a telecommunications operator, you cannot dig up the network and move it somewhere else. Thus, investors need to be sure that those costs can be recovered

against a background of fair treatment and predictable regulation. In a market which will find more and more countries seeking investment in telecommunications infrastructure over and above what their own capital resources can provide, it takes little imagination to foresee a shortage of, and thus competition for, investment capital for telecommunications. We have, for instance, already seen privatisations worldwide having to be phased to allow the capital market to absorb them.

On a projection of current investment plans in developed economies, the World Bank has estimated that the developing countries will need to find an extra US \$60 billion annually just to keep the present differential of telecommunications development the same, let alone catch up. With predictions from the International Institute of Economics of additional growth in telecommunications turnover, investment and consumer benefit of some US \$1000 billion over the next 10 years as a result of our agreement, this may even seem conservative. I do not believe it at all fanciful to predict that many more countries, anxious to improve their economic competitiveness, will in future seek to give the maximum assurance to potential investors in telecommunications. If they make these commitments in the context of the WTO, where an aggrieved investor can seek concrete redress if their expectation of fair treatment is disappointed, they are likely to secure more of what may become scarce investment resources.

Major benefit to world economy

I therefore believe that the 69 countries who did make commitments in February, and who are now working to implement those commitments in time for the ratification deadline of 30 November this year, have taken a step which will benefit their economies and with it the world economy as a whole. A number of countries who followed the negotiations, but

were unable to submit offers before February, may still wish to add their names to the Agreement before or possibly just after it comes into force on 1 January 1998. I believe that it will also help in giving a clear framework for countries at present negotiating accession to the WTO of the sort of issues which they will wish to address when looking at commitments on telecommunications.

In the UK we have seen over the last 13 years the benefits that liberalisation has brought to consumers and actually also to operators. This makes me sure that many more countries will over the next few years see that joining a worldwide trend in telecommunications will bring them the same economic advantages.

Biography

Neil McMillan has been Director, International Communications Policy in the UK Department of Trade and Industry since 1991. From May 1994–February 1997, he was Chairman of the WTO Group on Basic Telecommunications, which negotiated the agreement on basic telecommunications completed in February in Geneva. He was Chairman of the European Committee on Telecommunications Regulatory Affairs in the CEPT between 1992 and 1995, and set up the European Telecommunications Office in Copenhagen in 1993. After studying German and French at the University of Exeter in the UK and at the Universities of Kiel and Regensburg in Germany, he joined the British public service, and has dealt with the steel industry, GATT textile negotiations and consumer policy. Between 1982 and 1985, he was Private Secretary to the Minister for Technology and Industry, and was closely involved with the privatisation of British Telecommunications and liberalisation of the UK telecommunications market in 1984. In 1985, he was seconded to the German Federal Ministry of Research to advise on telecommunications reform in Germany, and served four years from 1987–1991 in the UK Permanent Representation to the EU in Brussels, dealing with technology and telecommunications policy.

Operations and Maintenance Centre: Delivering Network Services

Network management, administration and maintenance activities have developed dramatically since the introduction of digital exchanges, with their ever-increasing capabilities for remote interrogation and control. The operations and maintenance centre (OMC) was developed to exploit the exchange capabilities and fulfil the need for fast and efficient centralised administration and maintenance. This article describes the development of the OMC, outlines its key role and achievements in the delivery of network services in pace with ongoing requirements, and investigates the technical challenges being addressed as part of its evolutionary design.

Introduction

The operations and maintenance centre (OMC) has been a key component in BT's network management portfolio for over 10 years and is at the heart of the network modernisation programme. It is a centralised operations and maintenance support system providing powerful facilities for the management of both local and trunk switched networks. OMC systems are deployed in nine special-purpose computer centres throughout the UK and are operated by users at 10 command centres, called *network operations units* (NOUs). Each OMC provides a single-point interface between its users (for example, exchange maintenance staff and

administrators), external systems (for example, customer service system (CSS), network traffic management system (NTMS), CallMinder and the intelligent network (IN)) and the public switched telephone network (PSTN) digital exchanges within its catchment area.

As well as providing ever more efficient and fast mechanisms for network administration and maintenance, the continued development of the OMC is key to enabling the reliable provision of new services and products in the shortest possible timescales. The ability to link to various external systems, and the automation of those links where possible, are key enablers in this (see Figure 1).

Figure 1 – The OMC: delivering network services




Figure 2—OMC attributes and statistics


Thus the OMC plays an essential part in providing BT's customers with easy service provision or modification, and ensures the subsequent reliability of those services. Within BT, the OMC provides an extremely efficient means of network management on a day-to-day basis, and assists in the fast deployment of new services and features developed by other major BT development programmes.

The business benefits outlined above are possible because of the in-built capabilities of the system and the business processes used to implement frequent upgrades. These attributes may be simply expressed in terms of power, quality and flexibility, as shown in Figure 2.

The power of the OMC lies in its exploitation of advanced technology and cutting edge software, together with the provision of extensive fast relational-database storage and communications facilities. The quality of the system is maintained by applying stringent security requirements and by ensuring reliability through fallback provision and thorough testing of all new and modified processes. The flexibility of the OMC results from its modular design and the business processes used to control and implement all upgrades. The design is such that most changes can be data driven, while careful business planning, release control, and a fast-track process for urgent changes, ensure that all business requirements are met in a timely manner.

Delivering Interface Capabilities

The OMC provides a single point interface between digital exchanges of the network, external systems and users that enables service provision and network management. In addition to over 24 types of OMC user, there are currently 14 external systems that utilise the facilities provided by the OMC. Figure 3 shows some of these systems in relation to the OMC.




A major benefit of the OMC is that it provides simplified and consistent *service perspective* facilities for its users and external systems, while performing complex *network view* control-and-monitoring functions via man-machine language (MML) communications with the digital exchanges and switches. Complex communications, processing and data management functions are performed by the OMC in order to fulfil this role.

Translation processes are necessary to convert between the various exchange MMLs and the users' service views (these being independent of the exchange types and releases). The translation processes are complemented by extensive

distribution and scheduling facilities. By this means, the multiple interactions necessary, for example, to provide a customer with a new service such as CallMinder, can be performed 'behind the scenes' after only one initiating request from the CSS or an OMC user. Requests can also be processed in batches and at off-peak periods to further improve efficiency. The operational efficiencies provided by the above interfacing functions alone repay investment in the OMC many times over.

When the OMC is required to make a new service available to customers, it is often only necessary to add or modify a translation component; the distribution and scheduling being reusable together

Figure 3—OMC interfaces



with all the other background administration and security facilities of the OMC. Thus implementation costs, including associated user retraining, are kept to a minimum. Variants of the OMC such as the service control point (SCP) manager and reusable switch manager (RSM)¹ are further examples of the reuse of OMC components, as described later.

Network perspective interfaces

The OMC system supports the two major types of digital local exchange (DLE) in the UK core telephone network: System X^{2,3} and AXE10. All System X exchange variants, the AXE10 trunk (AS36), and two types of NorTel switch (DMS100 and DMS250), are also supported. Support is also given from the OMC to the following:

- *CallMinder*—this is a call answering service (formerly called *NBCAS*) which allows the diversion of calls when lines are busy and provides a message-taking facility via interactive speech application platforms (ISAPs) connected to trunk exchanges throughout the UK. The ISAPs are under OMC control. The OMC is required to configure both the ISAP and the exchange in order to provide the *CallMinder* service for a customer.
- *Intelligent network* (IN)—this is an overlay configurable network above the existing network, designed to enable fast deployment of new customer services. It moves services and features, that would otherwise be provided on PSTN digital switches, on to a separate real-time computer system called the *service control point* (SCP). The OMC SCP manager governs deployment of new services on the SCP.
- *Cashless and FeatureNet*—these systems use a variant of the OMC to manage the digital switches they use to route appropriate calls

from within these overlay networks on the PSTN.

- *London local fibre network* (LLFN)—this offers key BT customers in the London area flexible services over pre-provided optical fibre. A planned LLFN-OMC integration involves merging LLFN software onto an OMC to improve performance of the network and to enable LLFN functionality to be made available nationwide when installed as part of a future OMC build.

Service perspective interfaces

The following external systems use OMC functionality for service provision and to remotely control and monitor the UK network:

- *CSS*—this handles customer requests for services and requires access to the OMC to implement these services in the exchanges. Currently, 70% of these transactions are automatically processed by the OMC, with a further 20% planned to be automated within the next two years. Because of the importance of the CSS-OMC link, the multi-protocol router network (MPRN) and other resilience measures have been used to ensure its maximum availability for the automated transactions.
- *Network traffic management system* (NTMS)—this is a centralised network traffic control system, based at Oswestry. The NTMS can issue a set of high-level control instructions for any number of selected PSTN exchanges to prevent, for example, congestion on a part of the network. It uses an automatic link with the OMC to effect these controls quickly and efficiently.

Further external systems include:

- transmission domain surveillance system (TDSS),

- problem management application (PMA), and
- access network fault manager (AFM).

Delivering Solutions—System Evolution

A key aspect of the success of the OMC is that it has consistently delivered solutions to meet the needs of BT's networks. The following paragraphs demonstrate this aspect in the history and present developments of the OMC, and describe the mechanisms by which these solutions continue to be delivered in line with ever-increasing demands on technology and timescales.

Hardware and software evolution

Early developments of the OMC resulted in the installation of the first *OMC2*, operating on a DEC VAX VMS platform, in 1985. It has since been continually upgraded to satisfy the need for additional processing power; to improve performance and to support the ever-increasing variety and power of digital exchanges and other external systems. Table 1 summarises the main functionality associated with each *OMC2* software build from 1992 to the present day:

The OMC currently operates on DEC VAX computers running the VMS 5.5 operating system. They interface with a UNIX Data Server (UDS) comprising a Sequent processor and hosting an Oracle database. Front-end processors (FEPs) handle external communications to the exchanges, external systems, user terminals and printers. These standalone processors are connected to the main VAX processors via an Ethernet.

Each OMC is tailored to meet the requirements of individual installations. This is achieved via the VAX hardware options of processor power, peripherals, cluster architecture and communication interfaces. Plans are in place to upgrade to an Alpha

Table 1 OMC Software Builds Since 1992

Build	Description	Date
K1	OMC2/CSS interface Tariff group charging	2/92
K2	AXE10 Phase 3C	6/92
L	Migration of OMC2 SRS to the UNIX platform	1/93
M	OMC2 support for FAN Maintenance reorganisation for network administration improvement plan Support for national code change ⁴	4/93
N	Upgrade to VMS 5.5 and front-end processor replacement	7/93
O	Support for Billing 90s (elapsed time charging) Support for CLASS network services OMC security enhancement	11/93
P	OMC2-NTMS interface AXE10-DXE4 DN re-parenting	4/94
Q	OMC2 support for CLASS services Support for NBCAS Support for Centrex working on System X	8/94
R	Billing 90s Phase 2 (account code services) OMC2-CSS interface enhancement Enhancement of IN services support	12/94
S	Support for ISDN DSS1 services Migration of CSS-OMC2 links to MPRN Provision of communications gateway for Cashless, FeatureNet	2/95
T	Phase 2 for FeatureLine on System X and enhancements to CLASS network services	5/95
U	Feature interworking (introduced to support marketing of network services)	8/95
V	Number portability development	11/95
W	FeatureLine on AXE10 OMC-CSS link improvements	3/96
X	ETSI(DSS1) SCP Manager Phase 1	7/96

platform in the near future, increasing processing capacity even further to meet operational needs.

System software

OMC application software is predominantly written in the C programming language and makes full use of sophisticated facilities inherent in the VAX/VMS operating system. System administration software is predominantly written in DEC/VAX DCL programming language. Database services are pro-

vided by the VAX utilities and layered products such as the ORACLE relational database management system on the UDS. Front-end software provides on-screen menus and forms for the OMC users who manage the network.

Upgrades and reuse of OMC software

The OMC is a full client-server system which is data driven such that most upgrades to the interfaces and facilities it is required to deliver

can be implemented with considerable efficiency and reliability. A component-based software architecture has been implemented whereby files that were previously discrete have been combined into functional components that can be assembled as required for a particular application. This development has enabled the reuse of OMC system components for applications apart from the PSTN. Examples where the OMC2 is or has been used as the basis for cost-effective development within related BT programmes include:

- *SCP manager*—this is used to control IN services in response to IN service requests fed to it. The SCP manager reuses OMC functionality but uses TCP/IP local area network (LAN) connections rather than the FEP.
- *Reusable switch manager (RSM)*—this utilises code from OMC2 Build S in a stripped-down form and interfaces with three new types of switch, the NorTel DMS100 and DMS250, and the Ericsson AXE10 AS36 variant. RSM control of these switch types has enabled provision of the Cashless and FeatureNet services.
- *LLFN-OMC integration*—the integration of LLFN functionality on to an OMC will give the LLFN manager platform a standardised upgrade path for future enhancements.
- *Transmission network surveillance (TNS)*—this system was a very early reuse of the OMC and has had many components rewritten over time. The system monitors all the problems with transmission equipment in the network.

The OMC enhancement process

The speed of change sweeping the UK telecommunications market makes it impossible to give a definitive statement of requirements over a

Figure 4—National OMC transactions volumes (April 1993–March 1996)

long period. The requirements up to a year ahead can be mapped out with a fairly solid baseline, enhanced with smaller developments as and when the business decides that these can be accommodated. Beyond a year, some business needs can be determined, but the growth in risk and uncertainty runs parallel to the explosive growth in competition and the increasingly rapid introduction of new and converging technologies.

Planning and delivering new releases


The OMC2 programme usually has four releases underway at any one time. Each release is typically eight months from its definition to national roll-out completion, with a new release every 10–12 weeks!

The OMC programme has a three-year quality plan and budget (QPB), which provides the approval envelope for development, support and infrastructure. This QPB identifies major programmes of work for the OMC each year; for example, Billing 90s, IN, Corniche, FeatureLine and CLASS. The enhancements on the programme are driven by facility requests (FRs) which specify the business opportunity or problem. These are reviewed frequently via a prioritised workstack, and all high-priority FRs are referenced by a client requirements definition (CRD) for the release.

When the requested enhancements are agreed, the associated FRs are specified and designed. These, often complex, changes are carefully engineered and tested before an operational pilot can commence. To ensure BT's operations continue without any adverse impact, two in-service pilots are undertaken of each new release. If both are successful, then national roll-out commences.

Upgrade performance and continuing service provision

A key measurement of success in the provision of new OMC functionality is that the software releases have consistently been delivered to the



time agreed. Service target statistics are also regularly produced to ensure performance standards are met, with additional statistics produced following each major release (see Figure 4).

Mechanisms exist to address problems quickly and efficiently. Each problem is given a unique reference number and is owned by a support engineer. A fix can be delivered to all sites by an emergency release procedure or, if that is not immediately possible, helpdesks are advised of a workaround or a mandatory notice is issued.

The whole OMC programme is complemented by a support organisation to ensure effective resolution of in-service problems should they arise. This cover is provided 24 hours a day, every day of the year.

Improving the OMC infrastructure

The main emphasis of the OMC programme has been to deliver new functionality as quickly as possible to meet essential business needs. Aggressive delivery schedules have been used to meet this need, placing demands on the development, operational and user communities which they have continually met.

Since it is also important to address improving processes, working practices and the underlying quality of the core delivered products and services which comprise the OMC infrastructure, an OMC2 Improvement Programme has been implemented. This is an ongoing programme of work which proactively addresses areas that cause pressure in either the development or operational areas.

Fast track

To further add flexibility to the OMC programme, a small enhancements process is run called *fast track*. This aims at introducing new functionality on the OMC in 30, 60 or 90 days from requirement to delivery.

With the growing demand for a rapid response to requests for network management solutions⁵, the OMC fast-track process makes it possible to implement new self-contained developments quickly and cost-effectively over the whole network without sacrificing the traditional high quality of engineering. This process represents a radical approach to the development and integration of network management solutions. It exploits best practices from OSS experience while allowing a rapid applications development approach to reduce significantly the time taken to deliver solutions to the market.

When an OMC fast-track request for change (RFC) is approved, it is because it will realise commercial benefits for BT's network operations, or a significant business advantage is to be gained by implementing the development on an in-service release. The process features a non-disruptive software life cycle commensurate with the small, but critical, amount of development undertaken. Although the functionality is delivered quickly, it is important that the released software is not of a lower quality than the mainstream released software. The processes associated with the development, testing, integration and installation of BT's normal high-quality of software are therefore maintained throughout the process.

A fast-track development is targeted at a mainstream OMC release but can be delivered earlier, on top of one or more in-service releases. This approach allows the OMC to treat the development as an in-service patch fix, using the infrastructure and procedures which already exist. After a fast-track development is delivered to the field, it is incorporated into the next mainstream release to avoid rolling it out as a separate delivery.

OMC fast-track achievements to date include the following developments:

- enhanced pure mode, whereby exchange MML can be sent to a large number of exchanges automatically;
- resources out-of-service (OOS) analysis, required to investigate how much equipment goes out of service each day and whether it is back in service the following day;
- improved retrieval routines for test and diagnostic procedures on digital exchanges;
- assistance in the roll-out of the CallMinder service in London via an automatic customer check; and
- improved AXE10 facilities, including automatic line testing.

Delivering New Services

With the constantly changing telecommunications market place, the OMC needs to provide facilities appropriate to ensuring that market needs are met in the following key areas:

- *Customer choice*—BT must be able to offer best value for money, a strong brand image and the ability to provide requisite products at the required time.
- *Flexibility*—BT must be able to respond to the increasing pace of

change if it is not to miss opportunities for competitive advantage.

- *Quality of service*—The service customers receive from telecommunications suppliers is a key differentiator. BT quality targets require that systems support delivery against commitments—on time, every time—with all that implies for data quality and issue escalation.

Major project areas currently concerned with the delivery of services to customers include:

- CSS-OMC link—increasing automatic functionality;
- ISDN;
- CLASS;
- FeatureLine;
- IN developments; and
- regulatory work; for example, addressing OFTEL requirements for number portability.

CSS-OMC link

CSS is the primary interface with customers for order entry and service provision, and in turn passes these on to the OMC to implement the service changes requested.

The OMC is continually evolving to provide more automated transactions for CSS. A fully-automated CSS-OMC link is being implemented in phases. Since the physical CSS-OMC link was established in 1992, developments have continued to increase its scope and functionality so that within two years the link is planned to provide automatic handling of approximately 90% of all transactions. Automatic transactions (current or planned) include:

- line provision and cessation, and start/stop of orders for single lines and PBX;

- provision of line-based features;
- temporary and restore out of service (TOS and ROS) for billing purposes;
- provision of ISDN2 and ISDN30 services;
- provision of Select Services; for example, call diversion, call barring, three-party call facilities, CLASS facilities;
- addition and removal of lines for FeatureLine business groups;
- provision of advanced services; for example, account code and CallMinder;
- handling of customer supplied data; for example, a call-diversion number; and
- support for payphones.

The scope of the CSS-OMC link is also expanding to include facilities such as:

- network intelligence infrastructure programme—the CSS-OMC interface is to be enhanced to allow the provision of new IN services, such as account code, in greatly reduced timescales;
- tactical billing outgoing call barring (OCB)—this permits OCB on those lines where payment is overdue while allowing the customer to retain incoming calls (this tactical approach increases the likelihood of retaining customers when the account is settled);
- support for new technology—this will provide automated provisioning on new network systems such as TPON, AVMUX, FAN, LA30, one-per-customer radio and HDSL;
- advanced PSTN services—this is aimed at the business market and

allows regional directory numbers (DNs) to be retained or allocated for businesses outside a region;

- feature interworking—this provides support for the marketing of network services packages containing CallMinder, call waiting, call diversion and call sign (formerly DRN) services;
- elapsed time charging; and
- call management information (CMI)—this provides customers with reports on the performance of calls, such as the proportion of successful/unsuccessful calls, reasons for failure, time to answer incoming calls and most commonly dialled numbers.

OMC and ISDN services

The OMC has always played a key role in the setting up and management of ISDN services (both ISDN2 144 kbit/s via 2-wire links, and ISDN30 2 Mbit/s links via screen cable or optical fibre) for customers on the main network. These ISDN services are currently being aligned with the European Telecommunications Standards Institute (ETSI) Digital Signalling System 1 (DSS1) with the result that additional areas of OMC support are being provided, including the following:

- test and diagnosis of faults—the OMC allows the initiation of line tests and the provision of subsequent fault diagnosis for all services including ISDN; additionally, a warning message is provided if the specified test will result in disruption of service to the customer (tests include error rate checking, line controller testing for ISDN30 and cable pair identification tone for ISDN2);
- collated fault reporting—the OMC collects fault reports relating to the ISDN services and makes them available in various presentations;

- provision and reservation of ISDN2 and ISDN30—supported via manual or auto allocation of equipment numbers;
- adding, altering and ceasing services—the OMC allows a new ISDN service to be allocated and an existing ISDN service to be modified or removed; and
- additional supplementary services for ISDN2 including D-channel packet handling and maintenance closed-user groups.

OMC and CLASS

CLASS is a range of advanced network services made possible through the development of calling line identity (CLI) facilities. CLASS forms part of the Select Services marketing programme, as shown in Figure 5, and is expected to contribute significantly to BT's revenue.

By taking advantage of the ability to record CLI information, CLASS takes Select Services a step further: Customers are provided greater control over their telephone contacts, helping them to meet the require-

ments of a competitive business world. CLASS services are listed in Table 2.

The OMC plays a central role in the provision of CLASS services by providing the platform through which these services are configured on the network. Implementation of CLASS services can be achieved via manual, automatic, and bulk-driven interfaces.

The OMC also collects statistics identifying equipment and DN's capable of supporting CLI, together with an indication of spare capacity. This information is obtained from an OMC database for System X, and from the exchanges for AXE10. OMC-collected statistics can also be used to identify which DN's have specified CLASS services.

Since the initial implementation of CLASS services, development within the OMC is focused on increasing the consistency of System X and AXE10 service levels, and on providing an expanded customer base for CLASS services across other line types and interfaces. Beginning with Build X, for example, CLASS services will become available to FeatureLine customers on System X exchanges.

Figure 5—BT's Select Services




Table 2 CLASS Services

CRN	Call return
CAD	Caller display
CIW	Calling line identity withheld
ACR	Anonymous call rejection
COT	Customer originated trace
RBWF	Ring back when free
RBWFI	Ring back when free inhibit
PDN	Presentation delivery number
DRN	Call sign (formerly distinctive ringing)

OMC and FeatureLine services

FeatureLine is the BT product name for Centrex facilities offered to small-to-medium business customers. FeatureLine provides services directly via digital exchanges, ensuring that users have access to state-of-the-art technology with each OMC upgrade. The product is supported on all System X and AXE10 exchanges, offering customers easy entry to PBX functionality without the up-front costs of a PBX switch.

FeatureLine offers a service package similar to that available on a conventional PBX, allowing up to 60 lines to be grouped together to form a cohesive business group (BG). The contributing lines in the group—or *extensions*—can include up to 10 multi-line groups or PBX lines. All

Figure 6—FeatureLine private network numbering facilities

lines share a common numbering plan.

Customers within a BG can call each other by dialling a two-digit number from any other extension within the group. Each extension is also provided with a full DN which can be used to call the extension directly, from within or outside the group.

FeatureLine customers are also offered a wide menu of service options, including private network numbering, hunt group working, and much more.


A private network numbering plan, for example, facilitates calls across a customer's BGs located across any System X exchange. Under this plan, customers dial an *inter-site access code*—usually 8—followed by the two-digit *site code* of the required BG and then the extension number (see Figure 6). Calls between BGs can still be made via the PSTN by first dialling 9—the PSTN access code.

The OMC operates as a central point for FeatureLine development, installation and product support. The installation office duty provides CSS with the facilities to install and support FeatureLine services. As each line is added to a FeatureLine group, all the existing facilities are removed and the selected package and options applied. The OMC also provides support for feature changes for DNs and extensions, separately for individuals, BGs, and hunt groups.

Ongoing plans call for continued OMC involvement in design, implementation and testing.

OMC and IN advanced services

The IN fulfils the requirement to implement any new service quickly and cost-effectively, over the whole network, without detracting from the operation of the existing network. The contributions to IN developments of the OMC and the SCP manager, which reuses OMC functionality, have been, and continue to be, extensive.



The IN forms an overlay infrastructure above the existing network that intercepts and interprets calls for advanced services and instructs the main network in their implementation. Concentrating the control of advanced services in the overlay network provides for ease of implementation and maintenance. For the most part the main network continues with its normal operations but provides the new services under intelligent control from the IN overlay.

The basis of IN operations lies in the separation of the logic of a service from the switching system (see Figure 7). IN moves the services and features running on a digital switch and places them in a separate real-time computer system called a *service control point* (SCP). In the IN environment, the switch asks the SCP what treatment to give a call, and the SCP provides instructions on what is to be done; for example, to re-route the call or play an announcement. This moves the intelligence out

of the exchanges, speeding up and reducing the cost of providing new services.


The digital switches have required considerable phased-in modification to allow communication with the SCPs. A switch with the capability of communicating with the SCPs is called a *service switching point* (SSP). SSP capability now resides within each of BT's digital main switching units (DMSUs).

The OMC and IN developments

Early groundwork for the IN included OMC-related developments in two areas which are not strictly IN services:

- *The CallMinder service*—this utilises an early 'intelligent peripheral' the ISAP, the development of which set the groundwork for a hardware platform that is reusable with different functionality, the SCP. Continuing ISAP development extends support to packages containing CallMinder in combination with previously incompatible switch-based services such as call waiting and call diversion. Future transfer of call control to the IN's SCPs will support the national launch of a full portfolio of feature interworking packages.
- *The advanced PSTN (A/PSTN)*—this allows businesses to maintain a local point of contact for their customers. The customer dials the local number and is re-routed, via ordinary digital exchange facilities, on to the digital derived services network (DDSN) for handling. The implementation of A/PSTN involved the configura-

Figure 7—IN components



tion of BT's digital exchanges for call re-routing and allows re-routing to the DMSUs with SSP capability.

The OMC SCP manager

The SCP manager is deployed as an application reusing OMC functionality and provides a platform for further support of IN services. The SCP manager supports the centralised administration of customer service data and holds a master copy of all service application data. As data is received by the SCP manager, it is posted down to each SCP, ensuring integrity across the network and preventing inconsistencies.

The SCP manager governs the deployment of new services on the SCPs, initially supporting the national roll-out of a new account codes service. With this service, the customer dials an outgoing number, then enters a billing code which appears on an itemised bill along with the listed call. This allows, for example, a solicitor's office to apply outgoing telephone charges to individual accounts, or a household to identify calls made by individual family members.

Associated with the SCP manager is the fast track national activity management system (NAMS) which manually interfaces with CSS to extract orders, populate a bulk order interface file, and transfer this data on to the SCP manager. In future, customer data transfers will rely increasingly on automatic CSS-SCP manager links.

OMC and number portability

With the 'opening-up' of the telecommunications market, business and domestic users have been provided with a much greater choice of supplier. This development has led to an OFTEL requirement that customers should be able to transfer between network operators, and retain their existing telephone number(s). BT has responded swiftly to this demand by introducing number portability, a facility allowing customers this option

when transferring to an other licensed operator (OLO) operating in the same locality, and to continue using that number if they return to BT. It also permits customers who transfer from an OLO service to BT to use their existing OLO telephone numbers. This allows for greater flexibility in the transfer of telephone numbers between operators, and has other advantages, including:

- the provision of a sound platform for winning new customers from OLOs, enabling the new service to be set up quickly and efficiently;
- the ability to retain a presence at an ex-BT customer's site, allowing the customer to transfer back to BT with minimum disruption; and
- presenting the customer with the appearance of a seamless interface between BT and OLOs when transferring service.

Further planned developments of number portability include the automatic provision of *managed migration* via the CSS, and enabling transfers for multi-line customers and single-line PBX directory numbers (DNs).

Delivering Improved Support—Real-Time Processing

Maintenance support facilities provided by the OMC contribute significantly to the efficient running of the network and consequently to customer satisfaction. Developments involve real-time interactions and processing to speed the resolution of faults by streamlining fault analysis, work assignment and problem resolution⁶. Examples are:

- the real-time alarm toolkit, and
- the black-spot analysis tool.

Real-time alarm toolkit

The real-time alarm toolkit was developed to provide increased

operational support for maintenance personnel using the reusable switch manager (RSM) to monitor FeatureNet. Two fast-track RFCs were implemented to provide this real-time alarm system, and because of its success on the RSM and FeatureNet, an extension of its use to all OMCs is under way.

The real-time alarm toolkit displays switch-originated FeatureNet alarms at the NOU. The system is provided via a PC which has access to the host RSM, and offers the following features:

- a single logon at the PC provides direct access to the RSM and all appropriate exchanges;
- real-time alarms are listed in a scrolling window (the alarm bar) colour-coded based on priority—alarms can be expanded to give full report details;
- users can mark individual reports as 'assigned' and 'unassigned', allowing all other users to monitor current report assignments across a group of exchanges; and
- new alarm information, task assignments, and clear reports are reflected back to all active PCs as they are created, in real time.

Within each RSM session, multiple windows can be opened in an enhanced transparent mode display (as shown in Figure 8). Individual screens and windows can be cut, pasted, and printed. All activity within the real-time alarm toolkit provides direct commands to the exchange itself, bypassing forms and menus.

Through a single logon, each user is assigned an individual profile of exchanges for which alarms are displayed. Users can add and delete exchanges from their list.

The alarm bar

As each user logs on, the system displays the last 1000 records for the exchanges being monitored on the alarm bar. New alarms are added to




Figure 8—Real-time alarm toolkit

the bottom of this listing, in real time, as they are generated.

The alarms listed in this window are colour-coded as to severity: critical, major, or minor. Colour coding further defines the alarms as they are assigned and unassigned, through use of shading. Colours are fully configurable by individual PC users, as are the fonts used. Users can also individually set whether an alarm bell will sound as new alarms are added, separately for critical, major, and minor alarms, or when the status of an alarm changes. Users also decide individually whether to show alarms assigned to users other than themselves.

The future in real time

The real-time alarm toolkit, developed expressly for RSM FeatureNet alarms, is now being considered for use with OMC core PSTN alarms and the Cashless overlay network. The prototype software has provided a baseline understanding of requirements and features, opening the door to future benefits in BT operations and service provision:

- reduced response time,
- easily coordinated exchange information, and
- greatly facilitated workload assignment and monitoring.

The expectation is to simplify fault analysis and consolidate work across

exchanges, and thus help to set the groundwork for servicing projected customer levels and further strengthen BT's position in an increasingly competitive business environment.

Blackspot analysis tool

Another PC-based tool being developed for use with the OMC is the blackspot analysis tool (BAT). This powerful tool uses standard Microsoft packages such as Access and Visual Basic to enable automatic analysis of fault trends and the presentation of results in a wide variety of formats. Threshold levels can be set to filter unwanted information.

Conclusions

The OMC is a key component in BT's ongoing network development. This article has shown how, from its early days when it 'unlocked' the potential within the then new digital exchanges for centralised management and maintenance, it has been continually upgraded to provide a most powerful and flexible tool with ever increasing capabilities applicable both within and beyond the confines of the PSTN. The technical challenges for the future are numerous as the competition to supply more and more customer services within ever shorter time frames becomes increasingly intense. However, the OMC system, the development team, and the development processes in place to support

enhancements continue to prove themselves capable of the challenge.

Since the OMC provides such a good foundation for further development, it can look forward to a long life as it is used to achieve the following objectives:

- to advance BT's market position by beating the competition on services,
- to be a prime key in down-streaming Breakout,
- to automate more processes and thereby increase BT's productivity,
- to provide a reusable chassis to reduce the number of systems engineered, and
- to evolve the core network toward INs.

References

- 1 AMBLER, H. G.; SCOBIE, C. H.; and BALDWIN, V. W. Switch Management: the Operations and Maintenance Centre. *Br. Telecommun. Eng.*, Oct. 1991, **10**, pp. 187-193.
- 2 WALSH, A. A. Architecture of System X, Part 4—The Local Administration Centre. *Post Off. Electr. Eng. J.*, April 1980, **73** p. 36.
- 3 STRICKLAND, L. F.; and HEWITT, M. A. New Operations and Maintenance Centres for Second Generation System X Exchanges. *Br. Telecommun. Eng.*, Jan. 1985, **3**, p. 286.
- 4 Network Administration for the 1990s. *ibid.*, Oct. 1990, **9** (special issue).
- 5 HELLEUR, R. J.; and MILWAY, N. R. P. Network Management Systems: Introductory Overview. *ibid.*, Oct. 1991, **10**, p. 168.
- 6 GARWOOD, G. J.; and ROBINSON, A. C. Work Management Systems. *ibid.*, Oct. 1991, **10**, p. 204.

Biographies



Andrew Dawson-Maddocks
BT Networks and
Systems

Andrew Dawson-Maddocks joined BT Liverpool Area in 1981 as an apprentice and on successful completion of his apprenticeship in 1984 concentrated on exchange maintenance. In 1989, he graduated from Essex University, achieving a B.Sc. in Electronic Engineering, specialising in computers, micro-processors and VLSI design. He moved to BT Laboratories and joined the OMC development unit, initially involved in verification, validation and testing. In 1991, he moved into the senior technical design role on the OMC. After returning from secondment with the intelligent network design team, in 1992 he was appointed project manager of OMC. Since 1994, he has had the programme management responsibilities for all of BT's switch management systems developments.



David Cooper
BT Networks and
Systems

David Cooper joined BT in 1981 after receiving a B.Sc. and Ph.D. in Physics from Liverpool University. He initially worked on the design and fabrication of single-mode optical fibres. He went on to work for several years on the design and fabrication of semiconductor junction lasers and optical amplifiers. He was involved in the transfer of this technology to the joint venture with DuPont (BT&D, now HP). Since 1982, he has been leading teams of software engineers developing network management software for BT's network. He currently leads a unit working on the OMC software controlling BT's switches, fault localisation and testing of transmission systems.



Colin Scobie
BT Networks and
Systems

Colin Scobie joined BT Fulcrum Communications, in 1981, working on materials requirements planning and production control systems. In 1987, he moved to Ipswich to join a project implementing network management systems for the Government telephone network. In 1990, he joined Development and Procurement and has worked on OMC and related switch management developments for the last six years. He is currently the switch management development programme manager.

Graham Walker

The Mirror—Reflections on Inhabited TV

Inhabited TV is a vision of future television services in which multi-user virtual environments deliver unprecedented levels of audience participation. Social chat and interaction are mixed with professional content and programming to create on-line communities. The Mirror was a ground-breaking collaborative experiment in inhabited TV, created by BT, Sony, Illuminations and the BBC. Six on-line worlds were available to over 2000 viewers of the BBC2 series 'The Net' in January and February 1997, and this article provides an overview of the project.

Introduction

Imagine combining the proven pulling power of professional broadcast television with the enduring appeal of audience chat and participation, and you have a vision of *inhabited TV*. The producer defines a sophisticated audio-visual framework, but it is the audience interaction and participation which brings it to life. Professional content mixes with social conversation in a rich graphical environment. A community develops around celebrity characters, staged events and unscripted encounters.

The Mirror was an early experiment in inhabited TV, which involved 2000 viewers of the BBC2 multimedia magazine series 'The Net'. A research project created by BT, the BBC, Sony and Illuminations, The Mirror comprised six multi-user on-line worlds which reflected the broadcast material. The worlds were launched on 13 January 1997 with the broadcast of the first programme, and closed after seven weeks with an *End of the World* party. This article expands on the background to the inhabited TV vision and discusses some of the experiences, data and anecdotes from The Mirror.

Shared Spaces

The inhabited TV vision is part of a wider belief in the importance of multi-user virtual environments, or *shared spaces*, as a new communications medium¹. People are represented in a three-dimensional environment by characters or *avatars*, and can move around, converse and interact in a common context of information and applica-

tions. Shared spaces provide a flexible interface, combining synchronous dialogue with asynchronous messaging, and human presence with abstract information in a unified communications environment. As the technology evolves, shared spaces will support applications ranging from on-line commerce to role-playing games and business conferencing.

The starting point of an inhabited TV 'programme' is a professionally authored framework analogous to the programme structure of traditional broadcast TV. The framework defines spatial and temporal structures for both a persistent on-line community and a range of *special events* which are played out against an evolving backdrop. The audience or *citizens* in inhabited TV are no longer passive couch potatoes, but can choose an appropriate level of involvement in the life of the community and are able to play an active role in the special events. Moreover, the potential for worlds with completely different physical and social rules opens up limitless possibilities for creative programming. Aspects of this vision are illustrated by a quartet of futuristic application examples in the inset boxes, while the remainder of the article concentrates on the current reality as explored in The Mirror.

Technical Background

The long-term vision for inhabited TV extends to high-speed terrestrial and satellite networks linking into set-top boxes or network computers in the home to form a global *virtual society*². However, delivering The Mirror as an early experimental project required an existing platform

Figure 1 – The Mirror entry portal

and infrastructure which could deliver interactive, multi-user virtual environments to a large community of citizens, and hence delivery over the public Internet to a PC was the only practical choice for the network and client hardware.

After a lengthy evaluation, the Sony Community Place (CP) development was chosen for the software platform². The CP software comprised a VRML2.0 browser; a server which handled the messaging required between browsers to maintain a consistent distributed state, and a facility for developers called the *application object* (AO) which enabled us to insert complex shared objects and behaviours into the worlds. The choice of browser meant that users required a PC with Windows 95, and a minimum specification of a 90 MHz Pentium was recommended for acceptable performance. This restricted the potential audience for The Mirror but, with continuing evolution in hardware performance, is not considered an issue for the longer-term vision of inhabited TV.

Virtual Reality Modelling Language (VRML) has emerged as the *de facto* standard for describing three-dimensional worlds on the Internet, and although proprietary solutions currently achieve superior graphics performance, for example in networked 'twitch' games such as 'Quake', it was appropriate for us to accept both the short-term limitations and the longer-term potential of a mainstream standard³. A further consideration was that the VRML standard had developed so rapidly that there was no possibility of satisfactory interworking between browsers and multi-user servers from different companies. However, it did prove possible to work with a range of authoring tools, confirming the advantage of an industry standard.

The VRML2.0 standard is a file format for interactive three-dimensional environments. However, while future developments may address



multi-user applications⁴, the current version makes no acknowledgement of issues such as shared state and consistency. The Sony CP development is a leading contribution to this research area, making them ideal collaborative partners for The Mirror. The BT interest in content, user issues and service management was complemented by Sony expertise in multi-user architectures and implementation, resulting in an intensive and fruitful dialogue. The collaborative project team was completed by the BBC and Illuminations, an independent TV, film and multimedia production company. The dialogue in this instance was around content for the worlds, and the BBC and Illuminations brought a complementary broadcasters' perspective to the design process.

Content

Development of The Mirror began in September 1996, and hence there were just over four months before the first programme of The Net was to be broadcast. Once the overall structure and outline content of six worlds had been agreed in parallel with the plans for conventional broadcast material, a period of intensive world authoring produced three largely-complete spaces for a beta trial in late November. Several designers worked on the development, using VRML creation tools ranging from text editors to sophisticated authoring environments. The

designers were supported by programmers with knowledge of the Java interface and other technical aspects of the Sony multi-user software. Although pooling of experience was an essential aspect of the content development, the mix of skills and backgrounds ensured that the worlds had very different visual styles and atmospheres. Content development was concluded in early December.

The worlds of The Mirror reflected the themes of the six broadcast TV programmes in their overall settings and individual audio-visual elements. Moreover, they were also designed to experiment with specific aspects of inhabited TV content, to explore which would be most appealing to both new and experienced participants. The worlds were linked by an entry portal (Figure 1), which highlighted a *World of the Week* corresponding with the broadcast TV programme. The graphical design of this portal was closely aligned with the mood boards and title sequence of The Net.

The six virtual worlds were built around the following themes: Space, Power, Play, Identity, Memory, and Creation:

Space: Based on aspects of navigation and space on a lunar terrain, the environment included alien creatures, some of whom responded to one's presence. Teleports were used to produce unexpected transitions, there were a




Figure 2—The cage in Space




Figure 3—Charles Babbage puts in an appearance in Power




Figure 4—The bouncy castle in Play




Figure 5—The End of the World clock hangs over President Kennedy's motorcade in Memory

number of visual illusions, and a cage (Figure 2) encouraged cooperation between visitors, since a trapped avatar could only be released by a friend on the outside. The ambient audio was closely linked with The Net, which included an item on the 'composer', Italian astrophysicist Fiorella Terenzi.

Power: Animated figures from the past and present of computing (Figure 3) were included in a hall of fame, which led visitors into a debating arena. The arena could be customised, with an option to modify the image at the rear of the stage and to include celebrity *super-avatars*, able to broadcast their chat to all of the audience. Additional functionality allowed the audience to record their votes, which were then visible on a scoreboard above the stage. Exploiting this special event functionality was a key aspect of The Mirror. A stirring drumbeat in the hall of fame was replaced by a background murmur in the region around the stage.

Play: An oversize play room filled with games and tricks, designed to promote cooperation and rivalry between visitors. Features included a rocket that required three people to launch, a shuffleboard with persistent scoreboard to foster competition, and a bouncy castle (Figure 4) which shook the avatars. As with all the worlds, audio clips from The Net were attached to objects—in this case larger-than-life toys—with the objective of prompting discussion related to items in the TV series.

Identity: Experimentation with notions of identity and the influence of the environment on people and places. The world changed between day and night, as did the characters and their surroundings. An X-ray machine identified new arrivals to those already in the world, a guided tour was on offer, and a garden with musical sculptures hinted at future instrumental possibilities. The atmospheric *Smokeys Bar* was a popular location, but was the only

area of The Mirror which was not shared—visitors arranging to meet in *Smokeys* were destined to be disappointed!

Memory: Significant events from the last few decades were brought alive along memory lane, which wound through an open landscape. There were snippets of technological, political and cultural history designed to prompt comment and discussion: President Kennedy's motorcade would drive along the lane (Figure 5), and Elvis made fleeting appearances. Audio clips and image flick-books suggested scope for streaming of broadcast audio-visual content within a shared space. A key feature was the clock counting down to the End of the World: three hours before the final shut-down of The Mirror at 2200 on 28 February, Memory world changed to a party setting complete with dance floor, and a beer tent.

Creation: Vibrant flora and fauna brought life to a world which provided visitors with a chance for



Figure 6—Carnivorous flora in Creation

'fifteen Megabytes of fame' (Figure 6). Creatures included frogs, a dragon and a turtle. User authoring will be an important element of shared spaces, promoting a greater sense of community involvement and ownership. A simple VRML2.0 authoring package, Spinner⁵, was supplied to citizens of The Mirror; and an art exhibition was held with exhibits downloaded into Creation world.

Identity and personal appearance is a critical aspect of shared spaces, and The Mirror included a choice of four avatars with clothing which could be coloured from an extensive palette—the first tentative steps towards fashion in cyberspace! Once customised in the changing room environment (Figure 7), the avatar was saved between sessions.

Including all six worlds and the avatar changing room, the content amounted to 2.4 Mbytes of VRML code, supported by an additional 4.6 Mbytes of textures and 29 Mbytes of audio files.




Figure 7—The avatar changing room

Special Events

Scheduled special events were an important aspect of The Mirror. Although the worlds were 'open' 24 hours a day throughout the experiment, the size of community was insufficient to sustain a continuous human presence in six distinct worlds. The Mirror worlds were designed as a rich backdrop for social interaction, and the events typically ensured an attendance of at least 20. Events are also an essential element of the inhabited TV vision, providing a structure for authors, producers and celebrities to influence the action in the space.

- **Party Night** Monday evening prior to the broadcast of The Net was 'party night' in the world with the same theme as that week's TV programme. This provided an opportunity for citizens fully to explore the features of the world under expert guidance from

members of the project team. Additional interest was added to the concluding parties with activities such as an 'avatar stacking contest' (Figure 8).

- **Debates** The super-avatar and voting functionality in Power was used for a number of special events. An initial trial debate between two members of the project team on the future of the Internet was an unqualified technical success, with the speakers able to respond to heckling and comments from the avatar audience. A second debate, between science fiction author Douglas Adams and Peter Cochrane, head of Advanced Applications and Technology for BT, attracted an audience of 45, who resoundingly rejected the motion that 'the book is dead'. Mastering text chat debating requires a new range of skills, and our experiences highlighted deficiencies in the interface, with multiple windows

Inhabited TV—Cup Final

As a member of the inhabited TV crowd, your avatar pulls up a seat in the private box reserved for members of your local supporters' club. While the commentator introduces the teams in the minutes before the kick-off, you swap views on yesterday's transfer announcement and then confirm the arrangements for Wednesday's on-line quiz night. Mixing social banter with rapt attention to the game, you are soon immersed in the atmosphere of the shared space. The box is ideally located, although during the half-time interval you elect to move it across to the other side of the pitch, shading your eyes from the rays of the setting sun. The experience almost matches up to a visit to the real stadium, and the local replay facility is an added bonus.



The 1996 world cup final in Memory



Figure 8—Avatar stacking on party night in Identity



Figure 9—Guests admiring a VRML fish at the opening night of the Art Exhibition in Creation

proving particularly confusing for novice users. 'Cut and paste' of pre-prepared arguments was an effective tactic!

- **Game Show** Although the stage in Power was designed for celebrity debates, it was realised that the same functionality could be used for inhabited TV game shows. In 'Word Jungle', pre-selected members of the audience were invited to assume control of the super-avatars, and in a variation on 'Call my Bluff', they proposed convincing definitions for an obscure word. The audience could then 'vote with their feet' and the winner progressed to the next round. The increased complexity of multiple participants resulted in a few technical mishaps, but it was an entertaining experience which sustained an audience of over 20 for 90 minutes.
- **Art Exhibition** Citizens were invited to submit images and VRML models for an art show in Creation (Figure 9). The 'works of art' were then dynamically downloaded into the world. The exhibition opened with the traditional first night viewing, which confirmed the enormous potential for social chat and personal involvement to sustain interest in otherwise mediocre content!
- **End of the World** The Mirror was brought to a close with parallel parties: a physical gathering at BBC White City and an on-line celebration on the dance floor in Memory (Figure 10). Personalised

VRML tee-shirts were awarded to leading citizens, a virtual conga saw more than 50 avatars snaking around the beer tent, and Auld Lang Syne was 'sung' in the closing minutes. The most distant party-goer was from Australia, and one unfortunate citizen logged on for the first time just 20 minutes before the end of the world.

Other special events included inhabited TV drama in the form of an enactment of a wedding ceremony.


Technical Implementation

In parallel with the design input, there was substantial technical involvement in the content development. The CP Java interface and application objects were used to introduce interactive and

shared objects. The behaviour of these objects was controlled using Java program scripts. The initial feedback from the beta trial highlighted performance as a key issue, and resulted in a number of optimisations in the areas of start-up or load-time and navigation/interaction. December was devoted to a testing cycle with performance targets for load time and frame rate on the minimum specification client machine.

The CP browser and the VRML worlds were distributed to participants on CD, and although content update via the WWW was an available option, it was highly desirable to ensure there were no fatal bugs in the pressed code. With complex environments, flexible multi-user interaction and a diversity of client hardware, the testing schedule could

Figure 10—Avatars on the dance floor at the End of the World party






Figure 11—Server configuration for *The Mirror*

only scratch at the surface of possible run-time states, but in the event only relatively minor bugs came to light after the CD had been pressed.

A further strand of technical work was the implementation and testing of the server and network infrastructure. Although the beta trial involved just 20 human participants, a key element was loading of the server and network using *dummy clients*, which could simulate the presence of many more users. Data from these tests suggested a maximum simultaneous on-line capacity of greater than 600 people across the six worlds. The final hardware configuration involved five UNIX workstations and three PCs, supporting the six world servers, the application objects, a WWW server and ancillary support and monitoring services. The servers were connected through a SunScreen firewall to BTNet, BT's public Internet service (Figure 11).

Another area of great importance to the success of *The Mirror* was the service management and surround. Although users can choose to conceal

their true identity behind an on-line avatar, for many applications it is nevertheless important that they are ultimately accountable for their words and actions in the shared space. In business applications it may even be appropriate to enforce consistent usernames. In the *Mirror*, citizens were free to choose and indeed to change their username, but an initial password-authenticated log-in linked them with a unique user-id. The unique identity was important for data analysis, and gave a facility to exclude people if their behaviour was deemed unacceptable for the space.

In addition to the three-dimensional VRML worlds, an extensive two-dimensional Hypertext Markup Language (HTML) surround was created for aspects such as registration, technical support, and user feedback. It was important to provide every possible assistance to users, while at the same time moderating their expectations by reinforcing the research nature of *The Mirror*. Further HTML pages were created to supplement the content of the spaces

themselves, including maps for Identity world, background material such as biographies and debate transcripts for Power world, and a weekly on-line newspaper.


Finally, there was a significant technical investment in the capture and analysis of data. In addition to routine logging of the network and server hardware, the authentication process and CP software generated log files of selected events such as world entry and exit, AO interactions and text chat. This was used to determine overall loading and performance of *The Mirror*, and in parallel with qualitative observations and questionnaires, was also used to assess user response and community development.

Experimental Results

The Mirror was introduced to the half million viewers of *The Net* in a three minute item on Monday 13 January 1997 in the first programme of a new series. The programme ended at 23.55, and there were more than 600 successful registrations in the first hour. Over the course of the next seven weeks the number of registered citizens of *The Mirror* rose to 2250, including 300 from outside the UK. Overseas citizens registered from countries ranging from Australia to Canada, and South Africa to The Netherlands. Approximately 1000 completed registrations were rejected, mainly because the

Inhabited TV—Soap Opera

In today's episode you take your own viewpoint on the latest events in and around the plaza, as you are guided by an inhabited TV script agent between key locations. This 'fly-on-the-wall' perspective brings you up-to-date with the domestic disputes, and in the bar-room brawl you feel really involved as your avatar is thrown to the ground. Later, you meet up with friends in the bar (thankfully now restored to its normal state) and check to ensure you're fully up-to-date with the key scenes. Discussion turns to the likely outcome of tomorrow's court case, which will be released in the midday update. As a registered citizen, it is just possible that you may be invited to take a place in the shared space jury, but then the odds are 12 in twenty million...




The Plaza in Identity

Figure 12—Registration profile, illustrating the impact of specific press and TV coverage

hardware fell short of the minimum required specification or the Internet ISP could not handle the required network connection.

The initial flood of registrations (Figure 12) confirmed the power of TV in generating a near-instantaneous response, although the variable delay introduced by postal distribution of the CDs ensured that the worlds themselves were not subjected to the same initial 'shock loading'. By the end of the first week, users were starting to appear in the spaces and the focus for the support effort switched from queries on registration to difficulties with installation and performance. All support was conducted by e-mail, and approximately 600 enquiries were answered over the duration of the experiment. The majority concerned minor registration or installation difficulties, poor performance on lower-end PCs, or were from users of unsupported platforms such as the Apple Macintosh. However, user feedback was instrumental in identifying a significant number of underlying technical and interface flaws, particularly with the authentication process.



Some were rectified, and others mitigated through the HTML service surround and publicised work-arounds.

Figures 13 and 14 are illustrative of some of the quantitative data. The user

hours data confirms the importance of special events in boosting usage, although a reasonable level of background activity was maintained throughout the experiment. A total of

Figure 14—Distribution of session duration in a single world





Figure 13—User hour data for all six worlds of *The Mirror*



almost 5000 users hours were recorded, and over a quarter of a million lines of chat were exchanged. A mean session duration of 15 minutes illustrates the commercial appeal of extended connect times in this type of service, while other data suggests a correlation between session duration and on-line population, thereby confirming the importance of the social interaction. Testing with dummy clients had predicted that the server configuration would be able to support 600–700 simultaneous users over a 2 Mbit/s link, and although the experimental loading fell well short of these levels, the available data was in-line with expectations. Availability of The Mirror worlds was close to 100%, and it is estimated that the total data transferred from the server was 40 Gbytes.

Qualitative observation and questionnaires were used to build understanding of user reactions to The Mirror. The project team provided significant support from within the worlds during the first couple of weeks, and this was judged essential in helping new users with basic interface and navigation queries. The ability to change name and appearance within the world was satisfying for new users, but it typically required several sessions to master the complex menus and multi-window interface. It is believed that initial frustration with the interface deterred a number of citizens who did not receive this personal tuition.

Speed was also a major issue for users. Typical load time on a minimum specification client machine was two minutes per world, with frame rates of 3–4 per second at the default window size of 390 by 290 pixels. Although well below expectations established by computer games, most citizens were willing to accept this performance and recognised the very different functionality offered by The Mirror. The slow frame rate detracted from a sense of immersion in the three-dimensional scene, and made for a relatively weak link between the text chat and visual windows. It is suggested that full-

screen frame rates faster than 10 per second should be a minimum design goal for future services.

An on-line Forum provided an opportunity for citizens to post news and views about The Mirror. The Forum was not moderated, and overall the comments were very favourable with several postings applauding the vision of the collaboration in exploring the current technical boundaries of inhabited TV. There were also calls for a 'Mirror anonymous' group for addicts and a threat from one citizen to withhold their BBC licence fee if the project really did come to an end!

Another key issue for users was the impact of *auras*, the grouping algorithm used to produce scaleable worlds. Although the spaces can potentially support hundreds of simultaneous users, the rendering and bandwidth demands of rapidly


disseminating all changes made by all other users are prohibitive. To deal with this issue, auras are used to place users in need-to-know groups of a dozen or so, which enables the system to restrict the amount of information sent to any one client. This approach maps well to user expectations—although it might be acceptable to 'see' 300 other avatars, a multi-threaded 300-way text chat is potentially confusing! Specific limitations of the current CP aura implementation are discussed in the next section.

A final comment on results concerns the clear emergence of 'characters' and, to a limited extent, a community in The Mirror. About 2–3 per cent of registrants became regular users, appearing on-line most days. Although they were also loyal supporters of the special events, it

Inhabited TV—Game Show

After a couple of weeks as a passive observer of the new-format inhabited TV adventure game based on the recently released Hollywood film, you decide that tonight your luck is in, and you enable the funds transfer capability of your avatar. With the first spin of the virtual wheel you are plunged into a mind-bending race against the clock as you attempt to reassemble the shattered puzzle—timely cooperation with a neighbouring contestant ensures that you both proceed to the next stage which is a test of reflexes in a shared-space shoot-out. Oh well, one of you had to lose! Maybe tomorrow you will spare your blood pressure, nurse your wounded pride and revert to spectator mode.

Shuffleboard with persistent low-scores in Play



was the social chat and other citizens which brought them back time after time. A wealth of VRML content had been created, but with only limited updates and user authoring, the interactive possibilities were nevertheless exhausted in just a few hours.

Discussion

The Mirror generated a wealth of quantitative data and qualitative experience and observations, and in this section we discuss a spectrum of our early findings. The goal was to understand the critical factors for a successful inhabited TV service, and to explore issues in production and delivery. The learning covered a mix of technical, content and user-interface lessons, and the complex interdisciplinary nature of the project meant that the issues are often inextricably linked.

Rapid prototyping, small cross-functional teams and flexible use of available authoring tools were key to the content development. The absence of clear guidelines and understanding of performance and optimisation issues was a source of regular frustration, and rigid speed targets from early on in the project would have helped. However, with browser, server and AO code all developing in parallel with the content, the absence of a stable platform limited the early testing opportunities.

Of the content itself, the most successful was that which engendered a strong coupling between the world and the avatars. Examples included the bouncy castle and the rocket in Play, where the interaction was fully shared and inextricably linked to the avatar position. The ability to leave personal memories and messages was also important, and although there were just a handful of entries for the art exhibition, it is probable that shared on-line content creation and editing would have strong appeal. The ability to customise the avatar also enhanced the coupling between user and avatar; and citizens quickly became creative in the use of the 'transparent colouring' option. However, in general social chat, visual cues were largely irrelevant, and the 'emoting' capability, whereby avatars could smile, frown or wave, was little used. Stronger coupling between chat and the avatars, such as the lip synchronisation used in the OnLive! shared audio spaces⁶, would be an important development.

The importance of special events and the pulling power of celebrities in maintaining the community has already been stressed. However, such events quickly built expectations among the citizens of a regular schedule and we should not ignore the high costs of content creation and management. Substantial costs were also incurred in the initial development of the six worlds, and although


acceptable for a research project, attention to both reducing costs and understanding an appropriate revenue model based on sponsorship, advertising, usage and interaction charges is an important area to address.

A final content issue for further study is in defining a new role for authors, producers and editors in inhabited TV. Although each of the six worlds received editorial approval from the BBC, in stark contrast with traditional linear broadcasting, the content is incomplete until the citizens play their part. The editor can sanction a framework, but his ultimate control is much more loose than in, for example, a live TV debate with audience participation. Moreover, every participant receives a unique experience, determined partly by the production team, partly by the other inhabitants, and partly by their own actions. There is consequently no unique broadcast stream that can be monitored and assessed against prevailing public standards.

Turning to more technical aspects, there were clear user demands for a more intuitive interface and greatly improved speed of navigation and interaction. This will inevitably come with advances in networks and, in particular, client hardware, but developers must ensure that their content ambitions are not perpetually one step ahead of platform capabilities. Experimentation with higher

Inhabited TV—Debate

This was a subject about which you really cared, and on which you felt qualified to air your views. It was therefore gratifying to have been selected for the primary inhabited TV audience. This would afford full interaction with the speakers, and you had ensured that your avatar was appropriately dressed for the occasion. The secondary spectator groups could also be informative, although the discussion was invariably restricted to a few caustic comments and you were only invited to vote on the main motion. Nevertheless there had been a turn-out of ten million at the end of last week's shared-space debate—so much simpler than in the days of telephone voting.



Super-avatars debating on the stage in Power

bandwidth cable and satellite links will be important, particularly in enabling high-quality streamed audio and video, as will developments aimed at the Web-TV or network computer as client. Multi-party audio will transform the interactions within the community, although it is suggested that the greater reliability of text chat will remain important for several years.

A final area for discussion is that of auras and the problems that they raise for application developers and users. The current CP implementation uses three configurable auras which determine visual, text chat and AO interaction groupings. For small populations, less than the maximum group size of typically 12, the auras cause few difficulties; users readily accept that they cannot necessarily see or chat with all the other avatars in a large space. However, for large densely packed gatherings, the allocation of groups is at best confusing and is often frustrating for the users. The unlucky thirteenth arrival may find themselves in a seemingly uninhabited world; grouping can become inconsistent such that only half of a dialogue between two other avatars is received; or two friends may agree to meet at a specified location and time, and yet be allocated to separate groups and hence be unable to communicate. Advances in this area are critical to the success of large-scale inhabited environments.

Conclusions

The Mirror was a pioneering experiment in interactive multi-user VRML worlds, and events such as the debate, game show and art exhibition were world firsts in inhabited TV content. Initial user and performance data point to a significant market opportunity as the technology matures.

In addition to the technical hurdles perhaps the greatest outstanding challenges are in understanding the economics and structure of compelling content. The Mirror

confirmed the appeal of social chat, and it can even be argued that anonymity acts as the alcohol of shared spaces, lubricating creative and open dialogue that challenges the conventions of physical gatherings. Users are also motivated by personal participation and involvement, both through a close coupling with their avatar and through the opportunity of '15 Megabytes of fame' afforded by user authoring.

Consolidating these lessons is only a start on the complex area of content structuring. Much creative and experimental work remains if we are to understand the essence of compelling content, and the features that will build on-line communities by turning surfers into settlers. Inhabited TV as a commercial service remains a long-term vision, but The Mirror is an important early experiment, from which enduring lessons can be learnt.

Acknowledgements

The Mirror was a collaborative project which drew on the expertise of many people within BT and the research partners, Sony Architecture Labs, Illuminations and the BBC. The contributions made by the following people are gratefully acknowledged: Laurence Bradley, Andrew McGrath, Paul Rea, Tim Regan, Jason Morphet, Marco Fauth, Mike Hinds, Paul Shumake, Richard Munn, Mark Learmouth, Stephane Dubedout, Natasha Gann, Chris Seal, Dennis Sheat, Alison Coe, Matt Shipley, Amanda Oldroyd, Matthew Polaine, Peter Platt, Charanjit Sidhu, Helena Athoussaki and Steve Hubbard from BT Laboratories; John Wyver and CJ Brazier from Illuminations; S. Matsuda, Yasuaki Honda and Rodger Lea from Sony Architecture Labs; and Robin Mudge, Surj Patel, Danielle Eubank and Jonathan Drori from the BBC.

References

- 1 BRADLEY, LAURENCE; WALKER, GRAHAM, and MCGRATH, ANDREW. Shared Spaces. *Br. Telecommun. Eng.*, July 1996, **15**, p. 162 and http://vb.labs.bt.com/msss/IBTE_SS/
- 2 Sony Virtual Society and Community Place, <http://sonypic.com/vs/>
- 3 The VRML Consortium, <http://vag.vrml.org/consort/>
- 4 Living Worlds, Making VRML 2.0 Applications Interpersonal and Interoperable, <http://www.livingworlds.com/>
- 5 <http://www.3Dweb.com/>
- 6 OnLive! <http://www.onlive.com/>

Biography



Graham Walker
BT Networks and
Systems

Graham Walker managed The Mirror project at BT Laboratories. Graham joined BT as a sponsored student and, after graduating from Oxford University in 1986, spent six years researching into coherent optical transmission systems. This work resulted in numerous publications, and the award of a Ph.D. from Cambridge University in 1992. More recently, he has been leading a group within Applied Research and Technologies, working on information visualisation and shared spaces. He has an MBA from Cranfield School of Management and is a Member of the IEE.

graham.walker@bt-sys.bt.co.uk
<http://www.labs.bt.com/people/walkergr>

Alwyn Lewis and Fred Westall

Whither Speech?

Convergence and the Future of Telephony

The authors explain why the commercial and technical forces of convergence are likely to overturn many design concepts in telecommunications – concepts that were very successful for so long. For example, the traditional design viewpoint is network-centric and focuses on the transmission and switching of bits. Yet bits are not, and never were, the stuff of the traditional product. Designing new services with such an outlook could be commercially perilous in an era of rapid convergence and keen competition.

Introduction

The rapidly decreasing cost of computing power and the even faster growth of the Internet is widely forecast^{1, 2} to drive a convergence of the currently separate computing, entertainment and communication industries. One manufacturer³ has already produced a computer with fax, voice mail, Caller ID, paging and duplex hands-free telephony. All sorts of new services are becoming affordable, such as teleconferencing, teleworking, teleshopping and video-on-demand⁴⁻⁶, raising new and sometimes conflicting issues. There is little agreement between independent observers about what might be the ultimate result of such a technological and business convergence.

This is the first in a series of three articles dealing with the technical issues behind two common questions—‘Where will convergence lead?’ and ‘What might convergence mean for telecommunications?’. The series will look forwards, backwards, sideways and around—and range from the prudent to the fantastic—to discuss the cultural and organisational implications for the future of telephony.

Will Speech Withers?

Some observers claim that all human communication will be transformed and telephony made obsolete, within a generation, because of personal computers and the Internet. Campbell, a Managing Director of Australia Telecom, implies the telephone might be killed⁷ by the

rapid growth of Internet, data and video traffic. Others say telecommunications will become a low-profit commodity, much like water or electricity supply⁸.

The mass media have just discovered the Internet, but it was conceived 28 years ago⁹ and has become a prodigious invention. It can put you in touch with people you need to contact, even though you know nothing about them. It has made a sub-culture out of serendipity (the making of happy, chance discoveries). Negroponte¹⁰ shows how a global computer network can have unique properties, wonderfully augmenting the power and range of human skills. But there are detractors, including some serious and hard-nosed users. Stoll¹¹ speaks from bitter experience—finding information equated to knowledge, meeting human contact without humanity and being drenched in electronic dross.


Unfortunately, the Internet is growing and evolving so frantically that it addles the brains of a few of those it fascinates. When we all have Internet connections, will we stop talking to each other? Should we forget how to walk when we learn how to swim? To resolve this turmoil, we need to go back to fundamentals and take a fresh look—reverting to infancy.

A Cry from the Cradle

Aaaa-ahaaa... mmnn-mmm....
mmnmAhaaa?

Countless human babies have made the same discovery—that simple sounds, joined together, can make

Figure 1—The dimensions of human cognition¹²



something new. To an adult it's just babbling, yet it marks the start of a mental and physical struggle that is long, arduous and mercifully always forgotten—a struggle to master that subtle, highly-structured and pivotally-potent ability we call *speech*. An ability that becomes effortless and so integral to our being that we accept its enormous power without a thought. Yet the rich complexity of our speech makes us unique among living things, lying at the heart of how we become what we grow to be.

Reflect for a moment how unlikely, even bizarre, the mechanism is. We can, at will, vibrate a tiny piece of our body. We can form, shape and colour these faint vibrations, impressing just a few thousandths of a watt of power onto a thin waft of air. Is it really likely that such insubstantial events could carry joy or

Truly, speech is the cradle of our humanity

sorrow, love or fear, laughter or pain across the vast gulf of consciousness between different minds—and do so in just a few heartbeats?

Speech is made in just such a way, and in much less than a minute can penetrate to our deepest emotions or fling our thoughts to the uttermost star. Truly, speech is the cradle of our humanity.

The Power of Speech

Vogt¹² has analysed the dimensions of human cognition (see Figure 1) to identify the key skills of his future breed of 'knowledge workers'. It is clear that eliminating the speech-related aspects and using only text, data and image-based features would seriously damage or even remove much of the power of human thought. One picture can indeed be worth a thousand words—vision is the highest bandwidth human input channel. But the mental pictures that form inside

our heads are more potent still. The most effective expression of those images is through the fast and flexible power of speech—our highest bandwidth output channel.

To a signal-processing engineer, speech is an awkwardly complicated thing¹³ and difficult to classify. It is sometimes ergodic but seldom stationary, being at times quasi-periodic and at others quasi-random. The one dependable property of speech is that it varies, not only with different talkers but also with the same talker at different times or in different surroundings.

The way our brains are wired for speech seems to involve multiple layers of complexity. Medical evidence, from patients with disease or accidental brain injury, provides clues about how these functions are arranged in the brain. Grammar, vocabulary, cognition (understanding) and linguistic (speaking) skills seem to reside in quite distinct areas of the brain. Literal and metaphorical understanding also seem to be separate abilities.

Any one of these functions can be seriously impaired without much impact on the others. For example, some stroke victims suffer *empty speech*, with perfect grammar and linguistic style but severe loss of, say, nouns, despite understanding the 'lost words' when they hear them. Children suffering from a condition known as *Williams Syndrome* show severely retarded intelligence, but can use erudite vocabulary in an exceptionally skilful way.

This evidence suggests that emulating human speech faculties in a machine might be impossible. But many applications of speech technology have substantial structure or prior knowledge about the task, which makes useful implementations practical.

What a Piece of Work is Man

The hardest challenge in further developing and exploiting the potential of speech technology lies in

gaining a deeper understanding of how speech works. Today's adolescents might well describe human skills in speech generation, reception and comprehension as 'pretty neat—yeah, wicked'. This phrase is itself a good illustration of the mutative variability of speech, since the message received depends on the receiver. Shakespeare put it more eloquently when Hamlet described humankind as: '...How noble in reason! How infinite in faculty! In form and moving how express and admirable! In action how like an angel! In apprehension how like a god!' (Act II, ii, 317).

Speech is more information-rich and revealing than is generally appreciated. Wiseman conducted a test¹⁴ on the relative accuracy of different ways of detecting lying. Observers who just listened to the potential liar speaking did significantly better than those who also watched his picture or those who only read his words. In this large-scale test, the observers with the most clues (those with sound and picture) produced the least effective detection result†.

The ability to invent metaphor is often considered one of the hallmarks of genius—and linguistic metaphor is, as any good politician knows, the most powerful and memorable of all. The power of speech is something that advocates of data and image networking, in their enthusiasm for new technology, occasionally seem to

† Detection by a total of 41 471 television, newspaper and radio observers was 51.8, 64.2, 73.4 per cent correct, respectively (50 per cent chance) in the 'Megalab Experiment' during BBC Television's 'Tomorrow's World', 25 March 1994.

forget. From the 'Grunters' of the Pleistocene to the 'Mass Mediators' of today¹⁵, communication by speech has fuelled the ascent of humankind. It is absurd to predict the death of the telephone, but plausible to predict its transformation into something that will look and feel quite different from the device we know today.

Electronic Puberty

The wild, wired disciples of the Internet have a point after all—there is drastic change in the air. The authors believe that electronic puberty is a good description of the change that the telecommunication industry has begun to experience on a global scale. All the ground rules have changed, parts of the body have acquired minds of their own, anguish is flavour of the month, and there is no clear idea of what happens next. Furthermore, new and exciting abilities have appeared, with an enormous power that is yet to be explored.

Adopting both the fervour and the bias of the evangelist, the authors are confident that the most powerful of these new abilities will be digital signal processing and speech technology. The authors also believe that processing signals will become just as central to telecommunications as transmission and switching are today. This implies nothing short of a cultural revolution in the hearts and minds of many telecommunication engineers. Good evidence to support these beliefs comes, strangely, from the earliest days of the telephone.

The Birth of Telephony

The notion of transmitting sound to a distant place goes back to antiquity. The world's first practical electric device called a *telephone* (from the Greek, meaning 'far speech') operated digitally. In 1861, Johann Philipp Reis, a Professor in the Garnier Institute of Frankfurt-am-Main, demonstrated an apparatus to his local Physical Society that he called a *telephone*^{16 p. 65}. It could transmit the

By courtesy of BT Archives



Figure 2—Alexander Graham Bell

pitch of speech sounds, but not their intensity or articulation. It had no true microphone, but used the vibrations of sound to switch an electric current rapidly on and off by a delicately-adjusted diaphragm-driven contact.

However, Alexander Graham Bell and Thomas Watson (Figures 2 and 3) are universally acknowledged as the true fathers of the 'articulating

© Science Museum



Figure 3—Thomas Augustus Watson

electric telephone' we know today. The earliest days of the telephone were inhibited by a lack of knowledge and by the limitations of the technology of the day. Nonetheless, those days were full of enthusiasm and inventiveness, with experiments in stereo and in the commercial distribution of entertainment.

Key events in the history of telephony are shown below.

Key Events in the History of Telephony

10 March 1876 In a Boston attic, Bell shouts 'Mr. Watson, come here, I want to see you'—the first intelligible sentence ever transmitted by telephone. Two days later, Bell wrote in his notes¹⁷: 'The effect was loud, but indistinct and muffled'.

1881 The British Post Office grants the first licences for commercial telephony in the UK, limiting service to a 5 mile radius and extracting a 10% royalty¹⁸.

1884 Clément Ader demonstrates binaural telephony, using two microphones placed some distance apart near the stage of the Grand Opera in Paris. Listeners in the nearby Exposition Hall use two Bell receivers, one to each ear, to experience spatial sound reproduction^{16 p. 68}.

1891 An undertaker in Kansas, Almon B. Strowger, patents the two-motion selector for automatic telephone switching¹⁹.

1892 The first automatic telephone exchange opens in La Porte, Indiana, using the Strowger system.

1893 The 'Electrophone Service' sells entertainment by telephone, from London theatres and churches. The service grows modestly until radio broadcasting brings closure in 1926.

1912 The first public automatic telephone exchange in the UK opens at Epsom in Surrey, using the Strowger system.

1927 The first transatlantic telephone calls are made, by radio between Rugby in the UK and Rocky Point in the USA.

1943 The first amplified undersea telephone cable links Anglesey in Wales with the Isle of Man.

1956 The first transatlantic telephone cable links Clarendville in Newfoundland with Oban in Scotland.

10 July 1962 Lyndon B. Johnson makes the first public telephone call via an active satellite in space and says 'You're coming in nicely..²⁰.

14 February 1989 The first satellite telephone call from an aircraft in flight is made, live on BBC Television's 'Tomorrows World', using the BT Skyphone™ speech codec.

Figure 4—Telephones from Gower-Bell to Slimtel

By the turn of the century, it was clear that the most important work in telecommunications engineering was to find effective and cost-efficient solutions for the twin problems of transmission and switching—to overcome attenuation with distance and provide easy connectivity. These tasks were to occupy generations of engineers, all around the world, for most of the 20th century. As a result, the network-centric view has become a fundamental tenet of the industry—that the disciplines of transmission and switching are the flesh, blood and bones of telecommunications. Any other view is heresy.

Recent History

Over many decades, the telephone has been transformed in appearance (Figure 4) and greatly reduced in cost. Modern electroacoustic transducers and integrated circuit technology have been incorporated, but without any fundamental change in specification or performance. In the network, transistors replaced

By courtesy of BT Archives



devices²² finally removed the limitations of analogue signal processing. This led to a dramatic renaissance in signal processing, which has already had a breathtaking impact on areas as diverse as communications, weapons, entertainment and transport.

DSP technology has enabled the economic use of a wide range of complex, adaptive or non-linear processing algorithms that previously had been too cumbersome or expensive to contemplate. Speech coding has reduced the cost of international telephony, and is the economic basis of value-added resellers such as

are the central battleground of the convergence debate. The outcome is likely to decide whether telecommunications will become a bit-transport industry, without any need or way to add value.

Specific Fallacy

Some already regard telecommunications as a bit-transport function. Malamud⁸ says that telecommunications companies are: '... no different than electrical, water, sewer or any other utility competing for the right to push stuff through a pipe'. The authors find this outlook bizarre. Do you tell your water company which specific drops of water, of all the many in the reservoir, you would like to use next? When you make a one minute telephone call, are you content to get 3.75 million of whatever bits first come to hand?

Telecommunications is profoundly unlike water or electricity supply, because of its specificity. Bits are not 'stuff' in the sense that water, electricity or sewage is. They are not even 'stuff' in the sense of a postal service, since in telecommunications an order of arrival different from that of dispatch destroys the message. Bits are not the 'stuff' of the product. Meeting human communication needs—in all their rich variety and with planet-wide technology-independent interoperability—that is the product. The supply of telecommunications is like the supply of water or electricity in the same kind of way that painting a picture is like painting a warehouse.

the network-centric view has become a fundamental tenet of the industry

thermionic valves during the 1960s, but the manufacturing tolerances of affordable inductors and capacitors still limited the practical complexity of analogue signal processing.

The most dramatic change in technology over the past decade has been digitisation, which enormously improved the economics and flexibility of switching systems. Semiconductor lasers and optical fibres revolutionised network transmission and copper proved resilient in the local loop²¹. At last, the major problems of transmission and switching found effective, though not final, solutions. But these solutions were not the only result of digitisation.

In the late-1980s, the precision and repeatability offered by affordable digital signal processing (DSP)

ConcertTM, of new mobile telephony systems such as SkyphoneTM²³ and the GSM service²⁴, and of voicemail services such as CallMinderTM. Speech recognition and synthesis now offer an affordable and effective way of talking with machines. Many more applications of digital signal processing technology are yet to come.

So we have travelled full circle and returned to those eager experiments in novel applications of telephony made in the 1890s. As a result, DSP technology is now on the threshold of claiming an equal place, with transmission and switching, as the flesh, blood and bones of telecommunications—challenging the traditional network-centric view. Unresolved questions about which industry will provide this DSP technology—when, where, and how—

Bit-Bare or Bountiful?

Bits are simple things, yet controversial. In the 'Negroponte switch'²⁵ everything which now comes over the air will come in bits by wire, while all that now arrives by wire will come in bits over the air. Lewis et al.²⁶ discussed two imaginary futures for bits. One was a 'bit-bare' world, with all connections mobile and the use of uncompressed signals illegal. The other 'bit-bountiful' world was the opposite, with mobile telephony illegal and any kind of signal compression laughable. The bit-bare future has grown closer, with increasing use of bit-rate compression for mobile telephony and network economy. There is also evidence of a bit-bountiful trend, with the growth of entertainment services on fixed networks. The future may turn out to be both bare and bountiful in bits.

Cultural Change

The pundits are equally polarised about the future cultural impact of telecommunications, predicting a world ruled by either Mephistopheles or Gabriel. The most dire prediction, of the censoring and monitoring of everything with zero privacy, would make Orwell's '1984' seem sunny and transform the Dark Ages into a jolly Utopia. The opposing view predicts a triumphant victory of the democratic, libertarian marketplace, where the Government disappeared after there was nothing left for it to do. Kinney²⁷ debates both views. Personally, the authors suspect that improving communications will continue to mildly foster democracy and that the politicians have more to fear than the people.

We seem set to enter a 'knowledge worker' world²⁸, where new technology alters the nature of work itself²⁹, 'work tribe' alliances blur national allegiances or lead to the death of the nation state²⁷, and society divides into information-access 'haves' and 'have nots'. The Internet is widely

cited as an engine of such cultural change, but with an impersonal or even de-humanised outlook—leading to a world of 'terminal junkies'.

The authors believe that telephone evolution, conveying more of the nuances of 'being there', will be a powerful factor in humanising the cultural changes of the information revolution. The telephone is so successful, yet has changed so little in form and function, that any possibility of substantial change is commonly neglected. The second article in this series will examine potential developments in the evolution of telephony.

References

- 1 WILLIS, P.; and DUFOUR, I. G. Towards Knowledge-Based Networks. *BT Technol. J.*, Apr. 1995, **13**(2), pp. 87–93.
- 2 STERLING, B. Dropping Anchor in Cyberspace. *Telecommunications*, Sept. 1995, p. 115.
- 3 Globalyst™ TPC, AT&T Global Information Solutions, 1995.
- 4 MARSHAL, I. W.; and BAGLEY, M. The Information Services Supermarket—An Information Network Prototype. *BT Technol. J.*, Apr. 1995, **13**(2), pp. 132–142.
- 5 LYNCH, T.; and SKELTON, S. Teleworking: A Necessary Change. *Br. Telecommun. Eng.*, July 1995, **14**, pp. 122–130.
- 6 McCLELLAND, S. Telework's Global Reach. *Telecommunications*, Sept. 1995, pp. 184–188.
- 7 Editorial: 'Telecom's phone revenues will fall to zero says Campbell'. *Exchange*, 1 July 1994.
- 8 MALAMUD, C. Viewpoint. *IEEE Spectrum*, Jan. 1995, p. 32.
- 9 CERF, V. G. The Internet's 26th Anniversary... We Think!, *Telecommunications*, Sept. 1995, pp. 182–183.
- 10 NEGROPONTE, N. P. Get a Life?', *Wired*, Sept. 1995, **3**(9), p. 206.
- 11 STOLL, C. Silicon Snake Oil—Second Thoughts on the Information Highway. Macmillan, 1995.
- 12 VOGT, E. E. The Nature of Work in 2010. *Telecommunications*, Sept. 1995, pp. 21–34.
- 13 SYRDAL, A.; BENNETT, R.; and GREENSPAN, S. Applied Speech Technology. Ch. 1, pp 1–45, CRC Press, 1995.
- 14 WISEMAN, R. The Megalab Truth Test. *Nature*, 2 Feb. 1995, **373**, p. 391.
- 15 WHEDDON, C. The Seven Ages of Communication. BT Lecture, Mar. 1995.
- 16 FAGEN, M. D. (ed.) A History of Engineering and Science in the Bell System—The Early Years (1875–1925). Ch. 3, Bell Telephone Laboratories, Inc., 1975.
- 17 BRUCE, R. V. Alexander Graham Bell and the Conquest of Silence, p. 181, Gollancz, 1973.
- 18 YOUNG, P. Power of Speech. Allen & Unwin, 1983.
- 19 YOUNG, P. Person to Person—The International Impact of the Telephone. pp. 11–12, Granta Editions, 1991.
- 20 SOLOMON, L. Telstar, Constable Young Books, London, 1963.
- 21 FOSTER, K. T.; YOUNG, G.; and COOK, J. W. Broadband Multimedia Delivery Over Copper. *BT Technol. J.*, Oct. 1995, **13**(4).
- 22 WESTALL, F. A.; and IP, S. F. A. (ed.). Digital Signal Processing in Telecommunications. Chapman and Hall, 1993.

- 23 LEWIS, A. V.; GOSTLING, C. D. *et al.* Aeronautical Facsimile—Over the Oceans by Satellite. *BT Technol. J.*, Jan. 1994, **12**(1), pp. 83–97.
- 24 BARRETT, P. A. *et al.* Speech Transmission over Digital Mobile Radio Channels. *BT Technol. J.*, Jan. 1996, **14**(1).
- 25 NEGROPONTE, N. P. *Being Digital*. Hodder & Stoughton, 1995.
- 26 LEWIS, A. V.; BRANCH, P.; BARRETT, P. and OGDEN, M. DSP in Network Modelling and Measurement. *BT Technol. J.*, Jan. 1992, **10**(1), pp. 173–189.
- 27 KINNEY, J. Anarcho-Emergentist-Republicans. *Wired*, Sept. 1995, **3**(9), pp. 90–95.
- 28 TOFFLER, A. *Powershift: Knowledge, Wealth and Violence at the Edge of the 21st Century*. Bantam Books, Nov. 1990.

Biographies



Alwyn Lewis
BT Networks and
Systems

Alwyn Lewis is an advisor in Advanced Applications and Technologies, BT Networks and Systems. He gained a B.A. from Cambridge University in 1971, and a M.Sc. from Essex University in 1972. After a spell with Plessey on formant-tracking vocoders, he joined BT Laboratories to work on acoustics, telephonometry and telephone design. In 1984, he became head of a team developing a duplex hands-free telephone with a custom VLSI chip-set and a DSP microprocessor. He then led the speech-coding team, helping develop and test DSP software for CallMinder™ and Skyphone™ facsimile and data services. His interests include beam-steering microphones, speech coding and enhancement, adaptive audio signal processing and the social impact of electro-technology. He is a Chartered Engineer and a Member of the IEE and IEEE.



Fred Westall
Brite Voice Systems

Fred Westall is the Engineering Director of Brite Voice Systems, formerly a manager with Advanced Applications and Technologies, BT Networks and Systems. He gained a B.Sc.(Eng.) from University College London in 1973 and a M.Sc. from the University of Manchester Institute of Science and Technology in 1975. After a spell in microwave development, he joined BT Laboratories to design speech-band modems and has been closely associated with digital signal processing ever since. In 1982, he became head of speech coding, developing novel speech-and-data multiplexers, and became manager of data products development in 1987, responsible for packet terminals and high-speed modems. Recently he was manager of the speech technology unit, downstream applications in the fields of speech coding, analysis, recognition and synthesis. He is a European Engineer, a Fellow of the IEE and a Senior Member of the IEEE.

Dave Machin and Phil Sheppard

A Computer Vision System for Natural Communication

Soon we will be interacting with computers in a more humanistic and natural manner. Low-cost video cameras, speech recognition, high-power desktop computers and smart software, will allow us to talk and gesture, instead of being confined to the mouse, keyboard and screen. Such an advance is especially important for video conferencing where the natural flow of human interactions is currently constrained by the limits of technology.

Introduction

Technological visionaries and science fiction writers dream of autonomous computers that can sense their environment and react accordingly. In the real world this ambition is a far cry from being fulfilled, except in a strictly constrained manner. To a user-interface designer, the goals may be less ambitious, with a broad requirement for systems to analyse the nuances of human gestures and interpret facial expressions to minimise the use of mouse and keyboard. Here we describe a computer vision system which goes some way towards achieving such an objective and discuss some potential applications. Several novel applications for human-human interaction using audio, video and virtual conferencing are also described.

Applications

Pointing by looking

While it is easy to conceive of computer vision systems to replace keyboard and mouse, we might anticipate low usability. For instance, a fixed gaze to a point on a screen

precludes reading. However, the addition of an ear would give more degrees of freedom. We could look at a telephone icon and say 'Call Phil Sheppard' or a document viewer and say 'display the document I wrote in June about reorganisation'.

Virtual conference

In a video multi-conference where 10 individual locations are interconnected, 10 collocated monitors and cameras are required to facilitate eye contact at each site. At the other extreme, in audio conferencing, it is difficult to visualise who is talking to whom. This promotes the notion of an artificial representation of the participants displayed in a virtual meeting room (Figure 1). To add a degree of realism, the computer vision system analyses gestures and facial movements to synthesise the animated virtual people.

Affective computing

Being affective^{1,2} infers an understanding of a user's emotional and subjective preferences in the training and adaptation of software entities. Software agents are novel because they are:

Figure 1 – Virtual meeting room



- *proactive*—taking the initiative to assist users by making suggestions and/or automating mundane tasks,
- *adaptive*—learning user's preferences, habits and interests as they develop, and
- *personalised*—customising their assistance according to learned experiences.

The disabled

There are a very wide range of applications for disabled people including artificial voices for those who can affect very little physical movement. Words can be spoken³ when the user looks at text on a screen and blinks.


Gaze correction

One disadvantage of today's video-phone is the location of the camera not being close to the eyes of the talker. This gives an impression of disinterest and detachment. However, transmitting the position of the eyes, along with the video image, allows the use of eye warping to give a better impression of eye contact (Figure 2).

The Prototype

A vision system⁴ has been constructed to analyse facial details without recourse to specialist image processing hardware.

Figure 2—With and without eye warping



Locating the eyes

The first stage in the facial feature location is an eye search. Once an eye has been located, the detection of other facial features is relatively easy.

The eyes are located by scanning model images (Figure 3) of left and right eyes over the top centre region of each video frame until a high correlation is found. After the first few frames, the search time is reduced by constraining the search to the areas immediately surrounding each eye. In the event of failure, the area is widened for each frame until a match is again found.

If multiple matches are found, due to background features with a similarity to an eye template, then a score is derived from the horizontal angle, correlation coefficients and separation. The eyes with the highest score are then used.

Head orientation

Knowing the positions of the eyes in relation to the head boundaries in an image dictates where a face is pointing. The skin and hair colour are used to

Figure 3—Model eye tracking image




Figure 4—Finding the sides and top of the head by colour segmentation of skin and hair

segment the face from the background. The skin is found in the horizontal and hair in the vertical stripe (Figure 4).


During the first few frames, after eye location, a skin image sample is taken from just below the eyes, and hair from above. The RGB levels of each pixel from these samples are used to calculate an average skin and hair colour. The Euclidean distance from the mean RGB distribution to each pixel to be classified gives the closeness to skin or hair. The sides of the head are determined by finding the ends of the skin class along the horizontal stripe and similarly with the vertical.

The distance of the head from the camera is derived from the ratio of the number of pixels across the head to the assumed head diameter. Simple geometry is used to estimate the face pose angle (Figure 5).

The Mouth

Knowing eye position and head width allows a good estimate of mouth location. Greater accuracy is achieved by passing a horizontal edge detector over the mouth region revealing the mouth as a set of long horizontal lines. This works even in the presence of facial hair which tends to lie vertically (Figure 6). The threshold level is calculated as a fraction of the grey level of the surrounding skin.

The presence of teeth in the image is detected using vertical line






Figure 5—Geometry head in relation to the computer screen and camera

detection. For conferencing applications, the lip shape is transmitted at 15 frames/second for a good representation of mouth shapes during speech.

Eye direction

The centre of the pupil and corners of the eye are also located by the same pattern matching techniques. Vertical and horizontal displacements between the corner of the eye and the centre of the pupil give a good estimate of the iris position.

Eyebrows

The resting positions of the eyebrows are measured and used as an offset to any position change. At present, some of the more subtle eyebrow movements have proved difficult to detect, although frowning is detected by counting the number of vertical lines between the eyebrows.

Hand gestures

Determining the position of hands in three dimensions involves two cameras as a stereoscopic vision




Figure 6—Segmentation of the mouth region by thresholding

system, followed by geometric computation. Skin colour segments the hands⁵ from the background.

Results

The system has been used in a virtual conferencing application with several human body models. One of these models was controlled by a text-to-speech system to act as a secretary to read electronic mail messages. The other models were controlled by gesture recognition systems.

To make the environment more natural, each model was adjusted to resemble the person controlling it (Figure 7).

The system performed well in a limited test environment with some difficulties encountered when the illumination conditions were varied. Techniques are being investigated to solve this and to improve the accuracy of the mouth tracking.

Conclusions

Tracking of head orientation and key facial features has been achieved in real time using a PC without specialist image processing hardware. The system has many potential uses including multi-party conferencing, user interface enhancement and affective computing.

Techniques for improving the discrimination between the subject and background are being further refined, as is the trade-off between robustness and frame rate.

Rapidly increasing processor power and ever improving image processing algorithms are opening up a wealth of new applications for computer vision systems of this type.

Figure 7—Texture mapped three-dimensional models. Left view with mouth full open and front view with mouth closed



References

- 1 PICARD, R. W. Does Hal Cry Digital Tears? Computers and Emotions; chapter in *Hal's Legacy*, Ed. by D. Stork. MIT Press, Cambridge 1996.
- 2 PICARD, R. W. *Affective Computing*. Forthcoming book to be published by MIT Press.
- 3 PAGE, J. H.; and BREEN, A. P. The Laureate Text-to-Speech System—Architecture and Applications. *BT Technol J.*, 10(3), July 1992.
- 4 MACHIN, D. J. Real-Time Facial Motion Analysis for Virtual Teleconferencing. IEEE International Conference on Automatic Face and Gesture Recognition, 14–16 Oct. 1996.
- 5 AZARBAYEJANI, A.; WREN, C. and PENTLAND, A. Real-Time 3-D Tracking of the Human Body. Proceedings of IMAGE'COM 96, Bordeaux, France, May 1996 and MIT Media Laboratory Perceptual Computing Section Technical Report No. 374.

Biographies



Dave Machin
BT Networks and
Systems

Dave Machin gained a degree from the University of Bradford in 1975, joining BT Laboratories to work on mathematical modelling of transistors for integrated circuits. After a number of years in this area he spent five years on a collaborative European project for automatic design validation of integrated circuits using electron beams. In 1990, he joined a team working on subjective testing of telephony services and developed a tool using digital signal processors to emulate international telephony networks. He has recently moved into the area of machine vision for teleconferencing.



Phil Sheppard
BT Networks and
Systems

Phil Sheppard graduated from Birmingham University in 1986 with a B.Sc. (Hons.) in Computer Science and Electronic Engineering. He joined BT Laboratories 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 Rapporteur. He currently leads a group working on a diverse set of projects to provide natural communication systems, including research into advanced non-intrusive measurement algorithms, conferencing systems (using virtual reality technology and image processing), development of BT's three-dimensional talking head and demonstrators for advanced information management systems. He has recently gained an M.Sc. in Telecommunications Engineering (developed for BT) from the University of London.

An Evolutionary Strategy for New Communication Services

The independent telecommunications companies of Finland, Finnet Group, are developing new communication services for both business and residential customers to retain a competitive edge. Ordinary telephone services are no longer sufficient. Business customers see telecommunications as a core-business process where the operator is providing segmented services, supporting core processes. Residential customers are now wanting new communication services, based on multimedia technologies, following the Internet revolution.

Introduction

Finland is often considered as a country with an advanced telecommunications infrastructure that is supporting the construction of the information society. This development is supported by both the public and the private sector. Government and private enterprises have both taken a proactive attitude towards the use of new opportunities offered by advanced information technology and telecommunications resulting in three main initiatives:

- The Ministry of Education has strongly supported the networking of universities and, more recently, other levels of education. The universities have been developing not only their internal networks but also a network between universities using high-speed connections.
- The liberalisation of telecommunications services started in 1985 and has led to open competition in all areas of telecommunication since July 1994.
- Menu-based electronic information services were permitted to be developed early in the 1980s for banking and public services.

Main Players

The main players in the telecommunications market are Finnet, the group of independent telecommunications companies, the

biggest ones being the Helsinki Telephone Company and the state-owned Telecom Finland. The Finnet Group consists of 46 independent companies and five subsidiaries, which they own together. Subsidiaries are responsible for different nationwide operations and services and for connecting the parent companies (which used to operate on a restricted local area before liberalisation).

Subsidiaries

- Datatie Ltd. was established in 1985 when business-data services were liberalised. Datatie owns the nationwide fibre-optic backbone and is responsible for business data services. The backbone was built to handle fully-liberated services, that is the long-distance and mobile telephony throughout the country.
- Radiolinja Ltd. was established in 1992 to operate as a mobile (GSM) operator inside Finnet Group. The analogue mobile networks (NMT) are still run by Telecom Finland as a monopoly operation. The first European GSM call was produced in the Radiolinja network. The mobile penetration in Finland is among the highest in the world, currently over 30%.
- Finnet Nine Ltd. was established in 1993 to become a long-distance operator within Finnet Group from January 1994. The first four months of operation resulted in a 53% market share for the company.

- Finnet International Ltd. is responsible for international voice traffic and started to operate in July 1994.
- Omnitele Ltd is a consulting company owned by Finnet Group. With nine years of experience, Omnitele has gained a leading role in the Finnish telecommunications consultancy market and a good position internationally.

Data Network Services Evolution

Finnet data services started with a X.25 packet network and leased lines as did companies in many other European countries. Competition led to an alarming decrease in the prices for leased-line services, so there was a need to look for more advanced services very much like the example of the big global corporates inside monopolistic telecommunications markets. The only usable solution was a router network, which was originally meant only for private networks. At the time, the challenge was to develop a public router service using available equipment, which was technically feasible but not operating from the public service point of view. The demand for 2 Mbit/s and 10 Mbit/s connection bit rates was surprisingly high from the potential customers, both in the industry and public sectors.

The brand name LanLink service started to interconnect local area networks (LANs) in January 1991 with the real mesh router network in the backbone network instead of comprising leased lines between customer locations. The router in the customer premises was also included in the service because the management, simple network management protocol (SNMP), resided at the router level. The service is still growing as a result of wide connectivity, but it is open and has only Internet protocol (IP) level security.

The technology for closed company networks was invented soon

after, and 1993 heralded the provision of LanLink Frame Relay service from the Finnet Group. By the end of 1996, LanLink IP and Frame Relay had together over 3000 customer sites which meant busy backbones, so asynchronous transfer mode (ATM) transmission was implemented in the router backbone and in Frame Relay more recently. ATM is also available to customers as an alternative on the access network.

Information Services Evolution

At the same time as the network evolution, the value-added networks (VANs) have been developing. Public e-mail systems originated in the 1980s with proprietary systems and were later changed to X.400 systems. Initially, information menu services offered a character-based user interface to banking and public services and news groups and they quickly became very popular with many different service providers.

Because customers were available and used to VAN services, the natural evolution of VANs was Internet services and Internet service provision (ISP), which were implemented in 1995 under a brand name of *Kolumbus*¹. *Kolumbus* consists of full-scale Internet services to residential and business customers on top of narrowband (ISDN) or wideband (LanLink) connections. Today, the penetration of Internet in Finland is the highest in the world and has created a positive bit-flow balance (export – import). There is only one other country, the United States, with a positive bit balance.

In the Finnet Group it has been realised that the next evolutionary step is to provide broadband multimedia services for residential customers.

Becoming a Big Player in the Information Society

The traditional way of thinking of value creation is based on the

management of physical resources to create value and competitive advantage as presented in the value chain concept of Porter². In the value chain, information is seen only as a supporting function of the primary activities. However, according to Jeffrey F. Rayport and John J. Sviokla³ every business competes today in two worlds—a physical world of resources that managers can see and touch and a virtual world made of information. A company can create value in both of these worlds and to be successful the company has to understand the value-creation processes of each world and how they interact. The processes to create value are not the same in these different worlds.

The emergence of the virtual world is due to the fact that digitalisation of the telecommunication, information technology, media and entertainment industries has facilitated a new ways of gathering, organising and distributing information. Information in digital form is much more liquid than in printed or electronic (analogue) form. The ever-decreasing cost of computing, storage and communication capacity in the digital world is also increasing the possibilities of value creation in the virtual world. This revolution in technology is accompanied by a shift from the industrial world towards an information society.

The digitalisation of the above industries, especially the telecommunication and information technology industries, is bringing these industries closer and the boundaries are shrinking. For example, telecommunication companies are offering their customers access to router-based data networks and, on the other hand, companies in the information technology industry are offering point-to-point telephone services on the Internet. Thus, these industries are converging and a new industry called the information industry is starting to take shape⁴.

This new industry will provide solutions for managing long chains

of virtual-world processes. The services and products made available by the information industry can be offered straight to end-users or they can be a part of the value creation of other industries. Companies in other industries are today investing more and more in information technology.

As the converging industries have, until now, been separate and developed their own solutions to solve problems inside the scope of their individual industries, there is a need for a mutual framework to describe the roles and relationships apparent in the information industry. The European Telecommunications Standards Institute has developed an enterprise model which tries to describe and visualise the value creation in the information industry.

The information industry also needs an infrastructure to support it and many projects have started around the world to design it. In Europe this effort is called the European-Information-Infrastructure and the enterprise model is used to identify which requirements it will fulfil to support the information industry.

Finnet Group's data and information service development has been based on customer requirements and market reaction. The solutions have also been exploited in other customer fields. Nowadays, most information is already in digital form or can be more or less automatically converted. As a result of this, digital transfer, data storage and data processing can be utilised in almost any service. While the focus has been on the customers and the services they need and are willing to pay for, generic solutions have been developed. Cross functioning between industries has been retained for use again in the future.

A good example is the LanLink service as a generic service platform for business customers. Finnet Group has built step-by-step a broadband backbone network in Finland as an infrastructural

network for IP and frame relay virtual networking for business customers. Based on this successful evolution, Finnet Group is determined to use the same strategy for new services. As the development of international or European infrastructure is a slow process, the company has to build its own infrastructure. This work has to follow, as much as possible, the international development of the infrastructure standards. It can also provide valuable input to this work.

The emphasis is shifting from the development of business services (that is the information industry) to providing the services for marketing, trading and public services creation for end users. The shift is from information industry development to information society development. In technological terms this leads to the need for a public broadband network and a broadband access for end users. The explosion of Internet has been a breakthrough for this development. IP addressing and the birth of 'universal' browsers have been the enabling factors for the shift from business-to-business services to business-to-people, people-to-business and people-to-people services.

However, an universal infrastructure requires much more work. Security, electronic money and various legal issues such as intellectual property rights, taxation, customs, and electronic signatures still have to be solved. The architecture of the infrastructure will depend largely on the political decisions that will be made on these issues. So the structuring of a universal solution has just only started. In this situation a natural way to get further is to build public infrastructure islands to evaluate and validate the services and the technologies to be used in the future. For this task Finnet group has started two major projects, Helsinki ARENA 2000⁵ and Broadband Village⁶. Other major approaches described here have been made for

example in the publishing industry (Kärkimedia) and in education (Virtual Language School).

Helsinki ARENA 2000

Helsinki Telephone Company (HPY) has initiated *Helsinki Arena 2000*-consortium project in collaboration with the City of Helsinki. The project aims to develop a virtual Helsinki in the cyberspace and has received widespread support from business and public organisations. The project will culminate in the year 2000.

HPY and the other members of the consortium will create a three-dimensional model of the city on the information networks. Telephone numbers and other information linked to geographic information systems (GIS) will be included in the objects of the model. People will be able to move through this virtual city using their personal computers. The city's cultural, commercial and public services will be within the reach of everyone via the information and telephone networks. The aim of the project is to create a local multimedia network to which every home can connect. The project will combine the fully digitalised telephone network, HPY broadband networks and the Internet, as well as existing and new services in an integrated package easily accessible to ordinary people and based on the local urban landscape.

In addition to HPY, the consortium will consist of the City of Helsinki, the largest Finnish computer and communications companies, cultural institutes and art colleges. Providers of additional commercial and service input and other information, entertainment and media content will become members as the project progresses. The consortium has a very strong local market position. Each member will implement its own services independently on a free or commercial basis, and HPY will coordinate the integration of these projects.

The project will amount to at least 140 man-years, but there is scope for further expansion. HPY's own initial development contribution is valued at ECU 2.5 million. As the number of users grows, additional capacity will be needed on the information networks. Then an investment of hundreds of millions of ECUs may be required. No decision on investments of this scale has yet been made.

Helsinki Arena 2000 will be a remarkable electronic forum by the year 2000, when Helsinki is one of the Cultural Capitals of Europe and is celebrating its 450th anniversary. The consortium aims to build meeting points where people can meet and interact physically or virtually.

The project also heavily promotes PC-based videoconnectivity. Users will be able to experience a virtual Helsinki in which they can move

freely, looking at their surroundings, viewing public areas in real time, accessing electronic cultural services on-line, making purchases in shops, conducting business with officials, receiving information services, using toll or advertising financed entertainment services, making PC and videophone calls, holding videophone conferences, visiting amusement parks and casinos, and meeting members of clubs and associations.

Users will require an ordinary home computer to move through this virtual Helsinki. To use the most demanding interactive services, people will require a 200 MHz Pentium Pro standard computer, fitted with a video camera, a display accelerator, a sound card and an ISDN or broadband xDSL communications link. According to forecasts, such machines will be ordinary home PCs in the year 2000 in Helsinki capital area. An industry-standard

Internet browser with virtual-reality plug-ins and Internet call combined with video-communication plug-ins will also be required.

The city model and the rest of the system services will be connected to the user's own computer via the nearest backbone node. The city model will be drawn with ordinary architectural software and transformed on the Internet into a fixed three-dimensional model using virtual-reality modelling language. Doors, telephone-box directory information and a map of the city model will offer links beyond the basic services. Moving through doors into interiors will be like moving from one Internet page to another.

Video-on-demand and web-TV applications will be an essential part of the services. Earlier field trials⁷ have shown that the users are very satisfied with the picture quality using 2 Mbit/s MPEG1 transmission over asymmetric digital subscriber line access or rate-adaptive digital subscriber line access. The need for more symmetrical broadband access is also expected and hence the single-pair digital subscriber line access or the high bit-rate digital subscriber line access will emerge as the prices of these transmission systems are coming down in near future. Later, on a fibre-to-the-curb installation can offer high-definition TV quality using very high bit rate digital subscriber lines (VDSL) and existing copper subscriber lines. All digital subscriber line technologies are rapidly changing and developing⁸. Equipment prices, achievable bit-rates and extensions to operating distances over existing copper lines are developing towards the wide-scale use of winning configurations.

Helsinki Arena 2000 will develop telecommunications networks in an evolutionary way. Future development will be based on existing networks and standards. Interactive multimedia requires smooth, uncongested connections. In Helsinki, the existing HYP transmission network—the copper and optical-

Strolling in a virtual city

On my workstation I open a view of virtual Helsinki. I see Main Street (Mannerheimintie) and move to the Finlandia Conference Hall. I enter through the door and notice there is a public event underway. I go to the concert hall and even as I open the door I hear music flooding through it. On the edge of the hall, I see a camera point, which I click. A real-time view of what's happening in the hall appears on my screen.

After a while, I move from the concert hall outside into Hesperia Park. Here I see camera points under a 'Speakers Corner' sign. There are speakers on two of the cameras, and I listen with interest to one of them for a moment. He is speaking against technological advance.

I remember that I haven't finished filling in my tax declaration. I can't be bothered going all the way to the tax office so I go to the nearest telephone box, choose the tax office from the directory,

and am connected at once. I take the form I need from the form shelf, fill in my name, and continue filling in the form where I left off. I call a tax adviser and she answers the phone. She opens the video link. I obtain advice about my problem and continue filling in the form.

On completing the form, I remember that I promised to meet some friends. I move to a virtual clubroom, where even at the doorway I hear my friends' voices echoing around. On going inside I see all their faces on the walls except for Jarkko's, who doesn't have a camera on his computer. He has chosen a suitable tag to represent him. I still only have a 128 kbit/s link, so all the pictures besides the closest one are stationary images.

After a while I go and fetch a pizza. It isn't ready right away, but after I return to the clubroom the door bell rings and the pizza courier is at the door with a box. I charge the bill to my Merita Bank account.




Figure 1 – Helsinki ARENA 2000 network



Figure 2 – A view from the existing Helsinki ARENA 2000 3D model.




Figure 3 – Kärkimedia network for the major local newspapers in Finland.

fibre networks with ISDN, ATM and xDSL—will be used. The network will gradually be improved up to the year 2000 so that those households who wish can send and receive even television-quality moving pictures on their home PCs. The network architecture is shown in the Figure 1.

Today the platform network has three hubs in Helsinki and the first 30 end-user customers have been connected. The first version of the virtual city has been made. Figure 2 shows a scene from the cathedral square of Helsinki. The model will be used as one human interface for the services available on the network.

The model can be browsed by a VRML browser. It was built using a Lisp-based KCAD software from Arcus Software.

Publishing Industry

Another example of a service network is a network built for retrieving content for newspapers. The major Finnish local newspapers (26) have built a consortium for the delivery of advertising material to all the newspapers. The idea is to create a one-stop shop for customers, to rationalise the process of creating advertisements and possibly to use this consortium for targeted campaigns in various parts of Finland. The last issue was created to compete with TV-broadcasting companies having local advertisements in national programmes.

A network for this purpose was built using the facilities already existing in local newspapers. The network architecture is shown in Figure 3. Finnet constructed and now operates the e-mail server which is open to newspapers who publish seven days each week. The server receives all information from the customers and their advertising agencies, and provides billing information for the newspapers. The designed advertisements are held in the server and digitally transmitted to the newspapers in question. The data transmission takes place either through LanLink or Datanet (a competing network provided by Telecom Finland) or ISDN. As the network is in daily business use, it is possible to evaluate the need for new technology, in this case ATM, for the service. For this purpose a development project has been integrated into the network.

Today the network is used for new purposes. Newspapers realised the advantages provided by the network and now news and press photographs are delivered using it. Another interesting non-technological issue has been the cooperation (and competition) between Finnet Group




Figure 4. The concept of virtual language school

and Telecom Finland. As a result an infrastructure has been built for competing service (in this case also telecommunications) operators and the whole constellation has proved to be successful in business and gives advantages to all participants.

Virtual Language School

The goal of this project is to develop, demonstrate and pilot the language training implemented through interactive multimedia. The service will be independent of broadcast times, unlike television programmes. The students can take lessons whenever they want.

Training can take place at home, in the office or in the facilities of a teletraining centre. A student selects the multimedia lesson from the server and material consisting of video, audio, still images and text is provided. Guidance from the teacher is implemented through multimedia mail, videocalls or personal contact. The concept is shown in Figure 4.

Course material is located in the server, which a user accesses through a PC. Video files have been compressed (MPEG or AVI) before loading on to the server. Computer supported cooperative working (CSCW) tools (multimedia mail and desk-top videoconferencing) are used

in communication between a teacher and a student. PC access is implemented via telecommunications networks using ISDN and ATM connections.

Conclusions

The fast development of electronics and information technology is changing the development path of telecommunications. The evolution from mainframe computers to PCs and local networks with client-server architecture is now penetrating the traditionally-isolated industry of telecommunications and public networks. As this convergence is taking place, telecommunications development will be taking new paths. Customer-oriented service development and the revolutionary success of the Internet has narrowed the gap between traditional telecommunication and data networks. This evolution can be clearly seen in the development history of Finnish private telephone companies, the Finnet Group.

The new services currently offered by the Finnet Group take into account the continuation of this trend and a proactive strategy for this development has been created. The new services, the Helsinki ARENA 2000 network and services, the newspapers network and the virtual language school are good illustrations of this development. These services are based on the idea of a creation of network islands providing these new services and paving the way through interaction of these islands to gradual growth of a more universal broadband network, hopefully worldwide, in the near future. This evolutionary process resembles the creation of the worldwide Internet service.

References

- 1 <http://www.kolumbus.fi>
- 2 PORTER, M. E. Competitive Advantage—Creating and Sustaining Superior Performance,

The Free Press, New York, 1985, pp. 33-61.

3. RAYPORT, J. F.; and SVIOKLA, J. J. Exploiting the Virtual Value Chain. *Harvard Business Review*, Nov.-Dec. 1995, p. 75.
- 4 Overall Strategic Studies. Deliverable 1: Shaping the Information Society. EURESCOM Participants in project P311, 1995, p. 60.
- 5 <http://www.helsinkiarena2000.fi/>
- 6 <http://www.tpo.fi/BBV>
- 7 MÄNNISTÖ, E. (ed.); CLARKE, J.; DEFEE, I.; MURPHY, P.; and SAUVALA, S. Report of RACE DIAMOND Video-on-Demand Field Trials in Helsinki and Sligo 'Evaluation Report'. RACE R2105 DIAMOND Deliverable 35, 1996.
- 8 ABER, R. xDSL Supercharges Copper. *Data Communications International*, Mar. 1997, pp. 99-105.

Biographies



Heikki Sundquist
Helsinki Telephone
Company Ltd.

Heikki Sundquist received his M.Sc. degree in 1974, his Lic. Tech. in 1979 and his Dr. Tech. in 1982 in Metallurgy and Materials Science. Until 1982 he worked in the Finnish Academy of Sciences, Helsinki University of Technology and Salford University in thin film research. In 1982, he joined Lohja Co. to develop and manage manufacturing processes for EL-flat panel screens and terminals and in 1992 he joined Helsinki Telephone Company Ltd. where he is working as a research manager for the implementation of multimedia and broadband services.



Jari Karttunen
Omnitele Ltd.

Jari Karttunen received his M.Sc. in Electrical Engineering from Tampere University of Technology in 1989. He has been working in Omnitele since 1989 specialising in broadband networking covering access and backbone technologies including routers, ATM, xDSL and fibre technologies. He has been carrying out post-gradual studies regarding the influence of broadband to the telecommunications operator.

Authoring and Automating Content with its Associated Metadata for Multimedia On-Line Services

The range of multimedia services being introduced by BT is becoming increasingly extensive, especially those using the Internet. New, reliable, low-cost services must generate customer demand and sustain customer satisfaction. This article describes a framework for the authoring and distributed handling of content with associated metadata (information required to enable a service to be provided) accompanying an application. The use of a specification language to define metadata in conjunction with the tracking and verification of content are key aspects to minimising operational costs. The framework reflects the work being undertaken by the Digital Audio Visual Council.

Introduction

One of the technical directors from the computer games company, Sega, has been quoted as saying that 'the games platform is only an enabler; apart from price, three further things are required: content, content, content'.

This statement holds good for any interactive service. While the application may provide an interface and a 'look and feel', the content is what the consumer is seeking to access—and indeed most of what the consumer is paying for.

Content can be defined as the information which a service delivers. This content may be largely independent of the platform which delivers it, although the platform may limit the quality and richness of the content which can be supported and may also impose constraints on the format of the data. In this context the platform may be the World Wide Web, for example, or it may be an interactive TV service^{1,2}. The underlying commercial objective is to deliver content to a consumer, and so a potential service provider needs to study the experiences of a range of services and trials, regardless of platform, in order to prepare for potential problems.

BT's interactive TV (iTV) market trial confronted BT with the challenge of loading tens of thousands of content items onto a central server.

Items included video sequences (which could be several hundred Megabytes in size), previews or trailers, synopses, menu screens etc. If individual data fields are included, such as pricing and certification categories, the number increases by an order of magnitude. This content, from over 100 providers, was often supplied in a diversity of raw formats and required some processing (such as audiovisual encoding to ISO/MPEG standards³ often called *MPEG encoding*) before it could be loaded onto the server.

The cost of transporting large numbers of video tapes and the resulting encoded files was significant and not without its risks for the iTV trial—although the 'bandwidth' of a motorcycle courier is actually surprisingly large. Transportation of encoded content in tape form also introduced extra costs and delays in the manual loading and testing processes on the final server where the content was brought together to form completed programmes with their various descriptions, screens, trailers etc. This also made the scheduling of programmes more difficult owing to the uncertainties of availability of the finished content.

CampusWorld^{4,5}, BT's Internet-based service to schools, has faced the challenge of transferring content from a number of external providers onto a server and incorporating the appropriate testing and release

The project established at the outset the importance of developing processes that are consistent, where possible, with the requirements of standards organisations

processes. Responsibility for testing has been placed with the external provider, where appropriate, to reduce the uncertainty of its availability when it is due to go 'live'. However, there are still many manual elements in the transfer and release process which could be further automated to reduce the operational costs of the service as well as allowing more rapid update of material.

BT's experience with providing multimedia services has raised a number of challenges relating to content delivery which need to be addressed, and these can be summarised as:

- translation of content into the appropriate format for a specific service,
- automated or assisted testing of content,
- tracking and scheduling of content (that is, lifetime management of content),
- rapid and cost-effective transfer of large numbers of content items from provider to delivery server, and
- appropriate linkages to contractual agreements and financial transfer mechanisms.

A Strategy for the Development of Content Tools and Processes

BT's considerable experience in developing multimedia services has provided the opportunity to develop some reusable management tools; for example:

- to track the status of the content items and to allow content providers to download data to the server database via network links, and
- to transfer raw video to an encoder via BT's FacilityLine

service^{6,7} and to transfer the resulting MPEG files to a server via data links.

It has been necessary, up to now, for each new service to develop bespoke processes and tools to handle the transfer and loading of content. A separate programme, the Interactive Multimedia Services (IMS) Futures Campaign was set up to address the challenges that arose when introducing multimedia services building on past experience with such services. Specific aims of the campaign included:

- identifying a set of processes that could be applied during the life cycle of a multimedia service, and
- developing a demonstrator that shows how the processes work in an integrated way.

A workpackage has been set up to address the above aims and its strategy has been:

- to consider a service as comprising an application that provides users with access to content, and uses information about the service and content—this information, metadata, is required to enable the service to be provided;
- to define the requirements for the provision of facilities to ensure that suitable open content authoring tools may be developed; and
- to specify the interfaces and processes for automating the transfer of content and its associated metadata from an authoring environment to a delivery server making provision for the proper management of content and metadata.

A variety of proprietary tools is already available; for example, for encoding content from its raw format into a form that is accessible by end users at their premises. However, such

tools are not designed to work in an integrated way and provide solutions only for specific tasks.

This article outlines the processes identified and focuses on the following:

- *Specification of content and its associated metadata* By using a content metadata specification language (CMSL) within the application authoring environment, a standardised metadata format is produced—important if the format is to be incorporated in commercially-available authoring systems.
- *Tracking and scheduling of content* Tracking is concerned with the earlier part of the content life cycle; for example, recording whether content items have been tested satisfactorily. Scheduling is concerned with the latter part of the content life cycle; for example, the time when content is to be transferred to a delivery platform.
- *Testing of content and metadata* This is one of the most difficult processes to automate because it must be thorough; for example, to ensure that content is in a format which can be interpreted by the application. Testing must be performed as early in the life cycle as possible so that errors can be corrected without unduly affecting the schedule.

The project established, at the outset, the importance of developing processes that are consistent, where possible, with the requirements of standards organisations such as the Digital Audio Visual Council (DAVIC)^{8,9}. The interface between a delivery platform accessible by end users and the platform used by authors of applications/content (the DAVIC A10 interface) was of particular concern.

Links are also being established with commercial organisations to

Figure 1 – Content automation process model


encourage them to ensure that the content authoring and automation tools they develop are consistent with the processes developed by the project.

Processes for Automating Authoring and Handling of Content

Process overview

The content authoring process is considered in terms of a life cycle consisting of the following phases:

- **Acquisition** The procurement of any piece of content. This would ordinarily be raw content; for example, a photograph or a video tape.
- **Transcoding** The process of translating any piece of raw content into a suitable format for the target application; for example, scanning of a photograph or MPEG compression of video. There may be cases when the acquired material is already in the correct format and hence this phase may be omitted.
- **Quality assurance** The process of ensuring that the quality of the transcoded material is satisfactory; for example, checking that the artefacts on compressed video are sufficiently low.
- **Validation, verification and authorisation** The process of testing the content (and its associated metadata) to ensure that it is suitable, correct and complete, for the target application. Once tests are completed successfully, content is marked as authorised for transfer to the delivery platform.
- **Delivery** The process of transferring the content from the content producer to the delivery platform over the DAVIC A10 interface.
- **Loading** The process of distributing the content from the delivery point for the application to the



application itself. This process makes the content 'live' and is performed by the content loader.

- **Expiry** The process of preventing customers from accessing the content after it has served its useful or contractually agreed 'live' time.
- **Removal** The process of physically deleting the content.

The content automation process can be split into two main paths: the application production and the content production, as shown in Figure 1. Application production, the responsibility of application providers, would generally be a low-churn process in which the basic application would remain static. Content production on the other hand, the responsibility of content producers, is concerned with handling large amounts of data changing rapidly—that is, a high churn—and it is essential that content producers have an integrated tool-set to facilitate their work.

Applications are authored in advance of much of the content that will subsequently populate the service. Content and associated metadata, specifying the requirements and constraints that must be applied to content components, must be created in a format that the application can utilise. Metadata, as stated earlier in this article, is the information required to enable a service to be provided. Metadata is specified by an application author, in agreement with the content

producer, in a format defined by CMSL described later.

Content can be 'packaged' into two types:

- content item elements (CIE), which comprise the physical file (created and converted by a proprietary authoring tool) and its associated metadata; and
- content items (CI), which represent a collection of CIEs and/or other CIs; for example, a movie CI may consist of a film CIE, a trailer CIE and a synopsis CIE.

Testing of the application can be completed before transfer to the delivery platform by using a subset of compatible content and metadata. The platform for application testing will be similar to the delivery platform, although for simplicity the content loader, scheduler etc. are not shown in the figure.

The application authoring process also creates a *content loader* which resides on the delivery platform. As the loader is application specific, it understands the internal structure of the application and can intelligently distribute the content and metadata to the relevant locations in the delivery platform. The loader decouples the specific requirements of an application from the generic toolset.

Metadata required for both CIs and CIEs would be collated via a *content metadata editor* (CME). The CME extracts the CI/CIE metadata template from the content and metadata




Figure 2—Content life cycle

specification and enables the user to enter the values for each field; in addition to the metadata the CME also enables the grouping or association of components to a CI by extracting an association's template from the specification. When a new CI or CIE is created, the CME requests the allocation of a unique identifier (ID) via the production tracker, which creates a new entry in its database. This enables the production tracker to manage the content's progression through the content production life cycle to the required schedule. The CME provides a content producer with a graphical user interface (GUI) and defines the data that needs to be entered.

Once an item of content has been successfully tested, it is authorised for release. The scheduler manages the delivery of content to the delivery platform by specifying the items that need to be packed into containers and transferred. Transfer of containers may be performed electronically or via some form of physical media. On reaching the delivery platform, the container is unpacked and the content loader initiated to distribute the content and metadata.

Content metadata specification language

A multimedia application invariably involves the development of a custom-built tool set and application loader, both for use solely with the application and content/metadata specifications concerned. The toolset uses the specifications to verify that the content and metadata are correct before releasing them. Traditionally, the tool changes for each new application. CMSL is a tool designed to overcome the need for different tool sets by being able to understand a common metadata format.

Basing CMSL around a simple non-procedural language offers

complete reuseability, flexibility, portability and human readability for defining all types of multimedia content and its associated metadata.

CMSL source files are parsed by an interpreter. The interpreter provides an application programming interface (API) which the other tool set processes, and possibly the application loader, can call to obtain details on the CMSL constructs parsed.

The CMSL comprises several major syntactic constructs, each of which is enclosed with its own unique keyword and common end marker; for example, **ATTRIBUTE_SPEC.....END**, **CI_SPEC.....END**, etc. Within the bounds of a keyword and end marker, *attributes* can be declared to specify the details of the major construct. Parameters can also be associated with these attributes to define maximum/minimum values, default values, optional values, number of duplicate attribute occurrences, etc.

Major constructs can also contain references to other previously defined major constructs. These can be in the same or another externally referenced CMSL file. This means that, for instance, a standards-based set of available metadata attribute specifications could be defined in separate files and referred to as required. For example, DAVIC might produce a set of metadata attribute specifications for different services.

Six major CMSL syntactic constructs have been identified. These are briefly described in Panel 1 overleaf together with a simple example.

Content tracking and scheduling

Fundamental problems

For any given service, one of the fundamental problems is that of coordinating the different components that make up a CI so that all the components are ready before the

required live date. This is particularly complicated for systems such as iTV, where the different content types take significantly different paths to the delivery platform. For instance, the MPEG compression of a movie may have to be initiated well before the associated movie description is even written.

There is another problem of late or faulty content. For services like iTV, the long lead times on video content (and by long lead times we mean 1–2 days) means that lateness of a video can impact on the service being offered to customers. There is a need to identify at an early stage what components of a CI are late or problematic, and make the problems visible.

In addition, there is a fundamental need to reduce human interaction in all the stages of the content life cycle to reduce operational costs, minimise the risks of human error, and improve accuracy. Although it is impossible to remove human interaction altogether, many of the phases in the content life cycle benefit from having minimal human input. This is particularly true for the delivery, loading, expiry and removal phases.

Tracking content through the life cycle

In order to track content through the content life cycle, two specific tools have been designed—the production tracker and the scheduler. Both perform very similar tasks in that they both schedule when particular events in the life cycle should be completed and keep a record of completed events (these tools are separate from tools relating to traditional post production special effects processes which are outside the scope of this article). Their differences lie in the fact that they each regulate different parts of the content life cycle. Figure 2 shows the content life cycle and how it is split

Panel 1—Major CMSL syntactic constructs

Note that CMSL reserved words are shown in BOLD characters and that, in keeping with many programming languages, comments are allowed if preceded by a double slash '//'.

Content item element types

These define the type and format of the content that can be used by the application. The example shows the constructs for a small image and audio file. These specifications are referenced from CIE specifications and must be defined before they are used.

```
//This content type specification states that the content is an image.
//It must have a format of either JPEG or GIF. Its maximum frame
//size is 100 by 100 and can consist of 256 colours maximum.
CIE_TYPE small_image
    IMAGE
    FORMAT JPEG, GIF
    FRAME_SIZE 100,100 MAX
    COLOURS 256 MAX
END
```

```
//This content type specification states that the content is audio. Its
//only requirement is that it is of type MPEG1.
CIE_TYPE mpeg1_audio
    AUDIO
    FORMAT MPEG1
END
```

Content item element specifications

These specify a content type to use, by referring to a previously defined CIE type, and a content file size.

```
//This CIE specification states that its content must conform to the
//defined CIE_TYPE of small_image and have a maximum image file
//size of 144 000 bytes.
CIE_SPEC product_image_normal
    CIE_TYPE small_image
    FILE_SIZE 144000 MAX
END
```

```
//This CIE specification states that its content must conform to the
//previously defined CIE_TYPE of product_background_audio. Its
//file size is not relevant.
CIE_SPEC product_background_audio
    CIE_TYPE mpeg1_audio
    //no maximum file size specified
END
```

Content item specifications

These are used to declare which CIE specifications and which other CI specifications are associated with this one. Some associations may be optional or occur a number of times.

```
//This CI specification for a footwear product states that CIE_SPEC
//product_background_audio is optionally associated; that is, the
//specification is complete with or without it.
//There must also be a minimum of one or a maximum of two
//associations of CIE_SPEC product_image_normal.
CI_SPEC footwear
    CIE_SPEC product_background_audio OPTIONAL
    CIE_SPEC product_image_normal OCCURS 2, 1
END
```

Attribute specifications

CI or CIE metadata specifications (see below) declare which metadata attributes a CI or CIE metadata record will include. Before they can be included they must have been previously specified using an attribute specification construct.

Each specification includes a descriptive title, a type, for example, numeric, boolean, string, etc., and a group name to identify groups of related attribute specifications. Also, depending on the type of attribute specification, other parameters are included; for example, integer ranges, default values, enumerations, etc.

```
//This declares an attribute specification called 'footwear_type'
//belonging to the group of attributes called 'footwear'. It has a title
//of 'TYPE OF FOOTWEAR'. The attribute's value must be a string
```

```
//of a maximum eight characters. This attribute can only have
//values of 'sandals' and 'shoes'. When a metadata record is first
//created, this attribute must be given the default value 'shoes'.
```

```
ATTRIBUTE footwear_type
    TITLE 'TYPE OF FOOTWEAR'
    TYPE STRING
    LENGTH 8
    DEFAULT 'shoes'
    STRING_ASSIGNMENT 'shoes'
    STRING_ASSIGNMENT 'sandals'
    GROUP_NAME 'footwear'
END
```

```
//This declares an attribute specification called 'size' belonging to the
//group of attributes called 'footwear'. It has a title of 'SIZE'. The
//attribute's value must be numeric and in the range 1 to 10. When a
//metadata record is first created this attribute has no default value.
```

```
ATTRIBUTE size
    TITLE 'SIZE'
    TYPE NUMBER
    RANGE 1,10
    GROUP_NAME 'footwear'
END
```

Content item metadata specifications

These are used to specify which of the available attribute specifications (described above) are to be used in the CIE metadata record. Besides the selected attributes from the available attribute specifications each metadata record has, by default, the same set of generic attributes. These consist of a unique ID, a record title, an abbreviated title, a description and a content producer reference. If the generic attributes are not specified in the CMSL metadata specification they must still be included in the metadata record created from the specification.

```
//This declares a CI metadata specification called 'footwear'. Any
//metadata records created from it will contain metadata fields for
//Unique ID, Title, Abbreviated Title, Description, Content Producer
//Reference, Size and Footwear Type.
CIM_SPEC footwear
    GENERIC
        UNIQUE_ID
        TITLE
        ABBREVIATED_TITLE
        DESCRIPTION
        CONTENT_PRODUCER_REF
    END
    //Selected CI Metadata Subset
    ATTRIBUTES_USED
        size
        footwear_type
    END
END
```

Content item element metadata specifications

These are essentially used in the same way as CI metadata specifications but specify a content item element metadata record instead.

```
//This declares a content item element metadata specification
//called product_image_normal. Any metadata records created from
//it will contain metadata fields for Unique ID, Title, Abbreviated
//Title, Description and Content Producer Reference.
//Notice that even though the generic attributes are not specified in
//CMSL they will still be created in the metadata record. A content
//file size metadata attribute is required in this metadata record.
//No previously specified attribute specifications are used.
CIEM_SPEC product_image_normal
    FILE_SIZE
    //Selected Metadata for CIE
    ATTRIBUTES_USED
    END
END
```

One of the most difficult processes to automate can be the testing of content and its associated metadata.

so it can be managed by the production tracker and the scheduler.

This split is necessary in order to allow content reuse across multiple applications. The content production phases (from acquisition through to validation, verification and authorisation) are not application specific. They are common to all applications, and hence can share the management functions supplied by the production tracker.

Obviously, the delivery, load, expiry and removal times are application specific. Coordination of each of the phases for each piece of content that comprises a CI is ultimately the responsibility of the scheduler of the application for which they are destined. The loading of content at the load time is the main aim of the scheduler/production tracker combination; hence, the load time dictates the timing of all the preceding phases in the content life cycle for each piece of content.

For both the production tracker and the scheduler there will be a reporting mechanism that will identify any content items or components thereof which are late and have not completed any phase on time.

Production tracker

Default times will be assigned to a particular application for every content production phase for each type of content that appears in the application's content and metadata specification. A GUI will allow a user to edit the times for any piece of content in the production tracker, but will not allow the validation, verification and authorisation time to be altered as this would directly affect the application for which the content was destined. However, it does allow the content production team to change its timescales for producing content.

The production tracker manages the content flow from acquisition through to validation, verification and authorisation. The production tracker provides the facility for registering the times each of these

phases should be completed. This input would be provided either manually, or automatically by the scheduler using the default times. Other facilities allow the production tracker to be informed of progress as content passes through the life cycle phases, by storing dates of completion for each phase.

Content components will be identified for which:

- the completion date of the current phase is less than that given by the user,
- the previous phase has been marked as complete, and
- the present phase has not been completed. This can be used to produce a 'to do' list for people working on content production.

There will normally be one production tracker per content producer; however, it may be decided to use different production trackers for different projects. It is a choice that will be influenced by the scope of projects taken on by the content producer, the size of the projects and the terms of agreements made with other contributing participants to the service.

Scheduler

The scheduler manages the content flow from delivery through to removal and, as in the case of the production tracker, provides facilities for scheduling when each of these phases should be completed. When each of these phases is completed by the delivery process (for the delivery phase) or the loader (for the other phases), the scheduler is informed of progress. These are automatic processes which require no human interaction.

The loader scheduler also produces a list of components for which:

- the completion date of the current phase is less than that passed to it by the calling process,

- the previous phase has been marked as complete, and
- the present phase has not been completed.

There will normally be one scheduler per application which will interact with one or more production trackers. This allows the use of multiple content producers for any application and provides input on when content items (and more specifically their components) are required.

Testing and releasing content and metadata

One of the most difficult processes to automate can be the testing of content and its associated metadata. Analysis of a wide variety of services has shown that the testing requirements of multimedia content can be reduced to three main elements:

- **Format** Is the content in a format which can be interpreted by the application?
- **Completeness** Are all the components of the content offering present; for example, for a movie, the trailer, pricing, description etc?
- **Identification** Does the item match its label (for example, the piece of content labelled 'Mickey Mouse' really is a movie featuring Mickey Mouse and not an adult movie.)

Although these elements could be tested at different stages in the process, it is important they are tested as early as possible to avoid further delays to dependencies. The fundamental requirement of the test and release stage is to provide an acceptable level of confidence that an application's content is fully defined and correct. One solution is to associate some form of 'certification' label with the content at the production stage which places the

responsibility for testing with the content producer.

Test

The test component of the test and release unit provides the compliance and completeness testing activities required prior to transfer of content and metadata. The test requirements described above can be further broken down into specific tests.

Testing the format involves checking that a physical item (for example, video) meets the requirements of a specific application. For example, an application may require all of its movie type CIEs be MPEG2 encoded or that image type CIEs are JPEG or GIF format with a maximum of 256 colours. Additionally, an application may wish to limit the physical size of items to constrain the amount of data downloaded by the application at run time, in an attempt to benchmark performance. By interrogating the physical item it should be possible to confirm that the actual item exists and that it has met these requirements as outlined in CMSL.

Testing for completeness involves verifying that the CI or CIE has the correct set of metadata associated with it, as defined by the CMSL block that the item was built to. For example, has the CI been assigned a price, does it have the correct physical CIEs to make up the required 'offering' etc?

Validation of the actual metadata values is performed as the user enters them via the CME. For example, the CME ensures that if a value is to be of type 'number', as defined in CMSL, then it is of no other type. In addition, CMSL allows more rigorous validation by providing the capability to define metadata attribute constraints. For example a movie 'Rating' attribute may be constrained to take values from the following set {U,12,PG,15,18} and the CME would enforce these constraints. The use of CMSL in this way provides two key benefits: firstly, it simplifies the data-entry process as the tool may provide picklists etc.

from which the user can make selections, and secondly, it ensures, at the earliest opportunity, that the value is syntactically correct.

A requirement for batch loading of metadata was evident during the iTV market trial where several content producers already held their data in legacy databases and needed some way of automatically transferring the data from one area to another. It is therefore important that batch loading tools are capable of providing the same level of validation as described in the previous paragraph. It is expected however that the application provider may, in conjunction with the content producer, use existing data models to define CMSL structures in which case there should be a minimum translation and a maximum reuse of data. In addition, the fact that the data has already been used by another process should provide some level of confidence that it is correct.

Additional tests include quality assurance, which is performed by a process external to the test and release unit, and involves testing and accepting the quality of a transcoded physical item; for example, checking the quality of an MPEG encoded video. This type of test is particularly difficult, not only because of its subjective nature but also because of the fact that full automation of the process is almost impossible. Similarly, testing for identification correctness is a labour-intensive task and cannot be easily automated. For this reason it would be typical to combine both of these tests and perform them at the same time.

The toolset manages the test workload via dates held in the production tracker. A 'Test_by' date is used to determine which CI/CIEs are to be tested on a certain date. The results of these tests are held in a result mask which is used as the 'certificate'. Should a test fail then this will be recorded in the result mask along with the reason for failure, so that when the toolset next attempts to authorise the CI/CIE, it need only perform the tests that

failed previously. Notification of failures is fed back to the content producer so that appropriate action can be taken.

Once an item of content has passed all of its tests it can be authorised for release and transferred to the target delivery platform either electronically or via some form of physical medium.

Release

The release component is responsible for packaging the metadata (content items and content item elements) and physical content into containers to cross the DAVIC A10 interface mentioned in the Introduction. Requests from an application for content are handled by the content producer's transfer process. This then communicates with the release process requesting the creation of the specific container.

A container comprises a header component, a metadata component and physical item components. The header contains information regarding the contents of the container and what actions should be taken with them. The metadata component contains the serialized metadata objects and the physical item components are the physical content; for example, MPEG stills. The transfer process is responsible for the security, encryption, and charging (if applicable) for the requested items within the container. The release process maintains a database of requested containers storing information about the contents of a container, the destination of the container, the date and time the container was sent etc.

When a piece of content is physically very large, for example, a video, it may not be practical or economical to send it electronically. To cater for this, each request for an item must be accompanied by a delivery mechanism (for example, tape, electronically etc.). If any content is to be transferred via a mechanism other than electronically, then a list must be created and acted on accordingly by the content producer.

Figure 3—Developing a complete service from raw content

Achievements

Validating the framework against service scenarios

An overall framework has been developed that has been tested against, and has been found to be consistent with, the scenarios for a number of multimedia services including the iTV market trial and the Campus World service. The following two scenarios are examples to represent extremes in the range of services that may be offered by a multimedia service provider:


- developing a complete service, and
- renting server space.

Separate figures for each service (Figures 3 and 4) use shading to show how the responsibilities for undertaking individual processes vary.

Developing a complete service

This scenario is one whereby a film production company would be responsible for providing only raw content and basic information about that content (Figure 3). The other participant, the multimedia service provider, takes on the roles of application provider and content producer. Specific processes which that organisation undertakes are:

- Application authoring and the specification of content and associated metadata.
- Content authoring which will include the main areas of tracking, testing and updating the content at will using the loader. The raw content supplied by the third party is converted into a suitable format for the application during the content authoring process (for example, MPEG encoding).
- Application testing whereby the application and loader are transferred to a test environment



where the application is tested and verified using content supplied from the content authoring process. When application testing is complete, the application is transferred to the live platform and content is updated at will using the loader which will:


- move content item elements to the appropriate delivery platforms, and
- update the application database to include the new content.

Renting server space

This scenario is one whereby the multimedia service provider is essentially renting server space to a third-party company (Figure 4). Specific processes which the third party undertakes are:

- specification of content and associated metadata with the development of the application and loader;
- content authoring with tracking and scheduling, and testing, and

Figure 4—Renting server space



It is now possible to downstream the achievements of the project into new multimedia services to provide benefits of reduced development costs and shorter development times

- transfer of the application and loader to a test environment when application testing is complete and content is updated at will using the loader.

Specific processes which the multimedia service provider undertakes are restricted to application testing whereby the application is tested and verified using content supplied by the third party.

Conclusions

Considerable progress has been made in tackling the challenges relating to content authoring and delivery processes. Specifically, progress has been made with the following aspects:

- writing the CMSL interpreter as a Java package¹⁰;
- encoding content into the appropriate format for a specific service—current development needs to be made consistent with the overall process model;
- testing content; and
- tracking and scheduling content.

BT's association with DAVIC has enabled it to identify which participants have interests relating to content authoring and management. Any resulting opportunities may be exploited to initiate dialogues with those participants directly as well as within the DAVIC forum and enable BT to gauge the extent to which project results may be used on a wider basis. For example, BT has suggested that there could be greater flexibility when allocating responsibilities for undertaking specific processes without having an adverse affect on the significance of the DAVIC A10 interface.

It is now possible to downstream the achievements of the project into new multimedia services to provide benefits of reduced development costs

and shorter development times while providing quality services that end users have come to expect. The increased automation of processes has contributed significantly to the benefits that have been gained, and has also contributed to lower service operating costs.

References

- 1 KERR, G. W. A Review of BT's Interactive TV Trials. IEE Colloquium on Interactive Television, London, 2 Oct. 1995, pp. 6/1–5.
- 2 WITHNELL, J. BT Interactive TV Delivering Multimedia Services to the Home. IEE Colloquium on The Impact of Multimedia Services on the Home Environment, London, 12 Jan. 1996, pp. 2/1–2.
- 3 ISO/IEC. Coding of Moving Pictures and Audio for Digital Storage Media at up to about 1.5 Mbit/s. Committee 11172-1, 1991.
- 4 For CampusWorld information pack contact 0345 626253.
- 5 For CampusWorld, <http://www.campus.bt.com/CampusWorld/>
- 6 MAY, D. D1-Quality Video Transmission Marks End of Road for Runners. International Broadcasting, Feb. 1995.
- 7 For FacilityLine, http://www.vbs.bt.co.uk/bt_bs/faciline.htm.
- 8 CHIARIGLIONE, L. Digital Audio-Visual Council—Rationale and Goals. *EBU Technical Review*, No. 266, Winter 1995, pp. 51–63.
- 9 For DAVIC, <http://www.davic.org>
- 10 For Java, <http://java.sun.com/allabout.html>

Biographies

Stephen Searby
BT Networks and
Systems



Stephen Searby joined BT in 1979 as a sponsored student. He graduated from Oxford University in 1983 with a B.A. in Engineering Science and joined the Local Networks division at BT Laboratories to work on an interactive video library for the BT switched-star cable TV network. After a period of work on knowledge-based coding techniques for human faces, he became head of the advanced TV systems group which included collaborative research work on coding and transmission of TV and HDTV. More recently he has been project manager for the delivery of server platforms and MPEG encoding for BT's interactive TV market trial and is currently managing teams delivering platforms for new Internet-based services as well as the design of future multimedia content automation systems. He is a member of the Institution of Electrical Engineers, the IEE professional group committee for television, radio and data broadcasting, the Society of Cable Telecommunication Engineers and is a Chartered Engineer.

Janie Taylor
BT Networks and
Systems



Janie Taylor graduated from London University with an honours degree in Maths and Management and joined BT in 1981. Most of her time at BT has been in video-related topics, primarily involved in the software design and

development of the BT cable TV video library and later in multipoint audio-visual services. More recently she has been involved with the early stages of the iTV market trial and now in the IMS Futures work.



John Salmon
BT Networks and
Systems

John Salmon graduated from Newcastle University in 1984 with a degree in mathematical computing, and joined BT in 1988 after working for EASAMS and Admiral Computing. He spent his initial period with BT working on security in the areas of security policy development and user authentication. This was followed by a period assessing the suitability of using ISDN for X-ray image transfer for the National Health Service. Later, he moved over to systems testing for the interactive TV trial, and is currently concerned with working on developing generic content authoring and automation processes.



Paul Foster
BT Networks and
Systems

Paul Foster started his career in BT as a Technician Apprentice in 1973 and then, with a BT Minor Award, gained an honours degree in Computer Science at Leicester Polytechnic. Since 1983 he has worked at BT Laboratories and has developed systems and software for photovideotex servers, Talking Pages and BT's directory enquiries teleworking experiment. For the past three years he has worked on multimedia systems, in particular airline information services and IMS Futures.



Royston Walker
BT Networks and
Systems

Royston Walker joined BT Laboratories in 1994 having obtained a B.Eng. in Electronic Engineering from the University of Hull and an M.Sc. in Telecommunication and Information Systems from the University of Essex. Since then he has been involved with the development of multimedia services working on BT interactive TV and BT Touchpoint. He is currently a member of the On-Line Services Unit investigating management and delivery mechanisms for multimedia content as a part of the IMS Futures Campaign.



Richard Gepp
BT Networks and
Systems

Richard Gepp joined BT as a Trainee Technician Apprentice at BT Laboratories in 1986. After graduating in 1993 with an honours degree in Electronics from the University of Warwick, he returned to the laboratories to work on the development of Facility Line. He has been involved in audiovisual compression for video on demand services since the technology trials in 1994 and was responsible for BT's encoding facilities during the interactive TV market trial. He is currently developing content management tools for future interactive multimedia services.



Sally Mossman
BT Networks and
Systems

Sally Mossman joined BT Laboratories in 1989 as a Trainee Technician, and in 1994 completed her British Computer Society qualifications. She has been employed in a range of activities, including a short spell in the Home of the Future team, software development within the radio science area, and work on ADSI in the Network Intelligence Centre in 1993. Since 1995, she has been in Products and Services Engineering working on service creation and most recently the IMS Futures Campaign.

Customer Access Signalling on Concert VNS

Customer access signalling is a complex part of a customer's virtual private network solution, especially in the global environment in which Concert Virtual Network Services (VNS) operates, given the wide variety of signalling types available. With convergence to one or two global standards still an ideal, much confusion can be avoided by better understanding of the range of options available. Concert has over two years' experience of using the range of available signalling systems to connect customer equipment to the Concert VNS network.

Introduction

Concert Virtual Network Services (VNS) is the world's first global intelligent voice service and offers a wide range of signalling interfaces to connect to the broadest possible range of customer equipment in each country. With PBX manufacturers continuing to release proprietary feature-rich signalling variants, and with PTTs slow to adopt international standards, this article explores the issues of 'incompatibility' and how Concert has used over two years' experience of connecting customers to evolve its approach to customer signalling.

Concert VNS

Concert VNS is the world's first global virtual voice network to be owned and managed as a single entity. It aims to provide voice services, including a range of access methods and features, to customers with sites throughout the world. At the same time it provides the cost benefits of sharing international transmission, and the service management benefits of a single supplier-owned network. Concert VNS currently has *points of presence* (nodes) in UK, France, Germany, Netherlands, Belgium, Spain, Italy, Switzerland, Sweden, Australia, Japan and throughout the USA. Coverage is also provided in other countries such as Austria and Luxembourg via Concert's Global Managed Platform.

Concert VNS Architecture

Concert VNS is based on an intelligent network (IN) architecture in

which the 'intelligent' call processing functions related to number plan and features are removed from the core service switching points (SSPs) and centralised in a database (DAP). SSPs are located throughout the world and US coverage is provided by MCI's Vnet switches, which also have access to the DAPs. The US switches act as both SSPs and points of presence (PoP). Outside the US, each SSP is connected to one or more PoPs which not only concentrate the traffic for that country to the SSPs, but also provide a wide range of access signalling interfaces to allow connection to customer equipment.

In the most widely-used configuration, customer PBXs are connected to the PoP using a private leased dedicated access line (DAL). In the USA, digital T1 (1.5 Mbit/s) access lines are used. In Europe, digital E1 (2 Mbit/s) and analogue (4-wire) lines can be used. The DAL conveys not only the voice, fax and data traffic, but also the appropriate signalling to allow the PBX to communicate effectively with the PoP.

Each PoP supports a wide range of access signalling systems, the aim being to support the widest possible range of customer equipment throughout the world. Concert VNS acts on the basic call set-up functionality which is common to all signalling systems, which offers the advantage of a consistent set of network-based features globally, regardless of the signalling used at each site.

World of Signalling

Signalling is one of the most complex elements in a network. A

very large number of signalling systems exist, analogue and digital, simple and complex, which allow networks, exchanges and PBXs to communicate.

Signalling has evolved from primitive AC and DC analogue versions to modern digital signalling systems. Whereas the earlier digital signalling systems (channel-associated signalling (CAS)) were simply digital versions of analogue signalling systems, more modern common-channel signalling (CCS) systems switch intelligent messages in a similar way to packet switched networks. Also, while some CCS systems are designed primarily for use in private networks between customer PBXs, offering a multitude of features (for example, DPNSS, QSig), others are designed primarily for connecting 'intelligent' PBXs to 'less intelligent' networks—these are the so-called ISDN protocols. Still others (for example, CCITT No. 7 or SS7) are used exclusively within core networks. All are in use today (Figure 1).

To further complicate matters, whatever standards exist have been so loosely defined that individual PTTs and equipment manufacturers have tailored them to their own specifications. Thus many PTTs who support ISDN connections on their networks have developed their own, unique variant of ISDN signalling

(for example, 1TR6 in Germany, DASS in UK and VN3/4 in France). PBX manufacturers in those countries therefore need to support the local ISDN variant. However, PBX manufacturers have also developed their own, proprietary, feature-rich PBX-to-PBX protocols for use in private networks, and continue to develop ever more sophisticated variants (for example, Siemens CORENET, Nortel Q.931). Needless to say, virtually all are incompatible.

While access signalling in the USA appears to have converged towards two or three standards, this is not the case in the rest of the world. In Europe, a *Euro-ISDN* standard has been defined by the European Telecommunications Standards Institute (ETSI), but is only slowly being adopted by PTTs. PBX manufacturers are slowly following suit. In addition, a group of PBX manufacturers has formed a group (European Computer Manufacturers Association (ECMA)) to define a feature-rich PBX-to-PBX protocol which will allow PBXs made by different manufacturers to communicate with one another. The result, QSig, is now implemented on many new PBXs, although feature compatibility still remains an issue.

With ETSI ISDN yet to be adopted by all PTTs and PBX manufacturers, and a variety of customer equipment, digital and

analogue, still very much in existence, any global network operating in Europe (and elsewhere) must be capable of supporting a range of signalling options, or risk problems of incompatibility.

Access Signalling Support on Concert VNS

The US portion of the Concert VNS network is well-established; signalling in the USA has converged towards two or three well-established types, and for the purposes of access signalling, Concert VNS uses the same Vnet switches in the US to which MCI have been successfully connecting customers for many years.

However, in order to deal with the wide range of signalling systems available outside the USA, Concert VNS supports the following in Europe and in Australia:

Digital:

- DPNSS
- QSig
- ETSI ISDN
- CAS (DC5 and R2 varieties)
- 1TR6 (German ISDN)


Analogue:

- AC15A
- CEPT L1
- E&M V

DPNSS is a well-known and well-defined standard in the UK. Originally developed primarily as a feature rich protocol for use between customer PBXs, its standards allow for its use for connecting PBXs to networks. Indeed it is the standard protocol for connecting customer PBXs to FeatureNet and International FeatureNet in the UK. On Concert VNS, DPNSS is by far the most often-used signalling system in the UK, and has also been successfully used in other countries.

QSig is a feature-rich signalling system which can be used throughout Europe between different

Figure 1—Access signalling types in Europe/Australia



manufacturers' PBXs. Based on the Q.931 standard, it is effectively designed to play an equivalent role in private networks as DPNSS, but with greater acceptance outside the UK.

ETSI ISDN is designed to converge the wide variety of ISDN access signalling used throughout Europe towards a 'Euro ISDN' standard. Although it is expected to be adopted by European PTTs over the next few years, it is not widely used at present and PBX manufacturers are only now beginning to implement it on their PBXs.

CAS is a cheaper, digital alternative to the more sophisticated CCS systems quoted above (DPNSS, QSig, ETSI ISDN). It supports few, if any, features, but this is perfectly adequate for connection to Concert VNS, since Concert VNS acts upon basic call set-up messages. CAS exists in many forms and variants. Concert VNS supports R2 CAS and E&M CAS.

1TR6 (German ISDN) is supported in Germany, and has been used successfully in many customer connections. However, 1TR6 will gradually be replaced by ETSI ISDN (EDSS1) in Germany.

AC15A and *CEPT L1* are analogue, AC signalling systems conveyed in-band on four-wire analogue lines. AC15A is used predominantly in the UK, whereas CEPT L1 (also known as *AC15D*) is widely used in Europe. Both are useful for connecting low volume sites, where the cost of a digital access circuit cannot be justified.

E&M V is an analogue, DC signalling system designed primarily for use within buildings, since the DC signal cannot be conveyed over any significant distance. As such, it is not normally used for connection to Concert VNS.

Problem of 'Incompatibility'

Given the wide variety of signalling systems, standards and variants, care is needed when

specifying which signalling system is to be used when designing a customer solution. For example, the often-quoted Q.931 could mean QSig, ETSI ISDN or any number of incompatible manufacturer variants; Concert VNS supports only ECMA QSig and ETSI ISDN, and these are not mutually compatible. CAS can exist in many different forms (for example, DC5A, wink-start, R2) with many different options built in (for example, DTMF or pulse addressing). 'E&M' could mean analogue E&M or digital E&M CAS. Both analogue E&M and E&M CAS exist in more than one variety.

The scope for confusion is compounded by differences in terminology and language in different countries around the world, and the fact that a manufacturer may support a signalling system on its PBX in one country, but not in another. Inevitably, this can give rise to frustrations when selecting or recommending an appropriate signalling type.

Nevertheless, experience of connecting customers has shown that in virtually every case, a mutually compatible signalling system can be found for connecting a PBX to a Concert VNS PoP. The only exceptions are much older PBXs which support only analogue E&M (analogue E&M is supported by Concert VNS, but since it is a DC signalling system, it cannot be transported long distances without conversion to digital CAS or analogue AC signalling, and there are converters on the market for this purpose).

All of the available Concert VNS signalling systems have been used at one time or another for connecting customer PBXs to Concert VNS and most leading makes/models of PBX have been connected (for example, Nortel Meridian, Siemens Hicom, AT&T Definity, Alcatel Opus, Ericsson MD110, Philips etc).

By far the most popular signalling system used is DPNSS. It is one of the most established well-known

and well-defined common-channel signalling systems in the UK. It is supported by most major PBX manufacturers in the UK, and some elsewhere. Therefore, in the absence of a well-established PBX-to-network signalling standard (for example, ETSI ISDN), the support of DPNSS on PoPs throughout Europe has been important in establishing confidence in the Concert VNS product. Demand for QSig is increasing in Europe and Concert's experience and confidence with this signalling system is rapidly increasing to match that of DPNSS. Rare initial troubles with call diversion to voicemail, relating to the use of QSig and DPNSS as PBX-to-network protocols rather than PBX-to-PBX protocols, have now been fully resolved.

More or Fewer Signalling Systems?

Overall, reducing the options for connecting to Concert VNS has certain attractions: reduced time in decision-making, better-focused resources and expertise, leading to improved delivery lead times and faster fault diagnosis and repair. Ideally, signalling would converge towards a single internationally-accepted standard. However, although ETSI ISDN should be a suitable standard in the future, it is not yet well-enough established to allow standardisation without having to turn away a significant amount of business.

Nevertheless, Concert VNS has no plans to support manufacturers' proprietary PBX-to-PBX signalling systems. Not only would the development costs of keeping up with the manufacturers be prohibitive, but experience has shown that PBXs which support such proprietary protocols also always support more widely used protocols, such as those supported by Concert VNS; therefore, the development costs and added service complexity would not be justified.

However, recognizing the particular demands of the French and German markets, Concert will continue to exploit the capabilities of its existing hardware to the full, continuing support for 1TR6 in Germany alongside ETSI ISDN (EDSS1), and planning the introduction of proprietary signalling in France.

Conclusion

Despite fears of incompatibility, experience shows that a suitable signalling system can virtually always be found for connecting customer equipment to Concert VNS. Many leading makes/models of PBX have been connected to Concert VNS using the range of signalling systems available. There is therefore no reason to increase the complexity and running costs of the network by significantly increasing the range of signalling supported in Europe and USA. For example, Concert VNS has no plans to support manufacturers' proprietary PBX-to-PBX signalling systems.

On the other hand, the time is not right for convergence to a single standard, such as ETSI ISDN. The experience Concert has gained with alternatives, such as QSig and DPNSS will enable the continued use of these protocols with confidence. Concert also recognises the need to be able to respond in the short term to the particular needs of some national markets and is planning to introduce a proprietary access signalling type in France.

In general, greater experience of connecting customer equipment is leading to reduced time in making design decisions, reduced time in implementing service and improved fault diagnosis. This in turn should lead to improved customer satisfaction and save the significant revenue which can be lost while trying to resolve signalling issues during implementation.

As users continue to demand new services and as technology

advances, signalling will continue to be one of the most complex areas of network design and implementation. The experience gained in over two years of connecting customer equipment to its network means that Concert is well placed to meet this challenge.


Biography



Sanjiv Angrish
Concert

Sanjiv Angrish has been associated with virtual private networks since joining BT International from Aston University in 1990. As Operations and Maintenance Planning Manager, he supported the launch of International FeatureNet and became involved in Cyclone, BT's project to build the world's first global VPN. As Management Systems Design Manager, he led a team in negotiating a procurement contract and conducted a study into how to integrate BT's global product set, an issue of considerable interest when Concert, the BT-MCI joint venture, was formed. In Concert, Sanjiv worked as a customer network design engineer, supporting bids and implementations for major multinational customers. In October 1994, Sanjiv was called to the USA to design and project-manage the implementation of the solution for the first global voice customer, in time for the launch of the Concert VNS product on 1 November 1994. As senior engineer, Sanjiv continued to support bids and implementations and was responsible for supervising a team of designers, before moving to the USA in July 1996 to start work in his new role as Product Development Manager.

MuseNet



MuseNet, BT's national contribution to the Year of Engineering Success, is an audio and videoconferencing network linking seven major museums nationally: The Science Museum, The British Museum, The Ulster Museum, The National Museum of Wales, The National Museums of Scotland, The National Museums on Merseyside and The National Museum of Photography, Film and Television at Bradford. Museums were chosen for the broadness of their collections and because they had elements within their collections which they could share and were complementary to each other. This was an important consideration in forming the MuseNet museum project because the museums involved had to be prepared to extend boundaries in terms of their internal cultures and organisation.

Participating museums are linked together through an ISDN6 network with facilities to enable simultaneous reception of events and programmes originating from one site and to enable all sites to interact through a multi-link. BT's latest VS2 movable videoconferencing unit has been linked via a video-projector to a large screen display which enables direct interac-

tive communication with audiences. Video cameras and document readers extend the range and scope for the presentation of material.

MuseNet was developed with the museum community to enable museums to become showcase sites for a series of interactive programmes, in aid of the Year of Engineering Success, reaching a diverse national public audience as well as the engineering community. It also enables schools and universities to link directly into museum educational programmes providing they have ISDN videoconferencing. A further benefit to YES is that national audiences can be reached from the museum sites should they choose to promote their own events through MuseNet.

Museums are a content-rich community and are already centres for sophisticated and appreciative audiences. MuseNet adds value by multiplying audiences and consequently delivering greater rewards for time invested in educational programmes. It allows curators and educationalists directly to engage communications technology as a creative tool that facilitates visual and verbal messages which stimulate real-time discussion. Audience participation is a major feature of a MuseNet event. 'Tomorrow's World Live' at the National Exhibition Centre, Birmingham, from 19-23 March 1997, demonstrated how fascinated children become when they can speak to a character on screen and have their questions answered! The museums contributed 16 events via the BT stand and successfully caught and held the attention of a roving audience.


Experience of MuseNet highlights videoconferencing as an informal medium with different production values to television where time planning and rehearsal are key to success. Videoconferencing over ISDN6 looks similar to large-screen TV but can support inexpensive

productions without depleting intellectual content: substance and spontaneity take precedence. For the NEC, Phil Phillips from the National Museums on Merseyside, working alone, used MuseNet equipment to offer a 'mystery quiz' based on fossils, direct from his office desk in Liverpool. He skilfully investigated objects from his collection using zoom enlargement and microscopic detail via CD-ROM clips to engage his audience.

At a recent event held between the Ulster Museum and The National Museums on Merseyside called 'Life on Mars' (an event based around the recent discovery of life in a meteor from the planet), adult audiences were fascinated by being able to meet each other through the videoconferencing medium. The Vice Chancellor of Belfast University who attended the talk immediately made arrangements with the Ulster Museum for the University to use MuseNet for link-ups with other universities and institutions in the USA. In Scotland, communities on the Shetland Isles are keen to hold events through MuseNet to overcome remoteness.

Overall the MuseNet concept has started a process of change within museum culture moving institutions closer to their audiences. The MuseNet calendar is available on the Year of Engineering Success Internet site on: <http://www.yes-yearco.co.uk>.

MuseNet in use



Sue Nicholas, from BT's Strategy and Business Management division, developed the concept of MuseNet through to

realisation. Sue works on pricing new technology developments for emerging markets.

Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his regular column in the Journal with his vision of a three-click, one-second world.

Three-Click World

Only 15 years ago making an international telephone call could be something of an ordeal. The post-dialling delay of 15 seconds (remember those old mechanical dials?) was followed by concatenated electromechanical switching delays that could add a further 15 seconds or more. Similarly, the processing speed of computers and printers meant that you could wait for seconds or even minutes for screen fill or print out. How different it is today. Now we get irritated if we do not hear the familiar ring tone immediately we press the last digit on the telephone keypad, or when we have to wait seconds for a PC application to load. The generic problem is having to wait for a period that is too short to do anything else, but long enough to break our concentration. Delays of a fraction of a second can disrupt our mental agility and interactive creativity to an alarming degree.

In contrast to our recent past, we now have an abundance of bandwidth, storage capacity and processing power, with optical fibre, CDs, and power PCs. Moreover technology promises even higher levels of circuit density and clock speed at insignificant cost. We are thus approaching the realisation of a dream: to access everything, everywhere, anytime, within three clicks of a mouse and have screen fill and interaction within a second. For us to enjoy natural, and effective communication with people and machines, in real or virtual worlds, the need is for sensory delays of less than 100 ms.

Why should we foster such a dream? The principal reasons are twofold: firstly, we live in an accelerating world where we all have to do more in less time and communication delay limits our creativity and

output, and secondly, it can be done! Trying to interact with anything, or anyone, at less than natural human speed soon becomes counter-productive and extremely irritating. For anyone who has tried to communicate using a telephone operating over a geostationary satellite, which introduces over 300 ms of delay, it is obvious. An even more obvious experience is that of trying to access information from a CD ROM, LAN, or the Internet (the Information Super Footpath or the World Wide Wait). Here delay is endemic due to inappropriate protocols and layers of unnecessary and inefficient software. Even writing a letter, sending e-mail and manipulating simple documents now seems to require a Power PC to get delays down to a few seconds. The reality is that many PC applications waste increasing MBytes of RAM making the front end prettier, and providing unwanted and unused facilities, rather than making the process more effective.

In telecommunications the deregulated market may soon see the concatenation of digital mobile telephones (with an internal codec delay in excess of 120 ms), statistical multiplexers, ATM switches, satellite and cable links of numerous uncoordinated suppliers adding undefined transmission delays of random duration. This new regime of unpredictable delays will take us further away from realising another dream: matching man and machine to achieve effective and efficient communication and creativity. But perhaps more dangerous will be the prospect of economic routings chosen in ignorance of the final application. E-mail over GSM is no problem but speech with a total coding and a transmission delay of more than 0.5 s will be a disaster on a par with telephone calls over Internet.

For anyone trying to communicate effectively over a narrow-band video-conferencing circuit, where the codec and signal path can introduce delays of over 300 ms in the visual image, with a different, and disconnected, delay in the speech circuit—the problem is more immediately obvious.



Even with a high degree of practice and familiarity, such arrangements are a far from realistic replication of true human presence and communication. At a first meeting it can be almost impossible to communicate and work effectively with such delays. Furthermore, the lack of correct physical size, colour, definition, hi-fi sound emanating from the mouth, eye contact and body language, all detract significantly from the illusion of being there. If you know the people at the distant terminal, then your brain provides some compensation and business can be done. It works, but clearly, it is an illusion, not a real experience.

Experimenting with delay between human hand, eye and ear soon shows our critical dependence on precise coordination to achieve simple objectives. Much more than a fraction of a second between hand and eye and we have trouble writing. A mere 200 ms between lips and voice, and we are talking to a mannequin. At 300 ms we can experience severe coordination problems and confusion. Introduce fixed and/or variable delays between the three senses of sight, sound and touch of 100 ms or more and our ability to telecommunicate can be noticeably degraded.

It is interesting to note that sticking a pin into the end of our finger results in a message arriving at the brain some 30 ms later. This approximates well to the 30 ms period for sound discrimination of the human ear, movement detection by the eye, and a separation distance of 10 m at which we have difficulty hearing normal speech and discern-

ing facial expressions. Curiously, it also approximates to the time it takes a photon to travel between London and New York on an optical fibre. So in a telepresence future, a surgeon may wish to reach out and touch a patient in the operating theatre in New York while actually being in London. The effective distance between the brain and finger tips will then have doubled from 1 to 2 m, and without the synchronisation of sight, sound and touch, it will be impossible to interact in a coordinated way.

The only prospect of realising this dream relies on optical fibre to provide

a true super information highway with minimal electronic delays, bottlenecks, and bandwidth restrictions. It also requires a new attitude to software, which will have to become efficient, instead of prolific, for the processing and storage of information. We will also have to break out of a mindset that says bandwidth is expensive. It is time to trade bandwidth, processing power and memory capacity for human effectiveness—this is the most precious commodity.


Just watch children interact with machines and it is immediately apparent that they have an insatiable

desire for instant gratification—the shortest response time and the best graphics. Looking at professionals you see the same phenomena—a desire to be able to do more, faster. It really is time to start bending machines into humans—engineering a match between our natural abilities and that of technology. All the technology required is available today; we only have to adopt the right mindset and implement the solutions. In the meantime, I suspect our progress will continue to be frustrated by the delays of networks, computers and software configured for the past.

book reviews

WWW Unleashed 1996

by John December and Neil Randall



The Greek philosopher Zeno told his students the famous puzzle of Achilles and the Tortoise. Suppose that Achilles has challenged the tortoise to a footrace. Because he can run ten times as fast as the tortoise, Achilles agrees to give the tortoise a headstart of 1000 m. When the race has started and Achilles has run 1000 m, the tortoise has covered 100 m. When Achilles runs the next 100 m the tortoise is 10 m ahead. And so on. For ever. Achilles could get closer and closer to the tortoise, but never overtake it. The reader of *The World Wide Web Unleashed 1996* may feel a bit like Achilles as he reads each chapter.

At over 1300 pages and 50 chapters the book offers a lot of information to

the reader. But in a world where things change so fast that people talk of web years that last a few weeks, I wonder if the reader will have time to finish this book before the World Wide Web (WWW) has moved on.

Having said that, the authors have presented a comprehensive text with chapters that will interest everyone from complete novices to experts.

The book begins with an introduction to the WWW, with good descriptions of the internet, hypertext, and multimedia before describing how they all come together in the WWW.

Detailed descriptions are given of the major WWW browsers. It is a reflection of how quickly the WWW is changing that the National Centre for Supercomputing Applications (NCSA) Mosaic is given 40 pages, while Microsoft's Internet Explorer is too new to be mentioned. Instructions and tips for using your browser and to browse and search the Internet are given. This includes the WWW itself and those services 'at the edge of the web' such as file transfer protocol (FTP), gopher, telnet, wide-area information servers (WAIS), and news space.

A detailed and comprehensive guide to the WWW describes resources in business and commerce, entertainment and the arts, education, science and technology, publishing, and government information. The breadth of this catalogue must represent hours of web browsing and could save the

reader considerable time in finding the resource they want.

The authors next move from the reader's point of view to the publisher's. They also move from writing for the novice to chapters aimed at the more technically adept.

These chapters describe how to design, create, and maintain WWW sites. The reader is taught how to write WWW pages and to use a number of hypertext markup language (HTML) editors. Advanced WWW features such as forms, imagemaps, common gateway interfaces are outlined, as well as the extensions to Netscape that are becoming more accepted. Selecting, setting up, and administering a web server. A case study of the development of a corporate WWW site reinforces earlier lessons.

Sections on new technologies, such as Java and virtual reality modelling language (VRML), and topical issues, such as commerce on the WWW conclude the book.


The book comes with a CD-ROM loaded with Internet, multimedia, and web publishing utilities.

In conclusion, the book provides a good introduction to the WWW and useful software to help in creating WWW resources. The authors are well known authorities on the WWW. My only reservations are whether the book tries to be too comprehensive. The novice will want the first half, the

more experienced reader the second. But after reading the first half of the book, novice readers will probably want to buy a newer book that covers developments that have occurred since they started.

Published by Prentice Hall
ISBN 1-57521-040-1
£46.95. xxxix + 1392 pp.
Reviewed by David Alsmeyer

Java Starter Kit—Teach Yourself Java in 21 Days *by Laura Lemay, Charles L. Perkins*



The Internet has now passed into the public consciousness to such an extent as to become positively passé, although consciousness and understanding are, of course, different concepts. The so-called *killer application* which has made the Internet available to the computationally challenged is the World Wide Web. Servers running hypertext markup language (HTML) have been ideal as a mechanism for disseminating passive information to a geographically dispersed audience. Various tricks, such as server push, attempted to add the illusion of a true two-way dialogue, but HTML had not been designed with that in mind.

Enter Java from Sun Microsystems. Java is an object-oriented programming language designed to be small and portable across platforms and operating systems and now the subject of this review.

The Java starter kit come in two components. The book, *Teach Yourself Java in 21 days* can operate either in

isolation, or in tandem with a CD-ROM which contains the Sun Microsystems Java Developers Kit, Symantec Café Lite and a selection of example Java Applets as well as the source code of all of the examples in the book (once you have found the correct sub-directory). As you would expect, only a very small proportion of the CD is of any use to a Windows 3.X user. Installation (in my case under Windows 95) was a relatively painless experience.

The Teach Yourself Series is a well-established series with titles such as *Teach Yourself the Internet in a week*. This book (as with some of its companions such as *Teach Yourself Perl in 21 days*) is clearly aimed at current programmers looking to learn a new (and highly sought after) language competence.

As its name would suggest, the book is structured into three weeks of daily lessons. These start gently with introductions to object-oriented programming and Java in particular. It is then straight in to the language with comprehensive coverage of the building blocks such as variables and data types, working with objects and arrays, and creating classes and applications. Week two kicks off with how Java applets and applications are different. Much of the rest of the week deals with graphics and animations as a precursor to handling interactivity. In week three those topics vital to having something useful in the real world such as exception handling and multi-threading.


The book is very comprehensive, easy to read and well structured. There are some slight niggles. It would seem that the book was rather rushed to print judging by the number of minor errors and typos that more proof reading would have picked up. The publishers do, however, have a web site with an updated list of errata. Having examined this it would appear that most of those relating to the first printing (as reviewed here) have been corrected. Since one of the errors was so fundamental as to prevent the first example of Java written by the eager student from running this is pretty important!

In addition, it took a while to understand the relationship between the book and the CD. The book explains how to code and run applications direct with the Java Development Kit, whereas the CD installs the Café Lite development environment. A few minutes exploring soon cleared up what I needed to do, however. Once more this has been recognised by the author/publisher and better instructions are available on the Web Page.

In summary, I found this to be a useful and fun introduction to the Java language. No doubt the few slight problems on integration will be sorted out for the next edition. I would recommend it to anyone seeking to learn Java.

Published by Prentice Hall
ISBN 1-57521-097-5
£44.95. xxiv + 527 pp. + CD-ROM
Reviewed by Robert Temple

Distance Learning: Technology and Applications *by Daniel Minoli*



The first thing that strikes you about this book is its Americanness'. That's not to say it is bad, just that all the examples are American and so is the terminology, some of which is not explained, even though there are six pages of acronyms. For example, on page 1 the author talks about K-12 education. (I think it stands for Kindergarten to year 12 education.) The second thing you notice, is that the book is not about computer-based training (CBT) or interactive videos, but looks at the

role of the converging telecommunications and computing industries in the education market.

The book is in two parts. The first part, 'The Demand Side of the Distance Learning Industry', looks in detail at the three main drivers: schools, universities and corporations. Chapter one is an introduction to distance learning and the distance-learning industry. It considers the benefits of distance learning and how it can help the education and training institutions meet some of the challenges they are facing. It looks at the different technologies available; from video conferencing to groupware. It even analyses the costs of a typical distance learning classroom. The next three chapters look at schools, corporations and universities in turn. For each of these chapters, the author describes how technology is being used to provide solutions to meet their different needs. The author is not afraid to go in to some detail about the technology, and describes, for example, ATM-based videoconferencing solutions to frame relay networks.

Chapter five, the final chapter of part one, gives a brief overview of what is happening in every state, (yes, all 50). Some of the overviews are very brief. Delaware, for example, gets an 11-line paragraph. Apparently the state has been slow in supporting the interactive distance learning (IDL) initiatives compared with North Carolina, where the State Services Network provides teleconferences and education to over 250 schools, community colleges, medical centres and public buildings using satellite-based networks.

Part two, 'The NII and Distance Learning' is made up of four chapters. Chapter six considers the role of federal and state government in creating the National Information Infrastructure (NII). From the Clinton-Gore administration to the regulatory role of the Federal Communications Commission (FCC). Chapter seven, describes the activities of the telephone companies. Both the local exchange carriers (LECs) and the interexchange carriers (IXCs); in particular AT&T, MCI and Sprint, are considered. Chapter eight discusses

the activities of the cable TV companies. It focuses on the top four providers: TCI, Time Warner, Continental Cablevision and Comcast. The author describes the challenges facing the cable companies and how they are responding. The last chapter of part two considers the role of the Internet in distance learning. It describes everything from logon services and e-mail to the World Wide Web and videoconferencing.

The book finishes with three case studies. The first looks at implementing high-performance distance learning systems in schools. The second describes New York Universities 'virtual college'. The third describes how Westcott Communications, Inc. educates, trains and informs its employees.

Anyone researching into the subject of distance learning, in its broadest sense, will find this book worth reading. It would also be of interest to companies in the telecommunication and computer industries looking for business opportunities in the education market. Someone should write a similar book for the UK or Europe.


*Published by Artech House Publishers
ISBN 0-89006-739-2*

£59.00. xx + 352 pp.

Reviewed by Ed Kirkham

Managing a Programming Project, Processes and People

by Philip Metzger and John Brodie



This is the third edition of this best-selling book. It looks at programming

projects and explains how to manage them properly.

I planned to study the contents page of this book, read the first and last chapters with some care and then simply dip into the rest. No such luck, I was drawn into the book. For example, what to make of a line in the contents that reads: 'Getting at the possum'? You have to learn more.

Another line 'Compensating the achievers' is a title that says these boys know their stuff. For example, on page 90 they tell us 'weak managers have a tendency to push the salaries of both low and high achievers towards the middle'. Now why does that ring a bell?

The book is structured according to a 7 phase development cycle: definition, design, programming, system test, acceptance, migration and operations. Fair enough, and I'm all for any scheme that avoids the tired word 'requirements'. The authors call this cycle an 'overlapped waterfall', but the accompanying diagram is messy and unconvincing. I thought a salmon or two might have come in useful here to illustrate the strength of purpose it takes to stop the project in programming stage, perhaps, and struggle all the way back upstream to do the definition stage again.

This edition of the book (the first was in 1973) feels a little out of date in that there is not enough attention to object technology, Visual Basic does not appear in the index, the Web does not rate a mention, and Java is too recent. The INTERNET does get a name-check, but why the capitals? Nobody else does that.

While the book doesn't slavishly promote any of the popular methodologies, the book bangs the table hard for documentation standards. We're all inundated by paperwork, a lot of it useless, so the book declares bluntly that anyone on the project who produces a non-standard document should be fired.

The appendix has a nice model project plan, including a pretty credible set of document specifications. There are other examples of sound common sense: 'Change has to be managed, not eliminated ... you can't

book reviews

resolve questions of change until you and your customer agree that what is being discussed is a change'. Obvious? Maybe, but still worth restating.

There's a reference section at the back, and the authors are not afraid to show they have researched earlier efforts in the field. They tend to damage their credibility, however, by inventing a host of imaginary characters with humorous names like Walt Secondlevel and Peter Projectmanager. This gets irritating after a while.

The plentiful illustrations are lovingly collected and some are

quite funny. The figure illustrating the migration phase is a Breugel painting showing the blind leading the blind. Not cynicism, but realism.

The authors chip in their experiences, some of which make you wince. Take the problem of making sure dead-wood employees are not dumped on your project. One of the authors, failing to take basic precautions against this disaster, discovered four such dumptees had spent a whole day sitting round a table compiling a list of clichés.

How relevant is *Managing a Programming Project* to a large

telecommunications company? The authors use the preface to show they understand the challenge and inform us that 'Integration is a much bigger part of the job than it used to be'.

The book was most readable. I like that. Those of us who give up two or three evenings a week for career-enhancing reading prefer not to be bored on our time off. The book is also a good bet for the occasional quick dip.

Published by Prentice Hall

ISBN 0-13-554239-1

£35.95. xxiv + 372 pp.

Reviewed by W. Shane McMordie

telecom focus

MCI Shareholders Approve Merger With BT

On 2 April 1997, MCI shareholders overwhelmingly voted to approve the company's proposed merger with BT.

At the company's annual meeting, MCI shareholders voted 77 percent in favour of the merger, which will create a new company called *Concert*.

'MCI's legacy has proven that competition creates unprecedented benefits for customers and growth for our shareholders,' said Bert Roberts, Jr., MCI's chairman and designated co-chairman of Concert. 'Today, MCI's shareholders clearly acknowledged that an MCI-BT combination best positions the company to take advantage of the tremendous growth opportunities of a global communications market that will reach \$1 trillion by 2000.'

Roberts said that since MCI went public in 1972 at \$10 a share, the company's stock has split three times and appreciated nearly 3 000 percent. He also noted that MCI has invested more than \$18 billion building one of the world's premier global networks, helped pioneer long-distance competition that reduced prices for Americans an average of 70 percent since 1984, built the world's fastest Internet network, and introduced a wide

array of new services to the United States and global markets.

'The MCI competitive spirit that revolutionized the United States communications industry will be the driving force in Concert's ability to open markets worldwide to the delight of millions of users,' added Roberts. 'And, there is no bigger single market than the \$100 billion United States local market. MCI's number one priority is to offer Americans an integrated set of communication services which includes local calling.'

Roberts told shareholders that, when the merger is approved, Concert will have a stake in 25 communications services joint ventures, and distributor alliance agreements with 44 companies, representing all regions of the world. With combined revenue of \$43 billion, Concert will have an estimated 6 percent market share of the global communications services market, now valued at \$670 billion.

BT and MCI are also awaiting regulatory approvals from the FCC and Department of Justice in the United States and the European Commission. The companies have said they expect the merger to be approved by autumn 1997.

MCI, headquartered in Washington, D.C., provides a full range of integrated communication services to

nearly 21 million customers. Credited with opening up the United States long-distance market for competition, MCI is now leading the charge to bring competition to the \$100 billion local market, offering American consumers for the first time the freedom to choose their local carrier.

BT Shareholders Back Merger with MCI

Shareholders are backing BT's proposed merger with United States telecommunications company MCI to form Concert plc.

At an extraordinary general meeting at Wembley Conference Centre, 99 per cent of shareholders who voted backed the merger after Chairman Sir Iain Vallance told them: 'If BT is to develop and grow, it will do so most successfully with MCI, in Concert. Together, we can bring benefits to our customers and shareholders that we simply could not deliver alone.'

'We have the chance to make Concert one of the great companies of the late 20th century, and to create one of the first great companies of the 21st century.'

The merged company will combine the substantial financial resources and global position of BT with the growth momentum and market expertise of MCI.

BT and MCI will continue to sell business and consumer services under their own brand names in their respective home markets through separate operations.

Concert to be Listed on New York Stock Exchange

BT announced that it will be seeking a listing on the New York Stock Exchange (NYSE) for the new shares in Concert plc to be issued to MCI shareholders on completion of the proposed merger with MCI. These shares will be traded as Concert American Depositary Receipts (ADRs).

BT ADRs are currently traded on the NYSE and dealing in the new ADRs will begin on the NYSE when the merger closes. Ordinary shares will also trade on the London Stock Exchange and the Tokyo Stock Exchange. BT and MCI expect the merger to close in the autumn, following regulatory and shareholder approvals.

BT, MCI and Portugal Telecom Announce Strategic Alliance

BT, MCI and Portugal Telecom have announced a strategic alliance addressing the rapidly growing communications markets in Portugal and Latin America. With the agreement, Portugal Telecom becomes the exclusive distributor for Concert Communications Services' voice products in Portugal, enabling it to offer the most advanced portfolio of global communications services to multinational businesses.

Portugal Telecom and MCI will also explore opportunities in Latin America's largest communications market, Brazil, where Portugal Telecom has an alliance with Telebras. Brazil, a \$14 billion market, accounts for nearly 40 percent of Latin America's fast-growing communications market. In addition, the companies will also seek other opportunities in the \$36 billion Latin American communications market which is expected to grow to over \$60 billion by the year 2000.

Sir Iain Vallance, BT's Chairman, said: 'This venture underlines the fast

changing pace of global communications today. We are delighted to be associated with developments that will ultimately change the face of the markets in Europe and Latin America to the advantage of millions of customers.'

Bert Roberts, MCI's Chairman, said: 'We are pleased that Portugal Telecom has selected us to offer customers the benefits of seamless global communications services. Portugal Telecom is an ideal strategic partner for MCI and BT as we bring the benefits of competition to customers around the world.'

An alliance with Portugal Telecom enables BT and MCI to expand their presence in Latin America. Portugal Telecom and MCI will initially target opportunities in Brazil, which is in the process of privatising its telecommunications market. Portugal Telecom and MCI will also explore opportunities to work together in Latin America.

BT, MCI Announce Alliance with Telefonica

BT, MCI and the Spanish telecommunications company Telefonica have announced that they have formed a strategic alliance which they say will reshape the global telecommunications landscape across the Americas and Europe.

The alliance will combine the strengths of the largest telecommunications services provider in the Spanish-speaking world with BT and MCI as they move towards the completion of their merger to form Concert. Plans include a Pan American joint venture involving MCI and Telefonica's international unit, TISA, and joint investments involving BT and Telefonica in Europe. The move comes just days after BT and MCI announced a strategic alliance with Portugal Telecom.

New Undersea Cable System Planned

AT&T has announced it has signed a memorandum of understanding with Telecom Italia, Telefonica de Espana, S.A. and Companhia Portuguesa

Radio Marconi, S.A. to plan to build *Columbus III*, an undersea fibre-optic cable network that will link the United States with southern Europe by July 1999.

Columbus III is expected to span nearly 11 000 km and cost between \$270 million and \$310 million when completed. The system utilises two fibre pairs using state-of-the-art wavelength-division multiplexing technology to transmit voice, data, and images with an initial capacity to transmit up to 10 Gbit/s. The system has a design capacity which is upgradable to transmit up to 40 Gbit/s, or nearly 490 000 simultaneous calls.

The agreement calls for the cable to form a collapsed ring connecting Florida and St. Croix, United States Virgin Islands, to Lisbon and the Azores in Portugal, Conil in Spain, and Mazara, Italy. AT&T, Telefonica, Telecom Italia and Marconi will own and operate cable stations in the United States, Spain, Italy and Portugal respectively.

The new system is important because the explosive demand for high-speed telecommunications in the region has nearly exhausted capacity on the United States-Europe segment of the Columbus II undersea cable network. That 12 000 km cable links Spain, Portugal and Italy with West Palm Beach, Fla., St. Thomas in the United States Virgin Islands, and Cancun, Mexico.

'Columbus III will carry telecommunications growth between the Americas, Southern Europe and the Mediterranean through 2010,' said Thomas McInerney, managing director of AT&T's international cable planning group.

Removal of Restrictions on Sale of Airtime

Mobile phone customers will gain from new moves announced today by Don Cruickshank, Director General of Telecommunications.

He announced a programme for scrapping unnecessary restrictions on how mobile phone operators can market airtime and distribute their products. For the two smallest

operators—Orange and One 2 One—licence modifications will be put in hand as soon as possible, and come into effect in around six months' time. The larger operators—Vodafone and Cellnet—will also, in due course, be freed from these restrictions, when the Director General is satisfied that competition has become fully effective.

Don Cruickshank said he would review the competitive position in the mobile market again with this in view, though he did not anticipate undertaking a review before 1998. He said: 'The mobile phone business has exploded over the past few years, with seven million mobiles now connected and that figure set to rise further. With Orange and One 2 One now offering services, as well as Cellnet and Vodafone, the market has become increasingly competitive. Yet all four operators are still subject to special regulation, in some cases more onerous than that applied to BT. In particular, the freedom of these operators to decide for themselves how best to distribute their products to the public is still constrained by a system of rules devised before any of them had launched their networks.

'As with all aspects of the telecommunications industry I want to peel away unnecessary detailed regulation because competition is always the best regulator of an industry. But I can only withdraw completely when I am satisfied the market is fully competitive. I have looked very carefully at how the mobile phone market is developing and I have decided that I can act now to remove restrictions from the two operators without market power—Orange and One 2 One. When the market has become fully competitive, the same benefits will be extended to Cellnet and Vodafone.

'As a result of the changes to the licences of Orange and One 2 One, which will take effect when necessary statutory consultation procedures are complete, these operators will no longer be required to provide airtime wholesale to service providers for resale. This will allow them greater flexibility to sell, distribute and

market their products. The existing requirements have become an unnecessary restriction on the commercial freedom of mobile network operators without market power, and removing them has to be good for customers.

'I will be keeping the mobile market under review. When competition has become fully effective, I will also be able to relax those rules for Cellnet and Vodafone.

'These changes do not mean that independent service providers will cease to have a role in the mobile market. I am anxious to ensure both network and services competition. I expect independent service providers to continue to play a useful part, through their ability to identify particular market needs and offer retail products tailored to them. But market processes will determine when and where they can add value.

'Although the rules are being relaxed I will act promptly to stop any practices which I judge to be anti-competitive or an abuse of market power'.

Oftel Investigates Price of Calls to Mobile Telephones

Oftel has issued a consultative document setting out the emerging conclusions of its investigation into the cost of calling mobile networks from fixed phones, mainly BT calls.

Introducing the document, the Director General, Don Cruickshank, said: 'Mobile phones are an increasingly important part of everyday life. They are no longer expensive toys for the few but a mass-market product used everywhere. But the cost of calling a mobile phone is very high—much more expensive than an ordinary telephone call. Business and residential customer groups have complained to me about this and I undertook last year as part of the review of BT's Retail Price Controls that I would investigate it.

'The underlying issues about cost allocations and regulatory procedures are rather complex. But the emerging conclusion is straight forward enough. Prices of calls from BT to the mobile operators look to be too high

mainly because they charge too much for delivering calls to customers on their networks.

'In my view the average retail rate to call a mobile network from a BT phone might come down by about another 10 pence per minute. The document also sets out some of the measures we are taking to make customers more aware of the charging issues. In addition, I hope that the industry in the case of each operator will be able to move to a single retail price for calls from each fixed operator to all mobile operators. This would help customers. In the light of responses to this consultation, I expect to take early action to see that prices of calls to mobile telephones come down.'

Use of New Telecommunications Services

Customers must be confident that their telecommunications equipment will be able to access new services offered by a wide range of suppliers, says Don Cruickshank, Director General of Oftel.

He has said it is crucial for competition and consumer choice that most services offered by a network operator are also available on other networks, and equipment could access services from more than one provider. Oftel's policy on this link between services and networks—interoperability—is set out in a new document.

Launching the document, Don Cruickshank said: 'We are setting out when telecommunications-network operators should be required to make their services available over other networks. We also deal with how customers' equipment can access the wide range of services offered by different networks and service providers.

'Getting this right is important for customers. It widens choice by making services from a wide range of companies available to customers. And it makes sure customers can take advantage of that choice. It also strengthens competition by allowing network operators to offer a wider

range of services, knowing that customers will be able to access those services.

'However, operators may be less keen to invest in developing certain types of innovative new services, if they know those services will have to be made available to competitors. I believe that the proposals set out today allow us to strike the right balance: encouraging greater competition and choice without reducing incentives to develop new services. I believe the regulator should only intervene to insist on interoperability where a company has market power and the development of competition might otherwise be stifled.'

Electioneering Goes Hi-Tech for Schools

As the election campaigning gathered momentum, young people around the country were able to have their say, thanks to the latest BT technology.

A 'BT Young People's Parliament' took place using BT videoconferencing to link youngsters and candidates at eight sites around the country—Birmingham, Cardiff, Cambridge, London Docklands, Glasgow, Leeds, Newcastle and Warrington.

The videoconferencing forms part of an on-going series to give young people across the UK the chance to put their views directly to the political parties, and experience the latest BT technology.

The BT Young People's Parliament also intends to consider European issues and will establish contacts with schools and organisations in a variety of EU countries.

Ian Paterson, whose BT department is sponsoring the initiative, said: 'As more schools obtain access to videoconferencing facilities, this parliament will provide an interesting and valuable application which will compliment and enhance other learning activities.'

Activities of the young people's parliament are designed to develop good citizenship and debating skills. It is producing an active and valuable community of video-users as well as promoting a recognised platform with

increasing recognition to voice the views and opinions of young people in Britain.'

Glimpse of the Future at the NEC

BT Laboratories transported some of their innovations half way across the country to give visitors to the *Tomorrow's World Live* exhibition at Birmingham's National Exhibition Centre a glimpse of the future. Speech synthesis technology to help people with voice and hearing impairments, virtual meeting rooms, museums and games such as Wireplay, were among the applications on show.

Technology was demonstrated that enables people to walk around inside virtual worlds created by computers and transmitted down telephone lines. Using videoconferencing-type technology, three-dimensional images are projected inside a giant 'Vision Dome' to enable work colleagues, school children, and other groups of people to share and visualise products, designs and other information on a larger scale than ever before. Also on show was BT's 'talking head', which combines computer-generated speech with human physical features. This could form the basis of the next generation of videoconferencing or virtual reality, as well as helping people with hearing and speech impairments to communicate more easily.

Security Technology to Reduce Cellular Fraud

AT&T Wireless Services has introduced a new security service called *Authentication*, a new, state-of-the-art technology that provides cellular fraud protection.

Authentication prevents fraud by using a complex security feature that contains a secret code and special number based on a cryptographic mathematical process shared only by the cellular telephone and the wireless network. When a person places or receives a call, the wireless system asks the cellular telephone to

'prove' its identity through a question-and-answer process, which occurs without delaying the time it takes to connect a legitimate call.

'In simple terms, Authentication is like a locked door that has locks on both sides,' said David Thaler, vice president and general manager of AT&T Wireless Services. 'The only way to open the door is to turn two identical keys in the two locks at the same time. If the correct keys are turned at the same time, the door will open.'

The estimate for fraud losses in the United States last year provided by the cellular industry association is estimated at over \$600 million. Authentication is based on the same concepts and principles that the United States government uses to protect and access computers and telecommunications networks.

Oftel Looks At BT's Global Plans

The progressive transformation of BT into a global telecommunications company and the implications for regulation in the UK is the subject of a consultative document published by Oftel. BT is becoming more focused on global markets and a more major player in local markets worldwide. Is any change needed to BT's licence so that UK customers can continue to be assured of a good deal for telecommunications services, matching up to the best available in the world?

The proposed BT/MCI merger has highlighted this issue, but this is only one of many of BT's overseas investments. At the beginning of 1997 BT was involved in over 80 alliances, joint ventures and other overseas investments across 23 countries.

Don Cruickshank, Director General of Telecommunications, said: 'Global liberalisation and the international competition which it encourages are a welcome development for customers of telecommunications services in the UK and worldwide. It is possible, however, that BT's moves towards globalisation may have an impact on the company's ability and willingness to meet its UK licence obligations.'

BT's record so far demonstrates, overall, quality and reliability at both wholesale and retail levels. But there may be areas in future where BT's own interests no longer coincide so fully with those of UK consumers. In such circumstances, how can Oftel make sure that UK customers continue to enjoy world class services?

The Secretary of State and I have a duty by law to act so as to secure that all reasonable demands for telecommunication services are met and that companies providing those services can finance them. This is an onerous test. It is against this strict test that I must consider whether customers need any more assurance than is already in BT's licence.

As the market becomes more competitive, the demand for world-class telecommunication services in the UK will be increasingly met by others as well as BT. But, although the signs are promising and many customers have choice, we are not there yet. In this transition period the maintenance by BT of a high-quality infrastructure is vital both for customers with limited or no choice, and for competitors. While competitors are rolling out their own networks, they rely for their success, in part, on the capacity and performance of BT's network.

In looking at the impact of BT's global market position, my judgment is not a commercial one. It does not second guess BT's view of the financial or other consequences of the BT/MCI merger or of any of the BT alliances, joint ventures and partnership agreements. Nor is it a market analyst's judgment. It has primary regard to my statutory duties, which necessitate a degree of caution in my approach.

There are a number of options open to Oftel. They are, in summary:

- Oftel would continue to monitor closely the company's performance in meeting its obligations and be ready to take steps under existing licence obligations or the Competition and Services (Utilities) Act to impose specific quality standards, etc. No licence modification would be required.
- Oftel would modify the licence to impose a general prohibition on the company taking any action which could prejudice its ability to meet its service obligations in the UK.
- The option above with an added requirement for annual review by the BT directors and certification by the auditors that the general prohibition was being complied with.
- Oftel would impose an all-embracing duty on the company to secure the UK network business—this would include securing that it had adequate management and financial resources, supported by detailed annual certification by directors and auditors.

Recognition For Invisible Workforce

The voices behind some of the UK's most successful companies are being recognised in a national award scheme to find the country's smoothest operators.

Two thirds of the UK workforce spend more than half their day on the telephone either developing new business, selling or answering customer queries.

Now this invisible workforce, which contributes a massive £10 billion to the economy, is to get its own 'business Oscars' in the form of the 1997 Telemarketing Awards, sponsored by BT and Marketing magazine.

Howard Sandom, BT's head of Telemarketing Communications, said: 'Not only is the growing use of the telephone in business contributing to the economy as a whole, but it's helping to bring industrial regeneration to the regions, thanks to companies seeing the benefits of setting up telemarketing operations in the Midlands.

'And all credit must go to the people behind one of the biggest business revolutions since the advent of the PC. Telemarketing has changed the way businesses operate, creating

new jobs and new wealth for the country.'

Reflecting the rise of the call centre market, this year, the Telemarketing Awards include a separate category for 'Best Telemarketing Call Centre' which looks at implementation, applications and business benefits of call centres.

The nationwide awards scheme aims to find the best examples of telemarketing at work and to reward the people behind them. The telephone has been central to the direct-selling revolution—the supply of products and services directly from company to customer over the telephone. Companies such as First Direct and Direct Line have shown how the telephone can be pivotal in transforming a business.

BT's New Conference Call Instant Service

The overwhelming success of operator supported audio conferencing has been extended, with BT planning to introduce a new ground breaking, automated audio conferencing service. The unique BT Conference Call Instant service, now on trial internally, gives customers throughout the world two new methods of establishing telephone conferences—either via a telephone or using an Internet web page.

BT's Conference Call Instant audio conferencing is a flexible and spontaneous service controlled by the user. With the 'telephone only' system, customers are guided by user-friendly voice prompts, and can take full control of the connection of parties to their conference. When using BT's new dedicated Internet page, customers can easily pre-program groups of contact telephone numbers and then instruct BT's Conference Call Instant to set up the call—all from their personal computer.

Andy Pearce, Business Manager for BT Conference Call Services said, 'As our UK and International customers reap the benefits of audio conference meetings, their demand for new and extended services, complementing existing services, has grown. Conference Call Instant

delivers two key requirements to our customers. Firstly it offers the ability to quickly and easily set up their own audioconferences. And those with Internet access can initiate audio meetings at the touch of the keyboard. In response to customer demand, instant audio conferencing is set to break new ground.'

BT's traditional conference-call services with its professional coordinators, providing convenient, personal and guaranteed connection, will be offered alongside these two new 'Instant' service offerings. The new system has all the benefits of conventional audio conferencing which includes savings of time and travel which results in more rapid and effective decision making. Both the telephone booking service and the Internet system were designed at BT's development and research laboratories in Martlesham.

BT's Conference Call service is part of BT's Visual Business Area (VBA), a world leader in the development and provision of audio and videoconferencing products and services. In addition to its Conference Call service offering, BT's Visual Business Area has an extensive portfolio of systems and services, enhanced by product alliances with other leading manufacturers.

BT Opens Pacific Land Earth Station

BT is all set to offer its customers a fully global service (excluding the polar regions) with the official opening of a new NZ\$35 million (£15 million) land earth station in Auckland, New Zealand.

Called BT Pacific, it will work in tandem with BT's existing land earth stations BT Atlantic (formerly Goonhilly), in the UK, and BT Indian at Eik, in Norway.

'This is a measure of our total commitment to customer satisfaction world-wide', said Rohan Channugam, General Manager BT Aeronautical and Maritime. 'Customers will be able to use BT anywhere in the world on land or at sea.'

The station will provide the earth link for the Inmarsat Pacific Ocean Region Satellite. It will serve customers in the Pacific region and will provide the full range of communications facilities including voice, fax, telex, data and high speed data capabilities, introducing applications such as videoconferencing, Internet access and e-mail.

BT's partner in New Zealand, telecommunications provider Clear, has assisted BT by providing the national and international links to the BT global voice and data networks.


BT Celebrates the Past with Colourful Phonecards

The inventor of the telephone is commemorated on a new set of BT phonecards launched on 3 March.

The life of Edinburgh-born Alexander Graham Bell features on a new set of limited edition cards. It was the 150th anniversary of his birth on March 3 and the colourful phonecards reflect the creative spirit of Bell, showing portraits of the inventor and some of the early telephones.

Four phonecards are on sale throughout the UK and a special collectors pack, featuring two additional cards can be obtained—subject to availability.

BT Scottish Manager Bill Furness said: 'Bell was constantly striving to improve our ability to communicate and Bell's invention literally changed the way we all live and work. We are delighted to be paying tribute to him through these colourful and collectable phonecards.'



Could You Write an Article for British Telecommunications Engineering?

Contributions of articles to the *Journal* are always welcome. **Anyone** (BT or otherwise) who feels that he or she could contribute a telecommunications-related article (either long or short), which may embrace technological, commercial and/or management issues, is invited to contact the Managing Editor, BTE Journal, Post Point G012, 8-10 Gresham Street, London EC2V 7AG. (Tel: (0171) 356 8022; Fax: (0171) 356 7942). Authors are advised to contact the Managing Editor before committing significant effort to preparing articles. Guidance notes for authors are available on request.

Field Focus

Field Focus is a feature in the *Journal* comprising short articles (up to about 600 words) on specific local BT projects. For example, novel solutions to field problems have, previously, formed the basis of very interesting articles. Anyone interested in writing a Field Focus item is encouraged to contact the Managing Editor, or their local IBTE representative (see list inside back cover).

BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by 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.

© 1997: 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 £50.00 UK, £55.00 overseas; private individuals (paid by personal cheque) £30.00 UK, £35.00 overseas.

Single copies and back numbers—companies etc. £12.50 including postage UK, £13.75 including postage overseas; private individuals (paid by personal cheque) £7.50 including postage UK, £8.75 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'.

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/97 \$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	Peter Shaw (in Ipswich)	(01473) 224488
Midlands	Mark Broughton	(0116) 253 4476
North East	Geoff Jenkinson	(0113) 237 8163
North West	David White	(0161) 600 2468
Northern Home Counties	Graham Lovell	(01473) 605343
Northern Ireland	Allan McDougall	(01232) 215301
Paris	Garry Platt	+33 1 46 67 27 76
Scotland	Raymond MacMillan	0800 671991
South West	Peter Shaw (in Bristol)	(0117) 972 3216
Southern Home Counties	Nick Smith	(01227) 474823
Wales	Ferg Brown	0802 760343

IBTE International HelpLine: +44 171 356 8008

IBTE On-Line (BT intranet): http://mhwww.boat.bt.co.uk/ibte_nat/ibtehome.htm

BRITISH TELECOMMUNICATIONS ENGINEERING

Journal Subscription Order (Not BT Employees)

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: **£50.00 UK; £55.00 Overseas**
Private individuals (to be paid for by personal cheque): **£30.00 UK; £35.00 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.

Contents

Editorial: Optical Horizons	1
Building the Road to Optical Networks	2
World Trade Organisation Talks on Basic Telecommunications	13
Operations and Maintenance Centre: Delivering Network Services	18
The Mirror—Reflections on Inhabited TV	29
Whither Speech? Convergence and the Future of Telephony	39
A Computer Vision System for Natural Communication	45
An Evolutionary Strategy for New Communication Services	49
Authoring and Automating Content with its Associated Metadata for Multimedia On-Line Services	56
Customer Access Signalling on Concert VNS	66
MuseNet	70

