

VOLUME 17 PART 4 JANUARY 1999

# BRITISH TELECOMMUNICATIONS ENGINEERING

*Included in this Issue*  
*Voice-Data Convergence*  
*Optical Wireless*  
*Professional Mobile Radio*



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



# BRITISH TELECOMMUNICATIONS ENGINEERING

## Editors' note

---

Members of the Institution of British Telecommunications Engineers (IBTE) should note that the IBTE Web sites now include indexes for *British Telecommunications Engineering* and the *Structured Information Programme*. Furthermore, we have started to place electronic versions of journal articles onto the Web sites too; currently, these include the articles that appear in this edition and in the October 1998 edition. It is our intention that future articles will be made available on the Web sites (<http://www.ibte.org>) as soon as the work preparing them for publication has been completed. This will enable us to bring articles to many Members a lot sooner. Readers should note that the electronic articles are currently available only to IBTE Members through password access to the Members' closed-user group area.

For further information about this facility, please call the IBTE HelpLine on 0800 028 0209.

Contributions of articles to *British Telecommunications Engineering* are always welcome. Articles can be on any technical or business topic related to telecommunications or related industries, and can be short or long. In particular, the editors would be very pleased to receive articles on field activities; for example, novel solutions to engineering problems.

Potential authors are reminded that the IBTE operates an annual award scheme to encourage excellence in the production of articles.

The editors will always be pleased to give advice, and guidance notes are available for authors. Anyone who feels that he or she could contribute an article is invited to contact the Managing Editor, *British Telecommunications Engineering*, Post Point 2D05, The Angel Centre, 403 St John Street, London EC1V 4PL. Tel: 0171 843 7623; E-mail: [paul.e.nichols@bt.com](mailto:paul.e.nichols@bt.com).

## Contents

VOL 17 ■ PART 4 ■ JANUARY 1999

### *Customer Telephony Integration*

**Voice-Data Convergence and the Corporate Voice-over-IP Trial** 218  
Andrew Catchpole

### *Telecommunications in the 21st Century*

**IPv6—The Next Generation Internet** 225  
Stuart Prevost and Tony Dann

**A Platform for Electronic Commerce Service Trials** 232  
John Wittgreffe

**Agent-Based Solutions for E-Commerce** 237  
Martin Owen and Jorge Núñez Suárez

**Wither Video?—Videotelephony in Perspective** 245  
Alwyn Lewis, Graham Cosier and Charles Nightingale

#### Cover Picture

BT Corporate Picture Library;  
a BT photograph

#### Board of Editors

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

**Managing Editor** Paul Nichols, B.SC.

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

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

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

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

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.

*British Telecommunications Engineering Journal*, Post Point 2D05, The Angel Centre, 403 St John Street, London EC1V 4PL ISSN 0262-401X

*Optical Technologies*

<b>A Review of Optical Wireless</b>	251
<b>What is it and what does it offer?</b>	
David J. T. Heatley, David R. Wisely, Ian Neild and Peter Cochrane	

*Mobility and Radio Communications*

<b>Professional Mobile Radio—Service, Market Position and Technology</b>	262
Paul Tattersall	

*Customer Service*

<b>Developing Customer Service Skills</b>	269
Trish Jones	

*Structured Process Improvement*

<b>Structured Greenfield Planning</b>	274
Ken Cobb, Mark Rose and Ian Lawrence	
<b>GenOSys—Automated Tool for Planning Copper Greenfield Networks</b>	281
Don Asumu and John Mellis	

*Human Factors*

<b>'The Eye of the Beholder'—Designing for Colour-Blind Users</b>	291
Christine Rigden	
<b>Evolution of the BT UK Core Transport Network</b>	296
Ian Hawker	
<b>Video-Based Virtual Clubs</b>	301
<b>Using Multimedia to Re-present Technology to Women</b>	
Graeme Balfour, Sylvia Wilbur, Hilary Johnson and Stephen Furner	
<b>Convergence Between Public Switching and the Internet</b>	308
Hans-Ulrich Schoen and Stephan Scholz	

*Regular Features*

<b>pcochrane@btlabs</b>	314
<b>Book Reviews</b>	315
<b>Telecom Focus</b>	316
<b>Forthcoming Conferences</b>	320
<b>Information</b>	Inside back

**Theme Editors**

<b>Computer Telephony Integration</b> Doug Chesterman	<b>Telecommunications in the 21st Century</b> Steve Sim
<b>Optical Technologies</b> David Smith	<b>Mobility and Radio Communications</b> Nigel Wall
<b>Structured Process Improvement</b> Keith Beales	<b>Customer Service</b> Mark Jakeman

# Voice-Data Convergence and the Corporate Voice-over-IP Trial

*The convergence of voice and data networks and technology has brought new opportunities and threats to traditional suppliers of telephone switching systems. BT has recently completed a corporate voice-over-IP trial, involving over 50 users, and using a new breed of soft-PBX switching systems.*

## Introduction

This article is part of a theme discussing computer telephony integration (CTI). The first article<sup>1</sup> defined CTI as the *logical* association between the voice and data worlds. However, since that first article technology has moved rapidly forward towards the physical convergence of the voice and data communications.

The first part of this article discusses one aspect of this convergence technology – voice-over-Internet protocol (VoIP) with particular attention to intelligent IP switching systems, or *soft-PBXs*.

The second part goes on to discuss the technical evaluation of a soft-PBX system in the corporate voice-over-IP trial run by BT from June to October 1998 and involving over 50 triallists making over 10 000 calls.

Finally, some of the value-add applications are described which are now being realised as a result of voice-data convergence.

## What is Voice over IP?

Unlike the traditional public switched telephone network (PSTN), which is a circuit switched network, VoIP uses a packet switched data network for transporting voice and signalling information.

Figure 1 shows an IP telephony device such as a personal computer (PC) with full sound capabilities and connected to a data network. Also required is suitable Internet telephony client software such as the freely available Microsoft NetMeeting application, although many others are available.

Voice communications is initiated to another VoIP client by making a call to the IP address of the distant client (if known) or by using one of the many Internet location servers (ILS) or gatekeepers, as shown in Figure 2.

Once the call is connected, basic speech transmission can then take place. VoIP works by sampling speech

Figure 1 – Voice-over-IP telephony device

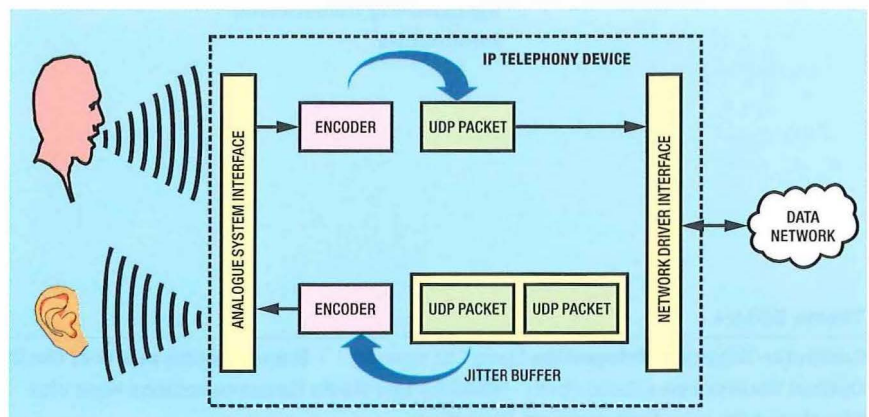



Figure 2—Simple gatekeeper routed VoIP call



(normally every 30 ms) from a microphone into data packets, which are transmitted to a distant receiver via the data network. In some cases, these speech samples are compressed to reduce bandwidth requirements using one of many coding standards.

When the IP packets are received at the distant end, they are buffered in memory before being decoded and played out through the sound system.

Generally, data network applications use higher-level protocols in the IP family. Most IP telephony applications use the user datagram protocol (UDP), which is most suited to real-time data such as voice but gives no guarantee that the packets will ever arrive at the destination. Signalling and call control information is generally carried using transmission control protocol (TCP) which, unlike UDP, does provide guaranteed delivery of packets over an IP network.

For a computer literate person, simple VoIP is relatively uncomplicated to set-up and use but still lacks the ubiquitous nature of the telephone network. A soft-PBX bridges the gap between simple VoIP and a commercially-viable fully-functional voice network.

### Soft-PBX Developments

A soft-PBX differs from a traditional PBX owing to two important characteristics. Firstly, call processing is performed by a software application executed on an industry standard computing platform. Secondly, the voice and signalling are carried over a data network rather than a separate telephone network—usually an IP network.

This, in theory, allows the replacement of conventional business telephone equipment and wiring by integrating data and voice traffic over a single local area network (LAN) or wide area network (WAN) infrastructure.

A soft-PBX system consists of the following functional components in addition to the data network itself:

- call processor,
- gateway,
- user terminal, and
- applications.

Unlike a conventional PBX, no switching element is required to connect the voice paths. Voice traffic is effectively switched within the existing data network as IP packets.

A unified communications infrastructure to the desktop allows common delivery of voice and data with the potential of reduced installation and maintenance costs. It also potentially lowers the cost of telecommunications when connecting remote offices. Since voice and data share the same data network, voice can easily be extended to a company's WAN. In practice the WAN would, in most cases, have to be upgraded to support the additional real-time bandwidth requirements.

Other potential savings can be realised by using speech-compression and silence-suppression techniques, which are normally cost-justifiable only in international networks. Audio compression means less bandwidth is required to transmit voice so bandwidth is used more efficiently.

All the traditional PBX features are provided by the soft-PBX but with the addition of multimedia applications such as data and video collaboration, which become much easier to incorporate in a soft-PBX environment and at a lower cost than ISDN-based solutions.

### Call Processor

The call processing function of a soft-PBX provides all the necessary call control and signalling normally

provided by the software of a traditional PBX. The call processing software resides on a server computer system connected to the IP network and gives basic call handling (for example, routing of internal and external calls) as well as supplementary features such as call transfer and call divert. It does not normally process the actual speech information itself, although exceptions to this include applications such as voice mail.

### User Terminals

A soft-PBX user terminal is generally based on an industry-standard multimedia PC and special IP telephony application software. Over the last year, however, new types of IP telephone have started to become commercially available—the *IP-phone* and the PC-based *soft-phone*, described below.

### PC-based soft-phones

Soft-phones are PC-based Internet telephony applications which exploit the PC's multimedia capabilities such as sound card, loudspeakers and microphone. Soft-phones can provide all the benefits of CTI applications, the major difference being that the soft-phone application also handles the audio streams (voice) as well as call control.

Soft-phones allow telephony users to log-on and acquire access to a company's telephone system from any remote site having an IP connection back to the call-processing server.

Soft-phones should be able to support all the features that a conventional PBX telephone supports and, in addition, can have even more user/system programmable features.

Figure 3—Corporate voice-over-IP trial architecture

One limitation of a PC-based client that could affect user acceptance would be the lower voice quality, specifically the higher audio delays within the PC. However, special telephony sound cards, such as the Quicknet PhoneJack, are available which provide enhanced audio capabilities with on-board hardware-based speech processing, full duplex speech and echo cancellation. Importantly, the card provides additional connections for a regular analogue telephone handset and headset if required.

### IP-phones

IP-phones are a relatively novel concept appearing during 1998. The main difference from a regular telephone is that an IP-phone is plugged into an office LAN, whereas normal PBX telephones require a separate telephone cabling system (although it is now common practice to use a structured cabling system for both data and telephony).

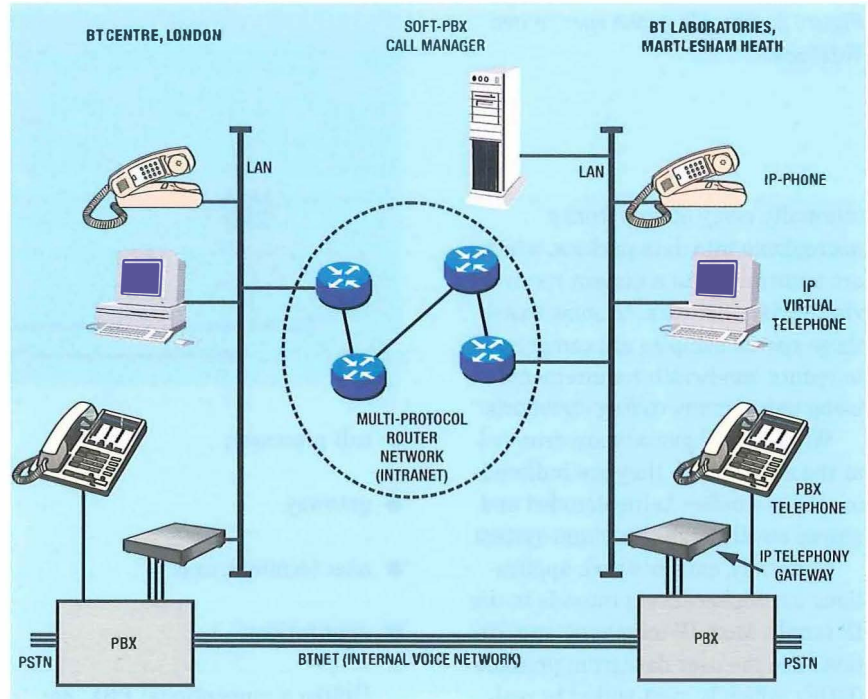
Also, unlike a standard telephone, an IP-based telephone can be installed anywhere in a corporate LAN/WAN network, while retaining its unique identity, without requiring a dedicated connection to the PBX. Setting up an IP-phone is more akin to setting up a PC's network configuration requiring host and server IP address details.

### IP Telephony Gateways

An IP telephony gateway allows IP telephone calls to be connected with the traditional PSTN or PBX networks. Gateways are also known as *PSTN-to-IP gateways* or simply *PIGs!*

Gateways can be either digital, connecting to the ISDN network (basic- or primary-rate ISDN), or analogue, connecting to regular PSTN lines.

Gateways can also be used to interconnect two or more PBXs via an IP network. In this configuration the use of voice-over-IP is a simple replacement transport system which does not require a separate call processing function. Many vendors now



provide these types of IP gateways, described later, to perform toll-bypass applications. In the future, PBXs will have IP telephony line interfaces as well as conventional line cards.

IP telephony gateways can be used in a soft-PBX system to make and receive calls to the PSTN or other conventional PBX systems.

### Corporate Voice-over-IP Trial

A number of potential soft-PBX solutions were shortlisted for BT's corporate VoIP trial. Requirements for the trial system were that it would have ISDN (primary rate) gateways, ITU-T H.323 standards-based PC clients and ideally IP telephones as well as offering as many PBX facilities as possible such as call records.

The following paragraphs describe the main components of the BT corporate VoIP trial system. Figure 3 describes the architecture of the soft-PBX system used in the trial.

#### Soft-PBX call manager

The soft-PBX call manager, which runs on an NT server platform, provided all the call processing and remote management functions required for the trial.

The soft-PBX server was administered through a World Wide Web (WWW) browser interface which allowed remote configuration and

management from any location within the BT intranet.

The call manager dial plan allowed flexible routings to be configured for internal and external numbers. For example: London triallists were allocated four-figure numbers, while other triallists were allocated six-figure numbers. All users were given the same internal number as their regular telephone number. The triallists were able to divert all incoming calls from their normal PBX telephone onto the soft-PBX using a private (ex-directory) number. Access to the PSTN and BTNet was available by dialling 9 and 00 respectively.

Call detail records provided a history of all calls made on this system and this could be linked to any billing system if required. During the trial, in excess of 10 000 calls were made using the system.

#### IP telephony gateway

A number of IP telephony gateways were connected to the Meridian PBXs at BT Centre in London and BT Laboratories which allowed calls to and from the PSTN or the BT internal telephone network.

The first gateways deployed supported only ITU-T G.711 uncompressed voice using 64 kbit/s. However, it is now possible to use voice compression at 6.3 kbit/s using the ITU-T G.723.1 recommendation. (Note: Voice is compressed only on

the IP telephony side; the PSTN is always 64 kbit/s.)

Analogue gateways are also available but these were not used during the trial. These gateways are useful for connecting to regular analogue telephone lines or to connect analogue devices such as answering machines, fax machines and modems.

The gateways used were all ITU H.323 compliant which allowed full interoperability with other vendors' H.323 gateways and H.323 PC client applications.

Management of all the gateways was controlled through the same WWW management interface as the call manager.

### IP-phones

The IP-phone is equivalent to a PBX feature telephone such as the Meridian telephones commonly used in BT. The main difference was that they were connected to the office LAN rather than a dedicated telephone wiring network. The trial deployed up to 50 IP-phones which gave users desktop connectivity to the soft-PBX.

Each IP-phone required a 10baseT LAN connection and a 48 V DC power supply unit. All telephones had programmable feature buttons, hands-free facility and LCD display that allowed the caller's number and name to be displayed. Each IP-phone was configurable to use either compressed or uncompressed speech.

Some of the telephones had LAN repeaters that allowed the desktop

PC to be connected to the IP-phone. This allowed a single LAN port to be used for both the IP-phone and the PC, although an IP address is still required for each device.

IP addresses were allocated statically or by using dynamic host configuration protocol (DHCP) which is a standard in the data network environment for configuration and management of networked computers. This potentially reduces in-service costs by allowing IP-phones to be moved and reconnected anywhere on the IP network with no configuration changes required.

### Soft-phones

The PC-based soft-phone is a low-cost alternative to the IP-phone. The soft-phone software was installed on a desktop (or laptop) PC equipped with a sound card and microphone. It combined the feature set of the IP-phone with a number of additional multimedia features such as data and videoconferencing.

The soft-phone currently supports only ITU-T G.711 uncompressed speech. This means that it is not yet suitable for remote working over dial-up connections. However, it is hoped to have this capability available later to give real 'road warrior' ability for the corporate user.

### H.323 telephony clients

Some users opted to use their existing Internet telephony applications such as Microsoft NetMeeting

that is freely available and uses the ITU-T H.323 recommendation for multimedia communications. By adding these users to the soft-PBX system they were able to utilise their H.323 client to initiate and receive calls to other trialists using the IP-phones. In addition, H.323 users were able to use the IP telephony gateways to make calls to the PSTN and to receive incoming PSTN calls using a direct dial-in (DDI) number.

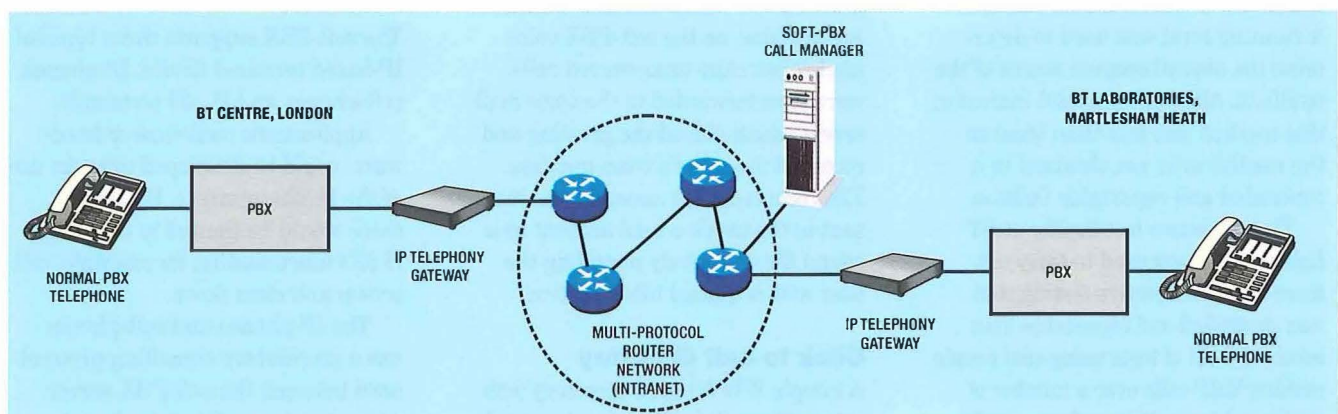
### Toll-bypass

Some IP telephony gateways are employed to link traditional telephony networks across an IP network, thus allowing network consolidation and potential reductions in PSTN charges. This type of service, commonly known as *toll-bypass*, can be used to route voice calls or even fax calls transparently across the IP network.

Figure 4 shows a similar toll-bypass system which was used during the trial for normal PBX telephone users but under the control of the soft-PBX call manager. Specifically BT Laboratories' telephone users were able to dial a unique access code on the Martlesham Meridian PBX to route calls to London and then on to any UK telephone number using the same corporate data network (multi-protocol router network (MPRN)).

Over 400 people at BT Laboratories were invited to participate in the

Figure 4—Simple toll-bypass call from a PBX telephone to another PBX telephone



toll-bypass service. Interestingly none of the additional 400 triallists were aware that it was a VoIP trial, just a new telephone service!

Alternatively the toll-bypass service was also available remotely by dialling a special access number. The call was automatically answered using a Meridian IVR system which prompted the user for a security number. If correct then the user was prompted for the required destination number. Once entered the call was then passed to the soft-PBX for onward routing.

By using either method the PSTN between BT Laboratories and BT Centre in London could be bypassed by using the VoIP network.

### Voice Quality using IP Telephony

Three different methods were used during the trial to determine the voice quality of IP telephony in a corporate environment:

- user opinion scores,
- subjective testing in a controlled environment, and
- objective testing using new BT test tools.

The first method required users to register their opinion of the voice quality after each IP telephone call. This process was automated by using specially programmed keys on the telephone to indicate call quality using the ITU listening scale of Excellent, Good, Fair, Poor and Bad. A running total was used to determine the overall opinion scores of the triallists. Although a useful indicator, this method was less than ideal as the results were not obtained in a controlled and repeatable fashion.

The subjective test facility at BT Laboratories was used to carry out more formal subjective testing that was controlled and repeatable. This involved a set of tests using real people making VoIP calls over a number of configurable conditions; for example,

different IP network conditions and different speech coders together with PBX, international direct dial (IDD) and GSM digital mobile calls.

To complement these subjective test methods, a new BT Laboratories test tool called *perceptual analysis measurement system*<sup>2</sup> (PAMS) was also used which was able to predict speech quality over the VoIP end-to-end connection. The tests were performed by injecting speech-like signals at one end and capturing the degraded signal at the other end. A quality prediction was then computed from a mathematical comparison between the original signal and the degraded signal. The same ITU listening scale, as with subjective testing, was then predicted. The results from the objective testing, using PAMS, was then compared with the performance of the corporate data network at any given time of day to check for any correlation between the data network conditions and the predicted quality of service.

Knowledge gained from the trial, in particular the testing discussed above, has provided essential information required to specify a commercially acceptable VoIP telephony solution.

### Applications

The soft-PBX provides several interesting applications which were investigated during the trial.

#### Voice mail

The voice-mail service allowed users to record their own personalised greetings, as with a telephone answering machine, on the soft-PBX voice mail server. Any unanswered calls were then forwarded to the voice mail server which played the greeting and recorded the caller's voice message. The recorded voice message was then sent to the user's e-mail account as a sound file, effectively providing the user with a unified inbox system.

#### Click to call directory

A simple WWW-based directory was set up that allowed anyone to search

for a VoIP triallist and automatically establish a call just by clicking on the hyperlink telephone number. The directory content was the same data as contained within the system configuration so was always up-to-date and correct.

#### Data conferencing

With a soft-PBX, the communications medium is no longer restricted to voice only. Both the IP-phone and soft-phone are able to add data conferencing to any IP-to-IP call, on demand, at any time during a voice call just by pressing an appropriate feature button. This is achieved by the IP-phone working in parallel with the PC-based soft-phone client. Data conferencing tools include:

- application sharing,
- white-boarding,
- file transfer,
- text chat, and
- videoconferencing

These collaborative working tools are useful for reviewing documents or showing slide presentations remotely. By adding a relatively cheap video camera to a desktop PC, videoconferencing becomes an inexpensive and simple-to-use reality. The soft-PBX now becomes a multi-media-PBX.

#### CTI Support

The soft-PBX supports three types of IP-based terminal device: IP-phones, soft-phones and H.323 terminals.

Applications (software or hardware) could be developed to make use of the H.323 interface. However, these would be limited to the basic H.323 functionality; for example, call set-up and clear down.

The IP-phones and soft-phones use a proprietary signalling protocol used between the soft-PBX server and gateways and terminals (tele-



phones) and provide much greater functionality. Using this protocol, or H.323, it would be possible to develop CTI applications with different graphical user interfaces (GUIs); for example, a BT-look-and-feel application. An alternative approach would be to use the Microsoft Telephony Applications Programming Interface (TAPI).

### Telephony API (TAPI)

TAPI was developed by Microsoft and Intel and has been available for several years.

The TAPI software architecture allows, in theory, any TAPI application to work with any TAPI server provider; for example, a TAPI compatible telephone.

Recently, TAPI 2.0 was released which attempts to address functionality required in a corporate PBX environment and comes as part of Windows NT 4.0 platform. TAPI 2.0 is the latest iteration that is built in to the client and server operating systems and provides support for third-party CTI applications<sup>1</sup>. These applications can then work across a number of telephony systems including PBXs.

### Potential Opportunities

Having a suitable CTI interface would allow the basic soft-PBX concept to be enhanced. Examples of these 'value add' applications are given below:

#### Internet call centres

Over the last few years, call centres have become a key technology area for many companies and organisations both large and small, local and global. However, the technology to date has been focused towards handling purely telephone-based customers. A trend is emerging for these centres to be *customer contact centres* where the communications channel will extend from audio only to include e-mail, WWW, video and data communications, effectively a multimedia call centre.

Example applications include the ability for people browsing the WWW to initiate a multimedia call to a call centre associated with a company's WWW page by clicking on a CALL ME button. The customer and the call centre advisor can then talk and see each other as well as exchange information and data; for example, to complete a mortgage application form. One example is the 'Call Me' service on the BT Shop WWW site (<http://www.btshop.bt.com/>).

IP-based multimedia call centres are seen as one of the more significant value-add applications for a soft-PBX system.

However, there are three evolutionary phases of the Internet call centre development:

- Internet call me back,
- VoIP breakout to existing call centre, and
- soft-ACD (automatic call distributor).

The Internet 'Call Me' function has been demonstrated by many companies but so far all these solutions require a second line to be available or for a PSTN dial back to be initiated once the Internet connection has cleared. The former obviously is not financially acceptable in most residential situations and the latter means that data sharing is no longer possible.

VoIP breakout to an existing call centre can be used to route Internet initiated voice calls via a gateway to a conventional circuit-switched PBX/ACD.

#### Soft-ACD

A soft-ACD allows a customer browsing the WWW to initiate a multimedia call to a call centre associated with a company's WWW page by clicking on a CALL' button as with the example above. However, with a soft-ACD the customer and the call centre advisor can, in theory, talk over the data network using VoIP. This is achievable only because

a soft-ACD architecture does not use a traditional circuit switched PBX or ACD system. Instead the VoIP is used to provide end-to-end connectivity between the customer and the call centre advisor that allows full multimedia (voice, video and data) communications.

#### IP-Centrex

Centrex is a network-based virtual private network (VPN) supplied by telcos for delivering voice services to large corporations, often over distributed sites. Centrex could be looked on as being a *PBX-in-the-network'* service.

IP-Centrex is the same idea but in this case utilises a company's IP data network to deliver the voice services to the desktop. In the future, it may no longer be necessary to have a separate PBX on every customer's site. All call processing could be achieved by the soft-PBX (IP-Centrex) server within the provider's network. The VoIP trial was effectively IP-Centrex with all calls controlled by a central call processor.

With IP-Centrex, access to the PSTN can be provided locally on customers' sites or centrally in the network by the service provider. This allows sharing of the PSTN gateway resources and breakout of the IP network at the most economic point in a country or the world. This potentially allows IP calls to bypass normal national or even international telephone routes. Ultimately this could lead to the forming of many small but flexible soft-telcos.

#### Web video screen telephones

A number of companies are designing and developing advanced screen telephone technologies which already support IP telephony.

Target markets for screen telephones include:

- corporate multimedia desktop,
- screen telephone for small and medium enterprises to deliver advanced network services, and

- mass-market IP telephony appliance (that is, non-PC dependent).

Screen telephones could allow BT to provide full multimedia (voice, video and data) desktop communications and could provide high-functionality easy-to-use desktop access to all the business applications such as e-mail, corporate directories and mobility applications. Advanced BT applications could also be downloaded onto the telephone without the need for expensive firmware upgrades or software upgrades to desktop PCs.

Another potential market is the hotel business where screen telephones could eventually be provided in every room of a hotel. In theory, a hotel guest could swipe a corporate smart card and download corporate computer/telephony services to a hotel screen telephone. The real opportunity here is not just the terminal but the soft-PBX (or soft-telco network) to support these new IP-based multimedia terminals and the network centric applications which will be downloaded onto them.

### The Future

BT has gained much knowledge and understanding from the VoIP trial. Work at BT Laboratories will continue to investigate the technical opportunities for the convergence of voice and data networks and soft-PBX switching architectures.

Further work will also look into the opportunities of convergence applications. In the future, as bandwidth becomes plentiful, it may be that applications generate the substantive revenue rather than just call minutes.

### References

- 1 CATCHPOLE, A., CROOK, G., and CHESTERMAN, D. Introduction to Computer Telephony Integration. *Br. Telecommun. Eng.*, July 1995, 14, p. 98.
- 2 For further information on Perceptual Analysis Measurement System, contact Richard Reynolds, BT Laboratories, richard.reynolds@bt.com

### Glossary

- ACD** Automatic call distributor  
**API** Application programming interface  
**CTI** Computer telephony integration  
**DHCP** Dynamic host configuration protocol  
**DN** Directory number  
**G.711** ITU-T recommendation for coding of analogue signals by pulse code modulation at 64 kbit/s  
**G.723** ITU-T recommendation for multimedia communications transmitting at 5.3 and 6.3 kbit/s  
**H.323** ITU-T recommendation for packet-based multimedia communications systems  
**IP** Internet protocol  
**ISDN** Integrated services digital network  
**ITU-T** International Telecommunication Union (Telecommunications Standardization Sector)  
**IVR** Interactive voice response  
**LAN** Local area network  
**PBX** Private branch exchange  
**PC** Personal computer  
**Soft-PBX** Server based PBX using voice-over-Internet protocol  
**T.120** ITU Recommendation for data protocols for multimedia conferencing  
**WAN** Wide area network  
**VoIP** Voice-over-Internet protocol

### Biography



**Andrew Catchpole**  
BT Networks and  
Information Services

Andrew Catchpole joined BT as an apprentice in 1981 in the Norwich Telephone Area. In 1989, he transferred to BT Laboratories and is now a senior professional engineer in the Network Intelligence Engineering Centre. His main work has involved investigating computer telephony technology and applications for the desktop, office and call centre environments. He has also worked on a number of ideas for advanced multimedia call centres resulting in a patent. In his most recent work he has been investigating opportunities with IP telephony which culminated in the recent BT corporate VoIP trial. He was one of the first students to be awarded the Martlesham M.Sc. degree from the University of London.

Andrew Catchpole can be contacted at [andy.catchpole@bt.com](mailto:andy.catchpole@bt.com).

*Stuart Prevost and Tony Dann*

# IPv6—The Next Generation Internet

*The Internet continues to grow rapidly and networks based on the Internet protocol will be the predominant communications networks as we move into the 21<sup>st</sup> century. Limitations in Version 4 of the Internet protocol have led to the development of a new version of the protocol, IPv6. This article looks in detail at the next generation Internet protocol (IPv6), the issues surrounding its introduction, and experimental work conducted in BT's Futures Testbed and the 6bone.*

## Introduction

Over the past 20 years, the Internet has grown rapidly from its birthplace in academic and US Government networks into becoming a vital form of communication for the next century. Over the past three years, there has been a gradual realisation that networks based on the Internet protocol (IP), will be the network-layer protocol of choice for all future modes of communication. Commercial investment in all things IP, and the growing promise of reliable and predictable service levels in IP networks, have revolutionised the network infrastructure evolution strategy of all the major telecommunications carriers.

Currently, all IP networks use Version 4 of the Internet protocol. This was designed in the late 1960s without the prior knowledge of the eventual global ubiquity of IP. As such, there are some non-ideal features which, although not catastrophic, do hinder some of the future design goals of IP-based services. By far the most important of these is the size of the address space, which allows only a 32-bit address. This is not enough for everyone on the planet to be allocated a unique address, and IP addresses have become a rare commodity. Workarounds have had to be developed to enable what is done today; for example, network address translation (NAT), although breaking the basic end-to-end principles of IP, is becoming increasingly prevalent and may lead to future problems.

To address the shortfalls, much effort has been expended in the Internet Engineering Task Force (IETF) to design the next generation

of IP, IPv6 (IPv5 was deprecated in the 1980s). IPv6 will offer the opportunity for a vast range of equipment, including mobile telephones, hand-held computers, TVs and domestic appliances, to be connected to the Internet.

The draft standard is close to ratification and already numerous early implementations of the protocol are available for various host-system platforms and router products. A global experimental network has been built, called the *6bone*, with the objective of evaluating early IPv6 implementations and identifying major issues with the design and operation of IPv6 networks.

## IPv6 Header

Let us start with a comparison of the current IPv4 header to the proposed IPv6 version (Figure 1). A main aim of the IPv6 developers was simplification<sup>1</sup>, achieved by removal of all unwanted and unused fields. The IPv4 header contains 10 fields, two addresses and some options. IPv6 has reduced this to six fields and two addresses. Some have been suppressed, others renamed and two new fields added.

Another goal was to remove the support for options from the IP header which necessitated recalculation of a checksum for every possible IPv4 header type. IPv6 options were placed after the IPv6 header and before the data. With this the IPv6 header has become a fixed size so a checksum field is not needed. This also makes IPv6 headers suitable for hardware processing owing to their fixed size.

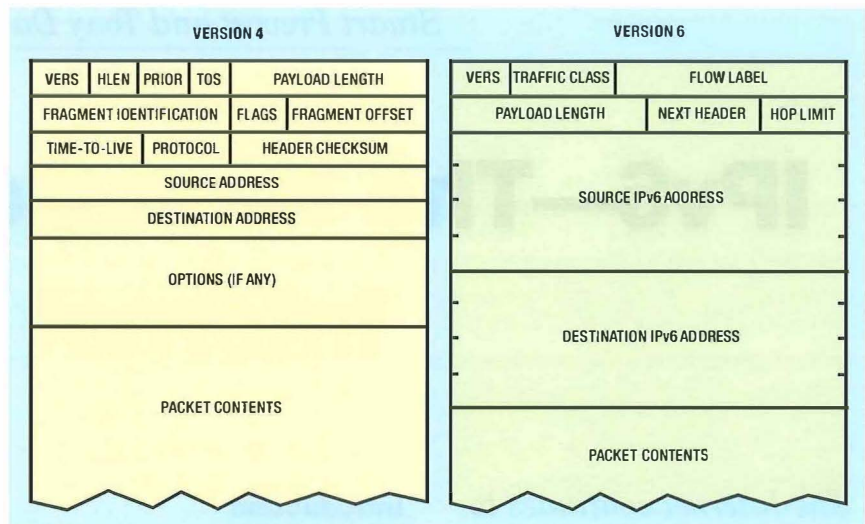
In detail, the IPv6 header comprises the following fields<sup>2</sup>:

Figure 1—IPv4 and IPv6 headers

- *Version*, a 4 bit field which identifies the version of the IP header.
- *Traffic class*, an 8 bit field that is used for distinguishing between different classes or priorities of IPv6 packets. Experiments are being conducted to establish the best way these bits can be used to offer various forms of *differentiated service* for IP packets.
- *Flow label*, a 20 bit field that will be used to identify a flow of IP packets from a source to the same destination.
- *Payload length*, a 16 bit field that indicates the length of the IPv6 payload following the header.
- *Next header*, an 8 bit field that identifies the type of header immediately following the IPv6 header. IPv6 uses the same values as IPv4; for example, UDP or TCP.
- *Hop limit*, an 8 bit field whose value is decreased by one for each router the IP packet is passed through. If the hop limit value reaches zero the IP packet is discarded.
- *Source address*, a 128 bit field that represents the source address of the IPv6 packet.
- *Destination address*, a 128 bit field that represents the destination address of the IPv6 packet.

The new IP header brings changes to the way intranets and the Internet can work. The new functionality that IPv6 introduces falls into four main areas:

- increased address space,
- plug and play,
- quality of service (QoS) and flow labelling capabilities, and
- security.



### Increased address space

The Internet has experienced exponential growth over the last 10 years which has led to predictions of the IPv4 address space running out in the future, especially with visions of domestic equipment such as refrigerators, washing machines, televisions, mobile telephones and hand-held computers containing an IP address.

IPv4 addresses are based on a 32 bit word which, in theory, gives IPv4 over 4 billion addresses. IPv6 increases the address space in the IP header to 128 bits to give the possibility of 4 billion × 4 billion × 4 billion × IPv4 address space.

This new address also allows for improvements in the routing structure of IP packets.

### Plug and play

IPv6 expands the basic IPv4 specification of one IP address per computer to multiple addresses which can be dynamically changed. IPv6 addresses can update themselves as you change

service provider, configuration methods can be turned on and off for complete control, and all are an integral part of any IPv6 implementation.

IPv6 allows hosts automatically to configure themselves using a protocol called *neighbour discovery*. Hosts or routers using this protocol can determine the addresses for neighbours connected to the same network, decide whether addresses are valid or even discover a router that is willing to forward packets.

There are two main configuration mechanisms, called *stateless* and *stateful*.

*Stateless* autoconfiguration allows IPv6 hosts to configure themselves using the in-built multicast capabilities of IPv6 and is achieved by the host joining a multicast group that identifies all routers on the local network. The router on the network sees the new host and sends the host the network prefix for the link. By using this prefix and the host's own Ethernet address, the host can build up a unique 128 bit IPv6 address.

### IPv6 Address Notation

IPv6 addresses are represented by eight 16 bit integers separated by colons. Each integer is displayed as four hexadecimal figures. All eight blocks go to make up the 128 bits of the IPv6 address.

#### Example of an IPv6 address

3FFE:2C00:0000:0000:0004:0369:34DE:98FA

IPv6 addresses can suppress leading zeros in each hexadecimal block. A double-colon convention can also be used once to represent a set of consecutive zeros. By using these shortcuts, the same IPv6 address can be represented as:

3FFE:2C00::4:369:34DE:98FA

*Stateful* configuration in IPv6 is a version of the IPv4 dynamic host configuration protocol (DHCP). DHCP is a client/server protocol that allows hosts to query DHCP servers for configuration parameters. The DHCP server then returns a valid IPv6 address to the host. DHCPv6 is different from DHCPv4 in that it allows multiple addresses to be assigned to a network card.

Another mode of operation is useful at the link level. When the computer is turned on, the interface associated with the computer initialises an IPv6 address known as a *link local* address. These link local addresses are available only on the local network and are used to communicate with other link local nodes. They are also used by both autoconfiguration methods to obtain a unique address for routed networks.

### Quality of service (QoS) and flow labelling capabilities

IPv6 will support quality of service and flows to support real-time IP traffic such as video streams. When the original IP was designed in the late 1960s no one envisaged watching videos on a PC. Now this is possible, and currently, IPv4 is being used to deliver some of these pictures with no inherent QoS mechanisms, which results in picture break up or loss of sound on congested links.

IPv6 has been designed to support 'real-time' traffic through the use of two fields called the *traffic class* field and the *flow label* field. These will be used to minimise the delay and delay variation of the IP packets.

The *traffic class* field will be used to assign different values to types of IP data, so video traffic could be assigned a value 1 and e-mail value 8, the lower values giving greater priority.

The *flow label* field can be used to identify a stream of IPv6 packets with the same destination and source address. By assigning a value to the flow of a video stream, a router will build up a table of the IPv6 flow

labels and destination addresses. In processing subsequent packets with the same flow label, a router simply forwards the IPv6 packet.

Use of both these fields is still under discussion and in experimentation within the IETF and academic community. This topic is currently very relevant to the IPv4 world within the differential services approach to network QoS.

### Security

Security is always a major concern to people using the Internet but IPv4 does not contain any security features. The designers of the new IP protocol decided that IPv6 should contain some inherent form of security, so a separate working group in the IETF has been defining IP-level authentication and encryption procedures that will be common to IPv4 and IPv6. The results of this will be that IP security will be retrofitted to IPv4, whereas they will be present from the start in all IPv6 products. The fact that security will be standard could be a key incentive for the adoption of IPv6.

The IPv6 specification has described two forms of security, which are the *authentication header* (AH) and the *encrypted security payload* (ESP).

The *authentication header* is used to ensure that the IP packet has not been altered during transmission and is achieved by a procedure that informs the recipient of a packet it has been sent by an authentic source address.

The *encrypted security payload* guarantees that only legitimate receivers will be able to read the content of the packet. This is achieved by encrypting the IP packet data using a key known only by the sending and receiving hosts. This key is then used by the destination host to decipher the IP packet; any hackers capturing the encrypted packets while they are transmitted will be unable to view the data contained in the IP packet.

### Addressing Structure

IPv4 addressing is 'flat' by nature. This means that IPv4 networks are not grouped together by organisational or geographic location. An organisation may use several different Class B and C networks spread across a wide area. With this mix of different network addresses used round the globe, every network address has to be advertised in the default-free routing tree.

IPv6's 128 bits allows for a new addressing structure<sup>3</sup> to be designed to solve the problems with IPv4. IPv6 enables a hierarchical addressing scheme, which will improve both routing efficiency and administration of the Internet.

The IPv6 address structure introduces a new topology which separates the public topology from the local site topology. This clear definition between topologies will allow individual sites to move Internet service providers (ISPs) without the need to renumber their internal networks. Two or more service providers can also be connected to any site. Such a multi-homed site can then decide which connection they want to use and when—a new service concept not currently possible with IPv4.

To understand these features fully, an appreciation of the structure of the IPv6 address format is needed.

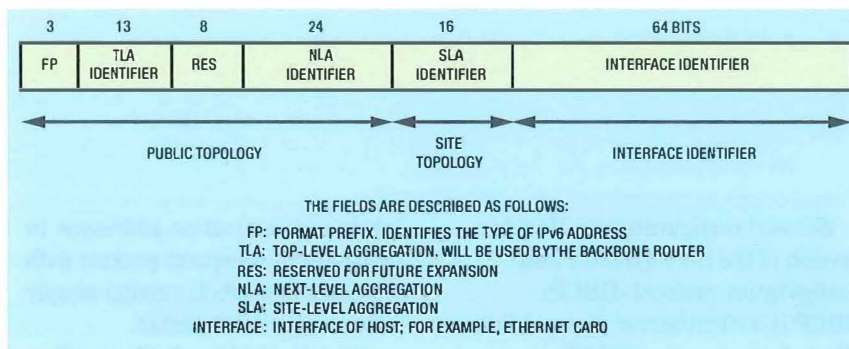
IPv6 addresses are basically larger versions of IPv4 addresses; in any IP address, a portion of the address identifies a network and the other portion identifies the host connected to the network.

IPv4 addresses were originally split into three classes, known as *A*, *B* and *C*. These classes represented the number of bits fixed in the address and identified the network portion of the IP address. These classes were later superseded by a more flexible method of dividing up the IPv4 address space, known as *classless inter-domain routing* (CIDR).

### IPv6 unicast address format

The IPv6 unicast address format is divided into six fields<sup>4</sup>, each repre-

Figure 2—IPv6 unicast address format



senting a specific part of the IPv6 address. The addresses are also organised into a three-level hierarchy. Figure 2 shows the 128 bits of the IPv6 address format.

The format prefix (FP) identifies the type of IPv6 address used and is specified by the leading bits in the address, which can be of a variable length. See separate panel on 'Assigned Address Blocks'.

The top-level aggregation (TLA) identifier identifies the top level in the routing hierarchy. Routers assigned a TLA interconnect to other TLA routers to form the default-free routing tree.

The next-level aggregation (NLA) identifiers are used by organisations assigned a TLA identifier to create an addressing hierarchy and to identify sites. The NLA identifier space can be divided up to create an addressing hierarchy appropriate to its network; the remainder of the NLA space is then used to identify sites.

The site-level aggregation (SLA) identifiers are used by an individual organisation to create a local

addressing hierarchy and to identify subnets. The 16 bit SLA identifier field supports 65 535 individual subnets (put in context, this is equivalent to the number of hosts addressable by an IPv4 Class B address block).

Interface identifiers are used to identify interfaces on a link, and are required to be 64 bits. These identifiers are required to be unique on the link, and usually based on the interface's link layer address. The interface identifiers used in the global address format are constructed using the IEEE EUI-64 format.

The above address structure describes the format for the global unicast addresses format, which accounts for only one-eighth of the total IPv6 address space. The remaining address space is mostly unassigned.

### Address formats

Address formats in IPv6 are used for different purposes depending on the circumstances in which the IPv6 host is placed. The main formats are described below.

- *Unicast addresses* are used by hosts to send information to the destination address specified in the IP header.
- *Multicast addresses* identify a group of destinations; information sent to a multicast group is delivered to all members within the group.
- *Anycast addresses* also identify a group of destinations. However, this process differs from multicast—instead of being sent to all members of the group the packet is sent to only one point, this point being the nearest member of the group.
- *Site local addresses* are used for addressing inside a site without the need for a global prefix. Routers must not forward any packets with site-level source or destination addresses outside the site.
- *Link local addresses* are used for addressing on a single link for purposes such as auto-address configuration, neighbour discovery, or when no routers are present.

### Assigned Address Blocks

The format prefix (FP) identifies the type of IPv6 addresses used and is specified by the leading bits in the address, which can be of a variable length. Currently, only four types of address are assigned, and three reserved. The table shows the different address types currently defined<sup>4</sup>.

Allocation	Prefix(binary)	Fraction of address space
Reserved	0000 0000	1/256
Reserved for NSAP allocation	0000 001	1/128
Reserved for IPX allocation	0000 010	1/128
Aggregatable global unicast addresses	001	1/8
Link-local unicast addresses	1111 1110 10	1/1024
Sitelocal unicast addresses	1111 1111 11	1/1024
Multicast addresses	1111 1111	1/256

### IPv6 global unicast address structure

The new global unicast address structure will be the main way IPv6 hosts connect to the Internet. The address format supports both the current provider-based aggregation and a new type of exchange-based

Figure 3—IPv6 unicast address structure

aggregation. Figure 3 shows how current long-haul providers will co-exist with the new IP exchanges. IP exchanges will be able to allocate IPv6 addresses to providers and sites, and by doing this an organisation will be independent from any single long-haul provider. Connecting to the Internet in this way an organisation will be able to change long-haul providers without having to renumber their site. A site can also be multihomed via the exchange to more than one long-haul provider without the need to have address prefixes from each long-haul provider. This new address format introduces number portability which is not present with IPv4.

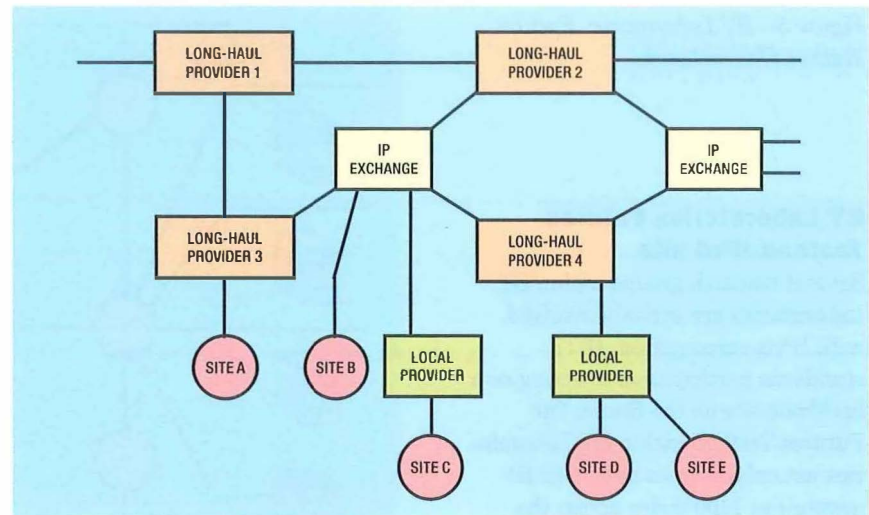
### 6bone

Research on IPv6 has been undertaken by the various working groups in the IETF, universities and leading-edge companies interested in the new version of IP. To make sure that IPv6 would be operational on the scale of the Internet an IPv6 test network, called the *6bone*, has been setup by the IETF.

The objective of the 6bone is to gain practical experience using early implementations of various IPv6 technologies. The information gained from this test network will go to form a set of documents detailing various mechanisms and procedures, ranging from advice on transition to native IPv6, sharing operational experience and maintenance of the global IPv6 default-free routing tree.

The 6bone was created in January 1996 by a few teams of IPv6 implementers who needed to test interoperability between their prototype systems. Because there was no IPv6 infrastructure in place to connect these sites, a method called *tunnelling* was used instead. Tunnelling involves using the existing IPv4 Internet to define two points over which IPv6 packets will be transmitted, the IPv6 packets are encapsulated into IPv4 and transmitted between these two points.

The 6bone is used:



- for testing new host and router implementations;
- for testing routing protocols such as RIPng and BGP4+;
- for developing new address allocation and management procedures; and
- for evaluating the new functionality of IPv6, such as the support for security and real-time information flows.

The fact that there are now over 332 sites across 39 countries connected to the 6bone proves that many interested parties are involved in

IPv6. These parties range from the end host companies such as Microsoft, router vendors like Cisco, to telecommunication research laboratories, ISPs and universities.

Out of the 332 sites connected to the 6bone, 51 are acting as backbone sites. Backbone sites are interconnected by tunnels to form the default-free routing tables using a protocol called *Border Gateway Protocol 4+* (BGP4+). This new version of BGP for IPv6 is used for routing the aggregatable global IPv6 unicast packets. All routers within the backbone must contain an entry in their routing tables for every other backbone site. Figure 4 shows the backbone routers on the 6bone and their interconnections.

Figure 4—6bone backbone<sup>5</sup>

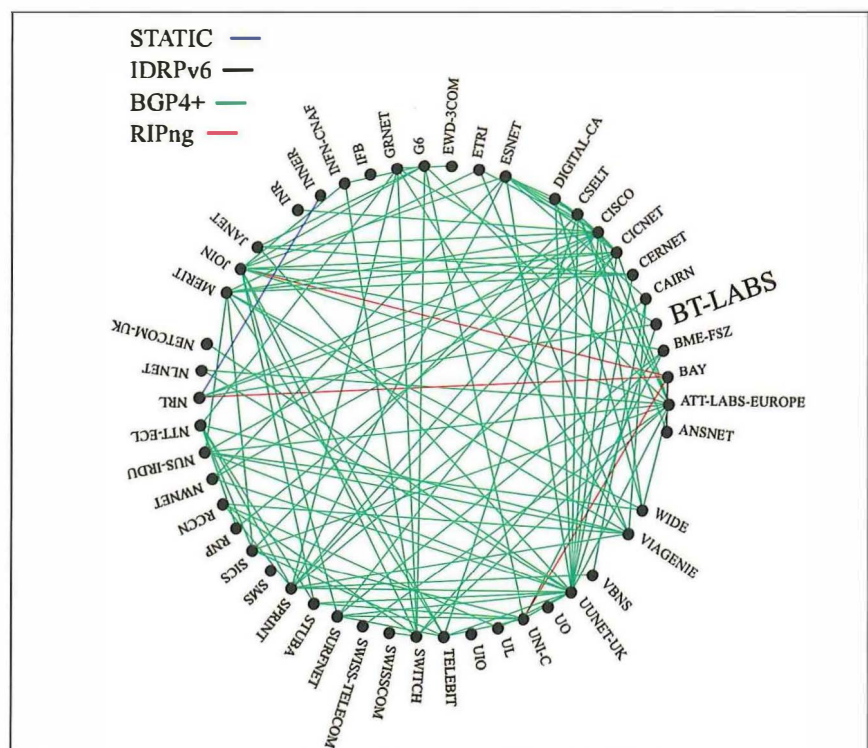


Figure 5—BT Laboratories Futures Testbed IPv6 network

### BT Laboratories Futures Testbed IPv6 site

Several research groups within BT Laboratories are actively involved with IPv6, ranging from IETF standards participation to acting as a backbone site on the 6bone. The Futures Testbed within BT Laboratories not only runs an advanced IP network to 1400 nodes across the site, but also runs an experimental laboratory looking at emerging IP technologies such as IPv6.

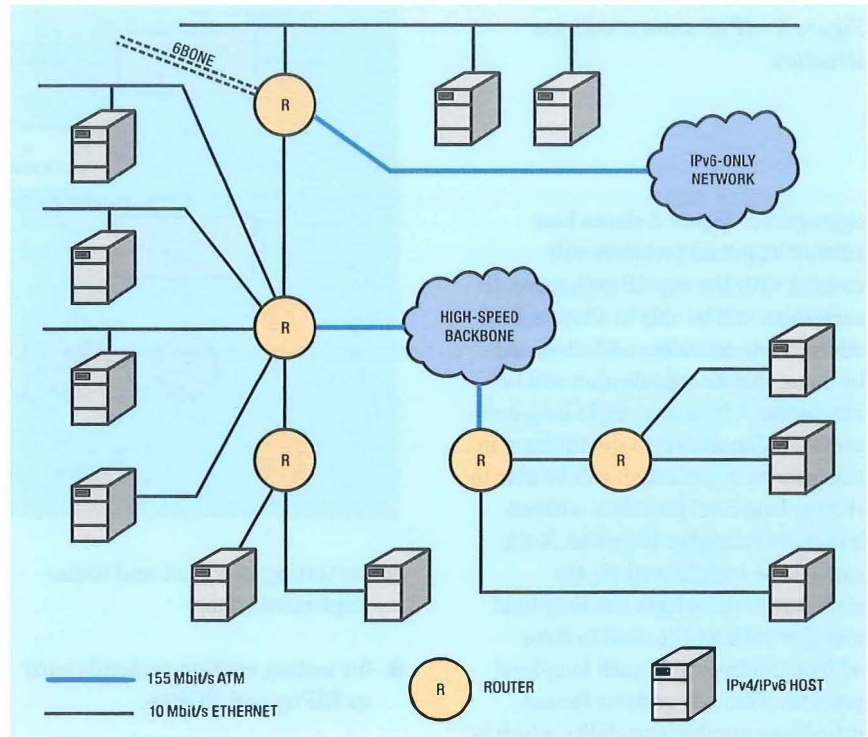
The Futures Testbed started work on IPv6 in January 1997, by first joining the 6bone using a single host. Since then the site has been developed into an IPv6 intranet (which interworks with the IPv4 testbed) and become a backbone site within the 6bone. The IPv6 intranet is a unique testing ground for various vendor networking and application implementations using IPv6. The purpose of the work carried out within this group is to gain practical experience in running an IPv6 network, to participate in the global effort evaluating IPv6 and to inform the business and raise awareness of IPv6.

The BT Laboratories Futures Testbed IPv6 site, as shown in Figure 5, consists of three experimental routers, with various end host implementations. The site tests the new auto-configuration methods of IPv6 and the interoperability of different implementations, as well as looking at the performance of IPv6 routing. This work is very important to BT as it gives early insight to the new protocol and its functionality, the advantages and disadvantages.

Applications that work on IPv6 are the basic applications that built the original Internet, such as telnet, ftp, ping, finger, sendmail, dns and tftp. Work is also being undertaken by collaborators such as Lancaster University who are looking at advanced applications that make use of the QoS functionality of IPv6.

### Transitioning to IPv6

To aid the transition from IPv4, it has been proposed that an IPv6



Internet will be deployed in parallel with the existing IPv4 Internet. A dual-stack strategy will be used where host computers contain both the IPv4 and IPv6 stack, and applications using the Internet will decide what IP protocol to use based on information provided by the domain name service (DNS).

The DNS system provides an easy way of mapping Internet addresses to IP numbers; for example, [www.b29net.bt.co.uk](http://www.b29net.bt.co.uk) refers to IPv4 address 132.146.239.148. A dual-stack host will have several records in the DNS server, one A record for each IPv4 address and one AAAA record for each IPv6 address. The DNS system will have to be updated in order to handle the AAAA records.

The transition period may take some years as organisations and users plan the upgrade of equipment and end host systems. One has to consider that eventually IPv6-only hosts will be deployed, with no IPv4 connectivity. In such cases some form of network address translation (NAT) will have to be placed in the IPv6 network to enable IPv4 networks to be accessed.

The NAT could be done by an IPv4-capable IPv6 router; the router would translate the data in the IPv6 packet to an IPv4 packet, with the IPv4 address being allocated from a pool of addresses stored in the router. When the IPv4 packet is returned to

the router, the data is placed in an IPv6 packet addressed to the IPv6 host. This process does have its drawbacks as each IPv6 packet has to be translated thus reducing the speed of communications and each NAT used would require a block of IPv4 addresses. One possible scenario could be that a company could change its intranet over to IPv6 only for its internal communications, making use of the increased functionality of IPv6, and using the IPv4 addresses it has freed up by doing this for its NAT routers. Then any communication outside its network to an IPv4 computer would be possible.

### The Road to IPv6

With all the perceived advantages of IPv6, one may ask why the world has not deployed it already. This is one of the key questions facing the Internet development community and it is a source of deep division. The problems of address shortage and flat addressing hierarchy are the areas of highest impact for IPv6. However, much effort has been expended to improve the existing IPv4 schema to encourage efficiency and scalability of backbone routing (by using an aggregation scheme called *CIDR*), and to minimise the need for public Internet address space (by using network address translation at the edge of the provider network). Many



people now believe IPv6 will not be needed for the foreseeable future. Furthermore, why would anyone consider going through the complexity of the migration process when it is not readily facilitated? There are also other major concerns. Firstly, all today's applications must be rewritten for IPv6. Secondly, current IP network technology is built for IPv4, and much performance optimisation is required for high transport efficiency of IPv6 packets. Finally, current operational support systems will have to be changed to allow support of the new features.

These factors have a major impact on the adoption timescale for IPv6. There is still a very long way to go with IPv4. At some point, driven by the ultimate need for global ubiquity of IP in all manner of everyday machinery, a gradual realisation of ease of deployment may develop and begin to drive the creation of IPv6 networks.

Until then, watch this space. The IP world is changing dramatically and there are major battles ahead. Strong market drive from particular segments and the increasing development of a whole range of novel IP functional services can radically alter industry perception in a very rapid timeframe.

## Conclusion

IPv6 has obvious advantages. It scales to global uniqueness. It has significant support. But, from an economic and management cost viewpoint, will it ever be feasible to migrate to IPv6? This is the big question. It would not be possible or desirable to arrange a single day when people switch over to IPv6—an evolutionary approach is mandatory and the method of achieving this is the subject of much debate within the IP community. An alternative question may be more apt; 'What are the disadvantages if we don't move to IPv6?' A clear view of this is a key objective for the decision to migrate for any major carrier, ISP and large company.

## Acknowledgements

The authors thank David W. Smith and Steve Sim for their comments when this article was being written.

## References

- 1 HUITEMAL, CHRISTIAN. IPv6 The New Internet Protocol (second edition). 1998, p. 11.
- 2 DEERING, S. and HINDEN, R. Internet Protocol, Version (IPv6) Specification. *IETF* draft-ietf-ipngwg-ipv6-spec-v2-02.txt, 6 Aug. 1998.
- 3 HINDEN, R.; and DEERING, S. IP version 6 Addressing Architecture. *IETF RFC 2373*, July 1998.
- 4 HINDEN, R.; O'DELL, M.; DEERING, S. An IPv6 Aggregatable Global Unicast Address Format. *IETF RFC 2374*, July 1998, section 3.1.
- 5 <http://www.6bone.net/6bone-backbone.html>

## Glossary

<b>BGP</b>	Border gateway protocol
<b>CIDR</b>	Classless inter-domain routing
<b>DNS</b>	Domain name service
<b>FP</b>	Format prefix
<b>IETF</b>	Internet Engineering Task Force
<b>Interface</b>	Interface of host; for example, Ethernet card
<b>IP</b>	Internet protocol
<b>IPv4</b>	Internet Protocol Version 4
<b>IPv6</b>	Internet Protocol Version 6
<b>ISP</b>	Internet service provider
<b>NAT</b>	Network address translation
<b>NLA</b>	Next-level aggregation
<b>RES</b>	Reserved
<b>RIP</b>	Route information protocol
<b>SLA</b>	Site-level aggregation
<b>TCP</b>	Transmission control protocol
<b>TLA</b>	Top-level aggregation
<b>UDP</b>	User datagram protocol

## Biographies



**Stuart Prevost**  
BT Networks and  
Information Systems

Stuart Prevost joined BT as a Trainee Technician Apprentice at BT Laboratories in 1987. He has spent most of his career in the Optics Research Division. His early work involved providing technical support for the team that developed the erbium-doped fibre amplifier, and later an optical neural network. After this he moved onto managing the unit's computer servers and joined the Futures Testbed in 1997. His present responsibilities include managing BT's IPv6 6bone site, evaluation of IPv6 and investigating novel networking technologies.



**Tony Dann**  
BT Networks and  
Information Systems

Tony Dann works within the Futures Testbed Group at BT Laboratories. He is leader of the Advanced IP Platform Technology team involved in the practical evaluation of new networking concepts such as high-speed routing, IPv6 and IP QoS. He holds a Ph.D. (University of Nottingham) and has held a Royal Society Fellowship for his work in Japan prior to joining BT.

*John Wittgreffe*

# A Platform for Electronic Commerce Service Trials

*Electronic commerce technology is evolving rapidly, with many vendors spurred on by the prospect of a multibillion pound market in the new century. To keep abreast of developments, BT Laboratories has created an electronic commerce trials platform, incorporating some of the latest e-commerce components. The platform enables rapid development of new e-commerce services by reusing the core e-commerce component set. Customer trials of new services can then be performed over the platform's live area, which uses a range of networks to reach consumers and business customers. This compelling combination of advanced technology, rapid development, and the ability to deliver live trials, is helping BT to craft its new e-commerce services.*

## Introduction

BT Laboratories' *electronic commerce trials platform* was conceived in 1995, with the goal of becoming a world-class platform in its field. By early 1996 the platform had already delivered the UK's first Internet gateway for processing credit card payments, bridging the link between a merchant's Web site and the acquirer. Since then the platform has been among the first in the UK to introduce a series of new e-commerce technologies, from advanced payment services, to on-line catalogues for the Internet, to more futuristic components such as intelligent trading agents.

New services can be configured rapidly by re-using the platform's functionality components, which are available on an internal testbed for development and experimentation. Finished applications are then put to the test in formal customer trials, using a separate high-specification operational platform. The platform now plays host to around 10–15 trials at any time, and is helping deliver some of BT's most innovative new products in e-commerce. The platform has claimed many recent successes, including on-line shopping in London with The Food Ferry, the BT Argent micro-billing service, and on-line payments for Wireplay.

This article gives an overview of the platform, including a discussion on the type of trials supported, and a summary of the platform's technical infrastructure.

## A Selection of Trials

Trials on the e-commerce trials platform are varied in nature,

ranging from new payment services, to novel Internet stores, to on-line catalogues for business procurements. The most convenient way to classify these trials is in terms of business aims rather than technology, where two distinct types emerge:

- a customer trial of a new e-commerce application or service, with a view to downstreaming to a product within one year, and
- a technology trial aimed at proving a more futuristic capability in advance of the market.

In our definition, a customer trial operates a prototype service in the marketplace and may include anything up to 1000 users, usually working in partnership with one or more businesses in the target sector. The trial may have a presence in the public eye and is intended to directly precede a product launch. By contrast a technology trial is relatively low key, available only to BT and collaborating organisations. The operational considerations for both types differ considerably, where the former requires a high operational specification, 24 hour support, and demands that all appropriate checks are in place for legal, regulatory, and intellectual property issues. By contrast, the latter primarily requires an advanced 'live' environment for putting state-of-the-art commerce technology through its paces.

To examine the formula for a successful trial in each case, it is appropriate to consider typical examples. For this article the Food Ferry customer trial and the three-

dimensional retail technology trial have been chosen. Within a few years, services similar to these are likely to be commonplace, both over the Internet and digital TV.

The Food Ferry service is an on-line supermarket trial operating over the Internet, running in the London area in 1998/99. Groceries are selected and ordered over the Internet (see Figure 1) using a high-end store system and are then delivered to the door by The Food Ferry. Before the trial, The Food Ferry already had an existing core delivery business operating using a paper catalogue and more recently a CD catalogue. The Food Ferry had a dedication to reaping the benefits of new technology and could work with BT daily on requirements. This combination made Food Ferry the ideal choice of partner for the initial trial. The location of the business was good from a market perspective, where the professional nature of the London clientele had potential for a high take up of the service, to provide a sizeable customer base for the trial.

Using the platform's capabilities, the application itself was designed and developed in just three months. With a development emphasis on reuse of scalable components, an existing catalogue database and an existing payment component were applied from the platform's portfolio. Further to this, a number of new components were developed for the new 'supermarket' application, including a flexible shopping basket, a new customer-handling component, and a component to aid back-end integration to a stock-control system. The development emphasis was primarily on ease of use in the marketplace and speed of operation. The trial was launched quickly. Once tested, the application could be moved directly from the trials platform's development area to the live area, where it was then available to the public and included within the 24 hour operational support.

From a market perspective, customers were phased in using



Figure 1—A typical screen from Food Ferry, BT's 'supermarket on the Internet' trial

groups of 30, then groups of 50, approximately weekly, to manage carefully any early operational problems. The Food Ferry handled first-line support, and the trials platform team handled second- and third-line support. The service has since proved highly popular with customers where The Food Ferry now complete around 10% of business on-line. During the trial, the platform's datamining techniques were used to evaluate the trial's strengths and weaknesses. This showed that the average on-line order was greater than the average paper catalogue order, and that customers stuck to the system with regular orders. The service has now been downstreamed to BT Syncordia who are using this in product solutions for a number of corporate clients.

Other customer trials live at present include the Argent micro-billing trial; Wireplay payments; Pro-Search, a business information agent working with British Gas; ProSum, a pay-per-usage document summariser; and the BT Labs Shop. Trials recently completed include BT PropNet, an on-line estate agency, the WEFAS on-line classified

advertising service, and a pay-as-you-go text->phone messaging service.

The second category of trial focuses on proving technology, rather than market acceptance. Components are typically more futuristic, possibly more error prone, and may not be in the public eye. For this reason hardware, servers, and components for technology trials are ring fenced from those for customer trials which require a higher quality of operation. The three-dimensional retail trial is one such example. The main purpose of this trial is to prove the use of three-dimensional catalogues in home shopping. For example, if you wish to furnish your home, the system allows you to select furniture from a catalogue on-line, drop this into a three-dimensional view of your target room, decorate your room to taste with carpets, wallpaper, paint, etc. and then walk around your room to view the full effect from all angles. In this case, one retailer was chosen to give the content due realism, namely Laura Ashley. The development team worked rapidly with Laura Ashley and with a specialist contractor CBL to create the new


Figure 2—Three-dimensional on-line catalogue

three-dimensional components that would appeal to the retail clientele (see Figure 2). Existing components from the platform were applied to help construct the on-line catalogue. After initial system testing, the service was made live to a limited number of customers from a cross-section of backgrounds. The key objective was to prove the usability of the new three-dimensional catalogue technology, before taking the technology to market. This trial is currently ongoing. Following this technology trial stage, the service may now go onto customer trial with a number of retailers.

Technology trials are not always intended to test user acceptance. By contrast, another technology trial ongoing at present is the European FACTS commerce task, where intelligent trading agents are being developed between collaborators BT, CNET/ONERA, KPN, and Broadcom. In this case a live trial is ongoing to prove a new European agent-agent negotiation standard, rather than a software technology. Here buyer and seller agents are employed to barter on price and conditions for business travel arrangements. An electronic travel agent works with various supplier agents to construct the best itinerary for travel based on convenience and cost. The output from the trial will be improvements to the agent negotiation standard, with a target of obtaining a specification suitable for products.

### Component Architecture

The e-commerce trials platform uses a component architecture to facilitate rapid and low-cost service development. These components are essentially building blocks of e-commerce functionality, used to support the different business processes in the e-commerce value chain. For each new application, the appropriate building blocks are selected and configured as required. For example, a new on-line store may require a



component to display an on-line catalogue, a shopping basket component, a component for ordering goods, a component for payment by VISA, plus a component to assist with back-end integration. By contrast, an on-line game may require only an on-line micro-payment facility for pay-as-you go, where the game's functionality is outside the scope of the platform.

The platform's components have been developed by four different BT development teams around the globe. The components incorporate both home-grown solutions and sub-components from external suppliers including Oracle, Netscape, Microsoft, and OMI.

An overview of the component architecture is shown in Figure 3.

To assist understanding, the commerce value chain is shown at the top of the diagram. From left to right, this includes pre-sales, sourcing, order placing, supply chain management, settlement, and post sales, with business 'collaboration' shown as the larger box encompassing every stage. The e-commerce components used on the trials platform are shown below, where the horizontal extent of the box indicates the approximate extent of the component in its support of the value chain. Bold text indicates that the components exist on the trials platform already, whereas plain text indicates that the component is under development on the internal testbed.

The 'payment' area of the value chain has been subject to the most development over the past two years. Payment services on-line now include VISA and MASTERCARD payments of £3 and over, a micropayment token facility for pay-as-you-go at less than 10p per unit, a small payments facility for sums between 10p and £3, and a micro-billing facility to allow accumulation of small payments to be resolved monthly by VISA. Each of these was developed at BT, initially in ART, and has since been down-streamed to a range of products including ARGENT (microbilling) and Wireplay. Other developments include cash on a smart card (Mondex) payment systems, and a global token exchange for monetary conversion between different kinds of electronic cash or tokens.

Components for 'ordering' and 'supply-chain management' include comprehensive on-line catalogue and storefront capabilities, with solutions of differing scale for small-to-medium enterprises (SMEs) and corporates. Additional components include an electronic data interchange (EDI) gateway, certification services, and capability to aid integration with legacy systems.

More futuristic components tend to be in the 'pre-sales' and 'sourcing' areas, where there is an increasing focus on the capabilities of intelligent agents for commerce. To this end, BT's award-winning Zeus platform is being incorporated,

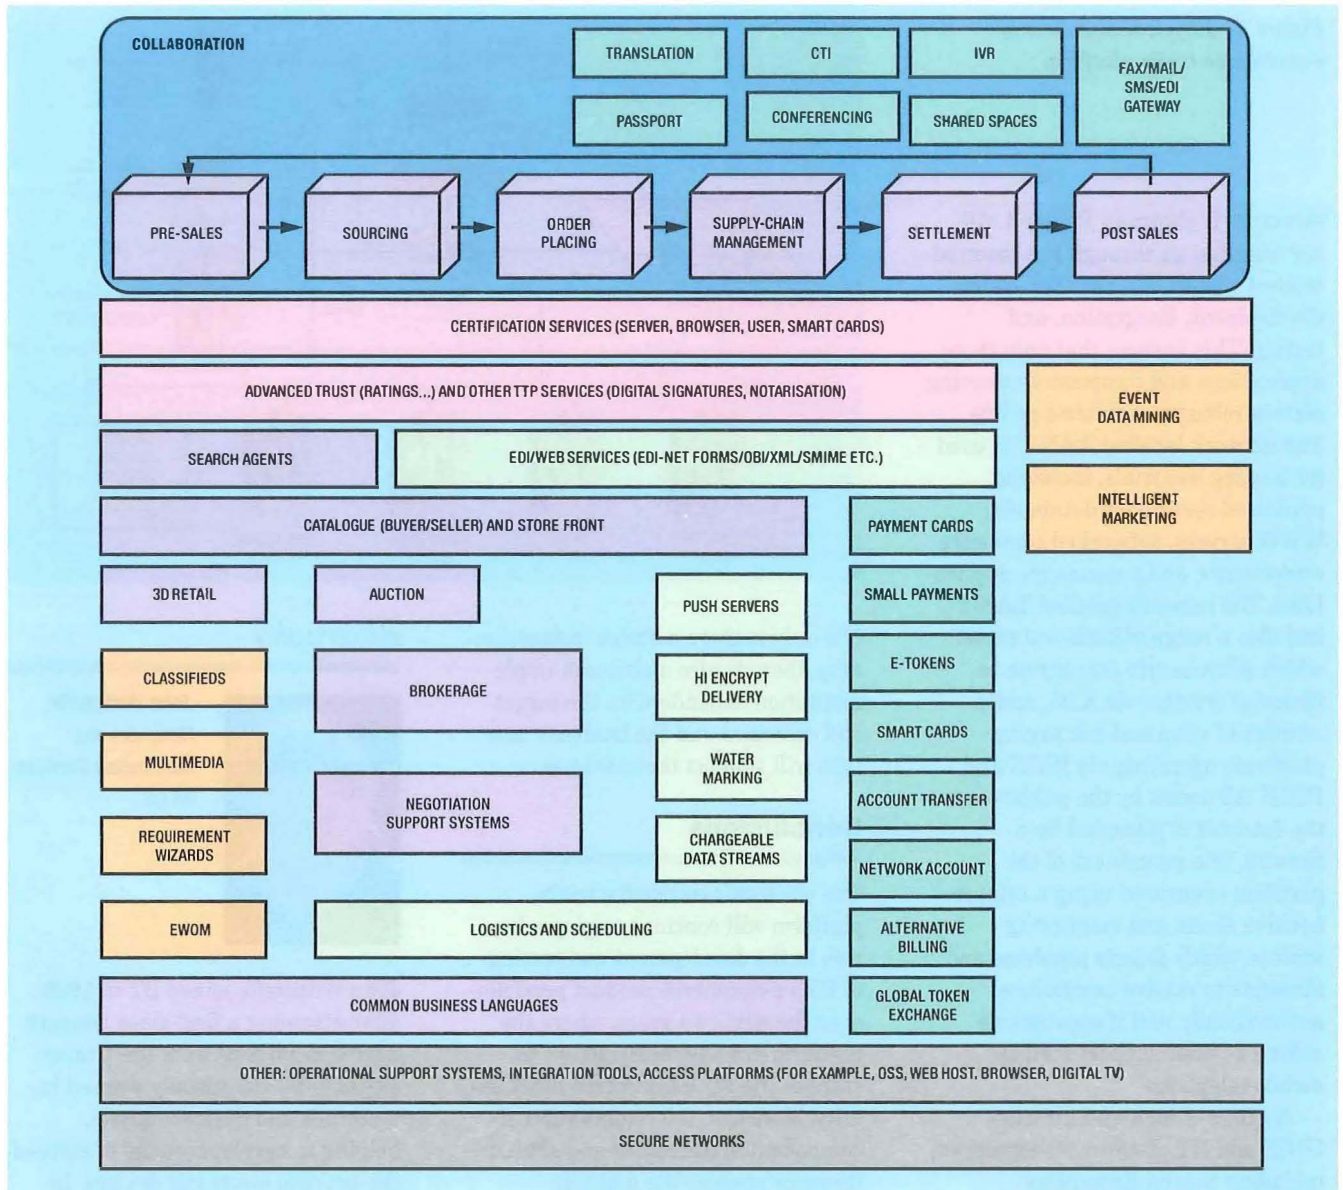


Figure 3—Component architecture of the e-commerce trials platform

where commerce agents are being developed to perform product finding, buying, and selling tasks. For example, an agent is already available to assist a buyer to find and select products from different competing suppliers. This agent was put to the test recently in the PropNet trial, where business accommodation could be selected from a range of competing on-line vendors, covering a combined portfolio of 10 000 properties. In addition, trading agents are already being used in the MATE trial, where agents act on behalf of buyers and sellers to automate the negotiation process within specified parameters and rules. This automation empowers even the smallest trader to negotiate between a largely increased number of traders, looking

to obtain the most suitable deal. For more information on commerce agents, see the related article in this issue of the *Journal* by Martin Owen†.

The collaboration area extends the platform from e-commerce to the broader context of e-business, with a variety of components to assist businesses in human communication during transactions. Messaging systems are largely organisation and application dependent; however, a core capability is provided to complement existing mail and

intranet capabilities. Examples include linkage to an interactive voice response (IVR) platform so that a voice interface can be used to access e-commerce systems over normal telephone lines, expanding use to those who have no data networking.

### Physical View

Key to the success of the platform is a high specification infrastructure of servers and networks, with 24 hour operational support. This ensures that trials can be operated at a high quality of service, while significantly reducing costs by sharing both operational support teams and infrastructure.

The simplified view of the e-commerce trials platform's physical

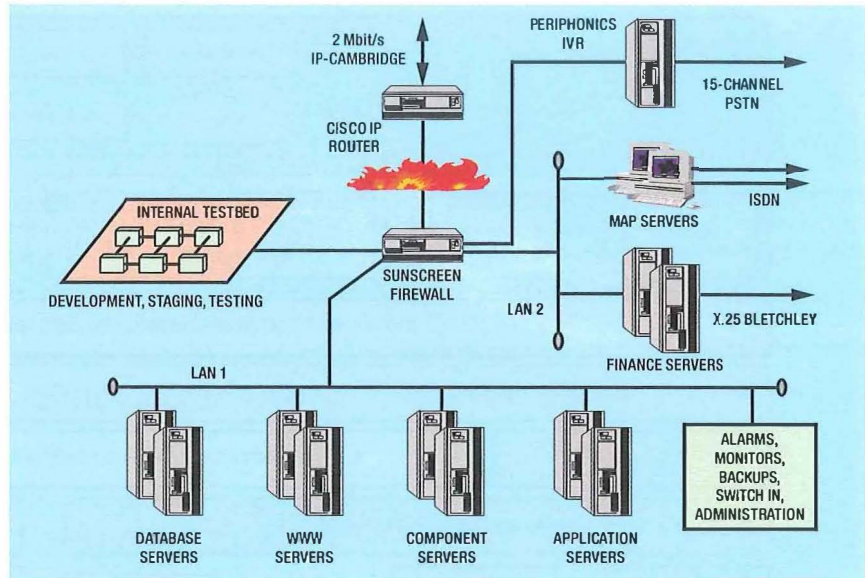
† OWEN, MARTIN; and NÚÑEZ SUÁREZ, JORGE. Agent-Based Solutions for E-Commerce. *Br. Telecommun. Eng.*, Jan. 1999, 17 (this issue).

Figure 4—Physical overview of e-commerce trials platform

structure is shown in Figure 4. All services first go through the 'internal testbed' shown left, used for system development, integration, and testing. This ensures that only those applications and components meeting certain robustness criteria go live. The network labelled 'LAN 1' is used for hosting live trials, including optimised servers for databases, WWW servers, networked commerce components, and e-commerce applications. The network labelled 'LAN 2' includes a range of back-end servers which allow secure connection to financial services via X.25, and a number of voice and messaging platforms operating via ISDN and PSTN. All access by the public over the Internet is protected by a firewall. The robustness of the platform is ensured using a comprehensive alarm and monitoring system, which detects problems and attempts to resolve any failure automatically, and if appropriate calling 24 hour support staff by mobile telephone.

A range of high-specification UNIX and NT platforms is employed, including Solaris Enterprise 450 servers incorporating UNIX on RAID5, and multi-processor Compaq 4000 servers for NT. Services are made live by a range of IP and non-IP networks including 2 Mbit/s Internet, X.25, ISDN30. Web servers include Netscape Enterprise, Oracle Web Application Server, and OMI secure server, database servers include Oracle 8.0.4 with OFA for additional file protection. Physical components span software languages C, C++, Java, PL/SQL, Perl; Web protocols HTML, HTTPS, XML, CGI, Servlets, ActiveX, CORBA ORB; web servers Oracle, Netscape, OMI, and Microsoft. The MAP servers provide text-to-voice messaging and audio conferencing over IP. The IVR platform enables access to commerce functionality using voice or tone commands over any telephone.

This variety ensures that the maximum flexibility is provided in configuring and operating new trials.



Not only is there a choice in functionality, there is also a choice in implementation dependent on the target end customer and the business unit that will support the customer.

## Conclusions

The electronic commerce trials platform will continue to play a key role in the development and testing of BT's e-commerce product portfolio over the next two years, where the platform provides early access to state-of-the-art e-commerce functionality. Moreover, the combination of component architecture and shared resource enables the multiple benefits of reduced time to trial, reduced development costs, and reduced costs for operation of a trial. The flexible approach from both a technical and business sense ensures that BT has the ability to experiment and craft the right combination of infrastructure and functionality for e-commerce.

A selection of the latest public trials to go live on the platform are accessible from a central Web page called *TranSend*, at: <http://transend.labs.bt.com>

## Biography



**John Wittgreffe**  
Networks and  
Information Services,  
BT UK

John Wittgreffe joined BT in 1989 after attaining a first-class honours degree in Physics from the University of York. He initially worked in materials and device analysis, helping to develop a range of state-of-the-art opto-electronic devices. In 1992, John moved to network management research, helping to create BT's operational support system enterprise modelling system TRIADS and its three-dimensional visualiser. In 1996, John developed an early interest in Internet technologies and in particular e-commerce, helping to form a new research group in this area in 1997. Since then John has technically managed and delivered several BT and Internet firsts including the 1997 General Election polling system, BT PropNet trial, the Food Ferry home shopping service, and Iceland home shopping. John is now the technical manager for the electronic commerce trials platform described in this article.

# Agent-Based Solutions for E-Commerce

*Agent technology is emerging as a major software engineering paradigm for the 21<sup>st</sup> century. BT has played a major role in developing tools for agent-based solutions. This article examines how agents can be applied to the domain of electronic commerce. The principal benefits and issues are outlined and viewed from the perspective of the end user, businesses and developers.*

## Introduction

The pace of technological change is speeding up. Nowhere is this more evident than in the arena of electronic commerce. Within the next three years, it is forecast that there will be a five-fold increase in UK e-commerce activity<sup>1</sup>.

The trading environments offered by electronic commerce will be radically different from the traditional market models in existence today. Users in this new rapidly changing market will require ever more sophisticated services to assist them. Buyers will need to identify and purchase the best from the myriad of products and services available on the Internet; suppliers will want to target potential customers in a more efficient manner, increasing revenue and reducing costs. Software, too, needs to keep up with rapid change and increasing complexity. The electronic commerce market will not tolerate slow response or delayed decision making. New ways of developing channels for interactive relationships with customers and building customer communities must be built. It is this business model which is forcing businesses to take note of software agents.

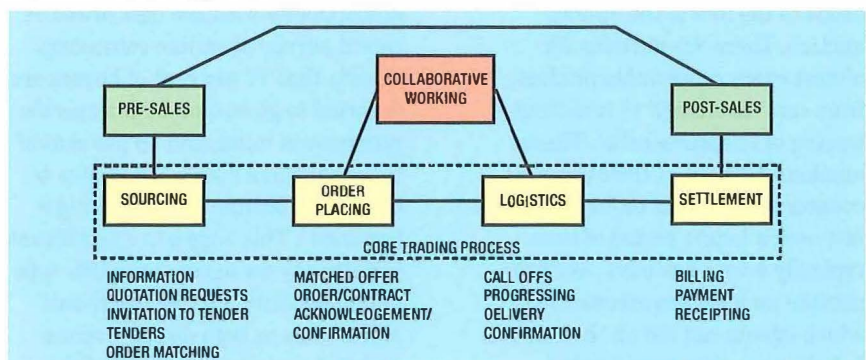
A *software agent* is a software entity that can act in an autonomous manner, can learn, is proactive and has the ability to interact with other entities, be they human or software based. Software agents are predicted to play an increasing role in electronic commerce as they exhibit much of the behaviour considered to be desirable or even essential. Agents endowed with these capabilities can significantly reduce the amount of human effort required, therefore providing real value to the user.

This article discusses how agent technology can be used to support electronic trading. It describes where agent-based solutions are being used and highlights other areas where agent technology is applicable. The infrastructure to support these services is also discussed including multi-agent-based trading platforms and other emerging technologies. The article concludes with a look at the technological and economic impact of the widespread use of agents.

## The Trading Process

The traditional commerce process can be represented by the framework shown in Figure 1.

Figure 1—Trading process



The core of the trading process consists of the following stages: sourcing, order-placing, fulfilment and settlement. The emphasis placed on these functions differs, depending on the market. In the consumer market, the emphasis is on the marketing functions, that is, the pre- and post-sales stages, with focus being placed on market matching and advertising. Companies typically attribute as much as 35 per cent of the costs of producing a product to marketing and sales costs<sup>2</sup>. In contrast, transactions in the business-to-business market typically involve establishing contracts that involve significantly larger sums of money and require more coordination between several enterprises. The processes typically involve multi-party negotiations and can last over an extended period of time. This article aims to describe the current and potential uses of agent technology from the perspective of both the consumer and the business.

### Consumer Market

One of the main trends developing in the consumer market is the rapid rise of market-matching tools. These tools facilitate the matching of buyers' requirements for particular products or services with appropriate suppliers. They can be broadly categorised into three areas: auctions, shopping robots (or *bots* for short) and shopping assistants.

#### Auctions

A simple example of a market-matching tool that has come very much to the fore is the on-line auction. There are auctions for almost every conceivable product, from cars<sup>3</sup> to energy<sup>4</sup> to real-time pricing of telephone calls<sup>5</sup>. These auctions differ from their physical counterparts in that on-line auctions last over a longer period of time, typically over a few days. Auctions provide an ideal environment in which agents can bid on their user's behalf, since they need limited




Figure 2—AuctionBot

knowledge of the strategies of other parties. The agents adopt a strategy that obtains the maximum profit and follows the rules of that market. One such example of this is offered by AuctionBot<sup>6</sup> (Figure 2) which enables users to define agents to act on their behalf, stipulating complex buying rules and even allowing users to specify their own auction protocols.

The agents provide added value for the buyer or seller in several ways: by eliminating the need to monitor constantly the process of buying (or selling), by keeping track of all the bids on the product that

increased transaction volume driven by shopping bots,' said Nicole Vanderbilt, group director, digital commerce at Jupiter Communications. Although some merchants may be reluctant to accept shopping bots which will pit them against their competition solely on price, Jupiter Communications believes that the benefits of bots far outweigh price concerns. On-line merchants must create unique shopping experiences and implement loyalty programs to differentiate themselves from the competition and to convert a portion of these one-time bot-driven shoppers

## *One of the main trends developing in the consumer market is the rapid rise of market-matching tools*

they want to buy or sell and by achieving the best price.

#### Shopping bots

A shopping robot is another example of a market-matching tool. It offers price comparisons to users wishing to make purchases over the net. Shopping bots are automated agents that search the Web for consumer-specified products and generate a list of merchants with the best prices. A recent survey of on-line consumers reports that 77 per cent of buyers are reported to go on-line with a specific purchase in mind, and 79 per cent of those shoppers visit several sites to compare offerings before making a purchase<sup>7</sup>. This suggests a significant opportunity for agents with this type of functionality. 'Online merchants will be able to reap direct revenue and cost savings as a result of the

into a loyal customer base. Typically, shopping bots, such as CompareNet<sup>8</sup> provide comparison shopping, allowing customers the facility to search across merchants based on a set of features, including price and product features; for example, processor speed and RAM. Acses<sup>9</sup>, a bot which finds best prices on books, is sensitive to user's geographical location and offers ranges of prices based on delivery time.

#### Shopping assistants

Auctions and shopping bots are useful in environments, such as in commodity markets, where the product can be characterised solely on price. However, in many markets there is a need to distinguish products on more than just price; for example, quality, service, warranties, brand and delivery cost/times. To add



to this complexity, many products are described by a range of characteristics that differentiate them from other products within the same market. In these types of markets, users have different sets of preferences over these products. In this case, agent technology can be used to provide digital shopping assistants.

A useful analogy can be drawn here to that of a trusted assistant. The assistant will need to understand and elicit the buyer's needs in terms of a precise set of requirements and preferences. They will also educate the buyer as they consider each of the characteristics of the product, informing the purchaser of best products, searching for appropriate suppliers, negotiating on the buyer's behalf and making the final purchase. A good example of where some of this functionality can already be seen is Amazon.com<sup>10</sup>. This site offers functionality which cannot be found in traditional book stores, delivering more information than its physical counterpart. Bookstores are limited by shelf space and reviews and comments are not available to customers while perusing books.

Purchasing advisers that can assist the user in providing this type of functionality will greatly enhance the benefits of the simple shopping bots<sup>11</sup>. However, in order to provide this added value, these new types of agents will need to elicit a complex set of user product preferences. They should be able to extract not only information about what type of product the users require, but more importantly their preferences about each of the product characteristics. Techniques for evaluating preference information about different combinations of product characteristics can be drawn from the domain of multi-attribute utility theory<sup>12</sup> and marketing, such as conjoint analysis<sup>13</sup>. Examples of this type of functionality are already emerging on the Internet, with services such as the decision aid system offered by PersonaLogic<sup>14</sup>. However, like many other e-commerce solutions it is client-server

based, only providing a front-end to the merchant's product description and therefore biased towards the needs of the merchant rather than the purchaser. Providing an infrastructure in which trusted third parties can advise buyers on purchases will enable users to obtain an independent perspective by removing any merchant bias. This is an important ingredient of any trusted purchasing assistant that is currently missing from many solutions.

Multi-agent based systems provide one approach, in which agents can be used to represent the various roles in the market place. Each agent has its own set of objectives or goals, which need to be satisfied. For example, merchants want to sell as many products as they can in order to maximise profit, while buyers want to obtain the best price for any product they purchase. This solution also allows new electronic brokerage roles to be dynamically created in response to demands in the market. These brokers create added value by providing specialised domain knowledge. The subject of brokers will be discussed in further detail later in the article.

Agent technology can also assist both purchasers and suppliers in other parts of the supply chain. Tracking agents can monitor the progress of an order and ensure that it is delivered in accordance to specified contracts<sup>15</sup>. There are also examples of agents automating initial sales enquiries and feedback using intelligent matching of customer queries with products from the suppliers database to tailor e-mail responses<sup>16</sup>.

### Targeted Advertising and User Profiling

A key element of the consumer market is advertising. Information technology has allowed companies to address ever-narrower market segments. This trend is moving towards the ultimate goal of one-to-one targeted marketing, a market of

one! Underlying the immediate benefits of this trend is the establishment of long-term relationships with customers, created by more subtle functionality, such as personalisation and customisation of Web sites. Budgets for broadcast media marketing will soon be re-allocated to interactive media marketing. The revenue generated from sales of Web advertising space has been forecast to be about 15 per cent of that generated from sales of newspaper advertisement space by the year 2002<sup>17</sup>.

Central to this trend is the notion of user profiles that include information about the user's previous transactions, preferences, user interests and personal information. The concept of user profiling within personal agent technology is used to assist in tailoring tasks such as searching and recommendations to the specific needs of the user<sup>18</sup>. This is a powerful technique which companies can apply to generate revenue from cross-selling of related or supporting products and services. Many of the current systems require significant input from the user. This approach is time consuming and susceptible to inaccuracies in the information being provided. However, several new products are coming to the fore that are based on learning the user profile in an unobtrusive manner<sup>19,20</sup>. This approach results in less data entry and is less prone to attempts by the user to falsify the data deliberately. However, it raises privacy issues as users may feel they have less control over what is being stored about them.

The issue of privacy is of paramount importance if user profiles are going to be utilised to the benefit of both the customers and suppliers<sup>21</sup>. To date, this area has been self-regulated. The Open Profiling Standards (OPS)<sup>22</sup> is proposing an open industry standard for on-line exchange of profile information between individuals, groups and businesses. This specification is currently being considered by the

## *A decentralised autonomous agent architecture provides a natural fit for enterprise-wide logistics*

Platform for Privacy Preferences workgroup within the W3C (P3P)<sup>23</sup>. Further work also needs to be carried out into understanding how the level of trust can be built up between user and agent, which goes beyond the traditional security by encryption<sup>24</sup>.

### **Business Process Automation**

The concept of businesses automating their trading process has been in existence for some 15 years with electronic data interchange (EDI). However, with specialised networks and expensive proprietary infrastructure, this has remained the domain of large corporations as it requires large economies of scale before it becomes affordable. EDI also tends to be based on pre-established contracts between customers and suppliers, resulting in a bias towards fixed supplier-customer relationships. Companies face global competition and customers who demand customised products at mass-produced prices. The Internet has reduced transmission costs and increased capacity. It has made it economically viable for small companies and even individuals to trade electronically. In line with the shift to greater customer power, businesses are re-engineering their processes to focus on their customer-facing information systems.

### **Supply chain management**

A supply chain consists of a collection of interdependent steps that, when followed, accomplish a certain objective such as meeting customer requirements<sup>2</sup>. Supply chain management encompasses the coordination of these steps, embracing a range of functions from logistics to knowledge management. As companies realise that competitive advantage can be gained by integrating supply chain activities, efforts will focus on coordinating these functions. This trend towards more efficient operations can be seen in the fundamental shift taking place from the tradi-

tional push-based model (build-to-stock) to a pull-based model (build-to-order), in which the customer initiates the supply chain. In this new model there is the need to support increased customisation, reduced lead times, and improved quality and lower costs.

Logistics plays a key role in supply chain management. Logistics is the process of planning, implementing and controlling the efficient, cost-effective flow and storage of raw materials, in-process inventory, finished goods and related information from point of origin to point of consumption for the purpose of conforming to customer requirements<sup>25</sup>. These are horizontal activities cutting across the traditional functional boundaries of purchasing, production, warehousing and distribution. Consequently, when organisations attempt to optimise the supply chain, current information systems tend to be limited to suboptimising individual steps instead of taking a holistic view that is required to gain the true benefits. A decentralised autonomous agent architecture provides a natural fit for enterprise-wide logistics, where it is necessary to balance the conflicting goals of minimum inventory with demand and production schedules. In this approach, the basic components of an enterprise system, including resources, processes and tasks are represented by agents. The agents coordinate the conflicting goals by using negotiation. This results in compromises between goals at a local level and the system level. Research efforts are already focusing on different aspects of logistics, including enterprise reference modelling<sup>26</sup> and manufacturing systems<sup>27</sup>.

Companies need to build more flexible processes that can adapt to the changing needs of business and take account of relationships between suppliers, partners and customers. Agent-based solutions offer several opportunities to support these requirements, notably in the area of work-flow management systems and

data integration. The objective of work-flow management systems is to replace the flow of paper documents by their electronic equivalents. They support business processes such as the order cycle in which a particular structure can be followed, thus reducing order cycle time and improving error and exception handling. Several Web-based agent tools are coming on the market which facilitate the integration of disparate heterogeneous information sources<sup>28</sup>. These tools can broadly be classified as automated-pull<sup>29,30</sup>, Web automation<sup>31,32</sup> interactive personalised catalogues<sup>33</sup> and information filtering<sup>34</sup>. Although these tools provide different levels of functionality, they provide a common foundation to leverage the knowledge management within their organisations.

### **Brokers**

Brokers create value within a market by:

- matching buyers and sellers;
- aggregating buyer demand or seller products to achieve economies of scale or scope and to reduce bargaining asymmetry;
- facilitating the market by reducing operating costs; and
- protecting buyers and sellers from opportunistic behaviour of other participants in a market by facilitating trust<sup>37</sup>.

There have been many predictions about the demise or otherwise of intermediaries, with new technologies such as XML<sup>36</sup> being used to disintermediate. It is the authors' contention, however, that electronic commerce will bring about many new broker roles owing to the dynamic requirements of a virtual marketplace. Examples of brokers can already be seen in the electronic marketplace (see Figure 3) on both the buyer and supplier side.

Figure 3—Broker roles in an electronic marketplace

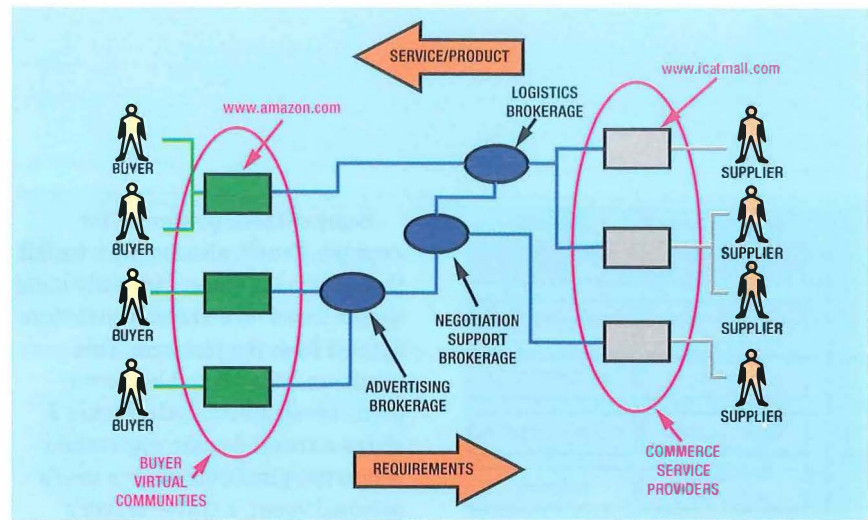
The provision of analysis and negotiation is seen as a key function in many of these roles. Electronic brokering would thus be facilitated by automating the negotiation process.

### Automated negotiation

As discussed in the previous sections, many of the systems on the Internet today, such as the on-line auctions and shopping bots, only allow prices to be compared between several merchants. They do not provide for any bartering, yet negotiation is an integral component of commerce. In traditional commerce, for example, consumers are often faced with negotiating a purchase, be it the price of a new car or moving house.

The subject of negotiation has been studied in a diverse range of fields, including social psychology, management science, economics, game theory and distributed artificial intelligence. The reason for this considerable interest in negotiation is that it plays a central role in our lives in terms of: resolving social conflict; establishing agreements in business; and, increasingly, coordinating interaction among autonomous software agents. As agents take on a greater degree of autonomy, agents have no direct control over each other and thus must persuade their acquaintances to act in particular ways. Negotiation is seen as a natural way in which conflicts can be resolved among these agents, precisely because they are autonomous.

Negotiation can be classified as either distributive or integrative<sup>38</sup>. *Distributive* negotiation occurs when the product or service being negotiated is characterised by one attribute; for example, price. In these situations, one party's gain is another party's loss. In contrast, when the product or service being negotiated is represented by a set of attributes, such as the product features, terms and conditions, the negotiation process is described as *integrative*. In this type of negotiation, the parties make concessions over the set of attributes,



with the potential for all parties to reach a beneficial agreement.

In the following description, the term *contract* is used to represent what is being negotiated between the parties. A contract can represent a purchase of a product, in which case, the negotiation would cover aspects such as the price, product features and warranty. Alternatively, the contract could represent a service agreement between a retailer and supplier, in which case the parties might negotiate over price, distribution and settlement arrangements. The negotiation process can be decomposed into three distinct phases: pre-negotiation, negotiation and post-negotiation. In the pre-negotiation stage, preferences for particular characteristics of the contract are elicited, analysed and revised. These preferences are expressed in terms of a party's utility. While the customer and supplier's utility of particular aspects of a product will inevitably be different, the goals of the customer and supplier may not be incompatible when the total set of product characteristics is taken into account. An individual's set of preferences defines his/her negotiation space. The negotiation phase consists of the actual bargaining. An agreement between the parties may occur if there is an intersection between the customer and supplier's negotiation space, although other factors may prevent an agreement being reached, such as psychological aspects. Each party uses its preference assessment as a means to analyse offers received from the other party and to decide on what counter-offer to make if appro-

priate. At the end of this phase, a decision is made as to whether agreement is possible or not. The post-negotiation phase is optional and may involve a mediating or arbitrating party. In this phase, both parties may submit their preference assessment and negotiated agreement to a third party, who might be able to suggest improved solutions.

A wide range of techniques have been applied to automate the negotiation process, including game theory<sup>39</sup>, genetic algorithms<sup>40</sup> and Bayesian learning<sup>41</sup>. This research will continue to receive increasing interest as companies strive to automate their supply chain to gain competitive advantage.

### Implementation Issues

The previous sections have described the type of functionality required, highlighting the need for a user-centric delegation. This section considers agent technology from the developer's perspective.

Software engineers are encountering increasing difficulties in automating business processes that can handle inconsistent or incomplete data in real-time with traditional engineering techniques. It is the authors' contention that problems with these characteristics lend themselves more easily to an agent-based programming paradigm. Agent-based solutions are key to the development of the next generation of e-commerce solutions enabling users to delegate high-level tasks to their agents. Figure 4 illustrates this shift in programming paradigms, demonstrating that agent technology

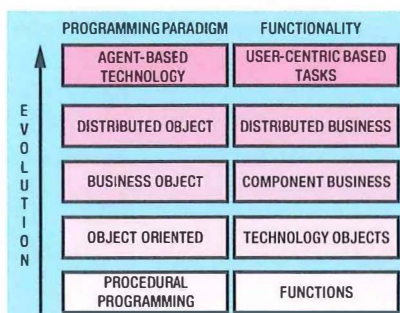


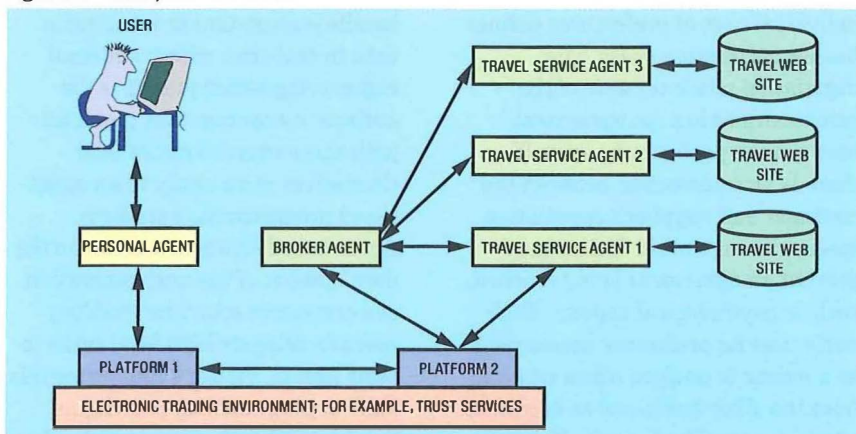
Figure 4—Evolution of programming paradigms

is a natural evolution, building on the distributed object technology currently being used.

### Agent platforms and tools

With the interest in multi-agent based solutions, there has been a rapid increase in the development of agent platforms. The role of an agent platform is to provide inter-operability between heterogeneous agent systems. Platforms achieve this goal by hiding the inherent complexity generated by different communication protocols, security mechanisms, directory systems, etc. from the developers. Some examples of the basic services provided by an agent platform include message routing and parsing, encryption and decryption of messages, directories of registered agents and agent life cycle management. Some platform implementers may choose to add more sophisticated functionality such as electronic payments, planning and scheduling, and reasoning engines.

Figure 5—Platform services



Some of these platforms, for example, Zeus<sup>42</sup>, also supply a toolkit that allows developers to easily build agents based on a skeleton structure derived from the platform. This results in a faster and less error-prone development cycle. Figure 5 shows a travel domain application where two platforms allow a user's personal agent, a travel broker's agent and a travel service agent to communicate. The platform protocols and languages used between the platforms and the agents registered with it may be completely different.

### Standards

One of the many advantages of agent technology is that it provides an elegant way to deal with the complexity of interoperation of heterogeneous systems. Since one of the most common approaches to

*Agent technology will have a significant effect on the way in which people use and view the Web.*

interoperation is standardisation, it comes as no surprise that several standards are either being defined or borrowed from other areas.

The most ambitious effort in this direction comes from the Foundation of Physical Intelligent Agents (FIPA)<sup>43</sup> who are currently standardising the following areas: an agent communication language, agent

management (creation, destruction, registration and mobility of agents across platforms) and software/agent interaction, security and human-computer interaction. Ontology services are also being studied which provide a service for representing and communicating knowledge between application domains. The result of this work will allow a FIPA-compliant agent to discover services on-the-fly, to migrate across platforms, and to send and receive messages securely.

Another standard that is receiving attention from the agent community is XML<sup>36</sup>, which provides a means to structure Web content. This is especially relevant if agents need to discover product and pricing information, compare it and manipulate or process it on behalf of the user.

### Technological and Economic Impact

Agent technology will have a significant effect on the way in which people use and view the Web. As more companies conduct their business over the Web, users will increasingly require more sophisticated functionality to manage the seemingly unlimited volume of information accessible over the Internet. Agent technology is seen as a key enabler, providing users with the capability to delegate high level tasks that not only search Web sites but also act on this information once the information has been collated. This trend will have the effect of making users less reliant on accessing Web sites directly which in turn may mark a transition to making the Web more transparent. The introduction of interactive digital TV will only accelerate this trend as more of the population gain access to the Internet.

The introduction of agent technology is also having an economic impact. As agents become more capable of discovering accurate product and pricing information, it will have an impact on the behaviour of the agents present in the electronic markets. It will become more difficult to exploit information, geography and other asymmetries. Studies conducted by IBM indicate that the impact of both shopping could result in price wars in which there are sharp discontinuous changes to economic stability<sup>44</sup>.

Although electronic commerce has existed for many years in the form of EDI trading, this technology only simplified the transactions between pre-established partners. It was a simple extension of existing channels and thus did not have a significant impact on the economic behaviour of the parties involved. In the business-to-business world, agent technology is facilitating greater automation of the supply chain, producing new trading models. The advent of electronic trading over the Internet has already significantly changed some economic sectors such as the publishing industry (for example, Amazon.com is already the second largest bookshop distributor in the US). A completely new aspect introduced by automated electronic trading is the fact that a competitor may suddenly appear on the other side of the globe through the formation of a virtual coalition of suppliers over the Internet. As a result of businesses gaining more accurate and up-to-date information, it is the authors' opinion that some sectors of the economy could mirror the chaotic behaviour of the financial market. Further work must be carried out in this area to investigate these effects.

## Conclusion

If businesses are to survive into the 21<sup>st</sup> century, they must be prepared to offer customers new services which can be obtained anyplace, anytime, anywhere. This article has shown how agent technology offers one

approach that is able to handle the complex nature of e-commerce. There is evidence to suggest that those businesses that have already deployed agent-based solutions within their organisations have gained a significant advantage over their competitors. However, there are still many complex issues such as trust and negotiation that need further investigation. As more businesses aim to adopt electronic commerce, it will drive the need to invest more research into these areas. In conclusion, agent technology provides a major new engineering paradigm that businesses must understand and use if they are to successfully meet the challenges of building the next generation of e-commerce solutions.

## Acknowledgements

The authors would like to thank colleagues within BT for their support in the production of this article, in particular Simon Steward. The authors would also like to thank Michigan University for permission to print a screenshot from its AuctionBot Web site.

## References

- 1 IDC. Electronic Commerce in the UK. IDC UK Research, Oct. 1997.
- 2 KALAKOTA, R.; and WHINSTON, A. B. Electronic Commerce: A Manager's Guide. Addison Wesley, 1997.
- 3 Aucnet. <http://aucnet.com>
- 4 KNE. <http://www.kne.com>
- 5 ArbiNet. <http://www.arbinet.com/>
- 6 Auction Bot. <http://auction.eecs.umich.edu/>
- 7 Jupiter/NFO. <http://www.jup.com>. Interactive survey, 1998.
- 8 CompareNet. <http://compare.net/>
- 9 Acses. <http://www.acses.com>
- 10 Amazon. <http://www.amazon.com>
- 11 STEWARD, S.; and VIDELO, I. Intelligent On-Line Purchasing. *Br. Telecommun. Eng.*, April 1998, **17**, p. 36.
- 12 LUCE, R. D.; and RAIFFA, H. Games and Decisions. John Wiley & Sons. Harvard University, 1967.
- 13 GREEN, P. E.; and SRINIVASAN, V. Conjoint Analysis in Consumer Research: Issues and Outlook. *J. of Consumer Research*, Sept. 1978.
- 14 PersonaLogic. <http://www.personalogic.com>
- 15 Fedex. <http://www.fedex.com/>
- 16 Brightware. <http://www.brightware.com/>
- 17 Jupiter Communications. 1998 Online Advertising Report. Jupiter Communications Research Studies, June 1998.
- 18 Firefly. <http://www.firefly.com>
- 19 AgentWare. <http://www.agentware.com>
- 20 OpenSesame. <http://www.cra.com>
- 21 Forrester Group. The Privacy Bomb. Forrester Group People and Technology Strategies, April 1998, **3**(2).
- 22 Open Profiling Standard. <http://developer.netscape.com/ops/proposal.html>
- 23 P3P. <http://www.w3.org/P3P/>
- 24 SKEVINGTON, P.; and REA, T. Engendering Trust in Electronic Commerce. *Br. Telecommun. Eng.*, Oct. 1998, **17**, p. 150.
- 25 Council of Logistics Management. <http://www.clm1.org/>

- 26 NICOLO, V.; MADENA, F.; COSSIO, A.; and PEPE, A. An Organisational Approach to Manufacturing Systems. Proceedings of 2<sup>nd</sup> Int. Conf. Autonomous Agents, May 1998, Workshop on Agent Mediated Electronic Trading, p. 88.
- 27 PARUNAK, H. VAN DYKE; BAKER, A. D. and CLARK, S. J. The AARIA Agent Architecture: From Manufacturing Requirements to Agent-Based System Design. Workshop on Agent Mediated Electronic Trading, AMET 98, p. 136.
- 28 MOUGAYAR, W. Future of Agent-Based Commerce on the Web. <http://www.cyberm.com>
- 29 Canopy/Junglee. <http://www.junglee.com>
- 30 Jango. <http://jango.excite.com>
- 31 AgentSoft. <http://www.agentsoft.com/>
- 32 WebMethods. <http://www.webmethods.com>
- 33 Tesseract. <http://www.tesseract.com>
- 34 GroupLens. <http://www.netperceptions.com>
- 35 SOLTYSIAK, S.; and CRABTREE, B. Knowing Me, Knowing You: Practical Issues in the Personalisation of Agent Technology. Proceedings of Practical Applications of Intelligent Agents and Multi-Agent Technology (PAAM), March 1998, p. 467.
- 36 XML. <http://www.w3.org/XML/>
- 37 BAILEY, J.; and BAKOS, Y. An Exploratory Study of the Emerging Role of Electronic Intermediaries. *Int. J Electronic Commerce*, Spring 1997, 1, p. 7.
- 38 RAIFFA, H. The Art and Science of Negotiation. Harvard University Press. 1982.

- 39 ROSENSCHEIN, J. and ZLOTKIN, G. Rules of Encounter: Designing Conventions for Automated Negotiation among Computers. MIT Press, 1995.
- 40 MATOS, N.; SIERRA, C.; and JENNINGS, N. Evolutionary Computing and Negotiating Agents. Int. Conf. Multi-Agent Systems. ICMAS 98.
- 41 ZENG, D.; and SYCARA, K. Bayesian Learning in Negotiation. Agent Theories and Language Workshop. Proceedings of European Conf. Art. Intell. (ECAI) 1996.
- 42 COLLIS, J. C. The Zeus Agent Building Tool-kit. *BT Technol. J.*, July 1996, 16(3), p. 60.
- 43 FIPA. <http://www.drogo.cse.lt.stet.it/fipa/index.htm>
- 44 IBM Watson. [http://www.ibm.com/iac/papers/alife\\_journal98/alife.html](http://www.ibm.com/iac/papers/alife_journal98/alife.html)

### Biographies



**Martin Owen**  
Networks and  
Information Services,  
BT UK

Martin Owen currently leads the work on agents and e-commerce in the Electronic Commerce Research group in BT UK. He gained an M.Sc. in Mathematics and Modern Control Systems in 1974. After a period in teaching, he was awarded an M.Sc. in Computer Science in 1991 from the University of Essex, focusing on artificial intelligence and formal methods. From 1991–94, he worked on requirements analysis and design of major network management systems. After a period working in formal methods and service creation, he moved to the field of agent technology. He has been involved in the initiation

of several European projects and has led several agent-related projects, including an activity in a EURESCOM project investigating the use of mobile agents. He commenced a part-time Ph.D. in the use of negotiation and agent technology at Queen Mary and Westfield College, University of London, in 1995.

Martin Owen may be contacted at [martin.owen@bt-sys.bt.co.uk](mailto:martin.owen@bt-sys.bt.co.uk)



**Jorge Núñez Suárez**  
Networks and  
Information Services  
BT UK

Jorge Núñez works in the Electronic Commerce Research group in BT UK and is currently investigating the use of intelligent agents in e-commerce. He received a B.Eng. in Telecommunications Engineering from the University of Vigo (Spain) in 1993. He also received a M.Phil. in Computer Science from the University of Brighton (UK) for his work on the application of multi-agent technology to complex real-life planning domains.

Jorge Núñez Suárez may be contacted at [jorge.nunez@bt-sys.bt.co.uk](mailto:jorge.nunez@bt-sys.bt.co.uk)

*Alwyn Lewis, Graham Cosier and Charles Nightingale*

# Whither Video?— Videotelephony in Perspective

*Television is very popular, but adding video pictures to telephones has not proved so appealing. This is the second article in a five-part series about the future role of pictures in telecommunications. The first article<sup>1</sup> outlined the development of television and discussed its contradictory reputation – clear and close yet reduced and remote.*

## Broken Dreams

The launch of videotelephony, as an enhanced replacement for the telephone, has proved unsuccessful more than once. The first public videotelephone service opened in 1936, on a dedicated link between Berlin and Leipzig, but was closed due to lack of traffic. Figure 1 shows an early prototype of a British videotelephone.

In the 1970s, AT&T spent an estimated \$500 million launching the Picturephone™, using analogue video technology. This service proved to be spectacularly unpopular<sup>2</sup> and became known as the *Edsel* of the Bell System, despite offering a better quality of picture than today's public switched telephone network (PSTN) videotelephones.

The reasons proposed to explain this unpopularity were:

- the limited bandwidth of contemporary networks,

- the small size of the pictures, and
- the high cost of the product and the calls.

Noll, who conducted a market survey<sup>3</sup> for AT&T in 1973, reported that customers felt no need to see, or even disliked seeing, images of each other on the Picturephone. Jakob Nielsen, a frequent user of modern videoconferencing equipment at Sun Microsystems, is also unenthusiastic<sup>4</sup>. He complains that it is too easy to ignore the people 'on the box' and that crucial but subtle aspects of face-to-face conversation can be lost.

## Less from More?

Is there a fundamental barrier to the effectiveness of videotelephony? Perhaps all instant pictures share the 'clear and close but reduced and remote' character of television. Might the accuracy and immediacy of a video image be a disadvantage? Is a conversation by telephone more subtly expressive and powerful than the same conversation by videophone? Does a moving picture reduce empathy by inhibiting the imagination? Do minds meet more easily when voices are disembodied?

Recent work at the Human Communication Research Centre in Edinburgh<sup>5</sup> suggests that, in some circumstances, this is so. In face-to-face conversation, we look directly at each other relatively infrequently. If a video link forces constant face-to-face gaze, the resulting distraction seems to increase the amount of speech needed to complete a task.

## Your Place or Mine?

Where, mentally, do telephone calls take place? Does the telephone bring the

Figure 1—An early videophone prototype, using a close-up view of the face (picture copyright British Telecommunications plc)



## *There is abundant evidence for the success of video communications in a growing range of unconventional applications in education, training and medicine*

distant person 'here' or is the local user taken 'there'? The universal appeal of telephony, and much of its power, may lie in its ability to take people to an imaginary 'somewhere else'. That place seems to mix the characteristics of 'here' and 'there', in territory that is more mutual than neutral. In football terms, the match is played at a friendly third-party ground.


The addition of pictures, and the sense of reality they bring, might change this situation. When a solo user joins a distant group by a video link, that user is likely to feel taken 'there' by the link. The solo user is playing an away match.

What happens in person-to-person videotelephony? Do the users change places when each is taken 'there', or do they meet in the middle? The authors suspect that both users feel the other is taken 'here'. Both teams are playing a home match, which is contradictory and might create a mutually incompatible kind of conversational behaviour. This reasoning implies that videotelephony might genuinely be a paradox, because adding real pictures could diminish the sense of imagined meeting. If so, inserting live images of people into synthetic backgrounds might increase the appeal of videotelephony.

### **Specific Success**

Yet other evidence is contradictory, because videoconferencing, using carefully-selected and high quality equipment, is increasingly successful and attractive for meetings and group collaborations. In Greenland, where travel is always expensive and often difficult, videoconferencing is used for legal, medical and educational applications<sup>6</sup>. It is also increasingly popular in business, where reduced travelling is a major incentive.

Videoconferences are typically more formal and organised than telephone calls, often with a chairperson in charge. Large and small companies have found that videoconferencing is an effective way of coordinating geographically-separated teams. The experience of the authors is that current



*Figure 2 – Frank Dobson watches a foetal scan at the Women and Children's Centre in Deptford Market, with a video link to experts in Guy's and St. Thomas's hospitals (credit: David Allen / Waymark)*

products provide natural interaction most easily when the users know each other. There is also abundant evidence for the success of video communications in a growing range of unconventional applications in education, training and medicine (see Figure 2)<sup>7-9</sup>.

The rapidly-diminishing cost of video technology means that a fresh look at the design of video communication equipment is now due. Ironically, the strongest resurgence of interest comes from outside the video and telecommunication industries. The computer industry has brought a different design

philosophy, where imaginative dreaming and multi-minded synthesis are given greater weight<sup>10,11</sup>. The success of this approach, in finding new markets and applications, suggests that users have not had a big enough say in the development of videotelephony.

### **Developing Standards**

BT has long played a major role in the development of video-related international standards<sup>12</sup>. Table 1 lists some current ITU Recommendations for videotelephony.

**Table 1 ITU Recommendations for Videotelephony**

Number	Title	Date Issued
H.261	Video codec for audiovisual services at p x 64 kbit/s	March 1993
H.263	Video coding for low bit rate communication	March 1997
H.320	Narrow-band visual telephone systems and terminal equipment	March 1996
H.321	Adaptation of H.320 visual telephone terminals to B-ISDN environments	March 1996
H.323	Visual telephone systems and equipment for local area networks which provide a non guaranteed quality of service	November 1996
H.324	Terminal for low bit rate Multimedia Communication	March 1996
H.331	Broadcasting type audiovisual multipoint systems and terminal equipment	March 1993



## Telephony provides a simple yet superb interface between sentiments and signals.

Current BT videotelephony and videoconferencing products include:

- *DVS100* PC-based desktop videophone and Internet data conference terminal,
- *VS1* self-contained videophone with ten-inch LCD display, for one ISDN line, and
- *VS3* room videoconference system with auto-tracking, for ISDN up to 2 Mbit/s.


### Meeting in Mindspace

Future developments promise truly interactive multimedia and 'virtual-world' applications of video pictures. Engineers, politicians or the public might 'walk' through a proposed building or 'use' a proposed product in a virtual space (Figure 3), interacting with computer-generated objects and each other. The Helsinki Telephone Company has started a collaborative project, called *Helsinki ARENA 2000*, to develop a three-dimensional computer model of Helsinki linked to video cameras, and to people, in the real city<sup>13</sup> that can be accessed via a home computer.

### Shared Barriers?

The effectiveness of these shared-imagination environments will be

*Figure 3—Collaborative working in the VisionDome (copyright British Telecommunications plc)*



much reduced if they are limited to singleton experience. The usefulness and value of these new kinds of communication depend on the subtleties of social and group interaction. The same barriers that inhibited the widespread acceptance of videotelephony potentially inhibit these new developments.

Understanding the human interface issues of videotelephony is therefore important because the appropriate use of pictures in telecommunications has a relevance stretching far beyond the current boundaries of telephony. Making an effective interface between the technical signals that represent pictures and the human sentiments of the users is an insufficiently understood but vital design aspect of videotelephony and telepresence.

### Co-evolution

Telephony provides a simple yet superb interface between sentiments and signals. The output voltage of a handset microphone is an excellent way of expressing the content and purpose of human speech. At the turn of the century, handset design was progressively refined as the social norms for making calls were being established. There was a co-evolution of the technical and social factors, creating a signal-to-sentiment interface which has become universally recognised as the metaphor for telecommunications—the telephone handset.

### Sleepwalking

Because of this success, the industry could sleepwalk into success, blessed with an easy-to-use product that everyone liked and wanted. Providing more of the same, at less cost, to more people was an immensely difficult challenge that occupied generations of engineers. Knowing why this should be done was never in doubt.

Telephones that once used carbon microphones and transformers now

use electret microphones and silicon chips. The signals have also changed, from analogue to digital form. But none of these developments has altered what the signals represent. There has never been any need to re-examine the link that was established, at birth, between the technical signals and the human sentiments in telephony. Until recently, the past was a good guide to the future for telephony.

Expanding the signal-to-sentiment link with video pictures has not duplicated the runaway success of the handset. The output voltage of a video camera is not necessarily a greater expression of the mental content and human purpose of a conversation than the output voltage of a microphone.

### What Next?

Telecommunications faces questions about what to do next, because new technology is abundant—in networks and terminals as well as in services and standards. The best example of the pace of this change is the stunning growth of the Internet. But the important questions are deeper and more fundamental than changes of technology. Technology is no longer the problem.

The hardest questions to answer are increasingly about purpose not technique, about why rather than what. The focus has shifted to gaining a better understanding of human needs, motives, abilities and reactions in both social and business interactions. Therefore, for more effective use of video in telephony and telepresence, a brief review of relevant human abilities is appropriate.

### Human Abilities

The human visual system, through a combination of optics, muscle control and brain processing, achieves a seemingly contradictory combination of high resolution and wide-angle yet rectilinear imaging. The peak

Figure 4 – Representation of the human eye's resolution, assuming a curved field

resolution of the eye (Figure 4) is about 60 seconds of arc, but we are usually unaware that our peripheral resolution is much lower. But our peripheral field is more sensitive to movement and to low levels of illumination than our central vision. The angle of view of human sight is about 160 degrees horizontally and 120 degrees vertically.

We are very well adapted to the identification of faces and facial expressions from surprisingly little visual information. The effectiveness of Baird's 30-line television system depended on this fact. In face-to-face conversation, distance is often related to intimacy. The subjective impression of distance in videotelephony is sometimes unlike face-to-face conversation, potentially creating inappropriate cues of intimacy or remoteness.

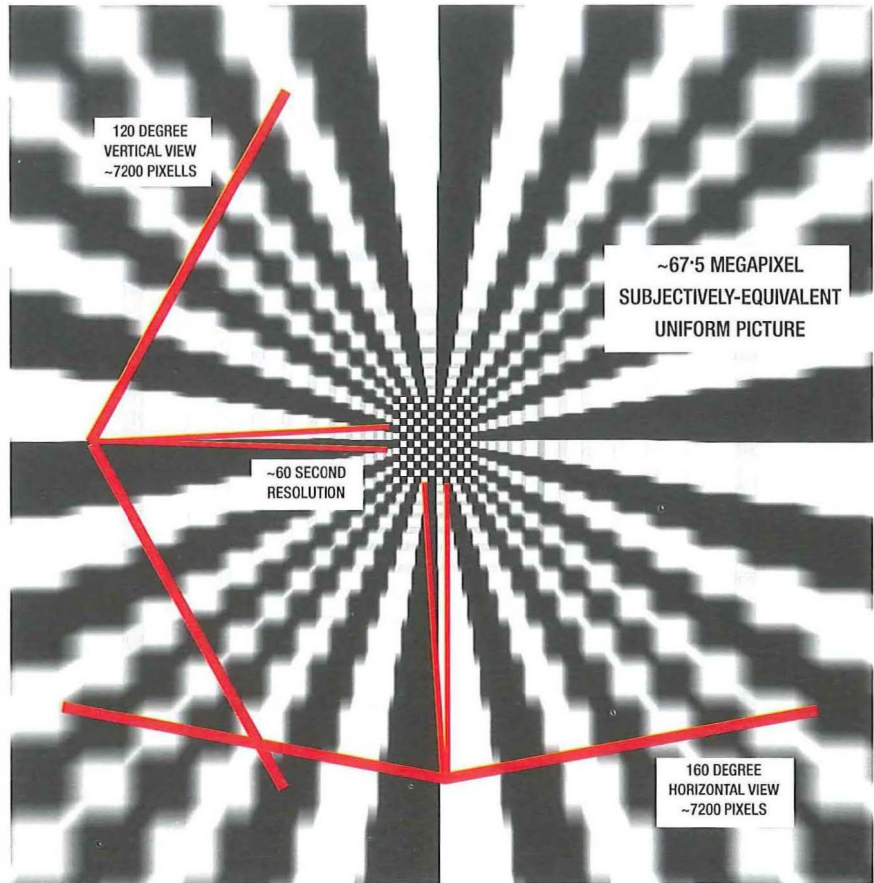
Human hearing has a much lower localisation accuracy than the visual system, but its sensitivity extends

*Gesture and body language can reinforce, diminish or even contradict the meaning of speech.*

behind, above and below the body. The visual and auditory systems interact in a non-linear way, so that sounds tend to be localised at visually likely positions and the perception of picture quality can be changed by the quality of sound.

**Seeing, Hearing and Meeting**

Human face-to-face communication is multi-modal, with markedly non-linear interactions between information in the different modes. Gesture and body language can reinforce, diminish or even contradict the meaning of speech. Cultural and social conventions in speech, gesture and posture are a subtle yet potent means of moderating and directing group interactions, especially during argument and debate.



An exchange of notes, figures or scribbled drawings is a common and effective way of structuring debate and focusing decision-making in face-

problem). However, this effect is less serious for pictures of people than it is for sound or text, because movement is an attention cue. Divided attention occurs naturally in face-to-face interactions with relatively little effect.

**In Human Terms**

Audio and video signals have very different technical and human characteristics, in conception, production and perception. Table 2 shows the properties of audio and video signals, when created and interpreted from a human perspective.

Speaking is cognitive and easy to control—calm phrases can accompany inner turmoil. It is much harder to consciously alter how we look—emotion tends to form our appearance and behaviour. Circumstances can lead us to choose our words with

to-face meetings. In group interactions, we can signal attention, reflection, agreement, surprise, doubt, enthusiasm, distraction, boredom and many more states of mind—all without saying a word. Indeed, two-thirds of all communication is said to be non-verbal<sup>14</sup>.

We can easily engage in a conversation while reading a book. But we find it hard to deal with the interpretation of multiple sensory inputs in the same mode (the divided attention

Table 2 Human Properties of Audio and Video Signals

	Transmission	Reception
Audio	Strongly cognitive Controllable	Cognitive Uncontrollable
Video	Mainly affective Circumstantial Mainly uncontrollable	Affective Cognitive Controllable

Figure 5—A technocentric view of videotelephony

care but can dictate our appearance or behaviour without choice. Listening is mainly cognitive, with emotional overtones, while looking can be cognitive, emotional or somewhere in between. It is impossible not to hear what is audible, yet it is easy to avoid seeing what is visible. This marked asymmetry of human properties suggests that the technical skills and terminology of speech communication are not necessarily useful in a video context.

### Putting Ourselves Across

If the information flow in face-to-face conversation is so complex, how is it that telephony works so well? The enormous success of the telephone owes as much to biology and brain power as it does to technology. When holding a handset, we change and focus our behaviour in quite subtle ways. We make a mutual pact with the distant person, to give a telephone performance, confident of how that person will experience our behaviour. Adapting our speech to the limitations of the telephone, in intonation, volume, choice of phrase and speed of delivery, is easy.

This psychological adaptation is an acquired skill. Shortly after inventing the telephone, Alexander Graham Bell asked the Chairman of an American Bank to use an early kind of telephone to speak to a clerk in a nearby room. The banker froze in horror, unsure of how or whether he might be heard, but not wanting to give offence by saying nothing. The serious and solemn gentleman shouted a stream of childish gobbledegook. The nervousness of some users, confronted with a videophone for the first time, may have the same cause.

### Eye of the Beholder

For relaxed and natural communication, we all need to know how we are coming across. Communication failure in telephony is often associated with the listener hearing


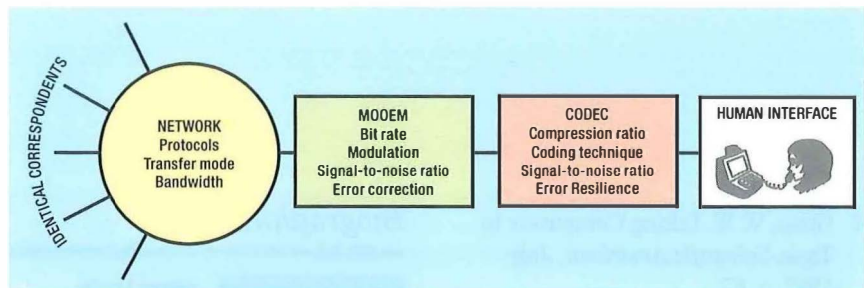


Figure 6—A human-centric view of videotelephony

something out of alignment with the expectations of the talker. Achieving psychological confidence in videotelephony—being able to enter the eye of the beholder—is much more difficult to do, because it depends on factors that are insufficiently understood.

The authors believe that is why videotelephony is a paradox—failing in general yet succeeding in specific applications. Videotelephones can be more powerful and effective than simple telephones, but they are nowhere near as simple to get right.

### Viewpoints

The important tasks in videotelephony have traditionally been the coding and transmission of pictures—the technocentric view of Figure 5. The authors think that the human interface issues, and the associated human behaviours, values and opinions, as shown in Figure 6 are the most important and hardest to understand aspects of videotelephony.

The technocentric view has diverted attention away from three essential kinds of skill and knowledge:

- in physical and psycho-physical optics—the cameras and displays;

- in pictorial culture—the artistic aspects of picture control and composition; and
- in psychology—how meaning and content are represented and conveyed by pictures.

### Coming Next

The next article in the series will discuss pictorial culture—the physical and artistic factors often involved, consciously or unconsciously, in the making, taking and viewing of pictures.

### References

- 1 LEWIS, ALWYN; and COSIER, GRAHAM. Whither Video?—Television and Telepresence. *Br. Telecommun. Eng.*, Oct. 1998, **17**, p. 158.
- 2 Bell Laboratories Record (complete issue devoted to the Picturephone), May/June 1969, **47(5)**.
- 3 NOLL, A. M. Anatomy of a Failure: Picturephone Revisited. *Telecommunications Policy*, May/June 1992, pp. 307–316.

- 4 GIBBS, W. W. Taking Computers to Task. *Scientific American*, July 1997, p. 67.
- 5 DOHERTY-SNEDDON, G.; ANDERSON, A. H.; O'MALLEY, C.; LANGTON, S.; GARROD, S.; and BRUCE, V. Face-to-Face and Video Mediated Communication—A Comparison of Dialogue Structure and Task Performance. *Jnl. of Exp. Psychology—Applied*, 1997.
- 6 SHANKAR, B. Breaking the Ice in Greenland. *Telecommunications*, July 1997, pp. 57–59.
- 7 <http://www2.state.ga.us/Departments/DOAS/GSAMS/> (distance learning).
- 8 <http://www.vision4.com/case2.htm> (health and safety training).
- 9 <http://www.mayo.edu/telmed/history.html> (telemedicine).
- 10 <http://www.infotech.co.nz/spreports/asbvid.html> (prisons in New Zealand).
- 11 <http://www.wpine.com/Products/CU-SeeMe/cu-success-stories2.html> (personal desktop videotelephony).
- 12 WHYBRAY, M. W. *et al.* Videophony. *BT Technol. J.*, July 1990, 8(3), pp. 43–45.
- 13 SUNDQUIST, H.; and KARITUNEN, J. An Evolutionary Strategy for New Communication Services. *Br. Telecommun. Eng.*, April 1997, 16, pp. 49–55.
- 14 BIRDWHISTLE, R. *Kinesics and Communication*. Univ of Pennsylvania Press, 1970.

The authors can be contacted as follows:

Alwyn Lewis: lewisav@boat.bt.com  
 Graham Cosier: cosierg@boat.bt.com  
 Charles Nightingale: nightic@boat.bt.com

## Biographies



**Alwyn Lewis**  
 Networks and  
 Information Services,  
 BT UK

Alwyn Lewis is an advisor in Advanced Applications and Technolo-

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



**Graham Cosier**  
 Networks and  
 Information Services,  
 BT UK

Prior to 1984 Graham Cosier worked in the area of electro-acoustics and published several papers on auditory human perception which contributed to the British Standard for hands-free telephony. In 1986, he became involved with the emerging digital mobile radio system in Europe and was elected chairman of one of the technical groups that standardised GSM. In 1988, he became head of the group concerned with the assessment of coding methods, channel optimising techniques and echo control systems for

increased utilisation of services using telephone networks. In 1994, he was appointed UK coordinator for ITU SG12 where he led the formation of the international standard for digital telephony. In 1995, he was appointed as the BT technical manager for the MIT Media Lab and was involved in establishing the UK Creative Art and Technology Centre. He is currently head of one of the units within the Applied Research and Technology Department at BT Laboratories where he leads several research teams looking into future telepresence systems, perceptual modelling, media environments, audio and vision based perception, gesture and affective communications. He is also the business support manager for BT Global marketing, championing a vision of 'through the screen telepresence'—the ultimate collaboration space, that will provide a 'better than real life' experience between remote locations with the vision that 'you don't have to be there, to be there'.



**Charles Nightingale**  
 Networks and  
 Information Services,  
 BT UK

Charles Nightingale joined the Speech and Technology Unit

in 1996, as a Senior Research Fellow. Previously he had worked for BT from 1971 to 1992, latterly as head of the Image Processing Section. He has a B.Sc. in Mathematics from London University, an M.Sc. in Applied Mathematics from Nottingham University, and a Ph.D. in Non-linear Network Theory from the University of Newcastle-upon-Tyne. He has many published papers, including some in signal processing, image processing, neural networks and in computational aesthetics. His current main interest is in image interpretation and perceptual organisation. He also has a strong interest in the human aspects of visual telecommunications.

*David J. T. Heatley, David R. Wisely, Ian Neild and Peter Cochrane*

# A Review of Optical Wireless

## What is it and what does it offer?

*This article presents a broad review of the field of optical wireless, from early experiments through to the high-performance systems of today. Emphasis is placed on understanding the benefits and limitations of optical wireless, all of which ultimately define the applications. Various systems are examined, each incorporating a different technological solution to suit the particular application.*

*Optical wireless technology has matured substantially in recent years and is now capable of delivering important benefits to modern-day telecommunications users.*

### Prologue

It is a new day and work beckons. While eating breakfast you catch up on news which is delivered to your laptop via an indoor infra-red wireless link. Several wireless systems exist throughout your home, ranging from short-distance point-and-shoot systems up to cellular systems that cover whole rooms. Your laptop automatically selects the link that offers the best performance at that time. At breakfast it has selected a short-distance link, the base station for which is built into the table (Figure 1). As you roam through your home taking your laptop with you, it remains linked to the network by seamlessly switching between different infra-red cells. All these cells are interconnected within your home and linked to the high-speed external network via a line-of-sight infra-red system to a communal 'pole' in the street (Figure 1).

As you enter your car it creates an in-car infra-red link with your laptop

and interrogates your diary for the day. Your car learns that your first journey is to your office in the city so it accesses traffic information via its GSM link and displays on the dashboard navigation system (or speaks) a recommended route for you to follow. On your way to work you pull into a petrol station. As you do so your car enters an infra-red cell and is instantly linked to a high-speed network. Your car uploads its performance data to an agency which will contact you later if your car needs attention. At the same time your laptop uses the in-car infra-red system to establish a connection between itself and the external cell. This done, it takes the opportunity to contact your office to update your diary and to up/download files from your desktop machine that it knows you will need. By the time you reach your office your laptop has exchanged gigabytes of data via a variety of wireless infra-red links and you are ready to face the challenges of the day.

Figure 1—Example optical wireless links

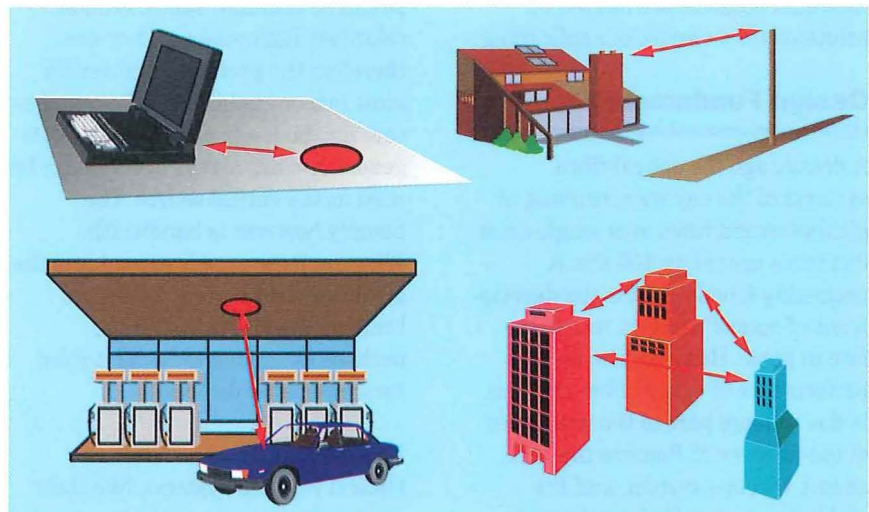


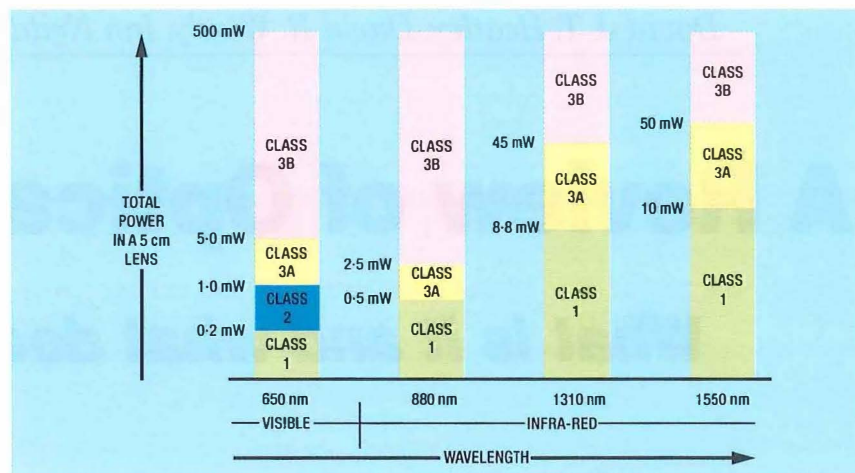
Figure 2—Laser safety classifications for a point-source emitter

During the day you make several personal calls on clients, all of whom have infra-red cells at their location, so you are on-line the instant you arrive. The buildings that you travel between are all located in the down-town business area and are interconnected via high-capacity line-of-sight infra-red systems that overlay the conventional cabled network (Figure 1). Consequently you are able to up/download vast amounts of data very quickly, thereby maximising your effectiveness while with your clients. As the day progresses your diary is continually updated whenever you enter a cell, and so on throughout the day.

Clearly this vision lies in the future; however, this article shows that much of the implicit optical wireless technology is in fact available today. In recent years much work has been done to understand fully the benefits and limitations afforded by optical wireless. This article calls upon work at BT Laboratories and elsewhere to present a broad overview of the field. The article begins by describing the underlying principles of optical wireless, paying particular attention to the intrinsic benefits and limitations. It then examines a broad selection of systems, some experimental, others now available commercially, that are tailored to a variety of operational requirements. Performances and deployment issues are compared and contrasted, and informed opinions are offered on solutions for a variety of applications.

## Design Fundamentals

A decade ago the optical-fibre systems of the day were running at gigabit/second rates over single-span distances exceeding 100 km. A promising foundation for the development of optical wireless was therefore in place. However, the high performance of optical-fibre systems is due in large part to the properties of the fibre itself. Remove the fibre, as in a wireless system, and the stable low-loss guided single-mode



propagation path is no longer available. Conveying light between the terminal stations in a controlled reliable manner then becomes a challenge. In addition to this, optical-fibre receivers, if applied to a free-space environment, incur a very high coupling loss which reduces the overall power budget of the system. For these reasons, and others, the transition from optical fibre to optical wireless is not straightforward and new design solutions are needed.

### Transmitter

The optical-fibre systems industry has spawned semiconductor lasers with broad bandwidths and high launch powers: features which are equally attractive to optical wireless systems, particularly for outdoor applications which call for high-capacity point-to-point links. However, for indoor applications, lasers pose a potential safety hazard because they are a point source emitter (see below on eye safety). Light emitting diodes (LEDs), on the other hand, are large-area emitters so can be operated safely, even at relatively high powers. They are therefore the preferred emitter for most indoor applications. To compensate for the lower powers that LEDs generally emit, arrays of them can be used as the optical source. The penalty however is bandwidth. Whereas laser speeds extend into the gigabit/second regime, LEDs are limited typically to 10 Mbit/s, perhaps extending to 50 Mbit/s for some specialist devices.

### Eye safety

Optical wireless systems, like their radio counterparts, can pose a hazard

if operated incorrectly; therefore, a laser safety standard<sup>1</sup> has been established in which optical sources are classified in accordance with their total emitted power. The principal classifications are summarised in Figure 2 for a 'point source' emitter such as a semiconductor laser.

Outdoor point-to-point systems generally use high-power lasers that operate in the Class 3B band to achieve a good power budget. The safety standard recommends that these systems should be located where the beam cannot be interrupted or viewed inadvertently by a person. Roof-top locations or tall masts are usual for this type of system.

Indoor systems pose a particular challenge because the safety standard recommends that they must be Class 1 eye safe under any conditions. Figure 3 shows that for systems employing laser sources, launch powers must not exceed 0.5 mW at the short wavelengths where most low-cost devices operate. Indoor systems that use lasers therefore find it difficult to achieve a good power budget. However, by using LEDs instead of lasers a much higher launch power can be used and still remain Class 1 eye safe. This is because LEDs are not 'point source' devices as are lasers, but are large-area devices, so if viewed the image on the retina of the eye covers a large area, and hence the power is diffused (Figure 3). Indeed, arrays of LEDs can produce substantial launch powers and yet be Class 1 eye safe. This and cost are the two primary reasons why indoor systems mostly

use LEDs as the emitter, either singly or in arrays.

Lasers operating inside the Class 3B band can, in fact, be rendered Class 1 eye safe by passing their beam through a hologram incorporated within the overall laser enclosure. The hologram breaks up the wavefront in the optical beam which causes the image of the laser spot on the retina of the eye to be diffused (Figure 3). The same hologram also shapes the beam to create a definable footprint; for example, square or rectangular. Using an in-house developed hologram<sup>2</sup> the authors and colleagues have demonstrated that a 40 mW Class 3B laser can be rendered Class 1 eye safe, and with certain refinements made to the hologram the same outcome could be achieved with a 100 mW laser. This technique, although proven in the laboratory, has yet to be taken up by the industry.

### Atmospheric loss

The power budget and overall performance of a point-to-point free-space link is strongly determined by atmospheric loss along the propagation path, which itself comprises free-space loss, clear air absorption, scattering, refraction and scintillation.

### Free-space loss

Free-space loss defines the proportion of optical power arriving at the receiver that is usefully captured within the receiver's aperture (Figure 4). A typical figure for a point-to-point system that operates with a slightly diverging beam would be 20 dB, whereas an indoor system that operates with a wide-angle beam could have a free space loss of 40 dB or more.

Clearly free-space loss is experienced by all forms of optical wireless systems, from short-distance indoor systems up to long-distance outdoor systems. However, the remaining sources of atmospheric loss described below are entirely the domain of the latter.

### Clear-air absorption

Clear-air absorption (the equivalent of absorption loss in optical fibres) is a wavelength-dependent process that gives rise to low-loss transmission windows centred on 850 nm, 1300 nm and 1550 nm (Figure 5)<sup>3</sup>. This process is essentially the same as that present in optical fibres, and hence the low-loss windows coincide with those in fibres. Consequently, the same opto-electronic devices can be used, which from a commercial point of view is crucial.

### Scattering and refraction

Scattering and refraction due to water droplets refers to the attenuating effect of rain, fog, mist and snow on the power reaching the receiver. This attenuation is continually in a state of flux so is most meaningfully quantified in terms of the percentage of time that it exceeds a given value (usually taken to be the point at which the error rate in the data is just acceptable). Over a period of a year, from December 1992 to December 1993, detailed measurements of

Figure 3 – The effect of different emitter types on the retina of the eye




Figure 4 – Schematic diagram of free space loss


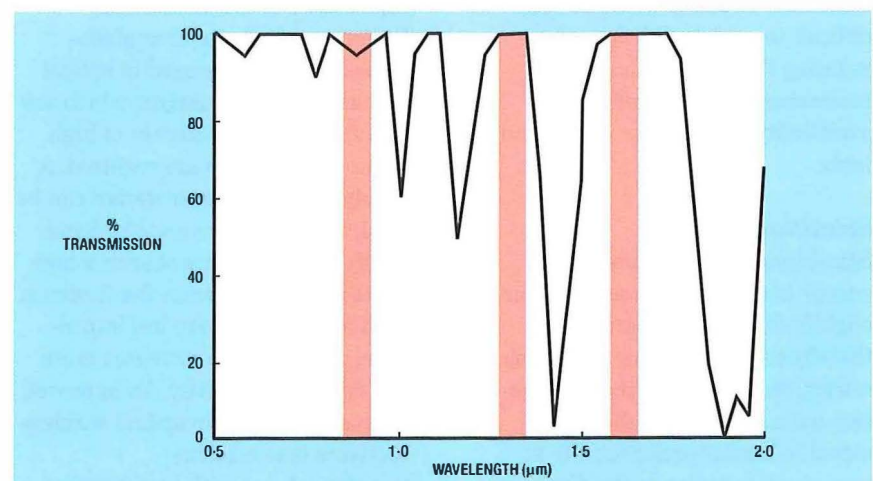


Figure 5 – Clear air absorption over a path length of 1000 feet at sea level



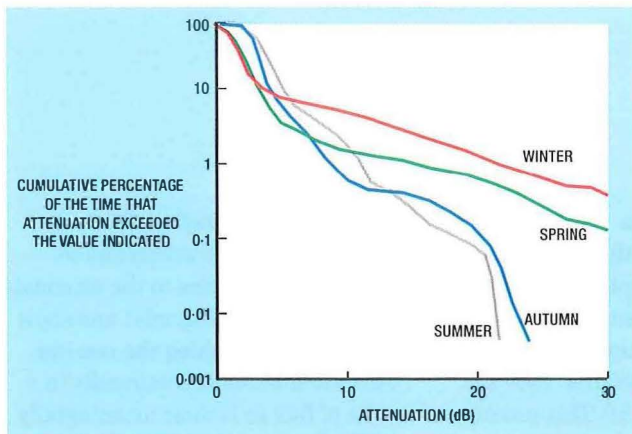


Figure 6—Seasonal atmospheric attenuation over a 100 m path for a 1 year period

atmospheric loss were carried out over a 100 m path in a rural area, close to the BT Laboratories site in Suffolk, England. Figure 6 shows some of the results obtained.

Not surprisingly the worst performance was achieved during the winter months owing to the high prevalence of rain, fog, mist and snow. Corresponding charts for an urban area show similar trends but with rather less variability between seasons and higher average attenuations.

A further study<sup>3</sup> was conducted using data provided by the UK Weather Centre, spanning a period of 9 years from January 1975 to December 1983. Figure 7 shows some of the results obtained for three major UK cities.

It is apparent that atmospheric attenuation is consistently low for all three cities for 99% of the time. Put another way, a free-space link spanning say 1 km in any of these cities would achieve an availability (that is, up time) of 99%. Figure 7 suggests that this figure could be improved to 99.5% given a suitable power budget etc.; however, 99.9% would be very difficult to achieve. Of course, by reducing the link length and/or increasing the power budget, the availability will increase correspondingly.

### Scintillation

Scintillation is the result of solar energy heating small pockets of air to slightly different temperatures, thereby creating regions of varying refractive index along the propagation path. This causes the optical signal to scatter preferentially at very shallow angles in the direction

of propagation, whereupon multiple signals, all phase-shifted relative to each other, arrive simultaneously at the receiver. This in turn causes the amplitude of the received signal to fluctuate rapidly by as much as 30 dB if conditions are unfavourable. The power spectral density of these fluctuations typically spans 0.01–200 Hz and hence can give rise to long bursts of data errors. Scintillation also distorts the wavefront of the received optical signal, causing the focused image at the photodiode to dance around the surface of the device. This requires a larger-than-normal photodiode to be used in order to ensure that the signal is never lost. Scintillation is generally not significant over distances less than 500 m but increases rapidly with distance.

### Receiver

Optical-fibre receivers mostly employ a transimpedance design because this affords a good compromise between bandwidth and noise, both of which are influenced by the capacitance of the photodiode. Because of the large-area photodiodes that must be used in optical wireless receivers, designs which are significantly more tolerant of high device capacitances are required. A purely transimpedance design can be used, but, to achieve a useful bandwidth in the presence of such a high photodiode capacitance, the feedback path must have a very low impedance, which in turn increases noise and reduces sensitivity. An improved approach common in optical wireless receivers is to combine transimpedance with bootstrapping,

the latter of which reduces the effective photodiode capacitance as perceived by electrical signals. This allows a relatively high feedback impedance to be used which reduces noise and increases sensitivity.

The circuit diagram in Figure 8 is typical of those used in optical wireless receivers. The feedback path (F) of the transimpedance portion is evident, as is the capacitive bootstrapping (B) at the input.

As in all optical receivers, fibre and wireless alike, their sensitivity is a trade off between photodiode parameters and circuit noise. For example, applications that require a good sensitivity will invariably use a small-area photodiode, which means that the aperture is small. Such receivers tend also to be high speed. Receivers for long-distance point-to-point systems generally fall into this category. Conversely, applications that require a large aperture must use a large area photodiode, whereupon sensitivity and speed are reduced. Indoor systems fall into this category.

BT Laboratories has developed a variety of receivers using PIN and

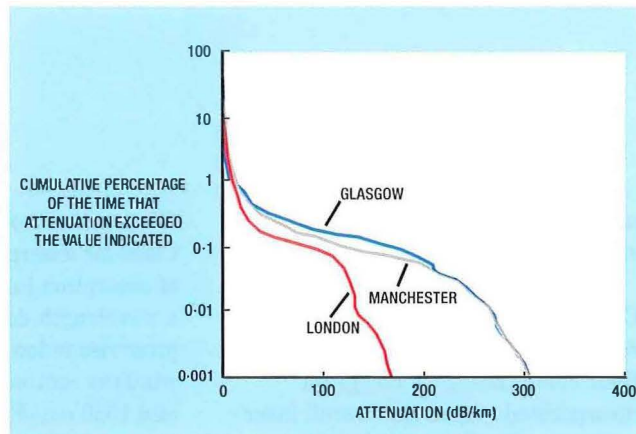



Figure 7—Atmospheric attenuation averaged over 9 years

Figure 8—Simplified circuit diagram of an optical wireless receiver





avalanche photodiodes (APD) of different dimensions. The circuits used are all based on the design in Figure 8. The sensitivities of some of these receivers are summarised in Figure 9, all for a 155 Mbit/s data rate. As predicted, the sensitivity improves (that is, reduces in numerical value) as the photodiode area reduces because of the correspondingly lower capacitance. However, small-area photodiodes incur a greater coupling loss due to the small aperture they present to the incoming beam, so a careful trade off between these factors is necessary to optimise the final performance.

Figure 9 also shows that a receiver with an APD gives a 10 dB sensitivity advantage over a corresponding PIN receiver, which is consistent with observations on optical-fibre receivers. APD receivers, however, are more costly and require high operating voltages, hence are predominantly used in specialist systems where performance is key. Systems in which economy is a priority, such as most indoor applications, favour PIN receivers.

The choice of photodiode, as well as influencing sensitivity, also impacts on operating safety. Most large-area PIN photodiodes are fabricated in silicon (Si) to keep cost low but they can only detect light at short wavelengths (up to about 900 nm). Systems employing such devices with laser transmitters achieve Class 1 eye safety by keeping the launched power below about 0.5 mW (Figure 2). Consequently, the

power budget is low. However a selection of large-area PIN photodiodes are now being fabricated in gallium-indium-arsenide (GaInAs) which enable receivers to operate in the long wavelength windows (1310 nm and 1550 nm). With these devices the associated laser transmitters can launch up to 10 mW and remain Class 1 eye safe (Figure 2). Indoor systems that exploit this can therefore deliver a good power margin together with a high capacity. Commercial large-area APDs in contrast are rather lagging behind. They are still predominantly fabricated in Si and hence are limited to short wavelengths and launch powers of about 0.5 mW for Class 1 operation. The greater sensitivity of APD receivers does of course compensate for the low launch power, but the PIN receiver is still the more elegant and cost-effective option for indoor applications. A limited selection of commercial large-area APDs in germanium (Ge) is available and hence can operate in the 1310 nm window where the Class 1 eye safe limit is raised to 45 mW (Figure 2). However, these devices exhibit very high leakage noise and can require operating voltages as high as 300 V to achieve an optimum performance, so once again the PIN solution is generally preferred for indoor applications.

### Interference from ambient light

Because of the wide field of view (that is, aperture) of a free-space optical receiver, stray light, in addition to the wanted optical beam, can reach the

photodiode. This is particularly the case with some indoor systems (see later on diffuse systems) in which the receiver can 'see' nearly everything above the plane of the photodiode. Ambient light raises the level of photonic noise in the receiver and hence can impair performance. Fortunately the properties of ambient light sources lend themselves to straightforward remedies.

First, by placing a narrowband infra-red filter in front of the photodiode the level of ambient light relative to the wanted beam is significantly reduced (Figure 10). Consequently, photonic noise in the receiver predominantly originates from the wanted signal, which is the optimum condition. Infra-red filters may be fabricated in glass or plastic, depending on the optical quality required and the application.

Secondly, the power spectral density of ambient induced noise extends from DC to typically a few tens of kilohertz, exceptionally a few hundreds of kilohertz, depending on the type of source. By conveying the data over the optical beam on a high-frequency subcarrier, or more commonly by applying a line code that contains no low-frequency components, interference is entirely avoided.

The strong DC content of ambient light also means that the dynamic range of the receiver could be impaired. By using the above infra-red filter and a receiver design which cancels or blocks any DC from the photodiode, this form of impairment is avoided.

Figure 9—Receiver sensitivity (at 155 Mbit/s) in relation to photodiode type and area

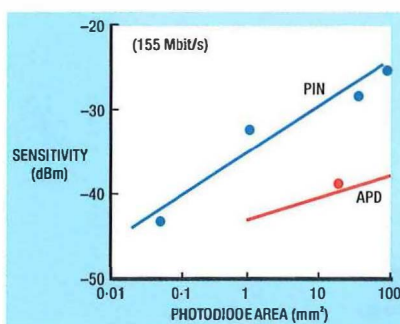


Figure 10—The use of an infra-red filter to reduce ambient light at the receiver





Figure 11—Experimental 4 km point-to-point system

## Long-Distance Systems (100 m–5 km)

Earlier in this article the strong influence that atmospheric loss has on spans greater than, say, 500 m is outlined. These systems must therefore have a good power budget to ensure error-free transmission, particularly so if high-speed operation is required (for example, 155 Mbit/s or higher). This can be achieved by addressing each of the three system elements: transmitter, propagation path, receiver.

A high-power Class 3B semiconductor laser can be used as the source, emitting perhaps 100 mW (+20 dBm). Alternatively, a lower-power device can be coupled to an erbium-doped fibre amplifier (EDFA) to achieve the same end result. BT Laboratories utilised the latter approach in a trial 155 Mbit/s link that spanned 4 km between Imperial College and University College in London<sup>4</sup>. The use of an EDFA dictated the use of a laser wavelength of 1550 nm, which enabled the system to be Class 1 eye safe while affording a useful launch power of 10 mW (Figure 2). The key elements of the system are depicted in Figure 11.

A good power budget requires that the overall propagation loss is minimised. Little can be done to alleviate the various atmospheric losses; however, steps can be taken to minimise free-space loss (Figure 4). This requires the use of appropriate optics at each end of the link to minimise beam divergence from the transmitter and maximise the aperture of the receiver. The BT Laboratories experimental system (Figure 11) achieved this by using an astronomical telescope at each end (Schmidt-Cassegrain, 20 cm aperture). The resulting beam diameter at the receiver was typically 2 m, which equates to a free-space loss of 20 dB. It was possible to reduce the beam diameter to only 0.5 m (8 dB free space loss) but then maintaining beam alignment became more difficult.



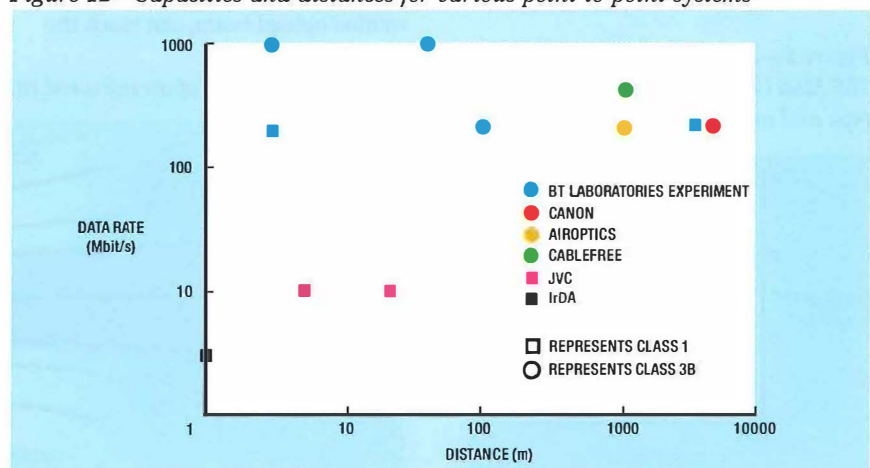
Beam alignment is a critical factor in long-distance point-to-point systems because it will drift with time owing to temperature changes and fatigue in the anchorage assemblies at each end. Some means of automatically maintaining alignment by mechanically steering the beams is desirable, particularly in commercial installations where maintenance intervention must be kept to a minimum or even avoided altogether. A long-distance system developed and marketed by Canon<sup>5</sup> embraces this requirement to the fullest extent. After crude alignment, each end of this system locks on to the other and maintains alignment even in the event of the transmitter/receiver units being jolted or moved off-axis by several degrees. Of course, such robustness comes at a premium.

The need for automatic beam alignment is taken to the extreme in orbital inter-satellite systems where path lengths of many thousands of kilometres are encountered. Such systems are being considered for high-capacity links between telecommunications satellites to re-route spare capacity and establish round-the-globe links without having to return to ground stations at intermediate points. Since atmospheric loss under these circumstances is zero, path loss is entirely limited by the optics employed.

Receiver sensitivity is the final factor in achieving a good power budget. Since the best sensitivity is usually delivered by an APD-fronted receiver (Figure 9), these are generally used in long-distance systems. The relatively high cost of these receivers is offset by the commercial advantage afforded by the greater reach. In any case, receiver cost tends to be a small factor in the cost equation for long-distance systems.

The data rate versus distance of the 4 km system and others developed by BT Laboratories are plotted in Figure 12, together with commercial systems from Canon<sup>5</sup>, AirOptics<sup>6</sup> and CableFree<sup>7</sup>. It is unlikely that long-distance point-to-point systems will afford reliable spans much beyond 4–5 km because of the extreme vagaries of atmospheric loss. Systems such as the Canon and BT Laboratories' ones are indeed capable of longer distances (for example, 8 km has been reported) but only when the weather, season and time of day combine to produce ideal conditions. This might equate to an availability of only 1–10% averaged over a year, whereas users require at least 99%. Indeed, some applications in telecommunications require 99.99% or even 99.999% (referred to as *five nines*). As Figures 3 and 4 show, obtaining that extra percentage point requires a significant improve-

Figure 12—Capacities and distances for various point-to-point systems






Figure 13—Optical telepoint

ment in the power-budget parameters, which in turn will be reflected in a higher cost. It is therefore believed that the largest market for long-distance point-to-point systems will involve sub-1 km spans, with the market for spans greater than 1 km largely confined to specialist applications. Of course, in countries such as the USA which enjoy a predominantly dry climate and contain large open expanses, the balance and size of these markets will be different.

### Short-Distance Systems (less than 100 m)

All of the design and performance factors described above for long-distance systems apply equally to short-distance systems but with one obvious exception—atmospheric loss has a much reduced effect, and none at all for indoor systems. The power budget is therefore determined almost entirely by the transmitter launch power, free-space loss, and receiver sensitivity.

### Point-to-point systems

Outdoor short-distance systems, used perhaps to deliver high-capacity links between neighbouring buildings, enjoy the luxury, as do long-distance systems, of utilising a high-power Class 3B emitter since these systems will be located away from people. A good power budget can therefore be achieved even with relatively simple components. This, combined with the low atmospheric loss, means that a higher free-space loss can be tolerated in return for the convenience of a larger beam diameter at the receiver. Setting up and aligning these systems is then significantly easier than their long-distance cousins, and, furthermore, automatic alignment and tracking is not required. Conse-




Figure 14—Optical telepoint locations

quently, overall complexity and cost is significantly reduced. Figure 12 includes the results of some BT Laboratories' experiments in this area. The 1 Gbit/s result over 40 m is particularly interesting as it shows the feasibility of using optical wireless as a high-capacity link between neighbouring buildings or as a temporary 'bridge' in the event of a cable failure<sup>9</sup>.

Indoor point-to-point systems do not differ from the outdoor variety in their operating principles; however, in practice their designs are invariably very different. Firstly, they must be Class 1 eye safe, which generally means that the optical source is an LED. This in turn limits the capacity to typically a few megabits/second. On the other hand, indoor systems need none of the weather proofing that outdoor systems require and only operate over short distances; hence they can be produced very cost effectively. A good example of such a system is manufactured by JVC<sup>9</sup>. Here a capacity of 10 Mbit/s can be conveyed over a distance of around 20 m (included in Figure 12). Such systems could be used to, for example, extend a 10 Mbit/s local area network (LAN) port to a different part of an office where no convenient port exists, or to link two separate offices via a corridor.

### Telepoint systems

Indoor applications, being fundamentally confined to short-distance spans, are uniquely appropriate to optical wireless systems that utilise wide diverging beams rather than the narrow beams of point-to-point systems. Such systems, sometimes referred to as *optical telepoint systems*

(Figure 13), afford several attractive features. Most importantly, each 'cell' created by an optical telepoint base station can be shared by as many users as can be sensibly accommodated within the cell.

The rationale behind telepoint is that users who wish to access a high-capacity link enter the cell. These cells are not ubiquitous as are those for mobile telephony, and neither do they overlap. Rather they are discrete entities and are located where users would normally require a high-capacity link. Once within the cell each user is free to roam within the cell, but not beyond. Some users may choose to enter the cell and then remain stationary; for example, as would be done in hot desking. In contrast, others may need to roam within the cell while maintaining a wireless link to the network; for example, hospital staff using some form of mobile terminal within a ward-sized cell. In either case, the users' terminal produces an up-beam that must always point at the base station. Achieving this in a roaming scenario is clearly a greater technical challenge than with stationary users.

BT Laboratories has experimented with several telepoint systems, creating cell diameters ranging from 10 m down to 0.5 m. The larger cells might be deployed in open offices and public areas such as foyers, libraries, waiting rooms, hospital wards, etc. The smaller cells (for example, 1 m or less) are particularly suited to applications in which a cell is dedicated to a single user; for example, a desktop or a telephone booth (Figure 14).

A good example of a commercial telepoint system, again manufactured by JVC<sup>9</sup>, generates a cell of


about 10 m diameter, which is sufficient space for up to say six simultaneous users. Each cell delivers a capacity of 10 Mbit/s (included in Figure 12) which is shared between the users via conventional LAN protocols. LEDs and PIN photodiodes are used throughout and line-of-sight paths are required between the base station and all the user terminals. As we will shortly discover, line-of-sight propagation is in fact vital if capacities more than a few megabit/second are to be reliably achieved.

The user stations in the JVC system are equipped with automatic acquisition and tracking which ensures that setting up and maintaining a link takes but a few moments. The mechanism by which this is performed is not suited to users who wish to roam within the cell. A later section describes alternative approaches to achieve that requirement.

### Diffuse systems

The converse to indoor systems that employ line-of-sight paths are those that employ diffuse paths (Figure 15). In this group of systems the beams radiate over a very wide angle and hence reflect off surfaces and objects in the vicinity such as walls, ceiling, floor, furniture, etc. In addition, the field of view of the optical receivers in the base and user stations is widened to such an extent that line-of-sight, as well as reflected, light can be detected with the same ease. Such an arrangement goes some way towards guaranteeing that a signal path is always present, regardless of obstacles or people

Figure 15—Diffuse optical wireless system



impeding the line-of-sight path. This arrangement also allows roaming to some degree. The penalty, however, of diffuse (that is, dispersive) transmission is a much reduced capacity compared with line-of-sight systems. This is entirely a consequence of the multiple signal paths reaching the receiver which cause classic pulse spreading and inter-symbol interference. Research has shown<sup>10</sup> that the theoretical capacity of a diffuse system is a function of many factors such as room size and geometry, the fabric and distribution of furnishings, the placement and orientation of the base and user stations, etc. Generalising, this work predicts an upper bound of about 25 Mbit/s in a room 10 m on a side, although higher rates have been demonstrated under particular conditions.

Interference from ambient light is a particular issue for diffuse systems because of the extremely wide field of view of the receivers. However, as outlined earlier, the use of infra-red filters and robust signal formats minimises any performance penalty that may arise.

A good example of a diffuse system is manufactured by Spectrix<sup>11</sup>. A single base station can comfortably irradiate a room 10 m on a side and deliver 4 Mbit/s, which of course is shared by the users within the cell. This system, like the JVC system described above, utilises LED emitters and PIN photodiodes to be Class 1 eye safe and to keep cost low.


### Roaming Architecture for Indoor Systems

In many indoor applications, it is desirable to deliver capacities far beyond the 4–10 Mbit/s described above for today's commercial systems. However, as mentioned earlier, to do so requires lasers to be used as the emitter because of their higher speed, and, to remain Class 1 eye safe, they must be operated at low power levels (Figure 2). In addition, propagation must be entirely line-of-sight. Given these constraints it

becomes necessary to reduce the diameter of the down-beam such that it illuminates only a single user station. Then, by designing the base station to produce several narrow down-beams simultaneously, with each affording a line-of-sight link to one of several user stations, a high capacity can be delivered to multiple users within the same cell. The beams would be re-used as users enter and leave the cell, the maximum number of users at any time being determined by the number of simultaneous down-beams that the base station can produce. In practice, this number would equate to the maximum number of users that can be sensibly accommodated within the cell. The beams could also be made steerable thus allowing a degree of roaming within the cell.

BT Laboratories developed an experimental system, operating at 155 Mbit/s, to investigate the feasibility of this concept<sup>12</sup>. The emitter in each transmitter was a laser array, and similarly the receivers used a PIN photodiode array (Figure 16). Device availability dictated that only a single base station and user station could be constructed, and the laser and photodiode arrays were one dimensional rather than the ideal two. Upon arrival of the user in the cell, pilot LEDs on the base and user stations enabled the down and up-beams to locate and lock on to the corresponding receivers. Thereafter alignment was maintained by monitoring the power detected by adjacent photodiodes in the array then realigning the beam as required by activating the appropriate laser in the array. At this point, a single laser in each

Figure 16—Concept behind the tracking system



array would be activated, the particular choice depending on the required launch angle of each beam. This process would be repeated for other users arriving in the cell, resulting in additional lasers in the base station array being activated. As a user begins to move within the cell, the detected beams would migrate from one PIN to the adjacent one in the array. This data, when fed into a tracking algorithm, would cause different lasers to be activated so that the beams remain directed at the corresponding receivers. This process would continue until the user was again stationary or had left the cell.

Of course the converse approach to a roaming architecture is to flood the floor area with stationary beams, with each cell wide enough to accommodate only a single high-capacity user (Figure 17). Roaming is then achieved by handing over from cell to cell, akin to cellular mobile telephony. BT Laboratories is investigating this technique for use in indoor locations such as finance centres and trading rooms where there are many mobile users, each requiring a very-high-capacity link.

These experiments have proven that very high capacities can be delivered to multiple users in an indoor Class 1 eye-safe environment. Commercial systems with these capabilities have yet to appear for reasons that include the immaturity and high cost of some of the devices, particularly the arrays for the tracking architecture.

Figure 17—High-capacity coverage using stationary overlapping cells



### The IrDA Standard for Very-Short-Distance Point-to-Point Systems (for example, 1 m or less)


An exciting development in recent years has been the arrival of very-short-distance 'point and shoot' optical wireless systems for laptop computers, PDAs, palmtops, printers, calculators, and, most recently, mobile telephones (Figure 18). Many of these products are already being manufactured with these systems installed as standard. These wireless systems are commonly referred to as *IrDA systems* in recognition of the IrDA standard that they embody.

The Infra-red Data Association (IrDA) was established in 1993 with the aim of producing an open standard for wireless data communication using mature commercially-available infra-red components. The resulting protocol provides a simple low-power low-cost reliable means of wireless infra-red communications for a broad range of computing, communications, and consumer devices. The IrDA now has some 160 member companies around the world, and IrDA systems are already resident on over 97% of new portable computers (this includes laptops, PDAs, palmtops, etc.) as shown in Figure 19.

The IrDA standard presently specifies the following parameters:

- data rates from 128 kbit/s up to 4 Mbit/s (raising this upper limit to several tens of megabits/second is being considered),
- maximum operating range of 1 m (increasing this to 10 m is being considered),

Figure 18—IrDA point-and-shoot concept



- 30 degree cone from the transmitter that equates to a 50 cm beam diameter at the receiver positioned 1 m away, and
- very wide receiver field of view.


All of these features ensure that a useful data rate is available over a wireless link that requires only the crudest of alignment.

IrDA systems have yet to have a noticeable impact. Many people who currently own laptops etc. are not even aware that they are equipped with an IrDA port, and of the few who do know this, very few of them actually use it. Of the many reasons why this might be so, the two prime candidates are the lack of marketing and the scarcity of locations which are equipped with IrDA base stations. A small number of businesses are beginning to install this technology in selected places of work so that their employees can 'hot desk', but these installations are no more than local trials. A more concerted and widespread push is needed if we are to benefit from the convenience afforded by IrDA.

### Comparison with Radio Systems

It is worth comparing the performance of the various short-distance (primarily indoor) optical wireless systems discussed in this article with their radio cousins because recent advances in the latter have narrowed the gap to the point that, in some instances, the optimum choice is not obvious.

Figure 19—Market penetration of IrDA systems in portable computers (source: Dataquest)



We have described how commercial indoor optical wireless systems can deliver 4 Mbit/s (diffuse) and 10 Mbit/s (line of sight) into a room of some 10 m on a side. The technology exists to improve on these rates and cell size, particularly so for line-of-sight systems, but standards and commercial systems that exploit this have yet to appear. Today's radio (or wireless) LAN systems can deliver 2 Mbit/s to multiple users over distances of typically 50–100 m depending on terrain, with the promise of 8–10 Mbit/s in the very near future. Standards are being produced for next-generation systems to deliver tens of megabits/second.

If not capacity, what therefore differentiates the two approaches?

Radio systems will almost always afford a greater reach and wider coverage than optical wireless systems because a higher transmitter power can be used and the receivers can take full advantage of sensitive heterodyning techniques. On the other hand, radio will always be a narrower bandwidth medium than optical, although this is not apparent in today's commercial systems because optical wireless manufacturers have yet to exploit the available bandwidth fully.

### Seamless Handover

It is clear that optical and radio systems are complementary in many respects. Optical wireless (telepoint in particular) can deliver a high capacity to multiple users within a small cell, whereas radio can produce a much larger cell, but, being shared by many more users, will deliver a lower capacity. It is therefore easy to envisage scenarios in which both types of system would benefit from being used in combination. For example, a large open plan office would deploy optical telepoint systems to establish several discrete high-capacity cells, each enclosing a small cluster of perhaps 4–6 desks, and additionally a single radio system would be deployed to cover the whole

office. Each desk might also be equipped with an IrDA base station.

Clearly, it would be desirable for the workers in such an office to access the network via whichever wireless link delivers the highest capacity to their laptops (or some other portable device) at that time. Furthermore, it would be equally desirable to have sufficient intelligence within these laptops to handover, automatically and seamlessly, from one wireless scheme to the other as the workers roam throughout the room, thus maintaining a link at all times. A typical sequence might then be: a worker at his/her desk accesses the network from his/her laptop via 10 Mbit/s telepoint or 4 Mbit/s IrDA; upon leaving the cell the laptop hands over to the office-wide radio cell which delivers a 40 kbit/s link (derived from 2 Mbit/s shared between 50 users); upon returning to a telepoint or IrDA cell, the laptop hands over and the original data rate is restored; and so on. The intelligence within the laptop, as well as managing the handovers, would also take account of the different capacities by tailoring the quality of service delivered to the user. For example, streamed television (vision and sound) at the desk via telepoint or IrDA might be reduced to sound only in the radio cell, then back to vision and sound upon return to a desk.

Of course, handover need not be confined to indoor wireless technologies. GSM can also be included to great effect. A typical sequence might then be: a worker at his/her desk accesses the network from his/her laptop via 10 Mbit/s telepoint or 4 Mbit/s IrDA; while roaming the office the laptop hands over to the radio cell which delivers a 40 kbit/s link; upon leaving the office, or even the building, the laptop hands over to the 9.6 kbit/s channel provided by the GSM network; and so on. Once again the quality of the delivered service would be managed automatically in accordance with the capacity of the link in use.

BT Laboratories is investigating several handover schemes that can support these scenarios. For example, handover between IrDA and GSM, and IrDA and DECT, have been successfully demonstrated. By building on this work, handover schemes that have the ubiquity envisaged above and in the prologue to this article will become practicable.

### Closing Remarks

This article has summarised the principal developments to date in the field of optical wireless, with particular emphasis on the capabilities and opportunities presented by today's commercial systems. It is clear that the field has reached a reasonable degree of maturity and the large-scale deployment of these systems is now feasible. However, for this to happen, the industry needs to take a proactive position in promoting these systems, and their radio cousins, if we are to realise the future that we aspire to. The onus is on us!

### References

- 1 IEC825:1993, Safety of Laser Products. Equipment Classification, Requirements and Users Guide.
- 2 EARDLEY, P.; WISELY, D.; WOOD, D.; and MCKEE, P. Holograms for Optical Wireless LAN's. *IEE Proc. on Optoelectronics*, Special Issue on Free Space Optical Communications, Dec '95.
- 3 SMYTH, P.; EARDLEY, P.; DALTON, K.; WISELY, D.; MCKEE, P.; and WOOD, D. Optical Wireless—A Prognosis. *SPIE '95*.
- 4 WISELY, D.; MCCULLAGH, M.; EARDLEY, P.; SMYTH, P.; LUTHRA, D.; MIRANDA, E.; and COLE, R. 4 km Terrestrial Line of Sight Optical Free Space Link Operating at 155 Mbit/s. *Proc. Free Space Laser Communication Technologies VI*, SPIE Vol. 2123, 1994.

- 5 CanoBeam (<http://www.usa.canon.com>)
- 6 AirOptics (<http://www.jolt.co.il>)
- 7 CableFree (<http://www.cablefree.co.uk>).
- 8 Eardley, P.; and Wisely, D. 1 Gbit/s Optical Free Space Link Operating over 40 m—System and Applications. *IEE Proc. on Optoelectronics*, Special Issue on Free Space Optical Communications, Dec. 1995.
- 9 JVC VIPSLAN-10 (<http://www.jvcinfo.com>).
- 10 GFELLER, F. R.; and BAPST, U. Wireless In-House Data Communication via Diffuse Infra-red Radiation. *Proc. IEEE*, **67**(11), Nov. 1979.
- 11 SpectrixLite (<http://www.spectrixcorp.com>)
- 12 WISELY, D. A 1 Gbit/s Optical Wireless Tracked Architecture for ATM Delivery. IEE Colloquium, London, Feb '96.

## Biographies



**David Heatley**  
Networks and  
Information Services,  
BT UK

David Heatley obtained B.Sc. and M.Sc. degrees in 1978 and 1981, respectively, and a Ph.D. in Optical Communications Systems in 1989. He joined BT Laboratories in 1978 to work on the development of analogue and digital optical-fibre systems for video and broadband services. In 1985, he was appointed to head a group responsible for the development of optical receivers for terrestrial and undersea fibre systems. In this capacity he was a

member of the team that received the Queen's Award for Technology in 1990. He is presently with the Mobility Research Unit and heads a group with special responsibility for mobile telecommunication systems and future technologies. During his career he has published widely on telecommunications and has received several 'best paper' awards. Aspects of this work now feature in a recently released book he co-edited with Peter Cochrane on future telecommunications systems. He is a Member of the IEE and a Chartered Engineer. His business web site can be found at <http://www.labs.bt.com/people/heatledj/index.htm>.

David Heatley can be contacted at [dave.heatley@bt.com](mailto:dave.heatley@bt.com).

**David Wisely**  
Networks and Information Services,  
BT UK

David Wisely graduated in 1982 and worked at the Atomic Weapons Research Establishment in Berkshire on laser damage and weapon hardening. He then spent a year at Reading University studying optics and joined BT Laboratories in 1988, specialising in micro-optic fibre devices. In 1992, he was appointed project manager for optical wireless research and headed up teams developing innovative concepts in receiver design, holographic diffusers and tracked architectures. His current research interests include cellular IP protocol design and mobile network performance modelling. He has published many papers on his work, one of which received a distinguished paper prize.



**Ian Neild**  
Networks and  
Information Services,  
BT UK

Ian Neild graduated from Hull University and joined BT in 1993. He initially worked on tracking systems for gigabit/second infra-red wireless links and later progressed to radio networks. In 1996, he moved to work on WLAN and IrDA data solutions for use in payphones and as a precursor to UMTS. He is currently on secondment to the Head of Research at BT Laboratories.



**Peter Cochrane**  
Networks and  
Information Services,  
BT UK

Peter Cochrane joined BT Laboratories in 1973 and has worked on a wide range of technologies and systems. In 1993, he was appointed as the Head of Research. A graduate of Trent Polytechnic (now Nottingham Trent University) and Essex University he is also a visiting professor to UCL, Essex, and Kent Universities and a Member of the New York Academy of Sciences. He has published and lectured widely on technology and the implications of IT. He led a team that received the Queen's Award for Innovation and Export in 1990; the Martlesham Medal for contributions to fibre optic technology in 1994; the IEE Electronics Division Premium in 1986; Computing and Control Premium in 1994 and the IERE Benefactors Prize in 1994; the James Clerk Maxwell Memorial Medal in 1995; IBTE Best Paper Prize and Honorary Doctorates from Essex, and Stafford Universities in 1996. His business Web site is at <http://www.labs.bt.com/people/cochrapp/index.htm>.

# Professional Mobile Radio— Service, Market Position and Technology

*Professional mobile radio (PMR) is widely used in industry where people need to move and communicate in closed groups as part of their work. This article discusses how PMR service should be thought of as part of a business process, the system functionality required to make this happen, the supporting technology and market considerations which will be a factor in PMR growth. The outlook for convergence of third-generation mobility to be provided by Universal Mobile Telecommunications System (UMTS) and PMR technology is also examined, concluding that convergence is unclear at present, but there is similarity in requirements.*

## Introduction

Professional mobile radio (PMR) services are widely employed by a range of corporate users, such as the police and fire services, the utilities (gas, water and electricity), in transportation and on specific sites such as airports and industrial premises (Figure 1).

A multitude of analogue and newer digital PMR technologies exist, which are all designed to support closed-group communications that are an integrated part of a business process. Closed-group communications is a key mode of PMR operation, whether this is for a 'public safety' organisation such as the police, a fleet of delivery vans, site security or a taxi company. A dispatcher usually controls closed user groups. This is a person or system typically based at a central location, who initiates and controls the actions of the 'group'. The dispatcher can

control several groups at one time, and is able to join or leave groups at will. The dispatch terminal can be very simple, as say in a taxi office operation, or be a very complex IT-based system, as would be found in many police control rooms.

Public cellular radio systems are currently designed to provide point-to-point telephony and data communication. This functionality does not provide an ideal solution to the needs for group working and other PMR functionality within businesses (Figure 2). The range of PMR functionality is discussed in more detail later, but performance aspects such as fast call set up are also of key importance.

The nature of call traffic on a PMR system is fundamentally different from that found on a public cellular system. Apart from group communications, the length of calls is usually very short compared to those on a public telephone system, with

Figure 1—PMR is widely employed by a wide range of corporate users






Figure 2—PMR service is aimed at group working

restricted scope of content. For voice communications, these are typically 4–5 seconds in length, compared with 90 seconds on public cellular systems, and usually in half-duplex mode. Most communications are made to and from a dispatcher at present (Figure 3).

The use of status data codes such as ON DUTY, can be sent in one byte of data, and is a very efficient way of communicating routine information within a business. The use of packet data communications on modern digital PMR systems enables efficient database enquiries to take place, whether this be a check of stocked components, a trace of a suspect vehicle, a building plan or maintenance instruction to a service engineer. It is in fact possible to meet some users needs solely through the use of data communications, only reverting to voice when an out of the ordinary situation occurs!


The primary business aim of any corporate user is to improve productivity by utilising a service which enables faster, simpler and more cost-effective communications. This involves the integration of a PMR system with the organisation and its management processes. These management processes have in the past been paper-based, but nowadays more commonly use information technology (IT), where the use of data communications are of paramount importance to operational needs (Figure 4).

**PMR Service functionality**

The functionality of modern digital PMR technologies has been defined through the requirements identified in business operations and the new possibility for more efficient operation through more advanced technology.

Although not all PMR users will want the full service set provided by modern PMR technology, it is useful to see what is offered and how the functionality can be used.

The following is a very abridged version of the terrestrial trunked radio



(TETRA) service set and demonstrates key PMR features and use.

**TETRA system operation modes**

**Voice + data (V+D) and packet data optimised (PDO) modes**

The primary operation mode for TETRA is currently the so-called voice + data mode, in which the radio interface uses separate control and

traffic channels to route data to terminals. Another TETRA option under development in the standards, known as the packet data optimised (PDO), is a data-orientated solution that can be used where speech is not necessary or appropriate. For example, many taxi companies these days do not use voice since they can make better use of automated dispatch systems that talk in data messages not speech. Use of data can

Figure 3—Cellular telephone and PMR traffic 'profiles'

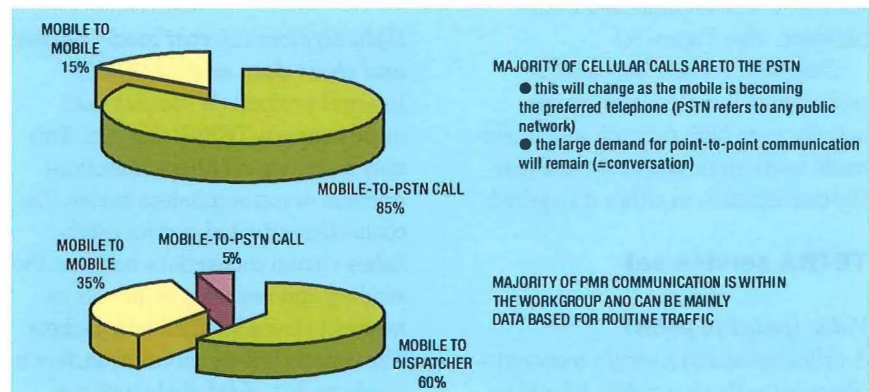


Figure 4—Professional use of wireless data




Figure 5—Direct mode communications

also be highly efficient and utilise less radio resource, therefore being more cost-effective for the user.

### Direct-mode operation (DMO)

Direct mode operation allows communications directly between handsets or from handset to mobile-mounted equipment, without the use of network infrastructure. This is very different from public cellular systems where all calls have to be set up via the network.

Not all users will need the direct mode facilities, but it is an essential requirement for public safety users, where communication must still be possible when access to the network is unavailable. It may also be used for radio coverage extension from the network; for example, when reliable communications to the network cannot be achieved into a building, a link can be set up via a direct-mode gateway mounted in a nearby vehicle, which is connected to the network. (See Figure 5.)

*Dual-watch* mode may also be used, where a user can listen for activity from both network and direct-mode users simultaneously, and join the conversation in either if required.


### TETRA service set

#### Voice (point to point)

A caller speaks to a single recipient—like most calls on a public telephone system.

#### Voice (group calls—many modes; acknowledged or unacknowledged)

Several terminals are linked together in simplex mode (only one talker at a time, while the rest of the group listens). The leader of the group is often a dispatcher; but not necessarily so. Group members are aware of the terminals connected if in acknowledged mode, which is normally indicated on terminal display. Groups are usually pre-defined and selected by operation of a simple terminal control. However the group can be defined dynamically and members can be added to the group while the call is in progress.



### Broadcast calls

A dispatcher or mobile user may broadcast a call to all users in a selected geographic area. This is a very useful work management facility; for example, a plumbing company may want to establish quickly who is available to attend an emergency call for repairing a burst pipe!

### Data services (circuit mode, packet and short data service (SDS))

Internet protocol (IP) packet data mode supports TCP/IP sessions. This may be set up in either connection-oriented or connectionless modes. The connection-oriented service establishes virtual connections between the sending and receiving terminals as needed to transfer X.25 data packets. The connectionless service transfers a single packet of data only without establishing the virtual connection. The functionality it offers is very similar to that of the IP.

Circuit mode is the establishment of an end-to-end bearer. This end-to-end circuit can be used unprotected or at various levels of protection. In unprotected mode, no error correction algorithm is applied and thus yields the highest channel data rate, while a high level of protection gives lower probability of errors, but lowers the overall useful data rate. The TETRA standard allows the concatenation of traffic time-slots on a single radio carrier to give a maximum data rate of 28.8 kbit/s. (TETRA voice + data system set up). Users, may of course apply their own end-to-end error-correction

protocols, but these would obviously work without detailed information of the radio-channel performance, which is internal to the TETRA system.

Circuit mode data is used in point-to-point communications, but development of an IP is underway so that groups may be addressed simultaneously. This is known as *user data protocol* (UDP).

SDS is a message service optimised for the exchange of short, user-defined or pre-defined (status) messages. These messages may be sent simultaneously with an ongoing speech call. User messages may be of three fixed lengths (16, 32, or 64 bits) or variable in length up to 2047 bits.

### Supplementary services


Several supplementary services modify or supplement the basic services. Examples of these are call prioritisation and pre-emption (where traffic grade of service may be controlled through the prioritisation, but allow emergency calls immediate access through pre-emption), talking party identification and dynamic group number assignment. These are in addition to telephone-like services such as call diversion, call waiting, call completion on busy, etc.

### System interfaces

#### Dispatcher interface

Most traffic usually passes through the dispatcher in traditional PMR systems. This was necessary because old technology only allowed calls to

Figure 6—Growth of PMR and PAMR users



be set up by a central dispatcher (like pre-automatic telephone operation). This mode was acceptable to users with a vertically structured organisation, where commands came from one source. Modern systems allow a different organisational approach, where routine commands and enquiries no longer need include the dispatcher, except for supervisory, special commands or more unusual enquiries.

The dispatcher interface does however remain a key and important part of a modern PMR system and must be able to support many simultaneous calls with a very high degree of reliability.

#### Line station interface

This allows the connection of a wired mobile, suitable for situations where radio coverage is not required; for example, a desk-based user.

#### Inter-working to PSTN, ISDN, PBX, IP networks

Public switched telephone network (PSTN), integrated services digital network (ISDN), PBX and IP network gateways allow users to make and receive calls from other networks. The provision of duplex operations means that calls to the PSTN and PBXs are just like present public cellular calls. PMR users are allocated their own national number identity, although incoming calls may be intercepted at an organisation's switchboard, in the same way that business PBX operation is often set up.

It is expected that mobility features—features which present PMR users with a virtual home environment, wherever they are and whatever terminal—will be developed eventually. There will however be restrictions on this for public safety users, for reasons of security.

#### System-wide roaming

The TETRA standard allows users to enjoy system-wide roaming (connection anywhere on the network) as opposed to limited area or regional services provided at present. How-

ever, an operator may choose to provide selective geographic service based on a scale of different call tariffs and classes of service.

Roaming on other TETRA networks is possible through an intersystem interface (ISI) agreed between operators. It is envisaged that European-wide service, using TETRA networks run by different operators, will be available in the future.

#### Data interface and application interface (API) on mobile

A comprehensive interface is provided on terminals for running independent data applications with call control. Many applications may be cited, but users typically want to have automatic vehicle location and central database enquiry facilities.

#### Charging records

TETRA manufacturers are making appropriate usage statistics available so that operators can deliver the usage and billing information in a manner that meets customer needs.

### European PMR Market Considerations

There is a multitude of disparate PMR systems being operated within Europe and worldwide. Only a few of these give nationwide coverage, while the majority are regional, town and site-wide systems. Many of these use old analogue technology and, despite limited functionality, are still providing good and cheap service. Therefore migration to modern nationwide digital PMR networks may not be rapid.

The PMR market is growing, albeit more slowly than it has for GSM public cellular service. The total number of users within Europe is currently around 7.5 million and this is expected to increase to approximately 11 million users by 2002<sup>1</sup>. These numbers may be split into PMR and public access mobile radio (PAMR) networks (Figure 6), where PAMR provides service to several users on the same network infrastructure. Although these user numbers

are small compared to public cellular, the traffic level and range of services used can be relatively high.

Much also depends on the availability of suitable radio spectrum, spectrum pricing and demonstration of cost-effective and reliable digital PMR network technologies. These are early days for digital PMR technology and additional spectrum is being made available only slowly by European countries for commercial PMR operation, although a fairly harmonised allocation spectrum for public safety use has already been achieved. The situation is better than it was for the introduction of GSM public cellular systems, as current digital PMR systems are narrow-band and therefore blocks of frequencies are not essential. However, fixed duplex spacing for up and down links means that the allocations must be paired.

Many potential PMR users currently use public cellular services which may also add to these projections. Nextel, who sell PMR service to more than 2 million users in the USA, provide an interconnect service, allowing public cellular like connection. The proportion of calls which may be connected in this way may be regulated in some countries, but does overcome the barrier of having separate PMR and cellular telephones and thus is an attractive package for users. There is, however, a threat from cellular operators who will be able to provide some key PMR functionality on cellular systems (see section on technologies below).

The European public safety PMR market position is fairly clear, with the users insisting upon essential services which can be provided only by thoroughbred PMR technology such as TETRA. The growth of commercial PMR service is much dependent upon the economy of scale (volume take up), the economy of service integration, the economy of business application and

capital efficiency. Several operators and manufacturers are keen to make this a successful business, and some keen competition is arising, so growth can be expected.

### Mainstream Modern PMR Technologies, Networks and Developments

Many analogue PMR technologies are currently used to serve the needs of PMR users within Europe. Perhaps the most successful of the more modern technologies is the so-called *MPT 1327* trunked system, which has many modern call features and makes efficient use of radio spectrum. The use of analogue technology has an inherent drawback as communications tend to become increasingly data oriented. These systems also lack communication security across the air interface, with many instances of the criminal fraternity and press intercepting public safety communications. The use of digital technology is helpful in overcoming these problems. This section therefore focuses on the more modern digital cellular technologies.

#### TETRA and TETRAPOL

The European community, through the European Telecommunications Standards Institute (ETSI) has always tried to agree common standards for interfacing communications systems guided by the views of operators, manufacturers and users. The agreement of GSM cellular radio standards has been an outstanding success, ensuring volume production across the European community and overseas. The PMR fraternity has tried to emulate the same success with the formation of the TETRA standard, which has open system interfaces and the comprehensive service set described earlier. One other digital PMR system has gained some ground within Europe—Matra's proprietary *TETRAPOL* system. This system has many good features, but is unpopular with many operators and other manufacturers because it is a propri-

etary solution. This may cause a divide within Europe and limit the success of TETRA systems both commercially and in ubiquitous service availability; for example, intersystem roaming.

Currently, 15 TETRA national and regional networks are under trial or roll-out, and this figure is expected to rise. TETRA does, therefore, seem set for success, but it will probably never be as cheap as current cellular service, as it will not have the same production volume. It is, however the only thoroughbred technology which will fully meet the needs of many professional mobile radio users. It is therefore difficult to compare the costs of public cellular service with PMR, as they are different services.

#### GSM PMR functionality

While the GSM system was designed for public telephone and data use from the outset, some features have been added which give it some key PMR functionality. These are known as the GSM *advanced speech call items* (ASCI).

The GSM ASCI features have been used in the GSM design for railways known as GSM-R for which an agreed frequency allocation at 900 MHz has been reached within Europe.

GSM PMR functionality falls into three areas:

- broadcast calls,
- group calls, and
- priority service.

This is clearly a limited service set compared to that of TETRA discussed earlier. It should also be noted that GSM cannot provide direct mode communications; however, not all PMR users will require this facility and the additional GSM functionality may be sufficient to meet many users needs.

There are also some fundamental differences in the relative performance of TETRA and GSM systems. A key performance metric for a PMR system is call set-up time. This is typically less than 300 ms for a

TETRA system, and around 1 second for GSM working in its PMR mode (3–4 seconds for normal telephony call set up). A call set-up time of around 1 second is probably not acceptable for a mission-critical operation, but may be acceptable to other PMR users.

The network architecture design for public telephony is designed for optimum performance against cost, based on an analysis of call destination and length. This is unlikely to be the same for PMR operation, so some compromise would have to be made if existing networks were augmented with PMR service.

GSM is therefore not a prime PMR technology, but does have the attraction of a volume market and lower entry costs for many users.

#### American PMR technology<sup>2</sup>

Several American digital PMR technologies are used throughout the world. Although an attempt has been made to produce and promote a common standard (APCO 25), in the same way as GSM and TETRA in Europe, this has not yet been successful and several popular proprietary solutions currently serve the market. The allocation of frequency also has a large bearing on the technology used and tends to preclude systems needing contiguous blocks of spectrum.

The more popular technologies are listed below:

- *Motorola iDEN* Widely used by Nextel for PMR and telephone interconnect service with upwards of one million customers. Coverage for 60% of the US population and a 45% market share. The next biggest operator has only 5% market share! Limited open interfaces.
- *Geotek FHMA* A modern technology using frequency-hopping multiple access (FHMA) to improve coverage at low speed and control interference from other users. Large number of regional licences, but technology not widely rolled-out yet.

Figure 7—Varying system resource by area

- **Ericsson EDACS** Enhanced digital access communications (EDACS) is a fairly old technology, but still going strong with new systems being installed in Eastern Europe etc.
- **EF Johnson 'D-Net'** A frequency-division multiple access (FDMA) with 12.5 kHz channel spacing.

### Future Digital Trunked PMR Technologies


#### UMTS

The Universal Mobile Telecommunications System (UMTS) is seen to provide full integration of fixed and mobile environments providing multi-service capability in all environments while users are unaware of any changes, perceiving they are still attached to their own home network (that is, the support of virtual home environment (VHE)). It provides dynamic bit-rate allocation and global roaming, and supports both personal and terminal mobility to obtain the 'Martini' concept of anyone, anywhere and in any medium! (See Figure 7.)

The UMTS system has, to date, been designed with public communications in mind and has an evolutionary path largely based on GSM developments. The system does, however, have features which have a broad similarity to those required for PMR service<sup>3</sup>.

One of these requirements within UMTS is the support of the VHE feature, whereby the visited network has a complete set of subscriber preferences, service subscriptions and restrictions, giving the impression to the user that they are still attached to their home network. For TETRA this may mean the identification of closed users groups, specific dialing codes, specific bearer resource allocation requests and priority call set ups etc.

Possibly the strongest of these is that the user, when making a request to the network for resource (that is, to make a call), indicates a set of service-related parameters and a destination.



The network uses these parameters to put the necessary resource in place to meet the customer's request, the parameters are known as the *quality-of-service information*. In the selection of network infrastructure and techniques many experts recommend packetised communication solutions since these can be configured to offer bandwidth on demand, user prioritisation and resource sharing.

Recent developments of Internet protocol are bringing group and real-time services to the marketplace. These service types are the basis of group services that will become normal practices in corporate local area networks (LANS). The key message is that, as networks and nodes develop, the access will become packet based and the communications solutions will be derived from the information technology (IT) domain, not the traditional telecommunications. These services are going to be developed much further.

#### TETRA technology developments

While the UMTS development is moving forward rapidly, significant third-generation developments are also progressing within the PMR world. TETRA developments of broadband service are worthy of particular note and this is discussed below.

Future implementations of the TETRA standard include the PDO proposal which is a different radio interface and traffic system utilising the existing mobility management and network database functionality. This provides a very useful model on an evolutionary path from TETRA voice + data and TETRA PDO towards a next-generation system.

As with GSM general packet radio service (GPRS), the PDO developments will bring an alternative network bearer to the marketplace. This can be seen in Figure 8 and linking into a packet network. For the PDO system to be effective there needs to be the appropriate dispatcher system and data network.

The Digital Advanced Wireless Services (DAWS) development<sup>4</sup> is showing that it is possible to implement a system as desired by PDO. In fact, the DAWS developments are based upon the PDO architecture, and the technology is being considered by the ETSI TETRA data working party for adoption as the technology solution.

Development of PDO is clearly on the same evolutionary path as GPRS, and it is always possible that the market can support only one system (since this makes the most of the volume of mobile terminal manufacturing) where the price is right. However,

Figure 8—Parallel networks for evolution

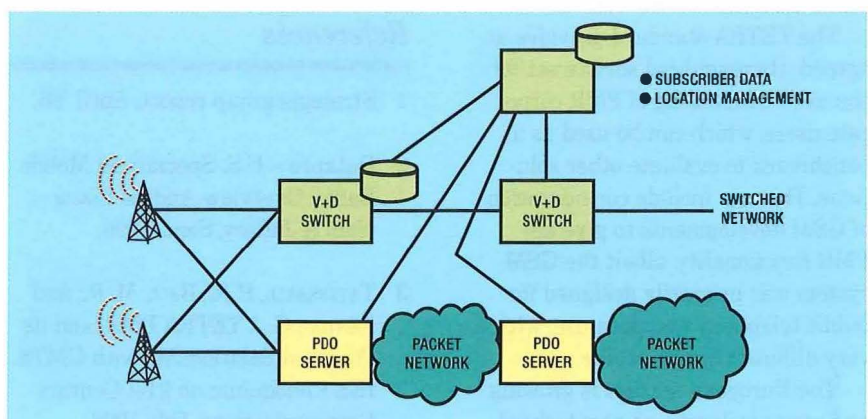
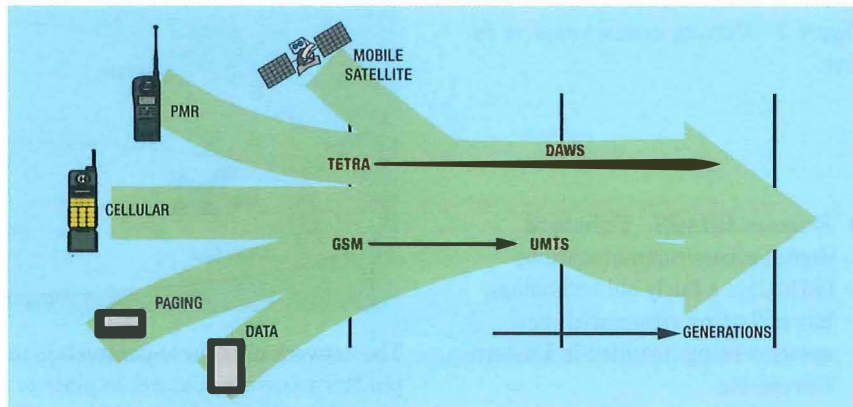


Figure 9—The way ahead



this is very unlikely to be the case since it implies the merging (or common access) to the subscriber database between the two systems. Whatever the evolutionary path the user requirements of both the point to point (for example, existing telephony) and group (for example, existing PMR) services can be met by packetised solutions.

A generalised view of TETRA, GSM, UMTS and other technology evolution paths is shown in Figure 9; whether these three technologies will actually converge to a common solution (UMTS perhaps) in the future will no doubt be subject to debate in coming years.

**Conclusions**

PMR service is widely used by public safety (police, fire, ambulance, etc.), utilities (gas, water and electricity), service industries, transportation, agriculture, as well as for specific site applications such as airports and site security. Typical applications are workforce management, surveillance, telemetry, database enquiry and updating.

PMR technology is designed to provide key business benefits to organisations when used as part of their business process. This can be thought of as ‘integrating voice, data, organisations as people’.

A multitude of existing analogue PMR technologies will continue to provide service for some years to come; however newer European and American digital PMR developments provide a much enhanced service set, with a leaning towards data-based communication.

The TETRA standard provides an agreed thoroughbred service set for the most demanding of PMR corporate users, which can be used as a benchmark to evaluate other solutions. This can include consideration of GSM developments to give key PMR functionality, albeit the GSM system was primarily designed for public telephony and data use, with a very different traffic profile.

The European market is growing and may accelerate, as new technol-

ogy such as TETRA becomes trusted and cost-effective solutions are implemented; however, the allocation of suitable frequencies and regulation in the use of public interconnect may restrict the potential growth. Since there is a substantial and active body of manufacturers and operators who are keen to make this a successful business, growth can be expected.

It is likely that PMR service will be provided by a few PAMR operators rather than the present situation where separate PMR systems are often set up for each organisation requiring service. This means that several groups of users will share a common network infrastructure, but work independently from each other in closed user groups, perhaps buying a different set and level of service appropriate to their needs.

It is not clear whether third-generation PMR systems will converge with public cellular systems and whether this will be within UMTS. Developments of the TETRA standard towards broadband mobile communications are well underway, and the UMTS concept of a VHE, which is similar to closed user group working in PMR, could be brought together. This will undoubtedly be a hot topic for debate in coming years.

**References**

- 1 Strategis group report, April '96.
- 2 Datapro – U.S. Specialised Mobile Radio: Overview Andres Llana Naqi A Jaffery, Sept. 1996.
- 3 TATTERSALL, P. R.; BELL, M. R.; and FENTON, C. J. TETRA PMR, and its Potential Relationship with UMTS. IEE Colloquium on 21st Century Communications, Feb. 1998.

- 4 RING, STEFFEN. DAWS (Digital Advanced Wireless Services)—A Progression of the TETRA Packet Switching Standard into Broadband Terrestrial Mobility. TETRA World Congress, Nov. 1998.

**Acknowledgements**

The author acknowledges the support and help of many of his work colleagues in the production of this article. Particular thanks go to Chris Fenton, Andy Cushing and Nigel Lobley who were able to give very useful perspectives on TETRA and UMTS service.

**Biography**



**Paul Tattersall**  
Networks and Information Services, BT UK

Paul Tattersall has worked on radio communication projects for the past 16 years, mainly on the implementation of GSM cellular systems, particularly with respect to radio performance. A keen interest in 1996 to contribute to pioneering work on designing a PMR system for UK public safety use brought about a new ‘tack’ in his career. This article culminates from further interest to understand better the PMR market and other technology, regulatory considerations and the rapid development of third-generation communication systems, in which he believes that PMR service will be a key and important part. Paul is a Chartered Engineer and a Member of the Institute of Electrical Engineers.

Trish Jones

# Developing Customer Service Skills

*Increasingly, customers' experience of how companies deliver services is becoming the differentiator that sets companies apart from their competitors. In addition, more and more people within companies, from all kinds of facets, are dealing directly with customers. In recognition of the important role customer service people play, a new institute—The Institute of Customer Service—has been set up to provide independent recognition and give value to the importance of what people do in the customer-service arena.*

## Customer Service—A Key Differentiator

Recognition of the fundamental importance of customer service and the vital roles and contribution of everyone working in this field has prompted 50 leading UK companies, including BT, to join forces as a Founder Council for a new institute—The Institute of Customer Service (ICS).

The ICS, launched in September 1998, is a new organisation whose principal purpose is to develop people and systems to raise service standards.

The ICS has two key objectives:

- to provide service leadership with a focus on commissioning research and promoting leading-edge thinking; and
- to provide people working in customer service with the means, through individual membership, of obtaining formal professional qualifications and recognition for their work.

It is becoming acknowledged universally that customers' experience of the service delivery—not the product—is increasingly becoming the differentiator that sets successful companies apart from their competitors. Customer perception on service has changed, with comparisons and judgements being made, not just within the same industry, but from all providers through day-to-day personal experiences; with banks, supermarkets, travel and transport, utilities, entertainment, etc.

Expectations are becoming greater and people in the 'front line' are increasingly needing to respond to the heightened expectations of their

*'a new era for our people if BT is to make its massive engineering and technology capability count'*

Ian Smith, Managing Director,  
BT UK Customer Service

customers. People who were previously in the back office and in technical functions find themselves through reorganisation and process changes to be the public face of their organisation. More often it is the attitude, behaviour, flexibility and tenacity of these people to manage their customer relationships effectively that will determine whether their companies will succeed or fail.

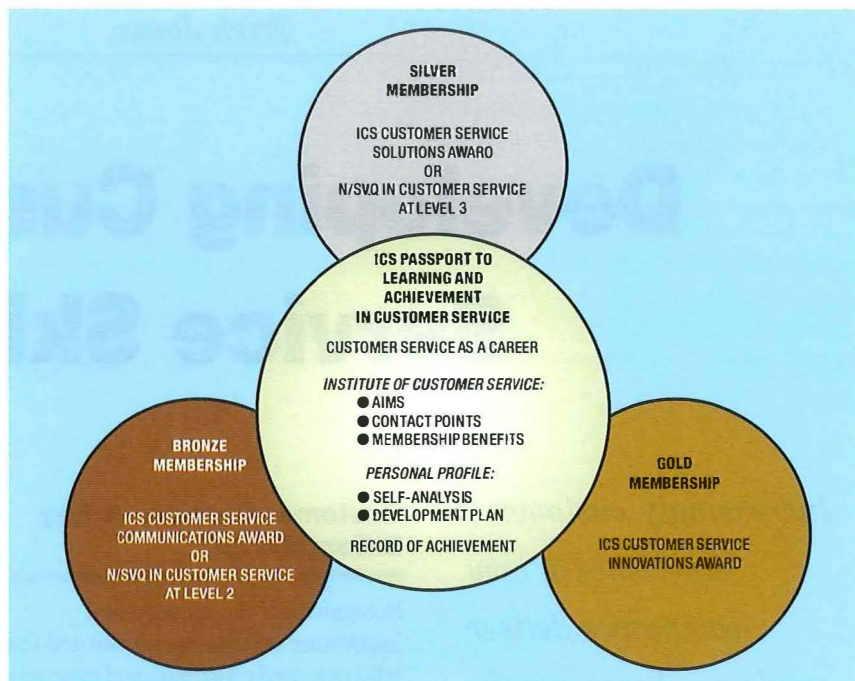
In BT, field engineers, operators and customer service advisors are at the forefront of this growing customer expectation, which is why BT has backed the ICS and contributed to help shape the Institute's strategic development including its membership and qualifications framework.

## Institute of Customer Service—Membership Framework

The membership framework provides the opportunity for customer service practitioners to become qualified members of their own institute. Becoming a qualified member means demonstrating skills excellence in managing the customer relationship including:

- empathy and understanding for customers and their needs;
- a high degree of IT skill—interacting with often complex software;
- commercial awareness, understanding the importance of

Figure 1—Institute of Customer Service practitioner framework



customers to the business, including the emerging demands of different customer types;

- the ability to solve problems, and the self-confidence, direction and knowledge to do so;
- a willingness to work as a team, supporting each other; and
- taking continuous personal and professional development seriously.

The Institute has its own award mechanism through three levels leading to membership as follows:

- *Communications* Award, leading to Bronze membership;
- *Solutions* Award, leading to Silver membership; and
- *Innovations* Award, leading to Gold membership.

Alternatively, membership may be gained through the S/NVQ route in Customer Service Level 2 for Bronze membership, or Level 3 for Silver membership.

At present, people can only work towards one of the Institute’s own awards if they are employed by one of the Institute’s Founder Council Members. Eligibility will be extended to Founder Corporate Members from October 1999 and to all other people from October 2000.

**ICS—Award Profiles**

The three levels of award provided by the Institute, *Communications*, *Solutions* and *Innovations*, have different requirements which have to be satisfied through four key areas:

- *You and your Customer*,
- *You and your Organisation*,
- *You and your Colleagues*, and
- *You*.

The requirements examine how skilled service practitioners are at managing their relationships in each of these areas and how well they flex these collectively to meet a customer’s requirements.

The panel opposite gives a brief overview for each award.

**ICS Approach to Assessment**

The Institute wants its awards to be stretching and perceived by everyone to be credible, worth having and maintaining.

It has designed its assessment processes with this in mind yet with an avoidance of bureaucracy and a focus on outputs; that is, on an individual’s own achievements and demonstrable skills, where words and actions are the same.

There are two distinct strands to assessment:

- *internal* working in partnership with a named coach from within the organisation worked for; and
- *external* working with an ICS assessor to provide an independent and objective assessment of attitude and ability to operate in the round; that is, the expertise of operating the award requirements in unity with each other, and using good judgement to balance the organisational needs in the best interest of customers.

The in-organisation assessment works with the individual, building a picture over time of how the award requirements are being consistently met. A coach is there to advise and support in this:

- by having a shared understanding of how the award requirements fit into the context of his/her own organisation and role being performed;
- by assisting practitioners to identify and evaluate their performance and sources of evidence; for example,
  - products of own work, customer records, notes;
  - feedback from customers, colleagues, line managers, suppliers;
  - self analysis;
  - key events dealt with and discussion around what went well, or not so well;
  - performance appraisals;
  - observation; and
  - development records and action plans; and
- by helping individuals to take ownership for their own development and progress with the awards.

The coach is also responsible for confirming that the evidence is real and valid, that an individual meets the award requirements, and is ready for external assessment.



### Communications Award

You will need to demonstrate that you can build relationships with customers—recognising them as people not labels—where confidence and trust is established early on. The highest standards of behaviour are expected at this first level of award as a foundation for the other two awards. Attitude will be one of the most important elements for success.

Knowledge and skills need to be demonstrated in the following framework:

- *You and your customer* You assess customers' feelings and adapt your own behaviour to put them at ease. You identify and understand their needs, and keep customers informed of progress. You are positive with customers.
- *You and your organisation* You understand the purpose of your work and how it fits with the needs of your organisation. You are confident and competent in using the organisation's forms of communication and you value and apply good service delivery practice. If things go wrong, you recover well; in so doing the customers' trust is restored.
- *You and your colleagues* You work well with others to the customers' benefit and you engage with colleagues and suppliers to get things done.
- *You* You are committed to developing your customer service skills and you are happy to take responsibility for your own personal and work related development by action planning.

### Solutions Award

You will have the qualities demanded of the Communications Award, but you will also have the wider experience and confidence to solve more complex customer problems. Many of these problems will be unusual, requiring one-off solutions and you have the flexibility to handle such issues without creating problems for the future. You inspire confidence in your ability, are proactive and alert to trends in customer needs. You are a team player and often coach others to improve performance.

Your knowledge and skills need to be demonstrated using the following framework.

- *You and Your Customer* You keep a constant eye on service issues, drawing conclusions from your own observations and from customer feedback. You are proactive in satisfying customer needs, tailoring in a flexible and creative way the organisation's products and services to delight customers. Your unprompted extra efforts and problem solving skills minimise complaints and turn potential problems into opportunities to build further customer loyalty.
- *You and Your Organisation* You develop commercial awareness through understanding the key business issues for your own work area and you know how your contribution impacts on your organisation's performance. You display a positive, enthusiastic and effective approach to implementing your organisation's change programmes.
- *You and Your Colleagues* You know how to adapt your behaviour to work effectively with different teams. You recognise the need for and build relationships outside the immediate work area. You coach others and encourage open and honest feedback.
- *You* You are reflective and analytical about your own strengths and development needs. You approach experiences as opportunities for learning and development. You actively follow through your personal development objectives.

### Innovations Award

As an experienced customer-service practitioner, you understand the importance of customer loyalty and retention to the future success of your organisation. You continually seek out new ways of improving customer service. When identifying new opportunities and generating new ideas you are innovative in understanding the customer base, the organisation's competitive environment. You are a motivator and influencer of others.

Your knowledge and skills need to be demonstrated using the following framework:

- *You and your customer* You recognise patterns in service breakdown, recommend changes in procedures, gaining agreement if necessary and implement your ideas where possible. You evaluate the impact of these changes. You are able to deal with unique customer requirements which are referred to you by others because of your skill and innovative qualities.
- *You and your organisation* You network effectively internally and externally. You know and understand your organisation's customer base, its different segments and its competitive environment.
- *You and your colleagues* You act as a coach and role model to others, and are seen as a clear point of reference. You inspire others through your motivation and innovative approach.
- *You* You are determined to keep fully up to date with the latest developments. You are eager to learn new ways of delivering excellent customer service.

This approach is designed to encourage individuals to drive their own development and achievements, emerging as more rounded professionals with a real depth of understanding about the complexities of providing good customer service.

The external assessment is there to test this depth of understanding, making sure there is not a surface approach to learning, and takes the form of a professional discussion, rather than a written test. Again the individual has the opportunity to drive this part of the assessment by outlining three-to-five key events identified as personal examples of providing good service, across a range of customers with differing needs.

Progress in the development of an individual's understanding and improvement in performance is also explored through a learning log.

This learning approach to achieving one of the awards is captured in the *ICS Passport*. The Passport is given to all subscribers of the Institute and contains everything an individual needs to help gain an award.

Completion of the Passport provides a record that demonstrates knowledge, skills and attitudes, and shows how learning from experience has been gained.

### **Benefits of Belonging to the ICS**

There is no other institute designed to create an opportunity for so many—3 million front-line people, extending to 4.5 million including the public sector, in the UK alone.

It will provide independent recognition and give value to the importance of what people do in the customer-service arena. Working through the award framework will raise performance and potential, and membership will demonstrate employability and credibility in the workplace.

The Institute will promote service excellence through courses, conferences and seminars, books and research reports. Members will

receive discounts to these important events and publications which will further promote continued personal development.

The Institute publishes a *Customer First* journal every two months aimed at those dedicated to customer service, covering latest developments and everyday service issues.

The Institute has 13 regions each run by a Regional Chairman responsible for providing networking opportunities through local groups. Other initiatives being developed are library and information services, Companion membership and sponsored awards for outstanding achievement and contribution to customer service.

Gold Members are entitled to use the designatory letters PMICS—Practitioner Member of the ICS—after their name, and Companion Members similarly CMICS.

The founder company members have a commitment to promoting and working with the Institute's awards to establish credibility.

### **Piloting in BT**

As a Founder Council Member, BT is piloting the new awards and their attainment and assessment processes in a number of ring-fenced groups involving over 120 people, until the end of June 1999. At this point, the results of the trial from across the 50 founder companies will be evaluated and the materials and procedures reviewed ready for a broader launch including 100 new companies.

### **A New Era for BT**

This initiative puts the UK at the forefront in developing customer service as a career supported by major employers whose prime purpose is customer service.

BT is not alone in developing strategies for expanding and seizing opportunities for growth around the world. Global market players need global telecommunications networks

and support, and investment decisions by other service companies are increasingly being influenced by the quality and service offerings of telecommunications providers. The UK is benefiting from the competitive telecommunications market which exists here, and consequently the UK is now the preferred European location for communication centres; for example, international call centres, shared service centres and behind the scenes processing functions.

BT has a major opportunity to capture this growing foreign investment into the UK telecommunications market and indeed many potential investors are already customers of BT or BT's international joint venture companies.

BT UK alone has 26 million customers with over 1 billion contacts per year.

BT invests £2 billion annually in the network and nearly £300 million in research and development. BT aims to offer everything from outsourcing and call centre technology through to managed desktop solutions. Concert offers customised solutions to meet the global communication requirements of international businesses. Aftercare programmes managed by account development teams are an essential part of maintaining customer loyalty and changing needs.

There is also the potential market of 300 million Internet users in just two years time, and BT is instrumental in providing innovative links for basic connectivity to advanced electronic business applications; for example, BT Highway, BT Connect to Business, and the even more recent BT Argent (a secure Internet payment service).

As BT develops and implements these services, the company is learning from its customers to understand their needs and create new solutions. Every experience of working with customers and suppliers provides an opportunity to be innovative and build customer loyalty and retention. The company's business depends upon it.

Technical skills alone will not be enough—a broader skill base has to be cultivated if BT is to take advantage of its massive engineering and technological capability. The ICS and its awards framework allows for that broader customer service skill to be developed where it matters—in the front line, the face of BT to its customers.

### *Acknowledgement*

---

Some of the material for this article has been drawn from 'The Pathway to Success—A Guide to Individual Membership', published by:

Institute of Customer Service,  
2 Castle Court  
St Peter's Street  
Colchester  
Essex  
CO1 1EW  
Tel: 01206 571716  
Fax: 01206 546688  
E-mail: [enquiries@instcustserv.com](mailto:enquiries@instcustserv.com)  
Web: <http://www.instcustserv.com>.

### *Biography*

---



**Trish Jones**  
BT UK Customer  
Service

Trish Jones has a background in strategy, operations and people development, and is currently on secondment from BT to the ICS. She has been instrumental in developing the ICS qualifications framework and has worked with a number of the Founder Company Members to pilot the awards across a diverse range of operations. She has been awarded the first Companion membership of the ICS in recognition of her contribution to the development of the ICS.

# Structured Greenfield Planning

*This article discusses the analysis behind new residential network greenfield planning policies and standards as the network planning part of BT's 1997 Greenfield Site Initiative (GSI). The analysis included both new design ideas and the application of appropriate new technology. Trials confirmed a 10% saving in on-site network-provision costs attributed to network design alone.*

## Introduction

Structured process improvement<sup>1</sup> concerns the integration of technology, process and systems to bring about cost reduction and/or revenue generation. In the greenfield planning arena there is much scope for such integration, a situation recognised by BT's Greenfield Site Initiative (GSI) Programme.

BT spends tens of millions of pounds per year on the provision of public switched telephone network (PSTN) services to greenfield sites. Nowadays much of this investment in new network is at risk due to competition, and some of these investments may yield little return on capital. In this situation any reduction in the planning cycle time or reduction in the initial costs is welcome. Reduced cycle time helps to ensure that planning does not have to be done too far ahead of implementation and therefore demand. Attention paid to whole-life costs ensures investment required is as low as practicably possible.

These reductions in cycle time and day-one investment costs have been fully effective after integrating the standards from the structured greenfield planning activity with the network optimiser tool (GenOSys<sup>2</sup>), which is the subject of an accompanying article in this issue of the *Journal*.

In order to establish the proposed standards that would yield lowest cost, desk-top trials were carried out where network designers (network planners) planned schemes to both existing and the proposed planning policies and standards.

Some interesting questions were posed by this study. Would there be any variation in scheme costs by

Zone? Where do the biggest variations occur? How large are they?

This article addresses these questions and presents some new ideas, embodied into a standards set, including:

- standards on lead-in provision,
- common trenching,
- duct tees distribution,
- road crossings without footway box terminations, and
- single pre-cast footway box usage.

The article also addresses layout methods for planning.

## Elements of Greenfield Planning

The analysis was conducted using the on-site telephony network consisting of underground ducted copper cables serving a new (greenfield) residential building site. To manage the analyses and to identify areas of high cost, the infrastructure was broken down into four main elements. Three of these defined the duct layer and the fourth defined the cable layer. This is shown in the scheme in Figure 1 and is defined below:

- *Spine* The main duct and jointing chambers installed on a main road leading onto a development site; the number of duct bores in a single trench can range between one and eight (dashed line in Figure 1).
- *Sub-spine* Duct and jointing chambers leading from the main

*in order to sell the plots all utilities need to be on line or at a point where the final connection can be made quickly*

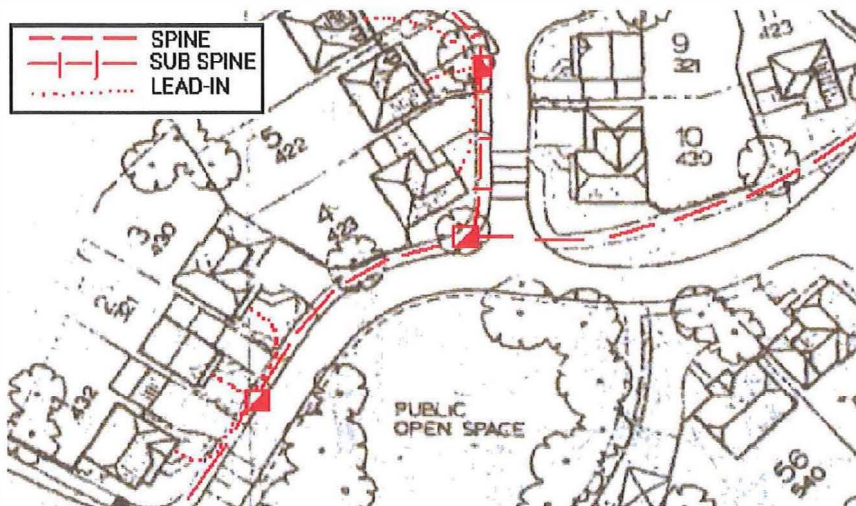


Figure 1 – Site plan definition of infrastructure elements

spine route to a distribution point jointing chamber, whereby the number of duct bores in a single trench was usually one or two (crossed line in Figure 1).

- **Lead-in** A single duct and cable from the distribution point jointing chamber to the customer's premises (dotted line in Figure 1).
- **Cable and jointing** All cables installed in the spine and sub-spine ducts; this was inclusive of standard joints and distribution-point joints. The definition does not include cable in lead-in duct as this was already defined in the lead-in definition.

An on-site jointing chamber or joint box is built as a footway box. The footway box enables access to the cable layer components. Typical uses for a footway box are for housing a cable joint, distribution point, or to enable cable installation personnel to turn cable through sharp angles during installation.

### Planning Considerations

Consider a proposed 20–150-plot building site of residential tenancies. In general, in order to sell the plots all utilities need to be on line or at a point where the final connection can

be made quickly. The most advantageous times to install all utilities' cable, pipework and duct occur at different times within the development of the site. All utilities are installed at different depths in line with the National Joint Utilities Group (NJUG) standard and their general routing follows the proposed line of the adopted roadside verge. BT plans its underground routing independently of the other utilities leaving the site developer to interpret and install the duct layer over the other utilities' networks.

Future application of lean methods<sup>3</sup> might point towards the site developer, or a single recognised planning authority, being entrusted to plan the utility networks to a series of specifications concurrently with the overall site plan. The specifications might then be embodied in systems and computer-based applications, and delivered just-in-time for implementation. At the time of the study this radical approach was not part of the GSI programme, but is currently being assessed.

For each residential building site, using national planning policies and standards, a duct network with joint-box positions is designed, spine cable and joint requirements are identified from the primary cross-connection point (PCP) to each distribution point

(DP) and a site jointing record is drawn up. It is recognised that the approach taken by each network designer is different.

The project was targeted at realising the most cost-effective initial design reflecting the necessary parameters both into national planning policies and standards, and, for incorporation into the network optimisation tool, GenOSys (also known as the *automated planning tool*). This work excluded the task of planning the off-site network necessary to support the greenfield site, but included the tasks necessary to complete the on-site work. The productivity analysis behind the work was based on current access quality task time synthetics since this enabled individual activities to be judged for time, manpower and cost. It was recognised that some aspects of the work would require the approval of other relevant bodies, such as the House Builders Federation.

For each greenfield network element this article considers the structured improvements leveraged from the analysis purely from a network designer's point of view.

### Customer Lead-in

The customer lead-in element is the most costly to provide since there is little shared infrastructure. Any savings identified in this part of the network have the most significant effect on the overall installation cost. A network designer is bound by routing the duct either down adopted roads, footpaths or within the curtilage of the served plot.

The analysis centred on three methods of installation technology and two different methods of trenching outlined in Table 1 and Figures 2–4. Each method consists of cable in a duct that is laid into open trenchwork:

- The radial duct system serves each plot with a single empty duct into which cable is drawn at a later date.

# underground duct tees lead-in technology is the most cost-effective over the range of plots of practical interest

**Table 1 Customer Lead-in Distribution Systems Analysed**

	Common trench	Non-common trench	Figures	Notes
CiD	X	X	2 and 3	
Radial duct	X	X	2 and 3	
Duct tees	X		4	Mixture of common duct and individual duct




Figure 2—Duct technologies and non-common trenching

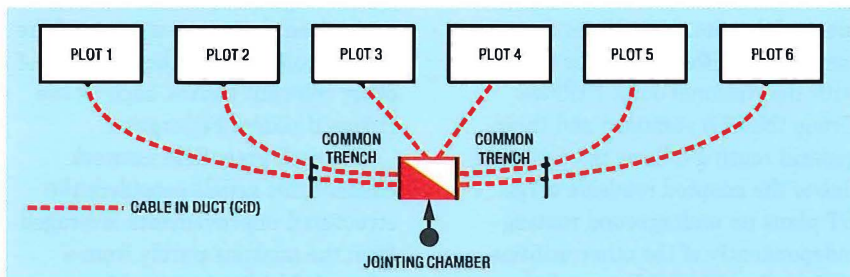


Figure 3—Duct technologies and common trenching methods




Figure 4—Duct 56 and tee technology

- Cable in duct (CiD), a relatively new technology, is where the lead-in cable is supplied to site pre-drawn into the duct on drums and is installed in the same manner as the radial duct system.
- The duct tees system (refer to Figure 4) is made up of a section of duct in which lead-in cables share

the bore, and a section of duct that serves the plot with a single lead-in cable. A duct tee component is used to join the two sections of duct together enabling the cable from the individual plot to be installed towards the distribution point.

- Common trench practice is where a single trench is dug and more

than one duct shares the trench for a maximum practical distance (for example, Figure 3).

The practice of directly burying the lead-in cable was not considered.

The analysis has shown the greater the sharing the greater the savings in productivity that can be made. This is borne out in Figure 5.

Figure 5 shows that lead-in costs of duct tees distribution and common trench cable in duct have similar per-plot costs over the practical range of plots fed by an underground distribution point. The cabling practicalities of an underground node are:

- A distance limitation of 50 m between node and served plot is one of the practical ability of a single engineer being able to install a cable through a populated duct.
- A multiple cable entry joint which can be installed easily, upgraded, and maintained without compromising its design functionality like ease of entry and no water ingress.
- The joint and cables must be able to be accommodated within the design of a footway box.
- Multiple duct entries through one wall of a footway box have the effect of compromising the integrity of the box thus necessitating a stronger box at a higher cost.

Figure 6 amplifies the two lowest-cost lead-in techniques: duct tees and common trench cable in duct. This shows that the underground duct tees lead-in technology is the most cost-effective over the range of plots of practical interest. The costs of per-plot duct tees were flat from between eight and 12 plots, with a jump after 12 served plots owing to a second duct within the same trench since the first was full. Hence the upper bound of tenancies

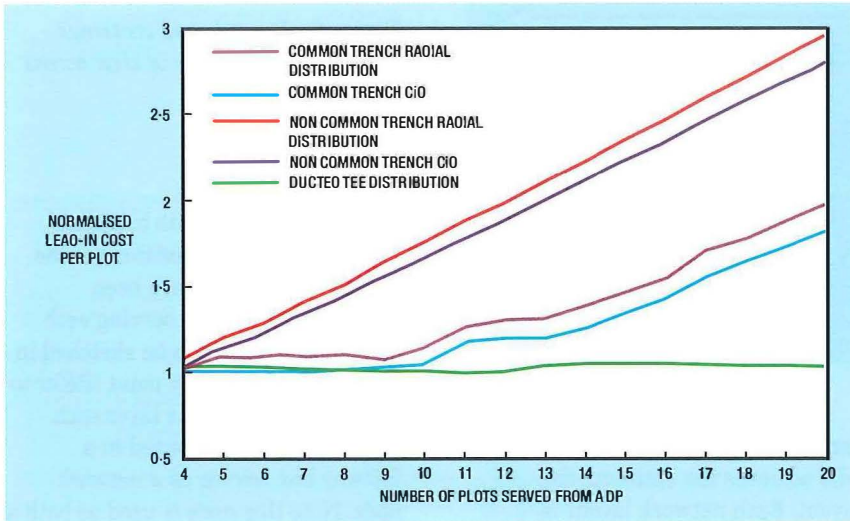


Figure 5—Cost against number of plots served from a distribution point for the five lead-in techniques described in Table 1


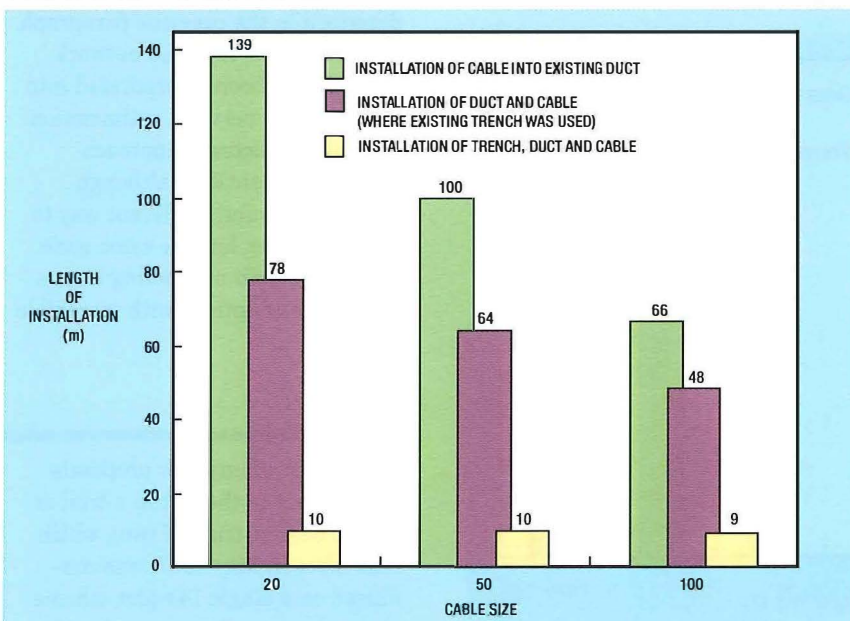


Figure 6—Cost against number of plots served from a distribution point for duct tees and cable-in-duct lead-in techniques

Figure 7—Graphical representation of maximum duct and cable combinations for a single Joint Box Type JB226



served was recommended to be 12 and the lower bound was eight. However, it must be emphasised that these costs have been analysed using a standard range of productivity synthetics, hence the use of cable in duct is still under review and on trial.

### Spine and Sub-Spine

The principal area where a network designer can make a significant influence is in reducing the number of footway boxes. This can be effected by ensuring all joints are sited within existing distribution-point joint housings and by placing distribution-point footway boxes at or near proposed road crossings and then routing the road crossing duct work at an angle across the road to feed the following DP footway box.

### Spine and sub-spine installation

Consideration was given as to the equivalent installation lengths of trench duct and cable, duct and cable, and cable alone for the value of a pre-cast footway box Type JB226. Footway box Type JB226 is not in use but at the time of the analysis it was under consideration. The equivalent lengths of 20, 50 and 100-pair cables are shown in Figure 7 indicating that the selection of footway box sites needs to be chosen carefully to minimise the number of boxes and to maximise duct usage.

### Spine/sub-spine road crossings

In many cases, even though road crossings necessitate deeper duct depths than under the footway, road crossings can be designed to link up existing footway box sites (DP sites too). Hence this avoids the addition of boxes to install cable around a right angle. The adoption of this methodology can not only reduce day-one costs but also lead to a reduction in maintenance costs. This can be achieved by slewing the duct route under the road using standard 18°




Figure 8 – Ducted road crossings using different angles of slew across the 6 m wide road

duct bends as shown in Figure 8. Figure 9 shows the normalised costs against each of the angled road crossings with 20 m of duct set under the footway.

Figure 9 indicates the cost of installing both a single-way ducted road crossing and two footway boxes either side of the road is double the cost of a comparative ducted crossing with a 20 m length of duct under the footway.

Through discussions with field staff, a ducted 90° road crossing (no boxes) was considered to be too tortuous a route to cable, yet 72° (and lower angles) was considered acceptable. A maximum route deviation angle of 72° for a road crossing was chosen because this was the minimum radius which a cable route can be deviated twice between boxes.

### Layout Methodology

Although a set of on-site network policies and standards can be

generalised, it is the network designer who achieves the optimum site layout. Each network layout is unique, in a way a work of art in terms of design. It must be functional in that it has to be able to deliver a telecommunications service network to its customer base at minimum capital investment and provide security—responsive to network change—and a low maintenance infrastructure at all levels (service, cable, duct etc.).

At the network design stage a structured approach to the layout is crucial. This has been recognised in areas such as the internal layout of integrated circuits found in every PC. It is for these reasons the network designer needs to approach the design at the micro level by roughly grouping together the customers into geographical nests of between eight and 12 and serving these from a single distribution point. There are maximum distance limitations between distribution point and

customer. Once the site has been grouped and rough positions of the distribution points have been identified, the spines serving each distribution point can be sketched in towards the site entry point (Refer to Figure 1). At the cable layer each distribution point, housed in a footway box, serves as a network node. Note this node is used as both a cable joint and customer distribution point. At this stage the designer can realign customers to distribution points until the network meets the functional criteria and policies and standards. Combining this approach with that of slewing duct under the

*At the network design stage a structured approach to the layout is crucial.*

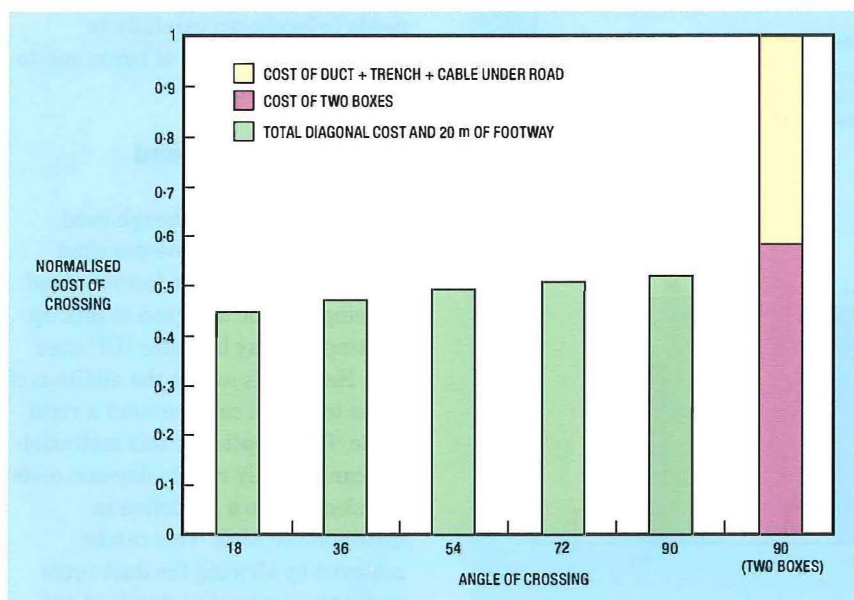
carriageway, the number of distribution points and extra footway boxes can be kept to a minimum. In this way the network designer coordinates various elements under their control to produce a site network which meets the functionality described in the previous paragraph.

Much of the reported network analysis has been incorporated into the GenOSys network optimisation tool<sup>2</sup>. The structured approach adopted by GenOSys, although applied in a subtly different way to the cable layer, has the same goals with the benefit of choosing from a vast range of options with incredible speed<sup>4</sup>.

### Trial

The quality of any new proposals must be put to the test in a trial or pilot. The first trial (of two), which took place in May 1997, was conducted on a single 144-plot scheme planned across the country by 13

Figure 9 – Normalised cost of different road-crossing scenarios shown in Figure 8





## *the savings attributed to the design of the on-site network alone were 10%*

planning offices. The trial established the variation in total on-site costs (Figure 10 and Table 2) and indicated the cost for each element. It was noted that the designs with least spine and sub-spine duct costs were the most cost-effective. This was attributed to a lower number of footway boxes. The design approach taken by each planning office was noted and two features are summarised in Table 2.

In the second trial (September 1997) new design ideas and new proposals for infrastructure components (for example, footway box Type JB226) were tried by 27 Network Design representatives on plans submitted and previously costed by themselves. The plans were for the same size schemes—about 50 plots. Prior to the trial the success criterion was set. A summary of the trial is shown in Table 3. The trial of new network design ideas and new technology identified a mean saving of 16%, however a 10% saving in value was attributed to the new approach to on-site network design.

### Conclusion

The article has described a structured approach to underground on-site metallic telecommunications infrastructure design for new developments (previously known as *greenfield* and *brownfield* developments). It has pointed up an important skill used to combine functional telecommunication service requirements with those of minimal installation and maintenance costs. The costs attributed to the lead-in are the most significant borne by the network provider and it has been shown (Figure 6) that the common use of resources in terms of duct and number of customers served from a single distribution point is beneficial.

Within the spine network the footway box represents the single most expensive item which can be commuted to either the installation




Figure 10—Regional variations in planned costs for a single residential scheme

Table 2 Variations in Key Planned Elements for the Single Scheme Represented in Figure 10

Description	Values
Total cost range represented as a percentage	100%–132%
Planned joint box range	20–30
Planned distribution point range	16–29

Table 3 Cost Summary of On-Site New Development Gains Due to Improvements in Planning and Technology

Component	Maximum	Minimum	Mean	Success Criteria
Day -one cost reduction of proposed against existing planning standards	51%	-2%	16.6%	6%–25%
Component	Existing Standards	Proposed Standards	Improvement	Success Criteria
Mean number of footway boxes per site	11.2	6.7	40%	Not defined
Mean number of lead-ins per DP	5.4	7.9	46%	Not defined

of 140 m of 20-pair cable or 10 m of trench, duct and cable (Figure 8). Footway boxes whose sole role is to turn cable through 90° at road crossings can be made redundant by slewing duct at 72° across the road.

Finally, it has been recommended that the approach to on-site network design be turned upside down. Thus the network designer should start by grouping customers together and assigning the group to a distribution

point, then work the network back to the on-site entry point. In this way each distribution point, housed in a footway box, can both serve up to 12 customers and be a joint where necessary.

Trials have shown that with the incorporation of the proposed policies and standards and with a change to the approach to network design, the savings attributed to the design of the on-site network alone were 10%.

## Acknowledgements

The authors acknowledge the work of John Foster, Kevin Clarke, Mick Stephens and Enda McSharry for their part in the analysis, and Ray Hooper for his helpful comments and without whom this task and paper would have been impossible.

## References

- 1 BEALES, K. Structured Process Improvement. Guest Editorial. *ibid.*, April 1998, 17, p. 1.
- 2 ASUMU, DON; and MELLIS, JOHN. GenOSys—Automated Tool for Planning Copper Greenfield Networks. *Br. Telecommun. Eng.*, Jan. 1999, 17 (this issue).
- 3 HOOPER, R.; and JONES, C. Lean Operations—The Route to Competitiveness. *ibid.*, April 1998, 17, p. 2.
- 4 MELLIS, J.; and ASUMU, D. Performance in Planning: Smart Systems for the Access Network. *BT Technology Journal*, Jan. 1999.

## Biographies



**Ken Cobb**  
Networks and  
Information Services,  
BT UK

Ken Cobb is a member of the Planning, Plant and Pilots Unit within BT Systems Engineering based at BT Laboratories. He is working on a number of initiatives addressing the planning of new developments both in the residential and small business sectors. Starting with a year on the shop floor at AVO he moved to GEC Marconi Communication Systems Ltd. and graduated with an honours degree in Electrical and Electronic Engineering from Bradford Univer-

sity. After completing his M.Sc. in Microwaves and Modern Optics at University College, London, he joined BT in 1981 working on the early installation of multimode and singlemode standard optical fibre systems. His role included liaising with BT's fibre and systems contractors and work on fibre characterisation and test methods. He took a post at BT Laboratories to develop the emerging portfolio of in-building fibre components and was involved in the integration of BT's blown fibre into the building. Subsequent posts have included work in the fields of Dropwire 12 verification, TPO, OTIAN™, video-on-demand trials and access fibre and copper planning consultancy.

Ken Cobb can be contacted at [ken.cobb@bt.com](mailto:ken.cobb@bt.com).



**Mark Rose**  
Networks and  
Information Services,  
BT UK

Mark is a member of the Planning, Plant and Pilots Unit within BT Systems Engineering based at BT Laboratories. Prior to joining BT in 1989 Mark worked for a company programming control systems for major chemical plants. Before joining the team at BT Laboratories in 1997 he worked for eight years in various access network planning groups. During this time he was awarded an HNC in Mathematics and Electronics from Cambridge Regional College and he developed a computer programme office targeting system (POTS) application for the Zone, which was a precursor to the current IDACS tool. Mark's first project at BT Laboratories was to manage and work on the project reported in this article. Mark wrote all the computer-based tools for data capture, analysis and presentation. He is now developing a PC-based 'tool for authorised

standard case' (TASC) application to be used by strategic planners to aid the decision between copper or TPO provision to a site.



**Ian Lawrence**  
Networks and  
Information Services,  
BT UK

Ian Lawrence is a manager within the Access Infrastructure Unit of BT's networkcapacity Unit. He joined BT in 1986 as a trainee technician on Network Provision and Repair. In 1988 he joined Customer Engineering specialising in the repair of small to medium business systems and network special fault investigation. In 1990 he moved to BT's Business Systems Technical Support Unit providing dynamic technical support for all CPE equipment up to 30 lines, and special network faults including electro-magnetic compatibility (EMC) and radio interference. During 1992, Ian took on the job of National Pilot Manager for BT's UK Access Division where he was responsible for the pilot trials of many of BT's new access technologies; for example, DACS, HDSL, Mk 2 Blown Fibre and LA30. He now leads a team responsible for owning, developing and implementing BT's access narrowband infrastructure planning policies and standards, which includes planning policy automation, the sponsorship of development work at BT Laboratories and BT Development and Training, and the implementation of BT's University of Planning. He is currently studying for an MBA with the Open University Business School.

*Don Asumu and John Mellis*

# GenOSys—Automated Tool for Planning Copper Greenfield Networks

*It is at the planning stage that the future risk to any network infrastructure build proposals can be minimised. Mistakes can have a massive impact on subsequent return on investment and hence continued profitability. Several complex techno-economic considerations usually need to be taken into account and often within aggressive timescales. These factors naturally lead to the need for tools that will facilitate at least partial automation of the planning process. This article describes GenOSys, a system which addresses the automation of several steps in the copper greenfield access network planning process.*

## Introduction

GenOSys (genetic optimisation system) is a copper greenfield service-delivery planning tool which evolved from a detailed study of BT's current processes for planning both copper and greenfield passive optical networks. Its primary objective is to automate the service delivery planning processes for copper greenfield networks, and, as this article points out, GenOSys has been defined by, and has in turn redefined, the greenfield copper planning process. In practice, the detail of greenfield planning differs according to operating zone, the individual planner's experience, patch geography and the interpretation of operating policies laid down in planning guidelines. The stages of planning from the onset of a proposal through to the creation of a detailed job pack which specifies every aspect of design and build of greenfield networks represent the front end of the network management life cycle<sup>1</sup>. It is at this point that the future risk to any network infrastructure build proposals can be minimised. Quite simply, getting it wrong at this stage has a massive impact on subsequent return on investment and hence continued profitability. The presence of competition, the need to turn proposals around in shorter timescales, the complexity of the underlying rules for design and build and the rate at which these rules change continue to be the main pressure points on a shrinking planner workforce. These underlying

forces represent a fundamental shift in how access greenfield networks are planned and subsequently built.

GenOSys is a design and planning tool which addresses many of the process steps involved in creating or modifying a copper greenfield network. By automating many steps in the process, significant and measurable productivity benefits for the planner, and hence the business can be achieved. It assures a consistency of design and build achievable because designs are based on a single and authoritative interpretation of current planning rules; no longer are rules subject to individual interpretation, misunderstanding or plain neglect. Furthermore, GenOSys, through the application of artificial intelligence is capable of simultaneously considering several thousand scenarios for infrastructure build, and then selecting the most cost-effective infrastructure design. It does so in a time shorter than a planner would take to justify and cost a single scheme. GenOSys has been successfully proven in test-bed and pilot use, and the business benefits projected to accrue from its field deployment are about £10 million per annum, or about 10 per cent of the total annual investment made in copper greenfield networks.

This article describes the present greenfield planning process and highlights the opportunities for automation. The architecture and engineering of the GenOSys system are discussed in terms of its requirements, evolution, build philosophy and build detail. The process of

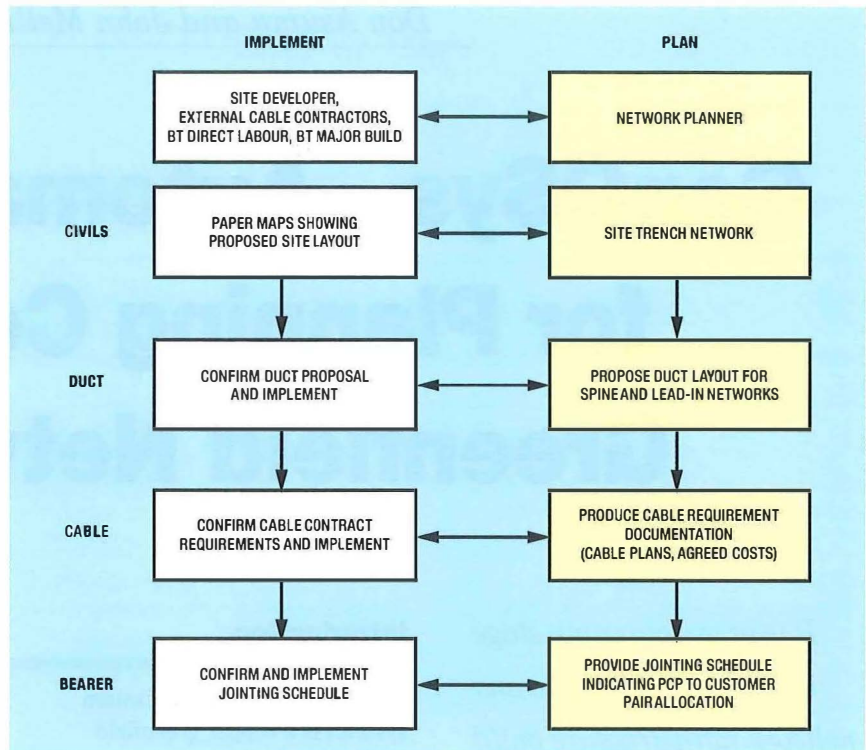
Figure 1 – A layered view of the copper greenfield planning process

detailed planning using GenOSys is described, further illustrating how individual process steps have been automated and where productivity gains arise. Finally the lessons learnt in this development are described along with the experience gained from deploying the system into test-bed and operational use across the business.

### The Generic Greenfield Planning Process

Most of the design activity carried out to create a scheme for a greenfield network is done between the primary cross-connection point (PCP) (made available at the entry point on the curtilage of the greenfield scheme) and each customer. This activity takes place across and within four layers. The site developer's initial geography and actions define the civils layer. For instance, the type and position of each customer, and the layout of site trenching. Within this layer is the duct layer. It consists of the duct pipes and intermediate boxes that serve to connect sections of ducting together. Enclosed within the duct layer is the cable layer. In this layer, groups of bearer pairs are aggregated into cable sheaths. Within the cable sheaths are the bearers themselves which deliver services to the connected customers. Figure 1 illustrates an overview of the process and the relationship to each of the layers highlighted above. Information and data flow between external contractors (site developers, cable contractors, major build) is represented by the double-headed arrows. The process begins with the site plans of the proposed site being received from a developer (typically on an A0 or A1 sheet of paper). Upon this plan, the network planner would lay out a duct network, and submit it for consideration to the developer.


If the proposal is acceptable, a cable network can then be produced. Additionally, a jointing schedule is prepared to specify how each



working pair on every cable is to be jointed back from each customer to the exchange. The duct network comprises of a spine network and a lead-in network. The spine network consists of 90 mm diameter duct laid between the entry point to the site—either a primary cross-connection point (PCP) or a box which is built or provided by BT—and access points which will subsequently contain the flexibility or distribution points (DPs) of the scheme. The lead-in network consists of 49 mm diameter duct and is used to establish connections between the DP housings and entry points to each of the site customers. Between the site entry point and each DP, the main spine duct comprises of sections of ducting which run between access points on the network. These access points are usually boxes that can either be pre-fabricated or built on site by the developer. These boxes enable junctions to be created on the duct network. Junctions serve two major functions in the duct network.

Firstly, they allow cables to be introduced into any part of the duct network. Secondly, they permit the introduction of points of flexibility or jointing into cables that pass through or are terminated at a box. Between the DP housings and each customer, the lead-in network consists of 49 mm diameter duct sections (Duct 56) and optional branch elements (duct tees). Tees enable sharing of runs of Duct 56 and obviate the need to run individual sections of Duct 56 between each customer and the serving DP<sup>2</sup>. Instead, cable can be introduced by rodding from each customer through an intervening duct tee up to the box that houses the serving DP. Figure 2 illustrates the rodding process. There are, however, engineering constraints on the use of duct tees which need to be balanced against the cost savings that they bring. In both the spine and lead-in networks, changes in duct section direction can be achieved by using a variety of duct bends. Furthermore,

Figure 2 – Duct tees in the lead-in network






Figure 3—Spine and lead-in topologies

direct connection of the spine and lead-in networks can be made using reducer elements which fit Duct 54 and 56 sections together.

Once a duct network has been agreed, a cable network can be introduced. The cable network also consists of two parts. The spine network is located between the PCP and each DP, while the lead-in network consists of each DP and its associated customers. In reality, because cables are manufactured in discrete sizes, with 100 pairs being the largest deployed size in the D-side network, a number of overlaid spine networks may exist. Each sub-network serves up to a maximum of 100 pairs. Beginning from the customer edges of the duct network, pairs are allocated based on the predicted level of demand per customer (dependent on factors such as the actual or potential level of competition). At this point, a decision about the number of customers that will be served (and therefore the total number of copper pairs) and the position of the serving DP needs to be made. DPs are usually sited at the most central box in the cluster of customers. This position tends to minimise the run of cables between the DP and each customer. Once all the DP pair counts and positions are determined, the design of the spine network cable can begin. For the spine network, the goal is to provide the smallest number of cables to provide service to the groups of DPs present on each sub-network. To achieve this, the DPs are grouped in such a way that the aggregate number of

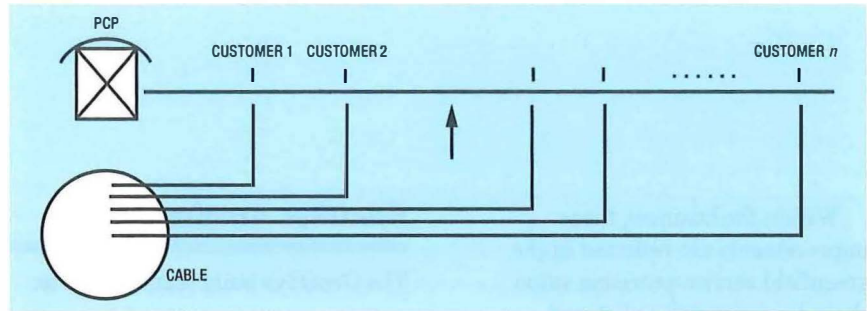


Figure 4—Allocation of cable bearers on basis of distance from PCP

pairs for each set of DPs does not exceed 100. Each cable that provides service to a DP group has its outer sheath exposed at the box which contains a DP, and then only the required number of pairs are jointed at the DP to the service cable in the lead-in network. Thus the cable spine network forms a logical bus. At the bearer level, of course, the configuration is a tree. Figure 3 illustrates this.

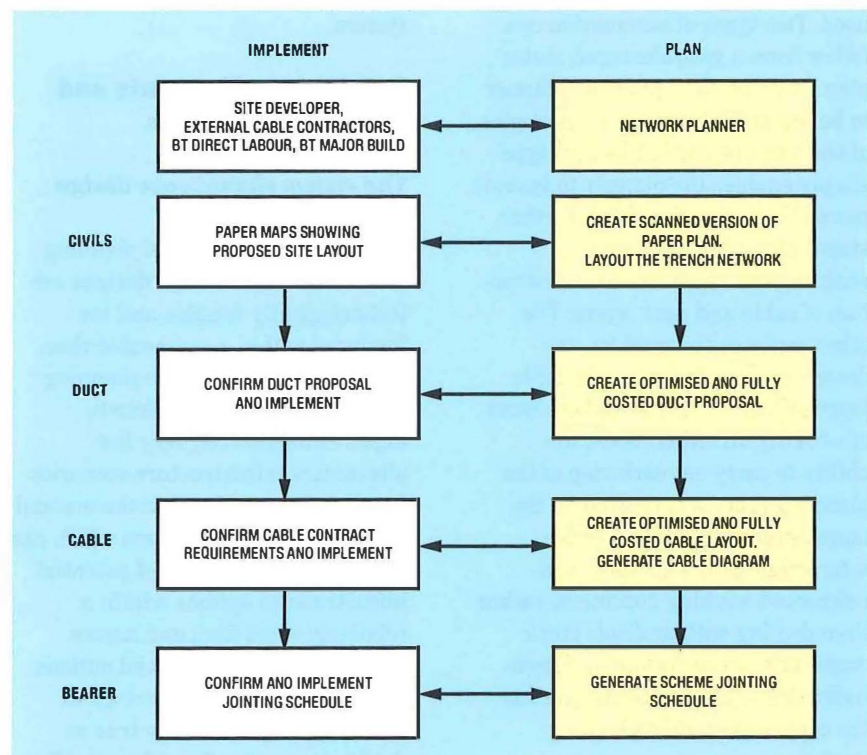
Upon completion of the cable network, a jointing schedule needs to be prepared. This document defines how connectivity at the bearer level is defined. In order to minimise the risk of damage, the pairs on the outermost part of a cable are allocated to the most upstream DPs or customers, with the assignment of the innermost pairs being to the

more downstream DPs or customers. Figure 4 is a simplified illustration of this process.

### Opportunities for Automation

Each of the process steps outlined in the previous section presents opportunities for automation. Figure 5 illustrates the process again, with an emphasis, on the right-hand side, on how automation can be applied. As far as the site developers, external contractors and direct labour units are concerned, the key change in the process is that the network planner suddenly appears to be significantly more responsive in dealing with proposals, in addition there is a greater consistency of final proposals.

Figure 5—Automating the planning process



Within the business, these improvements are reflected in the greenfield service-provision value chain by appearing as lowered network planning and provision costs. The basic unit of information exchange—the site developer's proposal—is first converted to a raster image by scanning. This necessary (but not sufficient) first step immediately provides the planner with an electronically flexible document upon which to carry out his/her design. A further step in enhancing the intelligence of the scanned document is a means by which it can be geo-spatially referenced to the real world. By this means, the planner is able to work with a true representation of the developer's site and be assured that positions and distances are correctly reflected on the raster image. Together, scanning and geo-referencing are the essential background for the representation of vector plant as a 'smart' overlay on the scanned raster image. With these key elements in place, the other stages in the process can be automated successfully. At this point, the term 'automation' has been rather loosely used. Two types of automation can follow from a geo-referenced raster plan. Both types require the planner to be aware of the possible topologies of the civils network. The first type simply enables the planner to become more effective in specifying further stages of design; for instance, enabling the rapid manual construction of cable and duct layers. The other removes the need for any design work at the duct and cable layers. GenOSys supports both ways of working. In either mode, the ability to carry out each step of the planning process is centred on the more effective model of flexible interaction with a spatially geo-referenced working document, rather than dealing with multiple static copies of a paper document. Essentially, this is the key to automating the copper greenfield planning process.

## GenOSys Architecture

The GenOSys build team set out to:

*Create a system which will autonomously design copper greenfield network infrastructures that satisfy engineering and cost constraints. The autoplanning system should further be capable of producing contract standard documentation and estimates of infrastructure cost as required by the site developer and any external contractors. The auto generated design should be cost-effective to implement, and meet any engineering constraints imposed by the type of network plant being deployed. Furthermore, any infrastructure designs created by the system should be produced with a minimum of intervention from the network planner.*

Given the opportunities which have been identified in the generic copper greenfield planning process, and the objectives which need to be met by the system, this section addresses the detailed requirements, desirable attributes, key components and business benefits of such a system.

### Detailed requirements and desirable attributes

#### *The system should save design time*

Planners work to a set of planning rules which ensure that designs are technologically feasible and are produced within a reasonable time. Given the pressure of job planning volumes and the time taken to explore alternatives, very few alternative infrastructure scenarios are usually considered in the manual planning process. A system which can explore a large number of potential infrastructure options within a relatively short time and assure consistency across explored options can yield considerable savings in time. This is particularly true as design progresses from the specifica-

tion of infrastructure to the detailed costing and documentation stages. Therefore, even when the planner must override significant portions of an autogenerated design, the time saved in recomputing costs and reworking documentation is essentially only a few mouse clicks away all the time.

#### *A feasible and cost-effective infrastructure design is required*

This is a central requirement for an automated planning system. It is a challenging task because of the number and kinds of parameters as well as the mix of engineering and/or policy constraints that must be factored into the creation of candidate infrastructures by the system. For instance, generating a cost-effective cable layer may lead to a duct layer network failing to meet engineering constraints. Changing the civils layer specification (done by the planner) may stop the system from meeting duct layer policy constraints. Many feasible designs may not be cost-effective. However, this type of task is well suited to the methods of artificial intelligence, and have been applied within the GenOSys tool.

#### *Optional and progressive automation is desirable*

Responsibility for the outputs of an automatically generated scheme must be assumed by a network planner. A highly experienced planner might for instance choose to specify every aspect of a proposal, using the system as a means to save time in the preparation of proposals. Alternatively, such a planner pressed for time might elect to have the system auto generate the duct and cable layers, and then override a number of design decisions made by the system. In this alternative scenario, it may be that the planner is completely satisfied with the outputs of the system and no further modification of the initially produced design is necessary. Therefore, a system that allows automatically

generated designs to be overridden or modified in the light of a planner's experience is desirable. Such a facility allows the proper flexing of automation against the experience of the planner. The challenge is to perform such actions seamlessly and in such a way that in whatever mode (manual, semi-automatic or fully automatic) the planner chooses to use the tool, exactly the same set of actions are required for interaction.

### **System component re-use is desirable**

In order to solve similar types of problems in future, it is desirable that the framework upon which the system is constructed lends itself to rapid re-use in related problem areas. For instance, the problem of primary network fibre provisioning is analogous to the problem of copper greenfield spine network provision. Similarly, the problem of service distribution from a central serving facility contains the same essential elements regardless of the medium used to deliver service (for instance radio or fibre). In such cases, the extent to which the similarities in

problem domains may be successfully exploited depends on the intrinsic nature of the framework used to express and implement the copper greenfield problem. In the most general case, the question is: how can the array of subsystems put together for the purpose of copper greenfield network planning automation be readily deployed in scenarios that require networks to be optimally planned and design and costing documentation subsequently generated? Indeed, at a low enough level of granularity, then can the subsystems which have been deployed within GenOSys be migrated to other problem areas?

Object technology has been extensively used in the design and subsequent implementation of all the constituent subsystems of GenOSys. It has subsequently been found that building systems in this way enables the inevitable change in customer requirements to be rapidly and effectively addressed.

### **Architecture detail**

Figure 6 highlights the major subsystems (modules or components)

of the GenOSys system as well as their relationships to each other. This architecture closely mirrors the requirements specified earlier in the article.

A key subsystem of this architecture is the communication object-oriented component architecture (COCA). This is a software model of the real-world access network infrastructure, abstracting essential properties of fundamental components as well as the layering and connectivity relationships between these components. COCA acts as a data pipeline in GenOSys, enabling infrastructure descriptions to be provided or updated by the planner, and enabling the automation subsystem to work on a given network description at the civils layer. The geo-spatial component enables the site plan obtained from the developer to be geo-referenced and displayed (after it has been scanned and converted into a raster image). Spatial geo-referencing is done by setting a scale, defining a real world point of reference (for instance, a known spatial position; for example, a PCP northing and easting) and

Figure 6—GenOSys—major subsystems

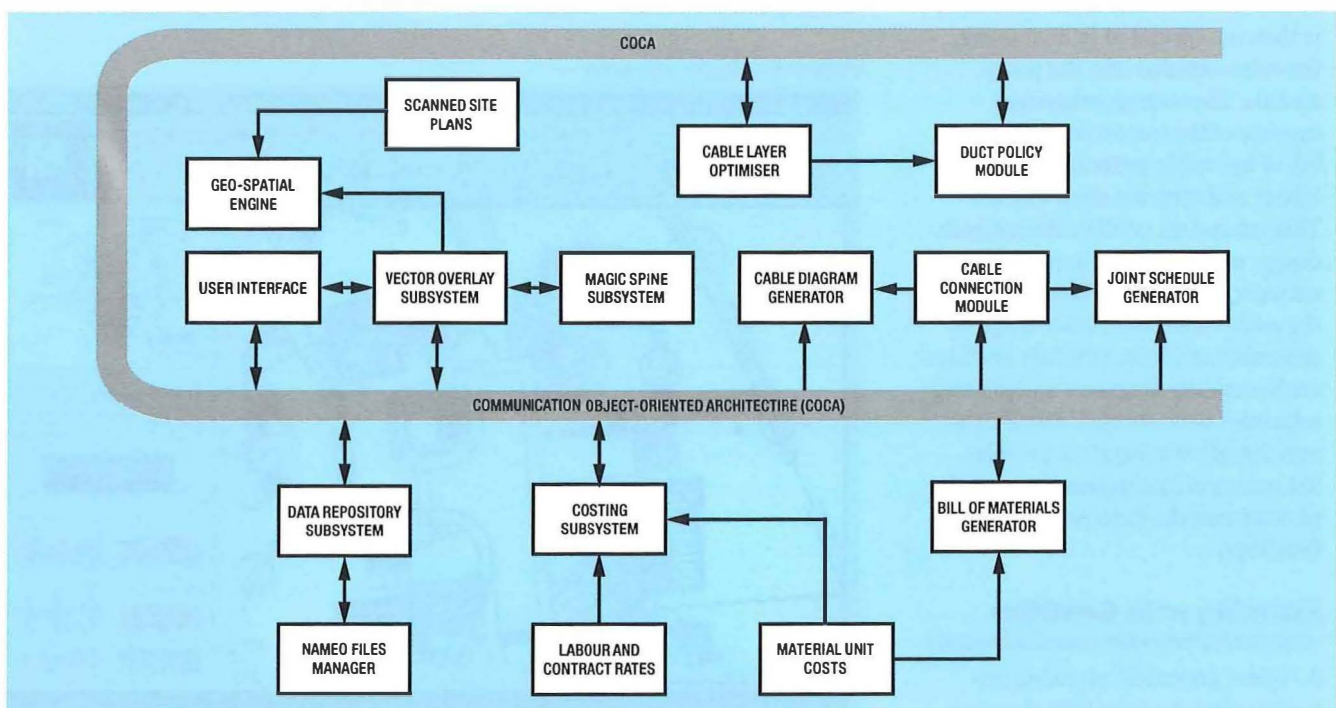


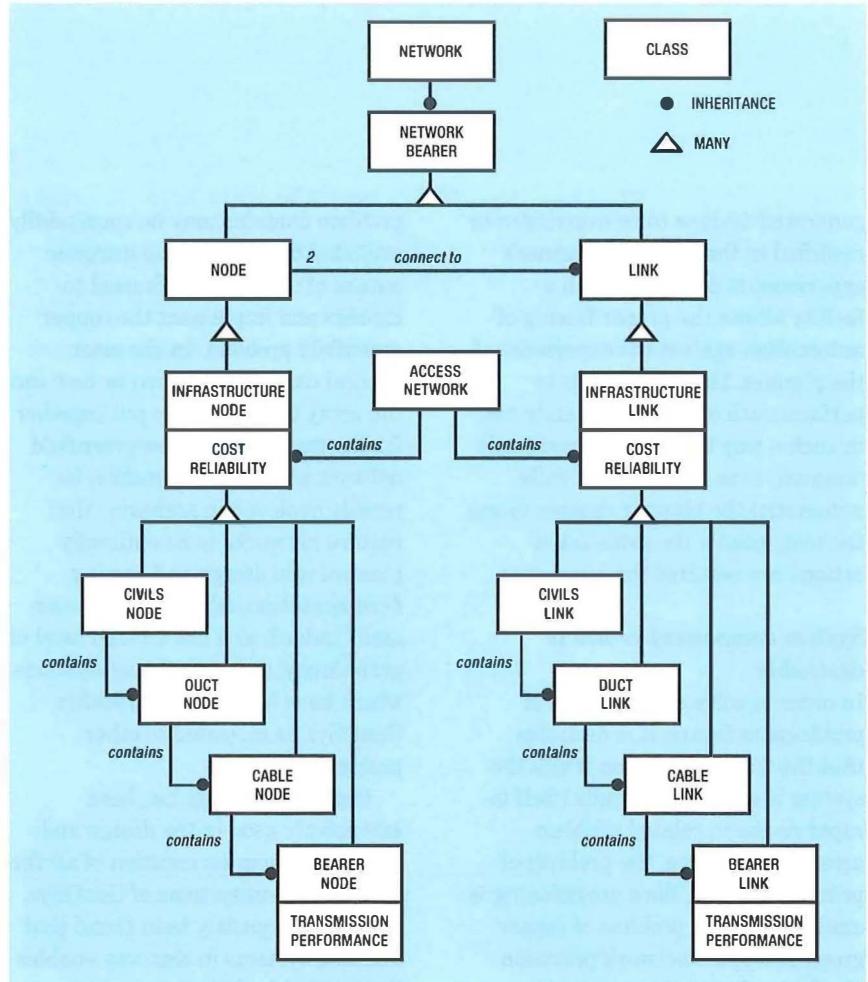
Figure 7—COCA

then setting the north orientation for the plan. Against the raster image of the site plan, the vector layer of the scheme can be overlaid using a mix of planner experience and automation. The vector layer is modelled using a set of meta classes (defined by COCA) which serve as a true representation of the plant items used in the access network. Figure 7 illustrates this abstraction. In concrete terms, the vector layer consists of items on the civils, duct and cable layers. The civils layer is entirely placed by the network planner and it consists of the geographical position, demand level (in voice circuits), and type of each customer together with the trench network that provides connectivity between each customer and the site entry point. The trench network consists of track sections along with the intermediate hole or pole positions that enable node access at the duct and cable layers. The duct, cable and bearer layers are discussed in the section on 'The Generic Greenfield Planning Process'.

The cable layer optimiser subsystem takes civils layer data from the COCA pipeline and applies a genetic algorithm to it in order to return data about the cable layer. This layer is then encapsulated in duct using the rules encoded into the policy module. The costing subsystem consists of the materials unit cost, bill of materials generator and the labour and contract rates module. This subsystem enables the network design to be costed at component or network level. On the bearer layer, the cable connection, cable diagram generator and joint schedule modules enable cabling diagrams and jointing schedules to be created. Finally, the user interface subsystem provides the means of interaction between the planner and the functionality of GenOSys.

**Planning with GenOSys**

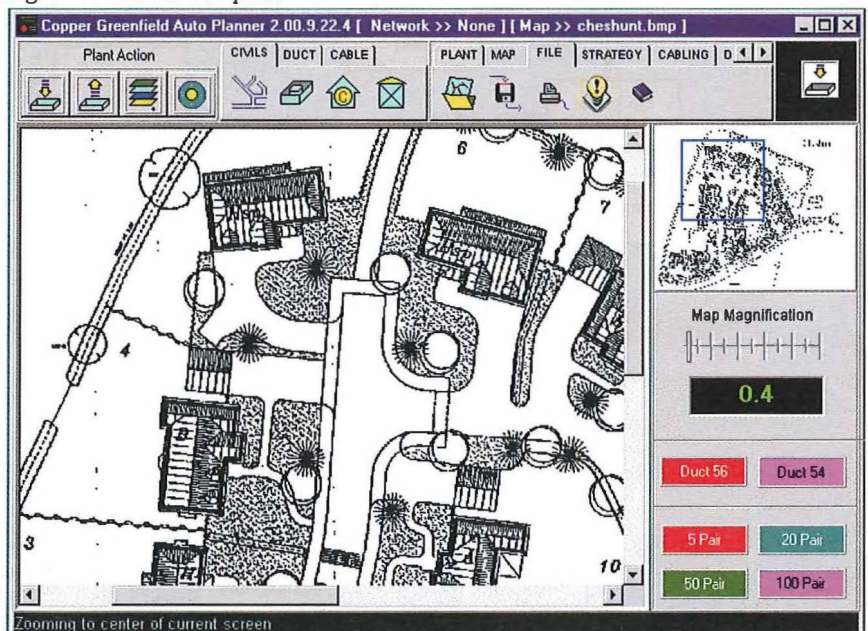
A typical greenfield planning scenario<sup>4</sup> using the GenOSys planning



tool would begin with a site proposal from a developer. This proposal would then be reduced to A4 (using a photocopier) if necessary and scanned to create a raster image. This raster image is then loaded into the main viewing window of GenOSys as illustrated in Figure 8.

This raster plan now needs to be registered onto the Ordnance Survey grid so that points on the plan correspond to their real-world positions. This is done in three steps. Firstly, as part of the GenOSys planning process a 5 cm line is placed on the developer's plan. This line is

Figure 8—Raster site plan








Figure 9—Setting up spatial geo-referencing

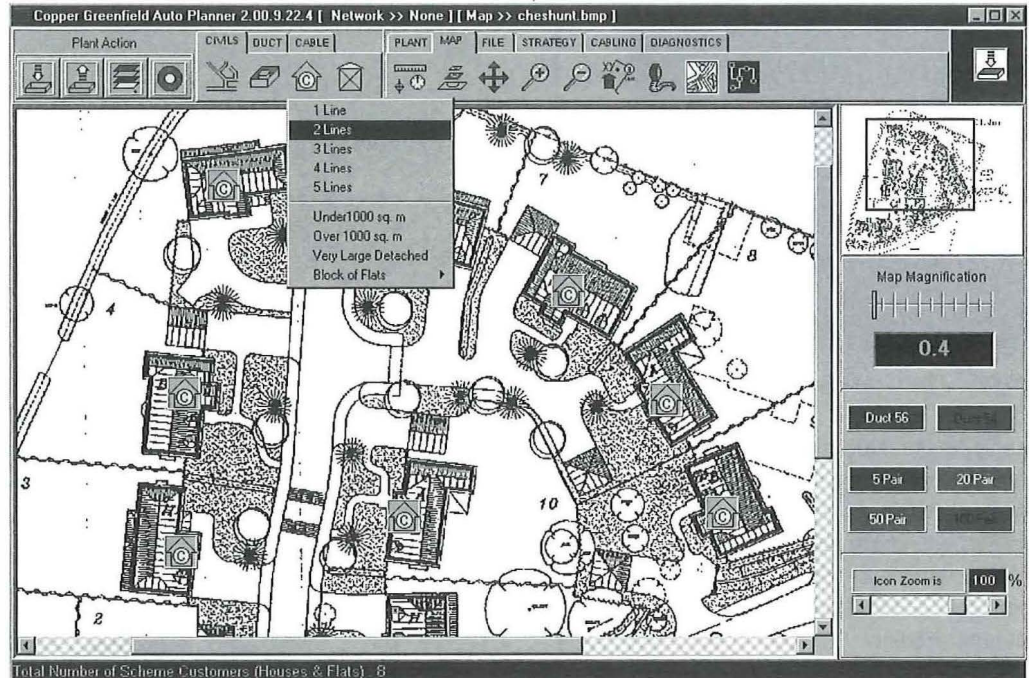


Figure 10—Spotting customers on a registered raster plan

overlaid in the software with a line that sets the scale of the plan. Secondly, the orientation of the north is set and finally, the scale of the map is set. Together these steps ensure that each piece of plant located on the map has a corresponding real-world geographic location. All these steps have been illustrated in the screen shot of Figure 9. Next, the greenfield planner would spot the customers on the proposed site, indicating the levels of demand (in 64 kbit/s channels) required by each of the customer locations. Options for both residential (flats and houses) and business customer levels of demand can be selected by menu option. Figure 10 shows how this is done. Next a civils layer (trench) network is laid out. The guidelines for trench links specify that they should follow the pavements on either side of the proposed roads on the site.

GenOSys enables this process by providing a *magic spine* civils link which enables the planner to lay a track on the pavement outside the housing units and then the system completes the layout of civils links and nodes.

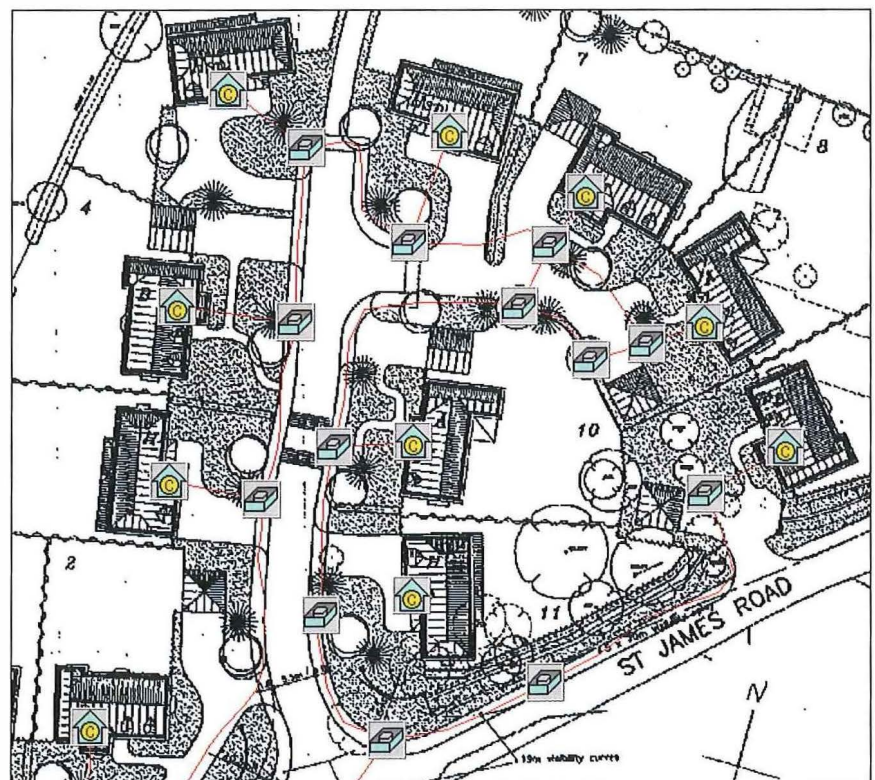
Civils nodes using the magic spine are placed closest to each of the customer locations and then connected as shown in Figure 11.

The positions of the nodes and links placed can easily be modified if

required, either by moving the nodes or by removing and re-laying links. Since civils links can be placed along any desired path, then a precise route can be followed in defining the trench. For trench nodes, each customer has a node left on the frontage area. Other positions for nodes are at road crossings and between long runs of trenching. If purely automatic planning is going to

be carried out by the planner, then a dense network of civils trenches can be laid down. This ensures that GenOSys has a wide range of alternatives to choose from when optimising for the best network configuration. At this point, the planner can either let GenOSys do the network topology optimisation for the duct and cable layers or carry on to finish the design of the duct and

Figure 11—Defining the civils network layout



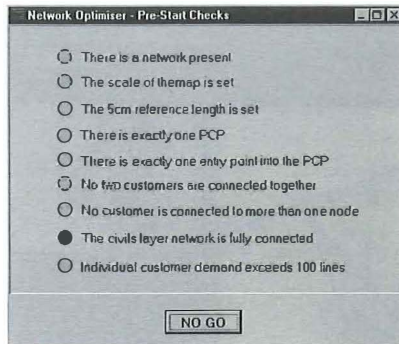


Figure 12—GenOSys pre-start checks

cable layers manually. Prior to any optimisation being carried out, several criteria must be met by the civils layer configuration. If any of these conditions (Figure 12) are not met then network optimisation cannot proceed.

For instance, full connectivity of the civils layer is essential. While this is currently detected in software, correction for gaps needs to be visually done by the planner. Figure 13 illustrates the view provided on the network to enable this process. If GenOSys performs the optimisation (Figure 12) of cable and duct layers, the planner still has the option of overriding the topology by manually placing items on the duct and cable layer. Placing links and nodes in the duct and cable layer follows the same drag-and-drop metaphor used for the civils layer. Each of the layers (civils, duct and cable) of a planned network can be independently viewed by selecting that layer into the main viewing window. Figure 14 shows how this is done.

Once a network has been planned, several key outputs can be created. These outputs can form part of a proposal that is passed to the building developer, or part of a job pack sent out to the field implementation team. A complete breakdown of the costs for the greenfield network build can be obtained on request as well as a bill of quantities which highlights the stores items required, their quantities and store codes.

Geographic duct diagrams (Figure 15) as well as schematic cable diagrams (Figure 16) can be produced from GenOSys. At the bearer layer, a jointing schedule is produced by the system on the basis of the distance of each customer from the site serving PCP.




Figure 13—Network overview

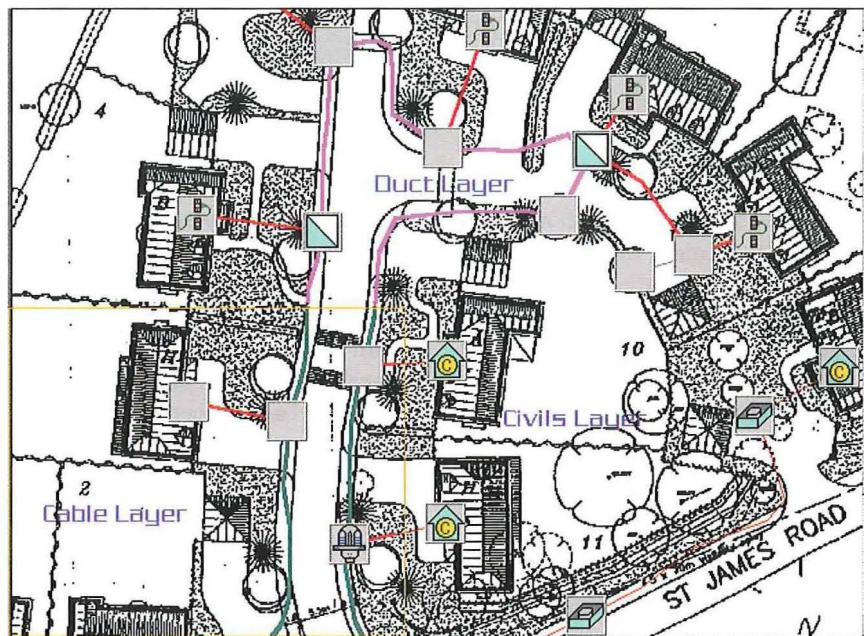


Figure 14—Layered network view

## Summary of Business Benefits and Lessons Learned

GenOSys was placed on a test-bed at Crayford, South London, for initial measurement of benefits in October 1997. At the end of this testing in February 1998, GenOSys had exceeded every single baseline measure of productivity, with average productivity increases of about 500 per cent being reported<sup>2</sup>. This translated into a complete clearing of the backlog of waiting greenfield proposals in south London. The system was then piloted in Southampton, Canterbury and Crayford in March 1998. Confirmation of the test-bed results led to the national roll-out of the manual version in September 1998. The fully optimising

version of GenOSys has just completed extended trials at Southampton. The reported savings in material costs are of the order of 15 per cent. This version is scheduled for national roll-out in November 1998.

Through the development and deployment of GenOSys, several important lessons have been learned:

- **Risk taking** Key technology decisions made by the product delivery team have been supported throughout, even when these decisions meant utilising products and applications not previously accepted or widely used within the BT development community. The freedom to sometimes make unconventional choices has contributed significantly to the success of product delivery.

Figure 15—Geographic duct diagram



set up and give the end users a chance to see how well this fits with their perception of what the final product will look and feel like. By surfacing enough stable functionality early in the product-build cycle, many course corrections can be made incrementally and continuously. This process can take place within a controlled test bed and/or enthusiastic product champions across the company. End users often use a product in ways not always possible to predict in advance. These ways of working can sometimes yield insights into how to make a particular feature or feature set easier to use. It can also enable problems not caught in normal testing to emerge. The process of early exposure also enables any misunderstanding of system requirements or specification to be exposed and handled. The importance of user involvement has also been identified in an earlier article in this series<sup>5</sup>. Such involvement concurs with known industry best practices such as rapid prototyping and concurrent engineering<sup>6</sup>.

- *A modular product architecture is highly advantageous* The ability to respond quickly to change in the short term and to evolve the basic product in the long term is a function of the modularity of its architecture. Several times, on the way to the final delivered product, there was the need to make changes to GenOSys either on the basis of its feature set or because unforeseen requirements arose. In each case, the modular architecture framework, based on object-oriented technology made it feasible to allow change within a reasonable time frame.
- *Use national roll-out as an opportunity to accept new product issues* When a software product of any reasonable complexity is shipped, there can be no absolute guarantees of functionality. This is well expressed in the end-user licence agreement statements of

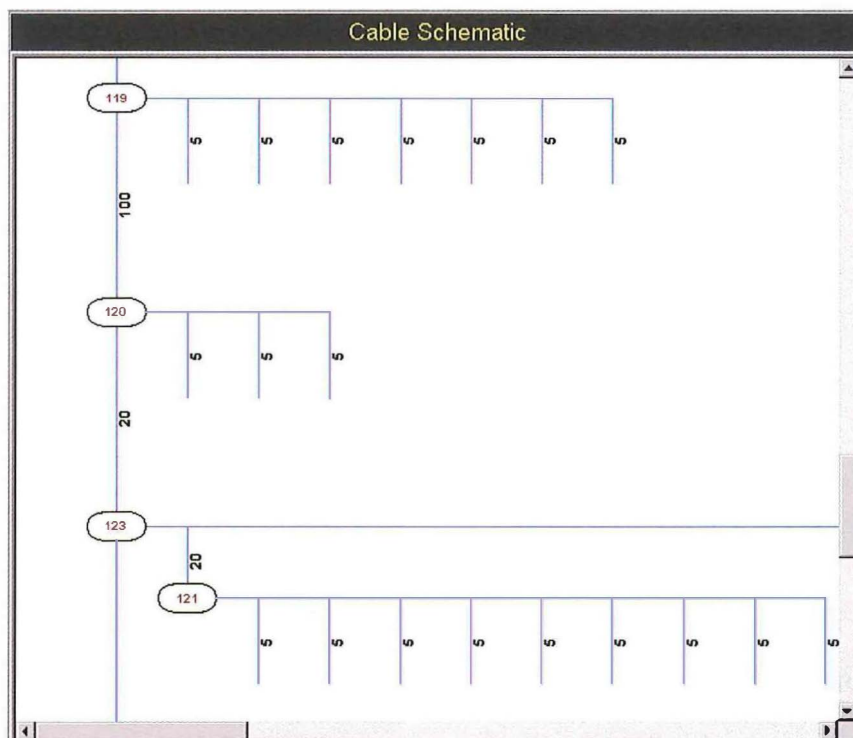


Figure 16—Cable schematic diagram

- *Clarification of administrative and product development responsibilities are needed at the beginning* The financial administration of the overall product delivery has to be integrated with the process of creating the product, but responsibilities for the other aspects need to be made clear from the outset. In the case of GenOSys, product development, resourcing and awareness issues were managed within the unit. Budget allocation and tracking, end-user training, user issues, road shows, test-bed and pilot administration were externally managed by the customer's representative.
- *End users are a key part of the development process* End users should be involved in the product-build process especially when the early prototypes have been stabilised. Many planners know what they want in a system, however, the initial difficulty is in expressing this in terms that can be translated directly into a product. Early prototypes enable 'straw men' to be

practically every shrink-wrapped software product on the market. In summary, fitness for purpose does not imply the complete absence of defects. As a product begins to be widely used, its true complexity makes itself known in the form of problems (major and minor) which were not caught in the stages leading up to roll-out. This does not imply negligence on the part of the design and build team, the subsystems and systems testers, test-bed or pilot users. It is simply an expression of the level of complexity of intellectual property that has been captured within the product. Therefore, roll-out and the subsequent issues and problems that arise as end users push the product's performance envelope represent a good opportunity to continue the evolution of the product. It is essential to maintain user confidence in the product by providing a means for them to express dissatisfaction about any aspects of it. For GenOSys, a Web site was set up soon after national roll-out. This site, along with a helpdesk is the primary means through which grievances can be aired by the end users and a response provided by the build team. This ongoing process assures that defects within GenOSys are clearly and openly acknowledged even as they are being resolved.

Finally, the GenOSys optimiser has brought about, through the reduced cycle time of the planning process, the prospect of planned greenfield sites on-demand. This will bring forecast and real demand closer together in time and reduce the need for costly rework. This could turn out to be a key gain in line with the ideas of structure process improvement.

### Acknowledgements

The author would like to thank Anthony Conway, Colin Birchenall and Danny Poon for the design and build of key subsystems of the

GenOSys tool; Steve Brewis who championed the product since its inception; Gilbert Wardrop and Derek Fred for their assistance in bringing the product to a successful completion.

### References

- 1 ASUMU, D.; and MELLIS, J. Performance in Planning: Smart Systems for the Access Network. *BT Technol. J.*, Oct. 1998.
- 2 BT ISIS Document NWK/LNK/C306. Access Narrowband Greenfield Planning Rules. 1997.
- 3 WARDROP, G. The Network Optimiser—Test Bed Closure Report. March 1998.
- 4 COBB, KEN; ROSE, MARK, and LAWRENCE, IAN. Structured Greenfield Planning. *Br. Telecommun. Eng.*, Jan. 1999, **17** (this issue).
- 5 BUTLER, DAVID. Implementing Software Tools for Planning Automation. *ibid.*, Oct. 1998, **17**, p. 126.
- 6 HOOPER, RAY; and JONES, CAROLE. Lean Operations—The Route to Competitiveness. *ibid.*, April 1998, **17**, p. 2.

### Biographies



**Don Asumu**  
Networks and  
Information Services  
BT UK

Don Asumu graduated from the University of Nigeria with a first-class honours degree in Electronics Engineering in 1982. He worked for Shell Petroleum Development Corporation as an instrumentation engineer for a year and then joined Schlumberger Wireline Services as a geophysical field engineer. He subsequently obtained an M.Sc. in Microprocessor Engineering and Digital

Electronics in 1986 from the University of Manchester Institute of Science and Technology (UMIST). This was followed by a Ph.D. in Artificial Intelligence also from UMIST in 1988. We worked as a research fellow and then a lecturer at the Manchester Metropolitan University until 1993. He joined BT in 1995 and currently leads a team involved in software delivery for future plan/build processes.

Don Asumu can be contacted at don.asumu@bt.com.



**John Mellis**  
Networks and  
Information Services  
BT UK

John Mellis leads the Planning, Plant and Pilots development unit within BT Systems Engineering. The unit specialises in the planning and launch of new network solutions for a variety of customers in BT, in the UK and globally. Current projects include the planning of new access, IP and data networks for BT UK, and European joint ventures, and the development of a range of expert-system-based network planning tools. He joined BT Laboratories in 1986, having previously worked on electro-optic signal processing with Nortel Technology (then STL), Harlow. From 1989–91, he was with BT&D Technologies (now HP Optoelectronics), commercialising optical amplifier and laser products. He received his Ph.D. from the University of St Andrews in 1983, for the mathematical modelling of CO<sub>2</sub> laser dynamics, and graduated with a first-class degree in Applied Physics (with Operational Research) from the University of Strathclyde in 1978. He is a member of the IEEE, sits on the technical advisory committee of the European Conference on Networks and Optical Communications, and is currently engaged in the BT Masters programme leading to a management degree in Telecommunications Business from University College, London.

*Christine Rigden*

# 'The Eye of the Beholder'— Designing for Colour-Blind Users

*Colour-blind computer users see things differently from most people, but this is seldom considered in the design of software or Web pages. This article offers a simple technique for simulating the differences, to help designers and developers avoid disadvantaging their users.*

## Introduction

Have you ever wondered how different the world could seem to someone with a colour-vision deficiency? Have you ever wondered whether they can understand your Web page, or follow the colour cues in your software?

If it is not your problem, you have probably never considered it. However, one man in 12 has some degree of colour-vision deficiency, or colour blindness (it is rare in women). The colour choices we make in software and Web design can make life more difficult for these people, if we are uninformed.

In good design, colour should never be the primary cue for information. The options should be clear without colour, and the colour is simply added as a means of emphasis.

Yet, colour becomes more and more a part of our everyday work, especially with computers increasingly supporting a greater range of colours. Often it is simply decorative, but colour cues are frequently used as the primary means of informing the user how to interpret information.

Published algorithms enable us to calculate how colours appear to individuals with certain types of colour vision deficiency. However, the mathematics is fairly complex, and designers and developers need an easier way to see how their choices impact colour-blind users.

This article offers a method for representing images as perceived with a colour deficiency, so designers can produce interfaces which avoid causing problems for their users.

Most colour-blind people have a deficiency with either red or green, and by focussing on the worst case scenario for these groups, maximum inclusiveness can be obtained.

'The Basics of Colour Vision' introduces how colour vision works, in the simplest terms. 'Problems with Colour Vision' looks at how colour vision deficiency happens. The section on 'The Colour Transformations' explains the theory behind the palettes, 'Using the New Palettes' describes their application, and 'Postscript on Colour Displays' looks at some of the technical issues with monitors. The 'Glossary' covers the basic colour-vision terms, and words which appear there are in **bold italics** when they are first used.

## The Basics of Colour Vision

We perceive colour via photosensitive cells in the eye, the rods and cones. The cones normally contain pigments tuned to receive wavelengths in three parts of the visible spectrum (Table 1). This gives us **trichromatic** vision. They can be referred to as **Red**, **Green** and **Blue**, although  $\rho$  (rho),  $\gamma$  (gamma) and  $\beta$  (beta) are sometimes used to avoid confusion with the common understanding of the colour names.

The wavelengths given are for peak receptivity, although they also receive light spread 100 nm or more either side of the peak, to a diminishing extent. Rods receive light at around 500 nm, but only as lightness.

The brain further processes the colour received by the eyes, and there are also psychological issues, but this

**Table 1 Comparison of Cone Names and Visible Spectrum Colours**

Name	Alias	Symbol	Wavelength <sup>1</sup>	Colour	Proportion
rho	red	<i>r</i>	590 nm	yellow-orange	60%
gamma	green	<i>g</i>	550 nm	yellowish green	30%
beta	blue	<i>b</i>	440 nm	bluish violet	10%

article deals only with the initial stage.

**Problems with Colour Vision**

In a person who is a *dichromat*, one of the pigments is missing. This is usually red or green, although sometimes blue. A quarter of colour-blind people are dichromats.

For the other three-quarters, either the red-sensitive cones have had their tuning shifted towards the green part of the spectrum, or the green peak is shifted towards the red wavelengths. All three sets of cones still work, so they are still trichromats, but because

of the shift in sensitivity are referred to as *anomalous trichromats*. The visual effect covers the continuum from normal trichromatic vision to dichromatic vision, so the common expression 'colour blind' is less appropriate to this group. (See Figure 1.)

Eight percent of Caucasian men are colour blind<sup>1</sup>. This is made up of 1% red-blind (*protanope*) and 1.1% green-blind (*deutanope*) dichromats; 1% red-insensitive (*protanomalous*) trichromats, and 4.9% green-insensitive (*deuteranomalous*) trichromats. Only 0.002% of men are blue-blind (*tritanopia*—tritanomalous condi-

tions are not known) and 0.003% of men are totally colour blind. This article does not deal with these rarer conditions.

The red and green forms of colour deficiency are caused by a recessive gene on the X chromosome, and are therefore more common among men. Only 0.4% of women have any sort of colour vision deficiency, most of them being the red or green forms.

**The Colour Transformations**


The work described in this article relates to the most common deficiencies with red and green. These deficiencies are also very similar to each other, in terms of colour perception. It also deals with dichromatic vision, as that is calculable and represents the 'worst case' scenario. If one designs for dichromats, the lesser degrees of colour deficiency will be well accommodated.

Interestingly, a red or green deficiency does not simply affect how one perceives red or green. Any colours that vary from each other by the amount of red or green they contain will also appear different. In addition, because protanopes are less sensitive to light at the red end of the spectrum, colours in this area appear darker to them. Deutanopes do not see this luminosity difference, as the other cones and rods compensate.

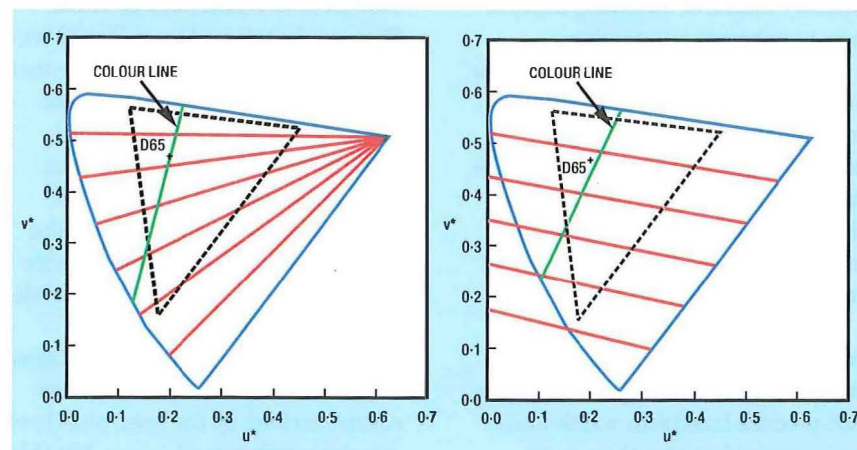
Another interesting aspect is that colours perceived by protanopes and deutanopes are broadly similar, with just a slight greenish cast to the *protan* palette, and an orangey cast to the *deutan* palette.

In colour science, the visible spectrum can be represented in a diagram known as a *CIE chromaticity diagram*†. In Figure 2, the outer triangular shape represents visible wavelengths. The lines that are roughly horizontal through this space (confusion lines) indicate the colours that dichromats get confused. Every

*Figure 1—Relation between trichromats and dichromats*



*Figure 2—The CIE  $Y_u^*v^*$  colour space, with confusion lines for protanopes (left) and deutanopes (right) (included with permission from HCI Resources Network)*



† The colours displayed on a computer monitor are a subset of this space, indicated by triangles in the diagrams.




Figure 3—The look of the Web-safe palette, and how it compares to the transformed palettes

colour along a given line looks the same as all the other colours along that line. The point at which the Colour line crosses them represents the perceived colour for that confusion line.

The position of any colour relative to these lines can be calculated. The relation between colours in a CIE chromaticity diagram and colours for given monitor phosphors is known, so RGB equivalent values can be generated for a monitor with known chromaticity coordinates<sup>3</sup>. It can, in theory, be done with as many colours as one cares to calculate (within the gamut for the monitor), but the work in this article is based on the standard 16-colour† web-safe palette<sup>4,5</sup>.

Values were calculated according to established algorithms<sup>3</sup>, and palette files were created from the

Standard

Protan

Deutan



Figure 4—The transformation when applied to the 16-colour palette

results in the \*.pal format (as used by Paintshop Pro). Figure 3 gives a visual representation of how the palettes compare† for the 216 colours. Figure 4 shows the corresponding transformation of the old Windows 16-colour palette.

A striking feature of these two transformed palettes is how similar they are. The transformations are all roughly divided into yellows/ochres, and blues. The protan palette seems to have the slightest greenish cast, whereas the deutan palette has a slight orange cast. Both red and green vanish from the palettes, as well as purples and oranges.

One can work out easily whether a given RGB value falls into the yellows or blues group. In general, if the Red value is greater than the Blue, it appears in the yellows group, and if Red is less than Blue it appears in the blues group. A colour generally appears neutral when Green equals Blue, more or less regardless of the value for Red.

The transformed Web colours are available as \*.pal files and/or \*.gifs which can be downloaded from the colour pages at <http://www.labs.bt.com/people/rigdencc/colours/>. Certain pages within the site also 'list' the Web colours visually, alongside their transformed colours, so that a direct comparison may be made. They are shown all together, and are divided into blues, yellows, and neutrals as well.

## Using the New Palettes

There are two basic approaches to using these palettes:

- apply them to images, to approximate how they appear to protanopes and deuteranopes; and
- use the understanding of the transformed palettes to select colours which work together safely for the Web site's navigational scheme.


For the first approach, the basic Web-safe palette is loaded or applied to an image or screen capture. Save the files into a convenient directory. In PaintShop Pro, choose **Colors - Load palette...**, and select *web-safe.pal*. 'Nearest colour' should be used for flat-colour images such as gifs, and 'error diffusion' for photographs†. (Other commercial paint programs have similar functions.)

Two copies are then made of the image. To one copy, apply *protan.pal* in the same way, but this time select the dialogue option 'Maintain indexes'. Apply the *deutan.pal* file to the second copy. The resulting images are a visual approximation of how it would look to a protanope or deuter-


† This will also show how the image will look when viewed on a 256-colour monitor, if it has not already been checked.

† There are 216 colours considered 'safe' to use on Web pages, because they are common to most of the computing platforms available. When other values are used, any system running 256 colours will substitute the specified colours for those it has available. Backgrounds will change colour, and flat colours in gifs will dither. Using the Web-safe palette gives better visual control over what the end user sees. The web-safe palette includes all colours (and no others) which have RGB values made up of 0, 51, 102, 153, 204, and 255 (or 33, 66, 99, CC, FF in Hex).

‡ Please note that printed colour may look quite different from screen colour, and these printed images will not be a totally accurate reproduction of how the colours look on screen. For more clarity, see the images on the Web site.



Using standard palette



Using protan palette



Using deutan palette

Figure 5 – Example transformations of a photograph

anope. Figure 5 gives an example of how this would work on a photograph.

It is important that the Web-safe colours are in their correct index positions, in order to correlate with the transformed palettes. If in doubt, re-map to the Web-safe palette supplied with the others.

For the second approach, the designer starts with the sort of colour scheme they are considering, based on the non-functional requirements of the site, aesthetics, and the general ‘feel’ being sought. They can then look at the corresponding colours in the transformed palettes, and use that information to help them choose the most suitable colours in the primary palette which will also ‘work’ in the transformed palettes.

For example, for the colour site where these files are kept, a primarily monochromatic colour scheme was chosen. By using a pale colour for the navigation panel, a colour could be chosen for un-followed links which was contrasting enough to be readable there as well as in the main page. To be easily distinguished from a green un-followed link, the followed links colour was chosen from among cooler colours, while still toning with the main colour scheme. Figure 6 illustrates how these compare.

The colour differences in the lighter shades are marginally more visible to anomalous trichromats, while the mid-tone and darker shades are more consistently confused.

Multiple colours will create greater difficulty in finding shades


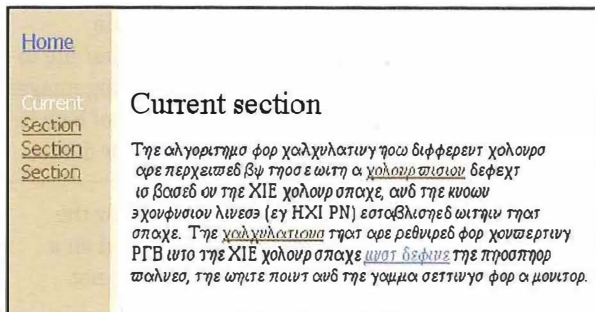
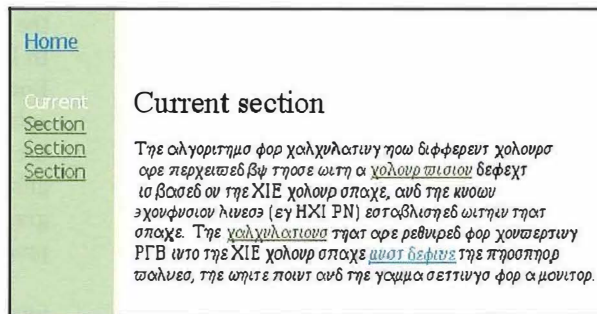
that are adequately different from each other when transformed. Early work suggests that only up to four colours may be chosen which can be clearly differentiated. When more colours are used, make sure their differentiation is less important.

Keep in mind cultural meanings also, when choosing how to represent information with colour.

The two essential points to remember are:

- Make sure that any text has a good luminosity (brightness) contrast with its background, for readability. Do not put a mid-tone against a mid-tone. Large areas of text for reading need to ensure greater contrast than text in a navigation bar.

Figure 6 – Sample colour selection based on a monochrome palette, showing how it looks normally, followed by how it appears with the protan (bottom left) and deutan (bottom right) palettes





- Make sure that the colours chosen for followed and un-followed links are far enough apart. It is safest if they fall on either side of the colour division in the palette, so that one colour is transformed into the blues group and the other one is in the yellows group.

Then double-check their effectiveness by taking a screen capture of your choices (as in Figure 6), and applying the transformed palettes to them.

## Postscript on Colour Displays

A prevalent concern about doing this kind of work in the past has been the lack of precision involved. The colours on every computer monitor are going to be slightly different, according to the gamma setting and the colours of the RGB phosphors. This makes it impossible to come up with a single 'true' RGB value for any specific real-world colour, and makes it difficult to manage accurate colour reproduction on screen.

However, in practical terms, RGB is used to define colour in software and on the Web, without concern for the chromaticity values. Modern monitors are becoming more similar than they used to be, as manufacturing becomes more standardised. Therefore, in spite of a lack of real precision, these transformations are a valid and practical means of ensuring that an interface does not cause problems for colour-blind individuals.

## References

- 1 HUNT, R. W. G. *Measuring Colour*. Ellis Horwood Ltd, Chichester, 1987.
- 2 TRAVIS, DAVID. *Effective Colour Displays: Theory and Practice*. Academic Press, 1991.
- 3 HCI Resources Network. *Colour Vision Deficiencies*. Retrieved May 1998 from the World Wide Web: <http://www.hcirn.com/>
- 4 WEINMANN, LYNDIA; and HEAVIN, BRUCE. <Coloring Web Graphics.2>. New Riders Publishing, 1997. Browser-safe colours also explained on Lynda's Web page at <http://www.lynda.com/hex.html>.
- 5 SIEGEL, DAVE. *Images on the Web—Use the Netscape Colour Cube*. Retrieved from the World Wide Web, Nov. 1998: <http://www.dsiegel.com/tips/wonk10/images.html>.
- 6 POYNTON, CHARLES. *Poynton's Color FAQ: 19. Is RGB always device-dependent?* Retrieved from the World Wide Web, June 1998: [http://www.inforamp.net/~poynton/notes/color\\_and\\_gamma/ColorFAQ.html#RTFToC19](http://www.inforamp.net/~poynton/notes/color_and_gamma/ColorFAQ.html#RTFToC19).

## Glossary

**trichromat** a person who has three colour reception pigments functional to some degree.

**dichromat** a person who has only two of the three pigments functional.

**anomalous trichromat** a person who is a trichromat, but one pigment is dysfunctional, or anomalous.

**pro-** Latin for 'first', and referring to red (first of the trio RGB).

**protan** relating to problems with red reception.

**protanopia** condition where the red (first) pigment is missing, 'red blind'.

**protanope** dichromat with protanopia.

**protanomaly** 'red insensitive' anomalous condition.

**deu-** Latin for 'second', and referring to green (second of the trio RGB).

**deutan** relating to problems with green reception.

**deuteranopia** condition where the green (second) pigment is missing, 'green blind'.

**deuteranope** dichromat with deuteranopia.

**deuteranomaly** 'green insensitive' anomalous condition.

**tri-** Latin for 'third', and referring to blue (third of the trio RGB).

**tritan** relating to problems with blue reception.

**tritanopia** condition where the blue (third) pigment is missing, 'blue blind'.

**tritanope** dichromat with tritanopia.

**tritanomaly** theoretically the 'blue insensitive' anomalous condition, but the condition is not known to exist.

## Biography



**Christine Rigden**  
Networks and  
Information Services,  
BT UK

Christine Rigden is part of BT's User Centred Design Group. Her work involves all aspects of designing interfaces for usability, but particularly developing guidelines for the conceptual design of software user interfaces and Web sites. She also offers interface evaluations and design tutorials. Christine joined BT in 1988 as a Drawing Office assistant, working in graphic design. She was sponsored to the BT Women's Bridging Course (Chelmsford) and subsequently to the University of Essex as a mature student, and graduated in 1994 with a B.Sc. (Hons.) in Information and Business Systems Technology.

# Evolution of the BT UK Core Transport Network

*To meet increasing demands for capacity, the BT UK transport layer will see the rapid introduction of advanced synchronous digital hierarchy (SDH) and wavelength-division multiplexing (WDM) technology. Unit costs will similarly reduce so that total costs remain within the investment envelope.*

## Existing Transport Network

The BT UK core transport network refers to the underlying plesiochronous digital hierarchy (PDH), synchronous digital hierarchy (SDH) and wavelength-division multiplexing (WDM) networks which support the various switching platforms (Figure 1) and routes capacity in 2–140 Mbit/s blocks between exchanges. Currently about 50% of network capacity is PDH which is continuously being replaced and upgraded by a software-managed SDH transport network. The current BT SDH network structure (Figure 2) is layered, with the long haul and regional networks largely deployed and the metropolitan and access SDH networks being rapidly rolled out.




Figure 1 – Transport and switch platforms

## Drivers for Change

A number of clear drivers for evolution can be identified:

- The BT transport network will require about  $\times 10$  increase in capacity within 5–10 years to support increasing volumes of traffic from broadband and Internet services. Broadband requires bandwidth in the range 2–140 Mbit/s, and both require (on average)

Figure 2 – BT SDH network

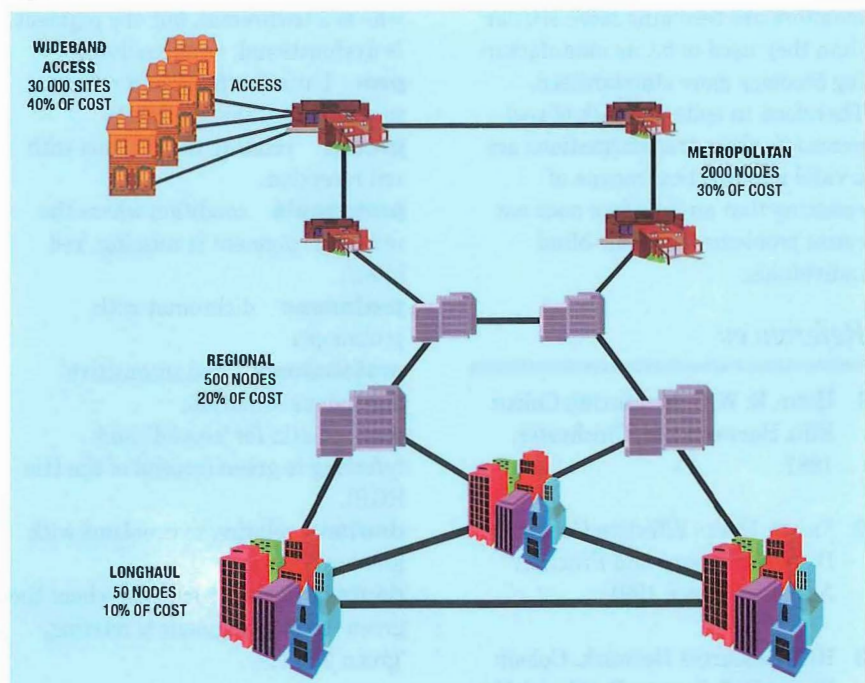


Figure 3 – Rapid technology introduction


holding times measured in hours rather than the 3 minutes typical of today's telephone calls. This is also evident for international circuits where (for example) TAT14, due to land in the UK in 2000, will require more than 300 Gbit/s capacity from the international backhaul network (which provides capacity between London, Madley and the frontier stations).

- To support these volumes, within the annual capital budget for transmission, unit costs (for example, cost per 2 Mbit/s circuit) must similarly reduce by  $\times 10$  within the 5–10 year time frame.
- BT will need to provide large chunks of  $n \times 140$  Mbit/s capacity at short notice for start-up companies. This is already happening in London so both high capacity and flexibility through software control are essential.
- Network resilience to cable hits (usually generated by utilities such as electricity and gas suppliers) is essential to restore service within milliseconds before the customer 'sees' a fault. This requires more advanced and cost-effective SDH SPRing and WDM technology, described later.

To meet this growth, BT will introduce high-speed technology through its suppliers, upgrading the SDH network from 2.4 Gbit/s today to 10 Gbit/s SDH in 1999 (Figure 3). This will be combined with 16/32 wavelength WDM systems providing  $16 \times 2.5 = 40$  Gbit/s in 1999 and  $32 \times 10 = 320$  Gbit/s capacity systems by year 2000. If the capacity increase is less than expected then this will simply provide more time to introduce the technology and new planning and operational practices.

### Advanced SDH Technology

We now consider a step-by-step approach to reducing unit costs



starting with roll-out in 1998 of advanced SDH technology. Within the core (layers 1–2), use of advanced SDH *shared protection rings* (SPRings) is reducing unit costs by replacing deployment of large SDH cross-connects by less expensive SDH add-drop multiplexer (ADM) rings (Figure 4). These have less flexibility than cross-connects but are easier to manage and offer high-speed protection. The payback period for SPRings is typically less than 2 years (Figure 5), compared with growing the mesh. SPRings are being extensively rolled out to carry private services and broadband ATM traffic.

A typical 2.5 Gbit/s SPRing (Figure 6) has eight working and eight spare 140 Mbit/s blocks (VC4s) per link which, in the event of a cable cut, are automatically re-routed around the ring within 50 ms using standby

routes from the 'pool' of spare capacity. Restoration from a cable break in an SDH mesh is slower; taking up to 5 min, owing to the extensive signalling needed between the element manager and the network control layer to locate the break, compute the restoration plan and make good using whatever spare capacity can be found.

SPRings have about 50% utilisation (ratio of live to total capacity), which can be increased by also carrying 'bumpable traffic' such as public switched telephone network (PSTN) traffic. This utilisation compares well with dedicated protection rings at 25%.

The unit cost reduction using advanced SDH products compared with the current SDH/PDH mix is estimated as 30% with 2.5 Gbit/s

Figure 4 – The price of flexibility




Figure 6 – Shared protection ring

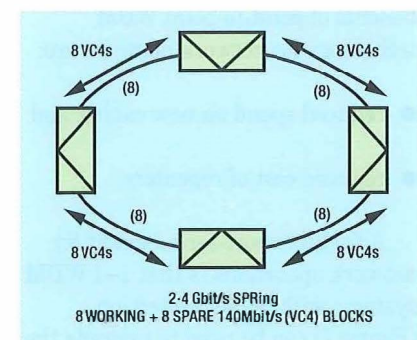


Figure 5 – SPRing 'payback' in two years




Figure 7 – WDM point-to-point system (protected)

SDH SPRings and a further 40–50% with 10 Gbit/s SDH SPRings.

Overall cost reduction using advanced SDH is approximately  $\times 3$ . (Target  $\times 10!$ )

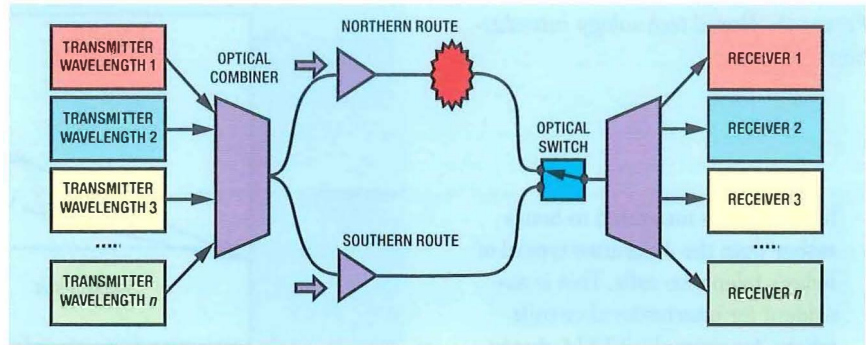
Clearly SDH alone is unlikely to produce the required cost reduction of  $\times 10$  within the 5–10 year timeframe and we must turn to WDM for further reduction in unit cost.

### Near-Term WDM Technology

In its simplest point-to-point form, WDM allows the capacity of existing optical cables to be increased by the number of wavelengths used (existing optical systems use one wavelength) at the expense of additional WDM multiplexing at each line system terminal. Also, optical amplifiers replace regenerators so that repeater costs are shared over each wavelength. WDM systems generally cost in (with respect to SDH) for line systems greater than about 80 km in length, although unamplified metropolitan WDM systems cost in at much shorter spans. Hence the immediate benefits of point-to-point WDM technology for network planners are:

- reduced spend on new cables, and
- reduced cost of repeaters.

A further short-term benefit for network operations is that 1+1 WDM systems with 50 ms protection (Figure 7) can be used to upgrade the performance of the legacy PDH network. One route carries live traffic while the other is on 'hot standby' ready to switch over if needed. Initially 1+1 WDM systems will be complementary to SPRings (which also have 1+1 protection) and deployed mainly at the network periphery or on special long routes. WDM systems currently use transponders to convert signals from legacy PDH and SDH client systems to a standard wavelength grid before launching onto the fibre. Transponders allow interworking of various



vendors' systems but add about 50% to the total cost of a WDM system. Hence there are strong drivers to reduce the use of transponders where possible using open interfaces being defined within standards committees.

The first two 16-wavelength WDM pilot systems (Figure 8) will be introduced in the BT inland network by Lucent (Carlisle–Belfast) and Ericsson (Newcastle–Edinburgh). These were chosen to increase cable capacity and currently carry a mixture of PDH 565 Mbit/s and SDH 622 Mbit/s and 2.5 Gbit/s systems. The Carlisle–Belfast route uses two submarine systems in the Irish Sea where capacity is particularly at a premium.

WDM may also be embedded into SDH, for example, increasing SPRing capacity up to 160 Gbit/s.

### A Managed Optical Platform

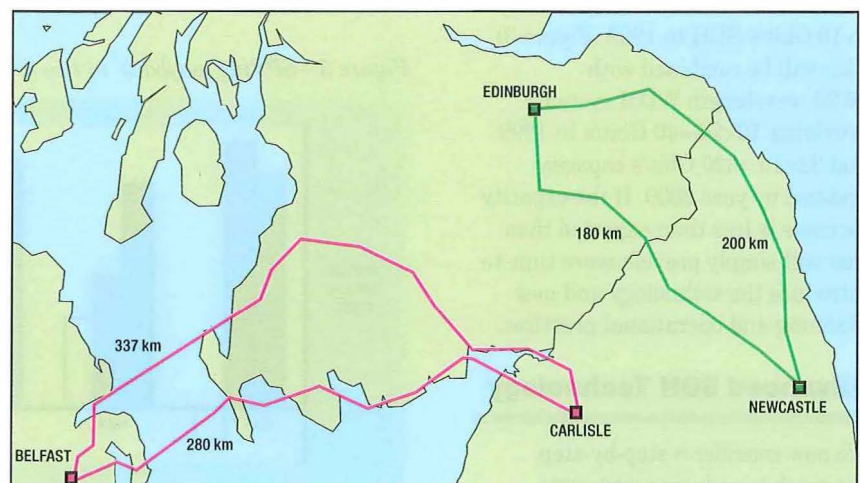
Whereas the WDM point-to-point systems have a specific use in reducing cable spend, the goal is to introduce a managed optical platform, containing wavelength add-drop multiplexing (WADM) and optical cross-connect (OXC) network capability. The new platform will support a large collection of client platforms including SDH,

ATM, Internet protocol (IP), etc. This ability to support multiple platforms transparently is unique to WDM and provides telcos with an opportunity for large cost reductions (and resilience to cable breaks and equipment failures) as traffic levels climb to  $\times 100$  current levels in 10+ years.

An example (Figure 9) shows a WADM system routing individual wavelengths (or blocks of wavelengths) between nodes on the ring and providing efficient 1+1 fast protection in the event of cable cuts or node failures. This is equivalent to a small optical network (up to about 800 km periphery on first product releases in 2000) and provides a relatively low risk way of introducing the WADM technology into the network because:

- All of the route lengths are well defined and designed to be within system limitations such as dispersion and crosstalk for up to 10 Gbit/s operation and about 32 wavelengths.
- Although the WDM technology can be described as 'state of the art' it does not contain any radically new technology (for example, large OXCs).

Figure 8 – WDM pilot routes



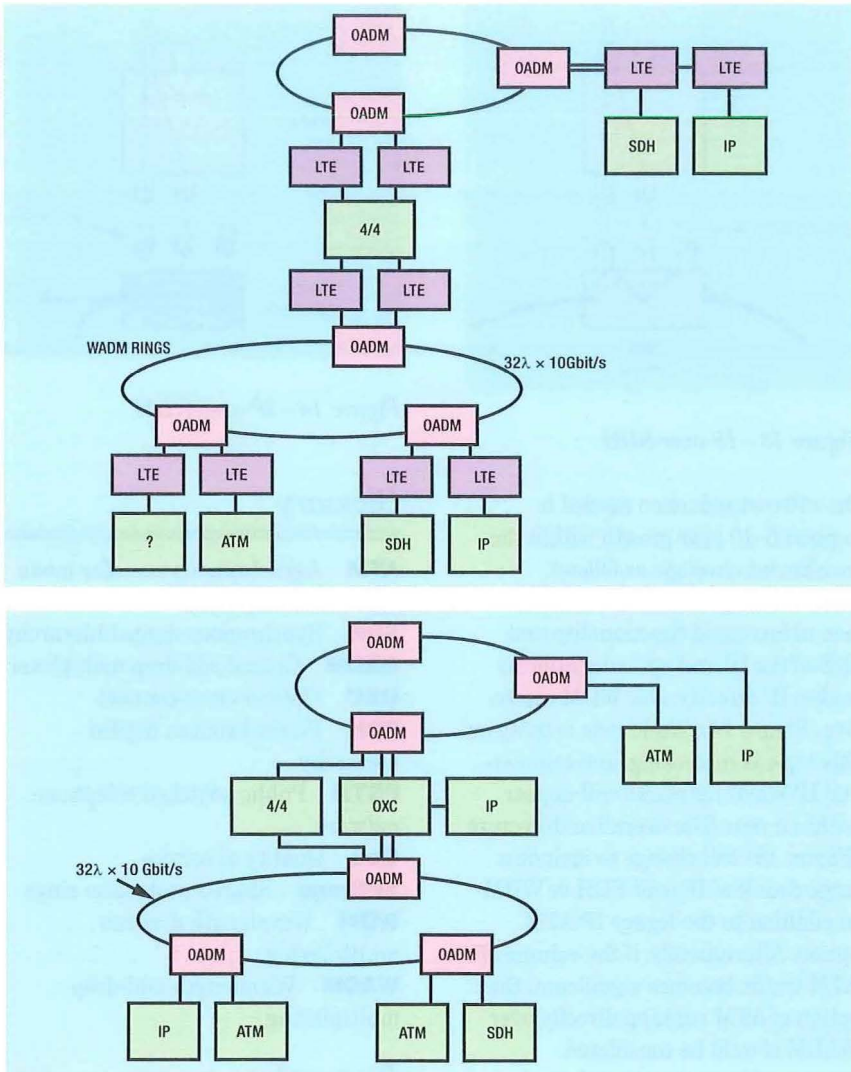


Figure 10—Building optical networks

The gateway between WADM rings is an SDH cross-connect which provides 140 Mbit/s traffic grooming. Later this can be supplemented by a small OXC allowing *express wavelengths* to be routed between rings (Figure 10). The OXCs are likely to use transponders to regenerate the optical signals between WADMs rings and reduce the impact of system constraints across the optical network. This network is effectively an optical multi-service platform supporting a mix of client platforms.

In the longer term (10+ years), when traffic levels have increased by  $\times 10$ – $100$ , optical networking can be expected to reduce unit costs for transmission by greater than 10 times compared with traditional PDH/SDH solutions, mainly because it can handle large blocks of capacity cost effectively. Within 5–10 years a 50% reduction in unit costs with WDM is assumed, so that combining cost reductions from

both advanced SDH and WDM gives an overall reduction compared with the current SDH/PDH mix of about  $\times 5$ .

Figure 11—Too many platforms!




Figure 9—Elements of an optical network

## Simplified Network Architectures

In order to meet the  $\times 10$  reduction in unit costs in the core we need to take a wider look at platform architectures to see if there are opportunities to reduce the number of platforms needed to carry broadband, IP and other services. Figure 11 shows an example of the current platform architecture where a typical service such as IP or frame relay may be transported over a stack of four–five other platforms. Generally this situation arises naturally as new platforms are rolled out over legacy platforms, but a simple cost analysis shows how cost reductions can be obtained by doing things differently.

Figures 12–14 show three options for transporting IP services. Figure 12 shows the use of the existing multi-service ATM platform to transport IP, which in turn runs over the SDH and cable/WDM platform. This is a reasonable option when IP volumes are relatively small and where good quality of service (QoS) (provided by the ATM layer) is needed from the beginning.

However when IP volumes grow to equal those of (say) PSTN the question arises ‘do we really need three platforms to transport IP?’ Why not route

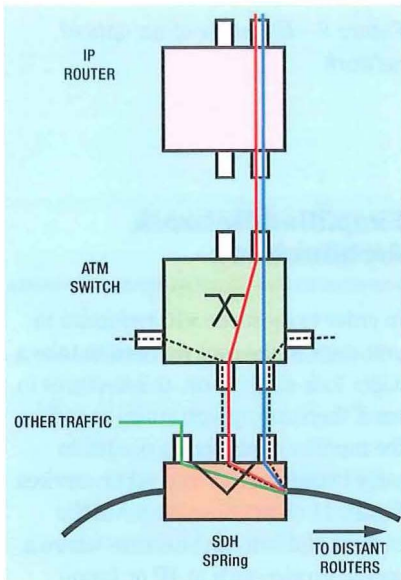


Figure 12—IP-over-ATM-over-SDH

IP directly over SDH (Figure 13), missing out the ATM platform and using increased functionality IP routers to provide QoS? The majority of telcos support this alternative, which becomes increasingly cost-effective as volumes increase (Figure 15). So, looking once again at unit cost reduction, we see that the combination of (1) advanced SDH, (2) WDM and (3) simplified platform architectures described above is expected to provide

Figure 15—Scalability of platform options




Figure 16—New network architectures


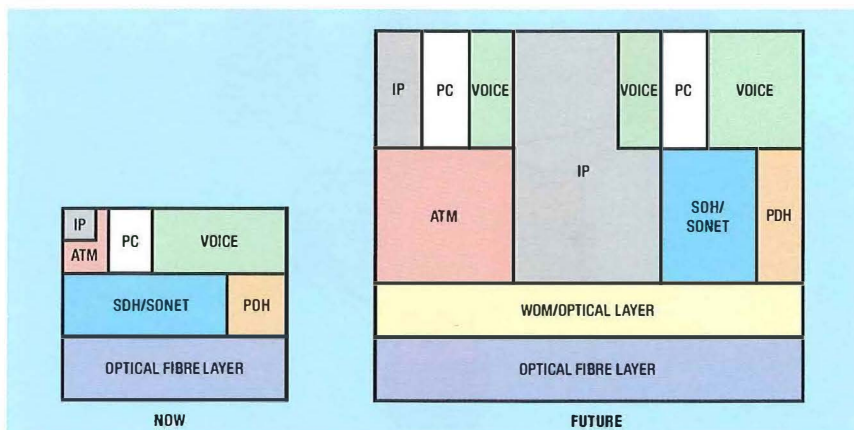


Figure 13—IP-over-SDH

the  $\times 10$  cost reduction needed to support 5–10 year growth within the investment envelope as follows.

In the longer term, the combination of increased functionality and QoS of the IP and optical platforms makes IP directly over WDM attractive (Figure 14). Worldwide activity on this topic is increasing, and commercial IP/WDM networks will appear within a year. The overall architecture (Figure 16) will change to include a large chunk of IP over SDH or WDM in addition to the legacy IP/ATM option. Alternatively, if the volume of ATM traffic becomes significant, the option of ATM running directly over WADM should be considered.

### Summary

In order to meet increasing demands for capacity by about  $\times 10$  in 5–10 years the BT UK transport layer will see rapid introduction of advanced SDH and WDM technology. To remain within the investment envelope, unit costs will similarly need to reduce in the same period, taking managed risks to introduce new technology quite rapidly.

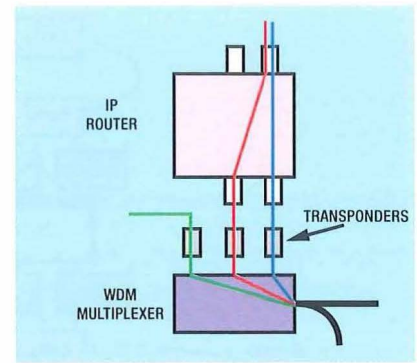


Figure 14—IP-over-WDM

### Glossary

- ATM** Asynchronous transfer mode
- IP** Internet protocol
- SDH** Synchronous digital hierarchy
- OADM** Optical add-drop multiplexer
- OXC** Optical cross-connect
- PDH** Plesiochronous digital hierarchy
- PSTN** Public switched telephone network
- QoS** Quality of service
- SPRings** Shared protection rings
- WDM** Wavelength division multiplexing
- WADM** Wavelength add-drop multiplexing

### Biography



**Ian Hawker**  
Networks and  
Information Services,  
BT UK

Ian Hawker is a  
group leader at  
BT Laboratories  
investigating

future transport networks on behalf BT inland and international backhaul networks. This includes cost analysis of emerging technologies, such as WDM and SDH products, and development of core transport network strategies. He joined BT in 1978 and worked on the development of submarine systems and 565 Mbit/s inland transmission systems. In 1988, he became involved in SDH transmission network studies including design of SDH networks and associated issues of network reliability. In 1990, he became leader of the Future Networks Transport team and is currently involved in planning of SDH and WDM networks in BT UK.

# Video-Based Virtual Clubs

## Using Multimedia to Re-present Technology to Women

*The Virtual Clubs project is providing a novel solution to the problem of how to attract more women into IT. The goal of the project is to engage young women's interest by changing their perceptions of what technology can offer them. Funded by a BT University Development Award, six secondary schools are linked to universities via desktop videoconferencing, using collaborative technology and Internet applications to create a virtual club. Young women are being enabled to acquire a range of new skills in creating and sharing multimedia information, thereby enhancing their perceptions of their own abilities in using technology. The project is creating a networked community of young women who might otherwise be excluded from the technological future.*

### Introduction

The *Virtual Clubs* project was conceived as a new approach to the problem of how to widen women's access to computing courses in further and higher education (FE and HE). The aim is to create an environment in which young women can be engaged with technology through the appeal of multimedia and the Internet, and through acquiring confidence in their own computer-related skills. It is intended that this will lead to a greater participation by women in computing courses at university.

The issue of women's low participation at all levels of IT has been widely debated over the last decade, but remains a persistent problem despite efforts by the Women into Computing (WiC) and Women into Technology (WiT) movements. Computing workshops organised by universities, for example, have had little impact on the numbers of women entering computing courses<sup>1</sup>.

Motivations for not giving up on the situation remain strong, however. Recent EU policies on social inclusion, the shortage of IT skills in the UK, and the ever-increasing role of technology in business and in the home, all provide good reasons for encouraging more women to take up computing careers. A speaker at a recent seminar noted the UK women's falling interest in entering a career in computing, at a time when the UK has the highest growth in IT-related business in Europe. Of British Computer Society members, only 9.2% are female, this falling to 4% as Fellows<sup>2</sup>.

The project described in this article takes a new approach to the

problem, based on a set of requirements identified from the literature and the authors' experiences with WiC workshops. The solution to the problem of involving more women in computing requires the following:

- engagement of women's interest in interacting with technology;
- a change in women's perceptions of their own abilities to interact effectively and creatively with technology;
- a change in women's perception of computers, their users and IT careers; and
- a change in knowledge about the capabilities of computers and computer-mediated communication.

The method adopted by the project to meet these requirements is to employ real-time multimedia conferencing, Internet software tools, and a long-term plan to help teenage girls develop a range of skills. The success of the project will be judged by measuring participants' attitude changes in relation to the above requirements, using questionnaires and other techniques.

Virtual Clubs began in October 1997 and the project is funded for 18 months by a BT University Development Award. The collaborating institutions are Queen Mary and Westfield College (QMW)—the lead partner, the University of Sussex, and the College of West Anglia†. Approx-

---

† The College of West Anglia, a further education institution, was formerly known as the Norfolk College of Arts and Technology

mately 50 students from six secondary schools† are currently participating in the project. Sites are connected over integrated services digital network (ISDN) for videoconferencing and Internet-based activities.

## Background

Explanations of women's low participation in IT have varied over the last decade. Initially ascribed to girls' low access to computers at school, the situation appears now to be more complex than originally thought<sup>1</sup>. Recent research indicates that the standard of female computer literacy is still far behind that of their male counterparts, despite increasing investment in computers, networking and Internet access in schools<sup>3</sup>.

Some studies suggest that girls' performances are held back by gendered expectations about their computer-related abilities, although these results are mixed<sup>4</sup>. Another reason may be girls' perceptions of computing applications and their users. For instance, media talk of 'anoraks' and 'nerds' populating the Internet may be creating a false image of the typical IT enthusiast that is unattractive to girls. The very language of computing itself is perceived to be masculine territory<sup>5</sup>. Terminology used routinely to describe the state of programs, such as 'abort', and 'crash', may be off-putting to female novice users.

### Women-into-computing workshops

Over the last decade, several UK universities, including QMW, have organised campus-based workshops for teenage girls, hoping to influence their study choices at A-level and higher.

† The schools participating are: Central Foundation School, London; Priory School, Sussex; Litcham, Methwold, Rosemary Musker, and St Clements schools, all in Norfolk

\* A list of Web sites specifically targeting teenage girls and women is given in the References section

Workshop programmes have aimed to engage participants' interest through a programme of girl-friendly activities, and to boost their self-confidence when interacting with computers.

The limitations of the workshop approach are:

- *Frequency and duration:* Workshops tend to be annual events, lasting only a few days at most; enthusiasm engendered at the workshop may quickly dissipate.
- *Context:* Campus-based workshops may be perceived by girls as 'outings', and unrelated to their normal, everyday activities at school.
- *Engagement with technology:* Workshop programmes normally include speakers, introductions to the university, and discussions, limiting time for actual engagement with technology.
- *Teacher involvement:* School staff, who are in the best position to motivate young women in schools on a daily basis, may not be aware of WiC issues and the aims of a workshop.

The workshops at QMW employed questionnaires, completed at the start and end of each event, to measure attitudinal changes. Data collected in these circumstances tends to be subjective in nature and difficult to validate. Most seriously, there is no opportunity to assess the persistence of attitude changes and impact on the number of women entering computing courses at HE level.

The limitations of the workshop approach, the resources required to organise them, and the lack of tangible results<sup>1</sup>, have led to new initiatives taking a different approach.

### Recent initiatives

In the US, the National Science Foundation's Partnership for Advanced Computational Infrastructure (PACI) programme is making efforts to

incorporate women and minorities into the high-performance computing community<sup>6</sup>. Several projects are ongoing, working with female and minority students at all levels. A particular theme of these projects is *persistence*; that is, providing support for women's progress from one level of education to the next.

Attempts to counter the perceived masculinity of cyberspace can be seen in World Wide Web sites such as Purple Moon, Cybergirl, Girls Clubhouse, and GIRL Tech\*. These initiatives can benefit women who have already acquired the skills and self-confidence to explore the Internet. They do not address, however, the problem of engagement, or how to involve women proactively in the future development of the world of technology.

### A New Solution: Virtual Clubs

The Virtual Clubs project takes a different approach, based on exploiting the benefits of multimedia conferencing. Multimedia potentially offers many advantages as a replacement for campus-based workshops, both in terms of persistent support and in the appeal to young people of interacting in multimedia environments. Project objectives are:

- to increase young women's interest in and engagement with technology;
- to change their perception of computers and their users;
- to provide a persistent, multimedia environment and support from a remote expert, to enable young women to acquire a range of computer-related skills;
- to change participants' perception of their own computer-related abilities; and
- to provide subjective and objective measurements of changes occurring in young women's attitudes.



An important element of the project is to develop a plan to enable these objectives to be carried forward after the end of the funding period. This is discussed later in this article.

### The virtual club concept

The concept of a 'virtual club' is being used to present desktop video-communication to teenage girls at school. This novel metaphor is meant to convey a more casual style of interaction than videoconferencing. It encapsulates the notion of an informal environment in which activities designed to reflect members' interests are organised, led by a club leader or facilitator. In the physical world, local clubs include programmes of events, cooperative development of skills, and talks from invited guests. The school club is a well-established component of the educational system in the UK, and provides a useful basis for students' participation in the project.

A virtual club links members in distributed sites, but the goal is to provide a computer-supported environment with features similar to a real-world club. The project's virtual club for young women is known to them as the *Video-Club*, and is an extension of each school's local computer club. The Video-Club's approach incorporates several key features:

- As video-conferencing takes place in the context of an informal school club, not the class-room, participation has a social dimension.
- Students participate in scheduled weekly sessions during which they have access to a club facilitator over a desktop videoconferencing system, and utilise shared applications to review progress and discuss new ideas.
- Students are encouraged to progress their ideas between scheduled sessions. To this end, details of activities undertaken during the scheduled sessions are

recorded in the Video Club's Web site† so that students can have easy access to the information and resources they need at any time.

- Students are encouraged to take increasing responsibility for the coordination of a Web design-related project and share ideas with other members. In this way the club will self-perpetuate in schools well beyond the funded 18-month period.

An important aspect of Video-Clubs is that participation is regular and ongoing. Weekly sessions with the club facilitator enable students to demonstrate their designs, discuss their ideas, request demonstration of new techniques, or ask for help with any difficulties encountered since the previous session.

### Multimedia conferencing support

Internet-based chat rooms and e-mail are popular with club members, but visual communication has proved to be essential for the purposes of the project. Currently, videoconferencing is enabled by a commercial desktop package over a 128 kbit/s ISDN connection. During conferencing, Windows-based applications can be shared using collaboration software.

The videoconferencing product currently in use is adequate for point-to-point connections, but multipoint conferencing over ISDN has not been used, for several reasons. The relatively high cost of conferencing via a multipoint control unit has been one factor. More significantly, fixed-format multiplexed images are not useful for realising the concept of a socially-oriented virtual club. A more flexible way of organising video windows, and support for multiple group activities, would provide a more appropriate environment.

An alternative solution would be to use Internet-based video applications supporting group conferencing. Internet tools, however, are not yet able to deliver the quality of service required for this specific application,

where high-quality audio is essential in parallel with shared applications and video.

### Development of an advanced group video application

The concept of a virtual club, where one or more groups of users are engaged in informal conversations and shared activities, has motivated a new research project called *Socialspaces*\*<sup>7</sup>. Advanced video encoders, supporting transmission of up to 32 simultaneous video channels, are providing an infrastructure for an experimental virtual club in which multiple groups of users communicate in real time. A prototype application based on the virtual club metaphor is being developed, and will be evaluated using the Video-Club students.

### Video-Club Activities

Each scheduled session typically involves a link between one school and a facilitator at QMW. The students use shared applications such as Web editors, browsers and graphics applications to demonstrate progress made since the previous session. This enables students to gain confidence from starting the session at their own level of understanding of their current situation. It also requires them to become familiar with the various controls available in the videoconferencing system, such as enabling application sharing.

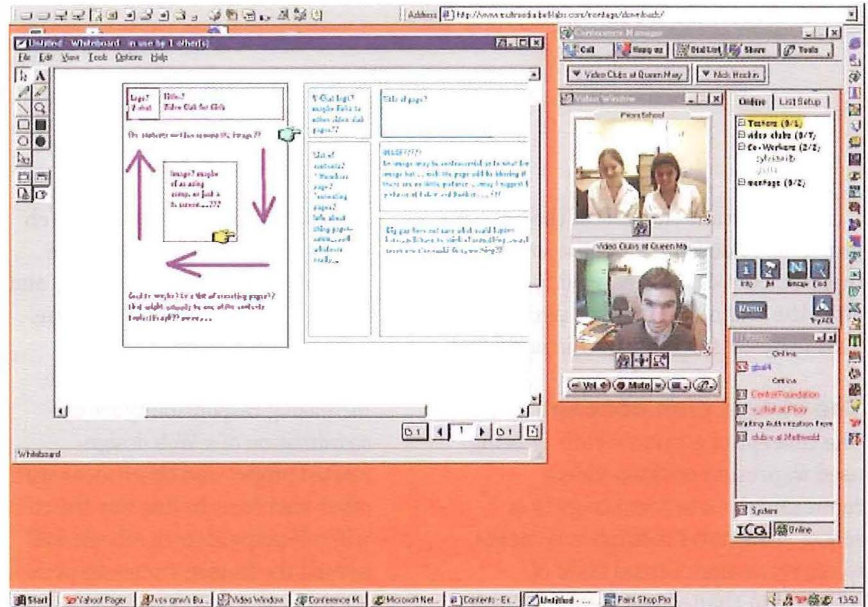
At this point the facilitator can help with any difficulties encountered, and students can allow the facilitator to take control of their document to demonstrate a technique. Once any difficulties are resolved, the students and facilitator can interact on the shared whiteboard to brainstorm ideas for further development. For

---

† The Video-Club's home Web page can be found at: <http://www.dcs.qmw.ac.uk/~vcs>

\* *Socialspaces* is a collaboration between QMW and Bell Labs Lucent Technologies

Figure 1 – A typical club video-conferencing session



example, students may ask how they can achieve a particular effect seen on other Web pages, such as animation. These requests can often be responded to immediately or scheduled for the next session. A typical session is shown in Figure 1.

Each student is initially encouraged to develop her personal Web page and experiment with design ideas for it. More recently, the girls in each school club have progressed to more substantial Web development projects of their choice. Projects now underway include designing an on-line gallery, Web-based school magazines and providing a topic-based guide to the Web tailored to the needs of examination subjects in their own schools.

As particular groups within one school may become involved in a project similar to that undertaken in another school, the need to collaborate has led to schools arranging conferences directly with each other. Such collaboration is positively encouraged, and over time is intended to become the primary form of club conference, with less need for scheduled sessions with the facilitator, although interest and support will nevertheless still be present.

The first generation of club members within each school is now beginning to pass on the benefits of their experience to a new generation. The Video-Club is starting to extend to a network of girls who are using technology confidently for communication and self-expression.

## Results of the Project to Date

Approximately 50 girls are involved in the project across the six schools. Numbers vary from school to school. Most of the students are currently in Year 11 (aged 15–16), but there are also a number in Year 10 (aged 14–15). Since September 1998, a Year 12 group (aged 16–17) has also joined the project in one of the schools.

An important by-product of the project has been the opportunity for IT staff in the various schools also to

develop new skills, learn from QMW how to configure their systems, and generally receive a level of support which might otherwise be lacking. The support of committed IT teachers in each school has proved to be crucial in getting the project going and maintaining momentum. Each school has its own approach to providing access to IT facilities in lesson time and at lunchtime or after school. Such practice is likely to impact on the time students have to experiment with technology and consequently may influence their wider perceptions of computing.

## Measuring changes in womens' attitudes to computing

The impact of the project is being measured by attitude surveys, which will provide subjective data over a period of time. For various technical reasons, some schools joined the Video-Club later than others. All participants completed a questionnaire designed to measure their attitudes to computing before their first videoconferencing session. The attitudes of a control group without exposure to the clubs were also measured. A second survey is carried out nine months after each school has joined the club and a final questionnaire will be completed at the end of the project.

The results of the first attitudinal survey reveal that there are substantial differences between each of the schools in terms of prior experience in using technology and, in particular, accessing the Internet. In one

school for example, no one had ever used the Internet (in school or elsewhere), yet, in another, many students had Internet access at home and communicated regularly with 'on-line friends' using instant messaging services such as ICQ, as well as e-mail and chat.

Throughout the course of the project, a range of such factors has emerged which may play a crucial role in forming students' attitudes towards computing:

- *Single-sex schools:* One school is attended only by girls, which means that students have no 'competition' with boys for access to equipment or attention.
- *Sixth-form presence:* The existence of A-Level or GNVQ courses in IT or computing within the student's own school presents a natural transition from compulsory to further education, particularly when a teacher familiar to the students is known to be teaching the course.
- *Awareness of staff:* Sensitivity to gender issues in computing can play a crucial role in students' perception of the subject but cannot be systematically observed or controlled. Perhaps even the gender of IT teachers and technicians may have some impact on student's perceptions.
- *Students' over-commitment:* High expectations of students and their

parents for participation in extra-curricular activities has led to some students becoming over-committed in their membership of school clubs. In consequence, these students have no time to progress ideas between scheduled Video-Club sessions.

- **Student ability:** In mixed-ability groups, some students may find club activities either too demanding or too trivial, impacting their engagement with the technology and their ability to interact with it.
- **Access to equipment:** Availability of technology affects opportunities for students to engage in club-related activities between scheduled sessions.

While each of these factors may impact on the findings of the attitudinal survey they are difficult to isolate or control, particularly as each of the factors manifest themselves in different permutations, variations and extents in each of the six schools. Consequently with such a small sample, it is not feasible to account for these in any objective, measurable way. These issues are nevertheless worthy of note and further exploration.

### Data analysis

Initial analysis of the results from two schools (12 subjects) who have completed the second survey indicates that the activities and opportunities offered through club membership have been attractive to students. Averaging across the two schools, members agree that using the Internet is exciting (80%), designing Web pages is fun (88%) and videoconferencing is enjoyable (81%). As a consequence, the number of students who have made use of the Internet several times a week has risen from an average of 8% of members prior to joining the club to an average of 74%.

Over the first nine months of the project, there was an increase in student's perceptions of their own ability to interact with technology. At

one school for example, 82% of members at the outset were unsure of their ability to undertake further studies in computer science. Now 75% believe they have the capability to do so.

This perceived increase in ability had been accompanied by a change in attitude towards further studies in computing beyond GCSE. For example, at the start of the project only 6% of members in one school expressed an interest in further study but, after 9 months, 38% were interested in studying computing after leaving school and 88% were interested in studying IT at A-Level. (This compares with interest in study at A-Level in English, 63%; Maths, 38%; Media Studies, 50%.)

These results are encouraging. Results from the four other project schools will further inform our statistical enquiry, and the findings of the third survey at the end of the project will provide a longer-term view of attitudinal change, including comparison with a control study.

The impact of the project on women's entry to computing courses will be measured by monitoring the participants' choices of study at further education (FE) and higher education (HE) at the appropriate time.

### Case Studies

The following case studies illustrate how membership of virtual clubs is raising young women's awareness of the relevance of technology to their personal goals and ambitions.

Nasima Begum (Year 12) is a keen artist and aims to display her work in a public gallery. She had rarely used the Internet prior to joining the video club. However, she has now started to develop an on-line gallery with a classmate. They seek video-mediated assistance from the club facilitator whenever needed.



Nasima is impressed by the fact that she is using technology she had previously only heard about on television.

The high quality of her scanned images and the ease of manipulation have impressed her. The skills she has acquired in Web page creation have given her the confidence to consider pursuing her artistic talents in the realms of Web design.



Violet (Year 12) has always enjoyed using her family's PC at home but had never used the Internet before. She has developed a strong interest in computing and is determined to learn more about graphics, appearance and advanced Web page design. She is keen to pass on her new skills to other young women.



Lucy (14, Year 10) is excited about their new Video-Club project, which is investigating on-line shopping. Since joining the club, she accesses the Internet regularly, and understands more clearly what it has to offer. Lucy has become familiar with the use of videoconferencing, and control of the software.

The case studies above are typical examples of the enthusiastic young women taking part in the scheme.

### Benefits of Multimedia Conferencing

The benefits of using desktop multimedia conferencing to link schools to university sites for video-clubs are as follows:

- the overhead of organising workshops and girls' travel is avoided;
- regular, weekly contact is easily achieved;

- use of software tools can be demonstrated during conferencing;
- the sessions are student-focused, giving them some sense of control and ownership over the technology and the media; and
- visual communication adds a new dimension to girls' interaction with technology.

Using multimedia presents technology in a new light to students whose experience of computing has been limited mainly to utilitarian needs such as word-processing. This re-presentation of computers as tools for communicating with others is an essential element of the project. Computing is no longer a private activity, but a means of interacting with the world beyond the school environment. Access to the Internet and construction of personal Web pages extends that interaction through the creation and sharing of information with others.

Typically, a cluster of girls share a single video-enabled PC. Grouping two or more girls informally round a machine in this way seems to help create the feeling of a club. This practice is also in line with a 1996 study by the University of British Columbia which found that girls tend to work better in clusters when interacting with technology<sup>8</sup>. In line also with educational psychology, it satisfies 'the need for learners to participate within rich communities of understanding, and to partake of the collaborative experience of learning'<sup>9</sup>.

Visual communication is particularly useful when several students share a terminal, as everyone can be seen even if only one person may speak at a time. Group participation eases camera shyness, and aids the introduction of new people to the club.

Arguably, the cost of resourcing the project—in terms of desktop videoconferencing software/hardware, ISDN connections and the facilitator's time—may be considered

too high for wide applicability. However, recent moves to improve technological support in schools combined with decreasing costs will make the approach feasible on a larger scale in the near future.

### **New Initiatives Associated with the Video-Clubs Project**

In addition to the Socialspaces project mentioned earlier, a new project at Sussex, the Emmanate project† will use girls from Priory School as focus groups.

Priory School is extending its work with the Internet, and is participating in the EU-funded Interreg Project. The aim of Interreg is to develop an integrated Web-based framework for distance learning. It will enable students to effectively develop, share and utilise curriculum and cultural information and resources, working both independently and collaboratively. The project incorporates approximately 10 schools from each of the three partner regions: East Sussex, The Somme and Seine Maritime (the latter two are in France). Each site will have networked Internet access (128 kbit/s ISDN or leased line) and its own Web server.

Priory School has also been awarded Beacon School status and is now providing Internet access to a cluster of local primary schools. This includes extensive staff training, technical support, and access to curriculum resources.

### **Conclusions**

It is too soon yet to draw firm conclusions concerning the long-term impact of Video-Clubs on women's participation in computing. The results of the attitude surveys discussed in the previous section, levels of participation in weekly club sessions, and the personal achievements of individual

† Emmanate is an EPSRC-funded project with University College London and BT to investigate mental models for networked applications

club members, are, however, encouraging. Further surveys will be conducted, and the study choices made by the young women participating in the project when they attain FE and HE level will be monitored and compared with a control group.

To enable the approach being pioneering to continue after the end of the project, an interactive resource centre is being created to support students' self-development. This will integrate software, useful Web links and techniques, and enable students to request assistance in new areas. The resource centre will enable schools to collaborate on projects, control their own progress, and generally empower their interaction with technology.

Finally, the fact that schools have been involved in the project and have been provided with new equipment, opportunities and staff development is already having a wider impact than was foreseen at the start of the work. The system has been demonstrated to great effect at parents' evenings and to official visitors, enhancing the way the schools are perceived. The increased self-confidence of staff in their IT-related abilities appears also to be raising the status of technology within each school. This in turn, will impact students and reinforce the goal of engaging more women in computing.

### **Acknowledgements**

The authors would like to acknowledge the valuable contributions to the project by the following people: Professor Hilary Buxton, Head of School of Cognitive and Computing Science, University of Sussex; Dr. Ian Wakeman, also from University of Sussex; Ms. Teresa Taylor, IT Coordinator, West Anglia College, Norfolk; and Ms. Holly Ward, of BT Laboratories, who contributed to early meetings of the project.

The Socialspaces project is a collaboration between QMW and Bell Labs Lucent Technologies, NJ, with equipment and funding from

Lucent Technologies. The authors would like to acknowledge the contributions of Dave Boyer, Mauricio Cortes, Marah Rosenberg and Michael Vernick, all of Bell Laboratories, Holmdel, NJ, to the Socialspaces project. Anjum Khan, a Research Assistant at QMW, is developing applications of Socialspaces, and these will be tested using Video-Club members.

## References

- 1 STACK, J. Women into Computing. *Interactions*, May–June 1997, pp. 17–18.
- 2 SCOTT, J. Women and Professional Practice in Information Systems. Report of IBM seminar. *Women into Computing Newsletter*, June 1998, p. 4.
- 3 DURNEDELL, A.; GLISSOV, P.; and SIANN, G. Gender and Computing—Persisting Differences. *Educational Research*, 1995, **37**(3), pp. 219–227.
- 4 WHITLEY, B. E. Gender Differences in Computer-Related Attitudes and Behavior: A Meta-Analysis. *Computers in Human Behavior*, 1997, **13**(1), pp. 1–22.
- 5 TURKLE, S. *Life on the Screen: Identity in the Age of the Internet*. Phoenix Press, 1995.
- 6 FINLEY, A.; REDELFS, A.; and MOSES, G. Education, Outreach, and Training. *Communications of the ACM*, Nov. 1988, **41**(11), pp. 74–75.
- 7 BOYER, D.; CORTES, M.; BALFOUR, G.; and WILBUR, S. Virtual Social Clubs: Meeting Places for the Internet Community. Poster paper in *Proc. of Multimedia 98, 6<sup>th</sup> ACM International Multimedia Conference*, Bristol, UK. pp. 40–41.
- 8 LEHMANN-HAUPT, R. Overcoming the Gender Gap. *Wired magazine*, Oct. 1998, pp. 10–11.
- 9 CROOK, C. Afterwords. Chapter 9 of *Computers and the Collaborative Experience of Learning*. International Library of Psychology, 1994, pp. 224–230.

## URLs

Web sites targeting women to use the Internet:

*Purple Moon*: <http://www.purplemoon.com/>

*Cybergirl*: <http://www.cybergirl.com/>

*GIRL Tech*: <http://www.girltech.com/>

*A Girl's World on-line Clubhouse*: <http://www.agirlsworld.com/>

## Biographies



**Graeme Balfour**  
Queen Mary and  
Westfield College,  
University of London

Graeme Balfour has been a research assistant in the department of Computer Science, Queen Mary and Westfield College, University of London since 1997, working on the BT-funded Virtual Clubs project. Before taking the M.Sc. in 1996/7, Graeme was a teacher of mathematics with responsibility for curriculum development. He developed substantial units of classroom materials, some of his work being published.



**Sylvia Wilbur**  
Queen Mary and  
Westfield College,  
University of London

Sylvia Wilbur is a senior lecturer in the department of Computer Science, Queen Mary and Westfield College, University of London. Sylvia has been a member of the department since 1983, and has been active in the areas of computer-

supported co-operative work and multimedia research since 1986. She has worked closely with BT on several collaborative projects. Sylvia was Chair of the Multimedia Group for the Esprit Working Group CO-TECH from 1990–93.



**Hilary Johnson**  
Queen Mary and  
Westfield College,  
University of London

Hilary Johnson is a lecturer in the department of Computer Science, Queen Mary and Westfield College, University of London. Hilary has been a member of the department since 1986. Her research interests include task analysis and modelling; developing task-related principles for the design of user interfaces and evaluation. She was the technical co-chair of HCI98.



**Stephen Furner**  
Networks and  
Information Services  
BT UK

As a consultant in Human Factors for BT, Stephen Furner has applied advanced usability engineering techniques to the research, design and development of BT products and services. He has led multinational teams of technical experts researching convergent technologies, and has managed national and international research and development projects for BT. He regularly reports the results of the research he carries out for BT at professional conferences, in technical journals and through the technical press and public media. His current research interests include computer haptics, social uses of computing and communications, self-positioning mobile data services and affective computing. He is both a Chartered Psychologist and an Incorporated Engineer.

# Convergence Between Public Switching and the Internet

*The Internet's exponential growth has a serious impact on established public switched networks with respect to user services, network performance and operators revenue. A solution to this conflict designed to protect the telecommunications operator's and service provider's tremendous investment in the existing network infrastructure is offered. An integrated Internet services architecture is proposed by Siemens that evolves the central office into the key network element of the seamless converged multimedia network of tomorrow.*

---

This is a Siemens' view of how they would move into the new IP world, but a number of other manufacturers are working on the problem and have different solutions.

## Introduction

### Internet—a challenge for the PSTN/ISDN

Today's public switched telephone network (PSTN)/integrated digital services network (ISDN) is the largest telecommunication network worldwide, interconnecting about one-sixth of the world population in the year 2000. The PSTN/ISDN represents an immense investment in infrastructure and carries all channel switched public traffic and substantial corporate voice and data traffic. With its nonstop processing network nodes it provides seamless intelligent services spanning the globe with extreme reliability and grade of service.

In the past years the Internet has developed to become 'the' global information network, attracting an exponentially growing community of Internet/on-line users worldwide by its wide variety of multimedia applications. Though still in its infancy with regard to reliability and guaranteed grade of service, the Internet is highly accepted as an information platform by end users and information content providers alike. The high investments being made in the Internet by the computer industry, network carriers and service providers may ensure that the Internet will meet the expectations for continued exponential growth predicted by analysts and the boosted performance claimed by the Internet industry.

Even today the Internet has serious impact on the established public switched network with respect to user services, network performance and PSTN/ISDN operators' revenue. This impact will

become more pronounced in the future as the Internet matures to provide services with guaranteed grade of service; for example, voice and video over IP. This may eventually turn the Internet into an alternative carrier network, challenging today's PSTN/ISDN operator's business by diverting substantial high-revenue traffic away from the PSTN.


Clearly this conflict has to be resolved such that the tremendous investments made by the telecommunications operators and service providers in the existing network infrastructure are protected. In consequence, the solutions which make a seamless migration to the network of tomorrow possible and best fit the telecommunication market's requirements will evolve the present telecommunications network infrastructure. This future network will provide services moving towards a true convergence by integrating the 'classical' network types (for example, PSTN/ISDN and public land mobile network (PLMN)) with the Internet. So it will be possible to benefit from the merits and innovations of both types of today's separated networks.

This article shows that by innovation with leading-edge technology the present telecommunication network elements can build the solid basis for a seamless converged multimedia network of tomorrow.

### Internet access from the PSTN/ISDN—issues and risks

The PSTN/ISDN is the major access to the Internet. The vast majority (90%) of the Internet end users gain access to the Internet from their PSTN/ISDN subscriber lines, using either analogue

Figure 1—Network structure today and tomorrow



modems or ISDN cards to set up data calls. Internet services providers connect their modem pools to PSTN/ISDN subscriber lines.

The channel switched call setup in the PSTN/ISDN is not adequate to handle the resulting packet data traffic over IP. In consequence high penetrations of Internet accessing lines result in serious degradation of the PSTN/ISDN network performance and substantial loss of PSTN/ISDN operator revenue:

- network performance degradation, due to:
  - long holding times of subscribers logged into the Internet (average 20 minutes compared to 3 minutes for normal calls) devalue the classical rules for network dimensioning;
  - blocked subscribers due to jammed concentration stages of central offices (COs); and
  - uneconomic use of 64 kbit/s channels to transport bursty IP traffic.
- loss of PSTN/ISDN revenue, due to:
  - long holding times at flat rate;
  - lost calls to Internet-busy subscribers;
  - local calls push away profitable long distance and international calls;
  - market demand for reduced tariffs for access to on-line services where there is no flat rate for local calls, because PSTN/ISDN does not add value;
  - migration of PSTN/ISDN services into the Internet; for example, fax replaced by e-mail and telephony potentially replaced by low-quality speech by voice over IP; and

—supplementary services provided by the central office are challenged by services from the intelligent terminal through to computer telephony integration (CTI).

### Internet access from PSTN/ISDN—two possible strategies

For the PSTN/ISDN operators there are two basic strategies to counteract these risks for their network and their business; either divert the Internet traffic away into a data network before it hits the PSTN/ISDN central office (CO) by using external xDSL devices; or control the Internet traffic (narrowband and, in the future, broadband) in the CO as part of the PSTN/ISDN traffic in an optimised way. Clearly the first option runs the risk of either installing an own IP overlay network with all consequences (for example, new infrastructure) separate operations and maintenance etc. or hand-over direct control of the subscriber line to an Internet access provider (for example, a cable network operator). As the Internet (potentially) develops into a full-service network, this may imply the complete loss of the subscriber; except possibly for low revenue POTS service.

The second option, besides ensuring direct control of the subscriber line, enables the PSTN/ISDN operators to reuse and to leverage their installed infrastructure for offering on-line services. The PSTN/ISDN operators thus have the opportunity to become Internet service providers (telecommunications company ISPs) to save and extend their existing customer base and gain revenue from subscriber services that utilise Internet technology. The technical realisation of the second option is now described in more detail.

### Integrated Architecture—Overview

A quick look at the network structures currently in service, shows parallels between telephony and Internet networks that call for a convergence as shown in Figure 1:

- physically, the access line for both services (telephony and Internet) is the classic copper telephone line ending with a line module in the local exchange (LE);
- logically the line ends up in a computer—LE for telephony and PoP for Internet, respectively. The computer interprets the callers destination wish (E.164 number or IP-address);
- both networks are hierarchically structured (according to similar hierarchical address schemes), thus both major backbone elements (the transit exchanges (TE) and routers) fulfil, functionally, the same tasks;
- both networks use the same mechanism (Sonet/synchronous digital hierarchy (SDH) and—if supplied—asynchronous transfer mode (ATM)) in order to transport bundled information over long distances.

Figure 2 gives an overview of a technical proposal how to integrate Internet technology in the EWSD Siemens' CO. The different steps are explained in the next sections.

### Integrated Point of Presence—A Platform for New Services

The interworking between the EWSD CO and direct connected or remote

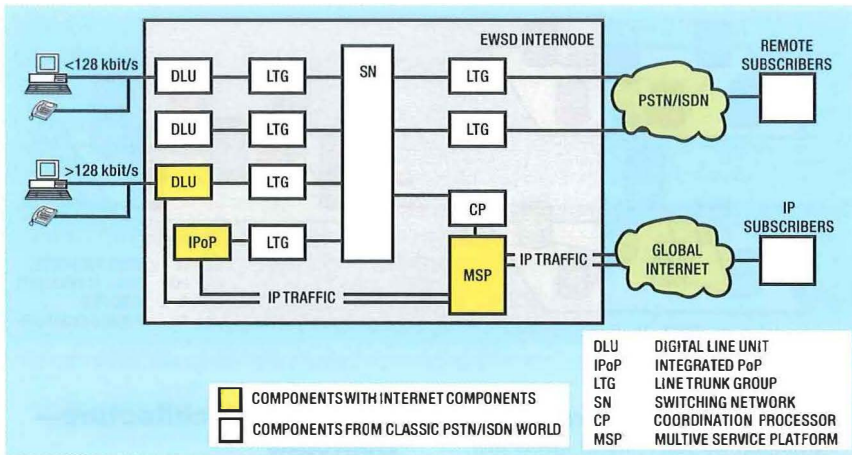


Figure 2 – EWSD with integrated Internet technology – architectural overview

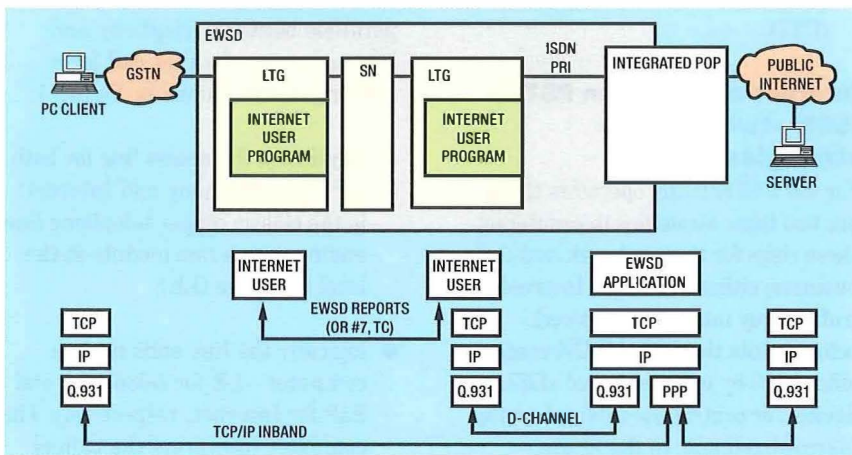


Figure 3 – Call control and IP interworking

dial-in subscribers (characterised by their IP addresses) is based on the integration of the point-of-presence (PoP) in the EWSD. The main task of the PoP is to terminate the modem protocol (for example, point-to-point protocol (PPP)) of the analogue or ISDN subscriber, who is communicating via IP with an Internet access point. In addition to classic Internet PoPs, a call-processing application is included that is able to send and receive IP packets from the internal router. On the other hand this application communicates with an EWSD Internet call-processing application via an internal communication protocol as shown in Figure 3.

For the communication between the CO and the Internet subscriber, the IP address of the subscriber is stored within the CO, referenced to the E.164 number of the same subscriber. When a message to the subscriber has to be sent, it is transported via an internal communication channel,

translated into an IP packet and routed to the IP address of the appropriate subscriber (using the IP/PPP stack). The same transport mechanism is used in reverse, to send information from an Internet source to the CO.

As shown in Figure 4, the architecture of the integrated PoP is chosen to be as open as possible, so that PoPs already in the field can be easily integrated – as far as the respective steering software is modified. Those modifications are already installed and running on the total control hub from 3Com and are planned to be realized on the Netblaser 8500 from ITK.

The benefits of integrating the Internet PoP into the CO are:

- physical integration of PoP facilitates network management through CO operations, administration and maintenance;
- bottlenecks are resolved on the network and at ISP-access level;




Figure 4 – HW-Architecture of the Integrated PoP

- the modem pool can be integrated into the remote concentrator (based on the remote switching unit concept);
- the offering of voice over Internet is facilitated (see next section);
  - telephone-to-telephone: voice compression and low-cost telephony,
  - PC-to-phone: voice-enabled Internet session, and
  - telephone-to-PC: call completion to a subscriber who is busy with an Internet dial-up session.
- the PSTN/ISDN operator can provide new subscriber features based on Internet technology (see next section):
  - SCI via webpages,
  - supplementary services for Internet calls (CWIB/CCIB), and
  - the PSTN/ISDN operator can offer ISP-based features supporting Internet access.

### Central Office and PoP – A Converged Platform for New Services

The integration of the PoP in the CO allows new services leveraging the close interaction of both technologies to be provided. Here the first set of services that has already been realised is described:

- *Call waiting Internet busy (CWIB)/call completion Internet busy (CCIB)* – the subscriber is informed during an Internet session that a call is waiting. This



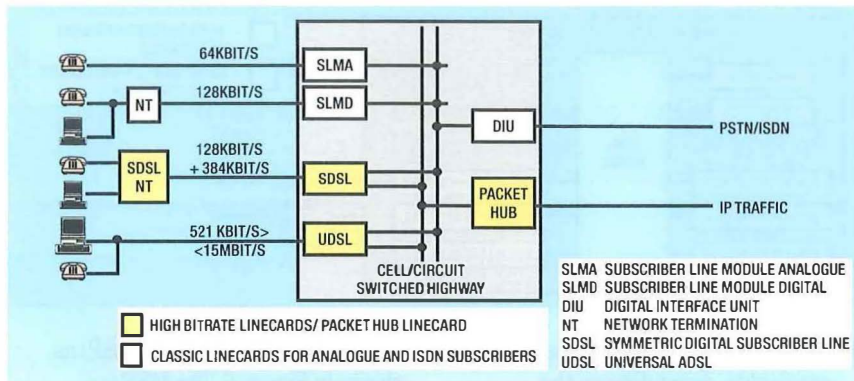


Figure 5 – IP-router based on the multiservice platform used also for CCS7-handling and ATM-switching

means the subscriber remains accessible even while surfing the Internet. If a call comes in while the subscriber is in an IP session a message pops up on his/her screen. The call can be accepted on the telephone in the case of CWIB. If the PC runs a VoIP client software (for example, NetMeeting and WebPhone) the call can be accepted directly on the PC via VoIP H.323 call set-up during the IP session. With CWIB the user remains accessible while surfing the WWW. For the operator the call completion rate increases.

- **E-mail waiting indication (EWI)** – the always-on feature EWI informs an analogue or ISDN user immediately—even if no Internet session is active—that an Internet mail message has been received by his/her e-mailbox. This indication is delivered to the subscriber's telephone line either by a stutter dial tone, an announcement or by display information if supported by the terminal equipment in compliance with ETSI message waiting indication (MWI) for classic telephony. The display indication may include valuable additional information like the number of messages received, originator identity, subject and time stamp.
- **Subscriber control via Internet (ISCI)** – allows a subscriber to control the services assigned to his/her telephone line by using the Web browser. The service provides a graphical PC-based user interface for telephone feature management. This user-friendly handling of telephone features without specialised and

expensive equipment is an obvious advantage for the subscriber. For the operator, the advantage is twofold: the ISCI feature itself can be marketed in the same way as other telephone services. Moreover, major operator benefit arises from the fact that ISCI significantly promotes existing analogue and ISDN telephone features. In particular, subscriber services which include control of telephone numbers (call forwarding, call screening, hot line etc.) achieve significantly higher penetrations, and higher feature revenues.

- **Improved Voice over Internet (IVoIP)** – IVoIP is a service that significantly enhances the capacities of a standard Internet telephone solution for calls originating from a regular telephone<sup>1</sup>. The EWSD InterNode fully integrates Internet telephony into the local or carrier gateway exchange to provide a truly low-cost alternative for the subscriber. Using the IVoIP service the subscriber can use the standard dialing procedure for VoIP calls, the subscriber does not have to provide authentication. Advice-of-charge or meter pulse information for VoIP-calls will be received and the billing will not start before the receiving party answers the call. VoIP calls appear on the regular monthly bill since IVoIP charging records are registered through EWSD in the same way as standard calls. Internet telephony is a good opportunity for classical operators to compete with alternative ISPs for new revenues. It provides a low-cost alternative to classic telephony, which increases

the number of long distance calls significantly. To make such a paradigm work, it is necessary to supply the service in a highly accessible fashion, which is exactly what is offered by IVoIP. In addition to increased user acceptance, the IVoIP service reduces operational costs because of full integration into existing traffic measurement and billing processes.

### Integrated High-Bitrate Linecards—Access to the Information Highway

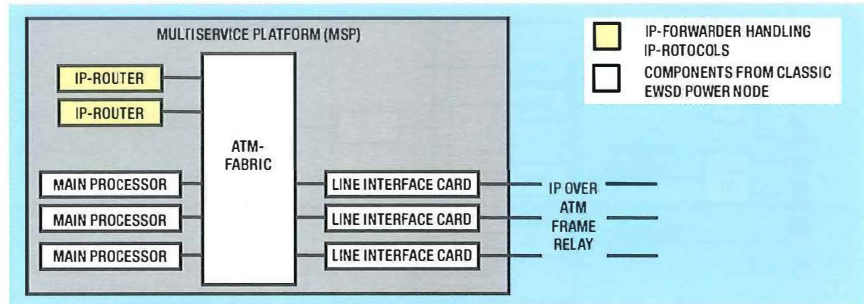
The subscriber demand for fast Internet access will promote the xDSL technology as a big opportunity to open millions of kilometres of installed copper wire for new services and applications. Based on this technology, the EWSD InterNode offers two new subscriber line modules: a universal asymmetric digital subscriber line (UDSL) module as a migration path for POTS and a symmetric digital subscriber line (SDSL) module as the migration path for ISDN technology.

Both modules are fully compatible with classic analogue (subscriber line module analogue (SLMA)) and ISDN line modules (subscriber line module digital (SLMD)) from the hardware and maintenance point of view. All subscriber line modules (SLMA, SLMD, SDSL and UDSL) fit into existing concentration units (DLUs) (see Figure 5) and are interchangeable thanks to universal line card slots. Thus, a migration from analogue or ISDN users to UDSL or SDSL is minimised in terms of operational effort. For the new modules the voice connection to PSTN/ISDN is identical as for the classic modules.

#### UDSL access:

The UDSL access provided by EWSD InterNode is based on the universal ADSL (UDSL) stand-

Figure 6—IP-router based on multiservice platform used also for CCS7-handling and ATM switching



ard—an industry standard promoted by PC, telecommunications, and network industry leaders with the following features:

- high performance: 1.5 Mbit/s @ 4 km loop length downstream, 512 kbit/s upstream; that is, 30 times faster than today's fastest modems;
- parallel voice and data support: analogue voice and data connections are supported simultaneously;
- plug-and-play solution: the UDSL compatible customer premises equipment only has to be plugged into the telephone jack—no network termination (NT), no splitter, no change in home wiring is necessary;
- end-to-end ATM support: the end-to-end service architecture is based on the PPP over ATM architecture. This allows services like high-speed Internet access, secured telecommuting or broadband content distribution;
- high integration: due to its highly reduced power consumption compared with full-rate ADSL a high port integration can be achieved: eight ports per module; and
- compatible CPE available: leading PC firms will provide UDSL PCs and modems and a simplified software setup of UDSL will be included in future versions of Microsoft Windows™.

#### SDSL Access

The SDSL module provides six B-channels in addition to the standard 2B+D channels of an ISDN-BA (basic access), which can be utilised application dependent, circuit switched or packet oriented:

- full ISDN-BA support: all standard ISDN features are still fully available;

- $n \times 64$  kbit/s support: the user can decide, depending on the application, to use the additional six B-channels (or parts of them) either circuit switched for high-quality service like videoconferencing or packet oriented (for example, for data transfer). In the first case all required B-channels are circuit-switched through the switching network (SN), in order to guarantee the quality of the connection;
- performance: full symmetrical 512 kbit/s circuit-switched or packet-switched application dependent;
- cost-effective and reliable technology: The reuse of major parts of well proven ISDN technology allows extremely cost-effective provisioning of SDSL and guarantees high reliability from the beginning;
- high integration: the low power consumption allows the SDSL, high port integration; that is eight ports per module; and
- NT and CPE available: The SDSL-NT as evolution of the ISDN-NT is provided with an Ethernet 10BT interface with normal POTS. Or an ISDN version with an  $S_0$  interface; that is all CPEs can be retained. The same NT can be remotely configured as either POTS or ISDN NT.

#### Integrated Router—Cost Effective Concentration of IP Traffic

##### Architecture

The IP routing and forwarding is realised on a module plugged into

the multiservice platform (MSP) as shown in Figure 6. The MSP is a central part of the EWSD Innovations architecture. It controls the high-speed CCS7-handling in the EWSD PowerNode and the ATM switching in the EWSD Broadband Node. The reuse of this MSP has the following advantages:

- cost-effective provisioning of IP switching in terms of first investment—only a few additional modules—scalable according to the needs;
- cost-effective provisioning of IP routing in terms of life cycle costs—one operations, applications and maintenance for different applications;
- the rigorous requirements for a carrier class design are automatically fulfilled—a must for CCS7 and ATM-switching; and
- perfect interworking with fast growing broadband traffic is inherently included—a key requirement for a future-proof architecture.

#### Features

With the integrated router EWSD offers—cost-effectively—the IP routing and forwarding functionality, with a wide range of powerful and distinctive possibilities:

- *Performance*—the IP switching is performed by a module directly connected with an ATM switching fabric featuring a strictly non-blocking throughput beyond 30 Gbit/s scalable in steps of 600 Mbit/s.
- *Bandwidth*—the integration into a multigigabit ATM switching architecture allows the support of the full range of high-speed

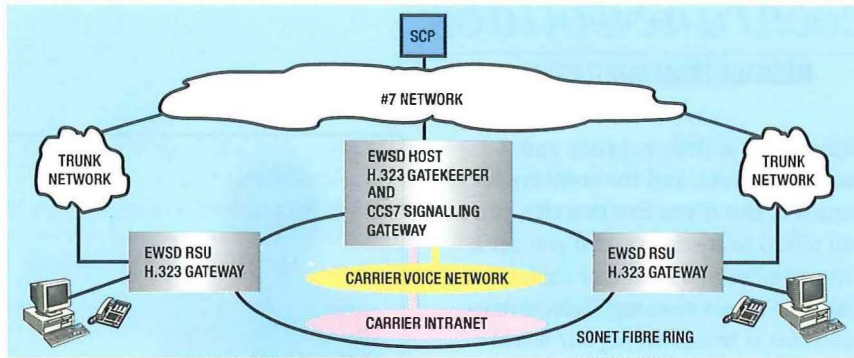


Figure 7 – EWSD InterNode configured as CCS7/IP gateway

interfaces, PDH/SDH/SONET and the ATM family from T1/E1 (1.5 Mbit/s) up to OC-12/STM-4 (622 Mbit/s).

- *Carrier class design*—the integration of the router functions in the general switching architecture guarantees the reliability, availability and serviceability features required for mission-critical networks.

### CCS7 Gateway

Minimised life cycle and operations costs are the key drivers of the consolidation of fully digitalised networks. In response to this requirement, the EWSD CO is being evolved into a distributed network of functionally specialised units: a centralised coordination platform providing a single point for administration, operation and maintenance and distributed remote switching units (RSU) that provide the full set of subscriber and network interfaces wherever connectivity is required.

This distributed switching platform provides the powerful basis that smoothly integrates the IP remote access points, VoIP gateways and wideband subscriber interfaces according to the InterNode concept as shown in Figure 7.

As a result of this combination of distributed IP and CCS7 switching, the coordination platform hosts the CCS7-signalling gateway and H.323 gatekeeper functions as specified by ETSI TIPPHON, while the RSU with integrated PoP functionality implements all functions of an H.323 media gateway. With this approach the EWSD InterNode turns into a full featured distributed CCS7/IP gateway fully reusing the complete range of SS7 signalling protocols

(ISUP, TUP, SCCP, TCAP...) of the EWSD switch.

### Conclusion

In this article an integrated Internet services platform has been presented that turns the EWSD switching system into an EWSD InterNode, the optimised gateway between PSTN/ISDN and the Internet. The key component is an Internet PoP constructed from cutting-edge IP technology that integrates smoothly into the EWSD InterNode and allows for interworking between PSTN/ISDN call processing and Internet services.

Through the presented concept, the EWSD switch of today evolves into an optimised Internet access point that fully supports the telecommunication operators' strategy to reuse their investment in the PSTN/ISDN infrastructure as they position themselves to become ISPs: the EWSD InterNode with integrated Internet service platform optimises PSTN/ISDN network performance by grooming Internet traffic and creates the basis for new subscriber services based on Internet technology.

The approach enables evolutionary extensions for broadband subscriber access starting from the installed narrowband infrastructure in the PSTN/ISDN—as user demand for Internet services balances telecommunications companies' investments. This way, the EWSD InterNode contributes to a convergence between PSTN/ISDN and the Internet towards a seamless multimedia network of the future.

### Biographies



**Hans-Ulrich Schoen**  
Siemens AG

Hans-Ulrich Schoen studied mathematics and physics at the University of Frankfurt and joined Wuppertal Gesamthochschule as a research fellow in 1979. He received his Ph.D. in Mathematics in 1983 and joined Siemens AG in 1985. Since 1996, he has been responsible for the call control of the EWSD switching system as Vice President R&D Application SW. Since 1999, Dr. Schoen has been chief executive officer of Zwut AG, the Siemens information communications network company in Poland. He is a member of VDE/ITG.

E-mail Address:  
Hans-Ulrich.Schoen@zwut.siemens.pl



**Stephan Scholz**  
Siemens AG

Stephan Scholz studied physics at the Technical University of Munich, joined the Max-Planck Institut for Physics and Astrophysics in Munich as a research fellow in 1986 and received his Ph.D. in physics 1989. He joined Siemens AG in 1990 and is responsible for the EWSD product planning and for North American sales.

*Peter Cochrane, Head of Advanced Applications and Technologies, at BT Laboratories, Martlesham Heath, continues his regular column in the Journal. In this issue Peter discusses the changes brought about by the 24-hour society.*

## 24-Hour Society

One Saturday morning in November 1998 found me at my screen in a hotel in Tempe, Arizona. Down the line came a request from a publisher for me to write a foreword for a new book on the 24-hour society. I read the contents, on and off that day, evening, and the next day on a flight to Los Angeles—all on my laptop. By mid-afternoon on the next day—a Sunday—I was sat beside a swimming pool at a hotel in Santa Monica, California, along with a bottle of Evian Spring water from France, the *Sunday Times* from England, and a bag of biscuits from Belgium. Because of my travel and work schedule I had to satisfy the request that day—at least to a good first draft standard. Only five years ago this would have been impossible. Today, for many, it is now the norm. We have become both customer and supplier, technomads on a shrinking planet where there is no escape, and we choose to work where and when we wish. If there was a 21st century bible of marketing and sales the opening paragraph would undoubtedly start with:

*'Customers don't want choice, they want what they want, when and where they want, and at a price and quality they dictate.'*

And how different this is to the 20th century version, which comes from a more sedate industrial, technological, and social era, and starts with:


*'Customers want choice, our choice, supplied at a place and time of our choosing, at a price and quality we dictate.'*

What is happening to promote such rapid and radical change? In short, there has been a transition from a world dominated by atoms to a world dominated by bits. If you live in a village you have to do every-

thing, there is little supplier and customer choice, and the economy is bounded. But if you live in a city, you can afford to specialise, and you get a vast supplier and customer choice linked to a vast economy. Telecommunications is transforming our world into a global city with people and intelligent things (machines, sensors, actuators, photocopiers, food dispensers and complete production plants) communicating over networks of optical fibre, copper, radio and satellites. The bit economy is with us, and with it comes increased competition, availability, and the disintermediation of markets.

If you want to buy stamps, or travel on a major trunk road and need a coffee, at half-past three on a Sunday morning, where do you go? Not to the Post Office or a restaurant, but to the filling station. If you want money at nine-thirty on a Saturday night, where do you go? Not to the bank, but to the cash machine at the supermarket. You want to order a book on a Wednesday evening, where do you go? Not to a bookshop, but on-line to <http://www.amazon.com>. Do you need the hassle of organising travel insurance for every trip? No! Just use your Master Card or AMEX and it comes with the ticket purchase. These are just a few examples of disintermediated services in a 24-hour world. And there is very much more to come. People in the United States have enjoyed 24-hour shopping for over a decade, Europe is only just waking up to the prospect, and characteristically it is the United Kingdom that is leading the way. While Europe worries about the information society—how to limit working hours and how to control bits—the Americans are creating an information economy that is free and unbounded. This is where the future lies, and it is why Europe will see its 'lunch eaten' by competition rising to the challenge of 24-hour working.

At the time of writing BT has over 200 companies inside the group, and over 60 joint ventures on the boil. The vast majority of these are outside the UK and its time zone. We



have become a global entity, a wired and virtual corporation with vast numbers of people working to maintain and expand a 24-hour, 365-day a year, spatially unlimited business far beyond the shores and thinking of Europe. For me as an individual, it means I have to be able to communicate and work from any location and at any time to respond to customers the instant they call. For many of us it means working unusual and unscheduled hours to appear, in reality or virtually, anywhere on the planet.

Now as customers what should we expect? Pizza and coffee at 04.30 on any day. New clothes and electronic hardware purchases at 22.00 on a Saturday night. Why not? IT is creating a new two-class society:

Class 1—Those who spend a lot of time trying to save money; and


Class 2—Those who spend a lot of money trying to save time.

Stone Age man only worked an estimated 15 hours a week. In this new economy most of us will be working 15 hours a day doing things that are far less risky than hunting for meat on the hoof, but sometimes, equally stressful. We need a 24-hour society to survive!

So, there I was still at the side of the pool on a Sunday afternoon—with my international press and consumables—job done. It was four o'clock and the pool was inviting, and no doubt the authors were in bed. But I still had time for a swim before night shopping for computer hardware, software, and a few items of clothing. But, then again, I was in California...

### Internet and Intranet Security

by Rolf Oppliger



The Internet was created to help foster communications among United States government sponsored researchers. Since then it has grown phenomenally, and the benign environment in which it once operated is now open to all the dangers that one can find in society.

The author says the aim of this book is to introduce and discuss security techniques and mechanisms that are available today to provide Internet and intranet security. Although tutorial in nature, it requires some familiarity with computer networks and cryptography.

Part 1 introduces us to the fundamentals that are necessary to understand the rest of the book. It starts with a history of the Internet from the advanced research projects agency network (ARPANET) through to the present day. The next chapter covers the technical bodies, documentation series, and the processes related to Internet standards. Each of the transport control protocol/Internet protocol (TCP/IP) stack layers are examined next, along with IP version 4 and IP version 6 packet formats, TCP header format, connection establishment using SYN and ACK, sequence number attacks, and SYN flooding attacks. One-way hashing, secret-key and public-key cryptography are covered next, along with a reasonably detailed explanation of the mathematics behind the RSA (Rivest, Shamir and Adleman) and Diffie-Hellman systems. Public key infrastructures and certification authorities are discussed, along with

the widely used ITU-T X.509 format for public key certificates.

Part 2, 'Access Control', addresses firewall technology. Firewalls prevent unwanted and unauthorised communications into and out of a protected network. They also provide a good place to collect information about system and network use and misuse. Firewall components such as packet filters and application gateways are studied in detail along with proxy servers, SOCKS, and SYN flood attack countermeasures. Different Firewall configurations (dual homed, screened host, and screened subnet) are examined, and the pros and cons of each discussed.

Part 3, 'Communication Security', addresses TCP/IP security protocols. The first chapter deals with the Internet layer, extensively covering IPSP (IP security protocol) and IKMP (Internet key management protocol). Moving on to transport layer protocols, we examine SSH (secure shell), the popular SSL (secure sockets layer), and Microsoft's PCT (private communications technology). We continue our journey up the TCP/IP stack to the application layer. Here we learn about security enhanced applications such as privacy enhanced mail, PGP (pretty good privacy), S/MIME (secure/multipurpose Internet mail extensions), and secure HTTP (hypertext transport protocol). The author concludes this part of the book by debating in which layer of the TCP/IP stack to implement security. Placing security in lower layers means that security can be implemented transparently to users and application programs, effectively killing many birds with one stone. Conversely it can attempt too many things simultaneously, and only application layer protocols (or above) can actually meet application-specific security requirements and provide corresponding services. The author's view is that a combination of security protocols, with each protocol selected for what it is best at, will usually be the best solution.

Part 4, 'Discussion', begins with a chapter on Electronic Commerce. Following sections are devoted to electronic cash, electronic cheques,

micropayments, and the secure electronic transaction (SET) protocol designed for credit card transactions. Next we get a very brief description of some of the security tools available, such as security administrator tool for analysing networks (SATAN), Pingware, Internet Scanner, and RealSecure.

I liked both the logical structure of the book and the way new information builds on previous chapters. The text is clear and concise, which helps when explaining some of the more complex concepts. However, the diagrams that sparsely punctuate the book are sometimes very basic. I felt that a more liberal use of detailed diagrams would have complemented the text. The coverage of subjects is, in general, comprehensive and thorough, with the right balance between too much and too little information. The author includes an extensive list of references to other publications at the end of each chapter. Hopefully this should allow readers to follow up on particular subjects of interest. Overall, I found the book useful and interesting and would definitely recommend it as an introduction to TCP/IP security.

*Published by Artech House*


*£55.00. xxi + 348pp.*

*ISBN 0-89006-829-1*

**Reviewed by Steve Smith**

### Newnes Telecommunications Engineer's Pocket Book (Second Edition 1998)

by Steve Winder



If telecommunications engineers cannot quite boast, with politicians, that a week is a long time in their

## book reviews

---

profession, they can at least console themselves with the thought that six years is getting close to a geological era. Since publication of the first edition of the *Pocket Book* in 1992 the pace of change in the industry has made an update long overdue. The real test of Winder's second edition is whether his revision has successfully reflected this change—indeed, whether such a vast field of endeavour as telecommunications engineering can any longer be adequately represented in this catch-all manner at all.

The answer has to be yes on both counts. While following the conventions of the 'Pocket Book' format (small physical size, rigorous structure, short discussions and emphasis on practical details), its scope is certainly epic. Although the discussion on the relative merits of different cuts of quartz crystals (the first topic in the 1992 edition) is among several items to have been omitted, the journey still takes us from relay contact materials to ATM switch architectures and most stops in between. Winder has recognised that these days most telecommunications engineers will not be doing serious business at the component level, and has successfully sought to distance himself from the physical layer while retaining sufficient material on basic principles. Thus, for example,

weighty sections on attenuators and filters have been quietly dropped along with the crystals, but fundamental transmission line theory is retained.

Winder has obviously wrestled at length with the material to arrive at a clear logical sequence. The book begins with a short discussion on the various types of signals which telecommunications systems are required to convey, and then moves on to terminal equipment. Sections follow on transmission media (copper, fibre and radio, including satellite) and then switching and signalling. A comprehensive section on multiplexing then covers FDM, PDH, SDH and WDM, followed by another on packet switched data. Finally, there are separate sections on test equipment, relevant organisations and standards.

The material in the sections is generally clear and sufficiently comprehensive to, as a minimum, let the reader identify the particular area where more detail might be required from a more specialised source. In most sections the coverage goes well beyond this. Diagrams are provided on a reasonably generous basis, although some more 'white space' in the layout would have been welcome in places. It is perhaps unfair to look to practit-

er's handbooks of this type to add yet more verbiage to current 'hot topic' debates in the industry—open network architectures and the implications of the datawave to name but two—and Winder has correctly concentrated his efforts on providing practical details of real systems. There is, for example, sound coverage of such 'advanced' topics such as xDSL technologies and TETRA.

One could have wished for a little more material relating to matters higher in the protocol stack—IP being the most notable example—and those interested in network matters such as network management and IN may feel a little hard-done by. Other than this, there is something relevant to all practising telecommunications engineers and the book will be particularly valuable to those looking for a basic but sound appreciation of fields outside their particular speciality. In short, an excellent handbook—but Winder cannot rest on his laurels for too long as the next edition will be required long before another six years have elapsed.

*Published by Newnes/Butterworth-Heinemann*

*ISBN 0-7506-3936-9*

*£15.00. viii + 284pp.*

*Reviewed by Bill Woodbridge*

## telecom focus

---

### Payphones in Prison

---

BT has announced that BT Payphones has won a £100 million, 10-year contract with the Prison Service for a new type of telephone system for prison inmates.

New telephones based on personal identification numbers (PINs) will replace the current payphones, introduced into prisons in 1987, which accept specially encoded phonecards.

The new contract is for more than 2500 telephones in 132 prisons across England and Wales. Each prisoner will be issued with a PIN number and account to use the telephone system. To make a call, the prisoner keys the telephone number required

followed by the PIN number. A computer management system then automatically checks whether the prisoner is allowed to call the number at that particular time, and has enough units in the account to pay for it. Prisoners buy units to credit their account.

The PIN telephone system has a special victim anti-harassment feature called 'GOTU'. This allows people with touch-tone telephones who are called by a prisoner either to accept the call or decide not to take it and to add their number to a list of those barred on the prisoner's telephone account.

BT PIN telephones have been on trial in two prisons—Buckley Hall, in

Rochdale, and Full Sutton, in York. BT Payphones will begin replacing the old cardphones with PIN telephones during 1999.

George Howarth, Home Office Minister, said: 'Providing prisoners with access to the telephone system is essential because it helps to maintain family relationships through the difficulties caused by separation and travelling for visits.'

The new system will afford us greater control and better intelligence on the illicit use of telephones by prisoners and is an important step forwards in helping to protect the public.

'PIN telephones should also reduce bullying within prisons by

removing the need for phonecards which can circulate as currency.'

### BT and Excite Link-Up

BT and US-based Excite Inc. have announced that they have reached an agreement for BT to buy a 50 per cent share in the Excite Inc. subsidiary, Excite UK. This \$10 million investment in the \$20 million subsidiary is the first time such a close and active relationship between a United Kingdom company and a major United States Internet portal has been forged.

The agreement will further develop BT's presence in Internet advertising and transactions—two major on-line revenue generators—complementing its expertise in Internet access. It also brings the benefits of Excite's considerable Internet technology and, specifically, its on-line marketing and sales expertise.

Excite UK will benefit through better access to one of Europe's most important and fastest growing Internet markets and BT's expertise in further localising its product to the UK market.

Excite UK, in providing Excite's Internet and other services in the UK, offers new and experienced Internet users a personalised and content-rich starting point on the Web. This includes 12 content areas covering subjects such as news, music and sport, plus advanced search technology, free web-based e-mail, communities and on-line shopping. Excite offers advertisers access to a significant number of Web users, as well as considerable technical expertise in areas such as advertisement display and on-line marketing.

In addition, Excite UK is marketing Excite Click, powered by BTClick, a combined pay-as-you-surf UK Internet access service and Web guide. This service requires no registration or monthly subscription and features immediately on the Excite web site ([www.excite.co.uk](http://www.excite.co.uk)).

BT offers a range of consumer Internet and interactive multimedia services including BT Internet; premium content through joint venture company, LineOne; current

and forthcoming innovative on-line services through BT Click+, multimedia kiosks and 'Open', the interactive TV service provided by BiB (British Interactive Broadcasting).

### AT&T Concert

BT has announced that AT&T has become a non-exclusive distributor of Concert Communications Services in the United States.

AT&T is selling the Concert portfolio of international voice, data and Internet protocol (IP) services under the AT&T Concert brand.

This arrangement ensures that AT&T and BT customers continue to be served effectively and continuously.

Concert is the world's leading provider of seamless global transborder communications services and has more than 4400 customers in 52 countries accounting for nearly \$2.75 billion in committed contract revenue. Concert services are available from 47 distributors worldwide.

### Oftel Plan's to Publish Complainants Information

Oftel has set out plans to publish information on complaints it receives from customers about telecommunications companies and services.

The proposals aim to make customers better informed when they are choosing telephone suppliers. They will also be a greater encouragement for telecommunications companies to sort out the root cause of the complaints.

Director General of Telecommunications, David Edmonds, said: 'This information is an important tool for telephone users.

'I want to ensure that it is accurate, presented in a form which is useful to customers, and fair to the telecommunications companies. That is why we are carrying out this consultation exercise.

'It is important that we get the views of consumer groups and input from the telecommunications industry at this initial stage. Everyone benefits from informed consumers.'

The Director General has a statutory duty to consider complaints and enquiries and around 40 000 complaints are received each year. From May next year, OFTEL plans to publish information about these complaints every six months.

Under the proposals, complaints will appear in two tables. One table would be the Top Ten complaints for the period. The other would be the top ten companies most complained about for the period.

### More Quality of Service Information Published

More telecommunications companies than ever before are publishing information about the quality of service they offer to customers.

Seventeen companies have provided information for the telecommunications industry's new comparable performance indicators.

COLT Telecommunications, Eurobell, Global One, Esprit and Norweb are publishing data for the first time in this publication. Worldcom International and Torch have pledged to publish data on their performance in future publications.

Welcoming the publication of the report, David Edmonds, Director General of Telecommunications, said: 'I am pleased that more companies have responded to Oftel's call and are publishing data for the first time. I look forward to seeing even more companies included next time.

'I also welcome the fact that this report contains extra information about operators' performance on service provision and fault repair. These are important aspects of their service.

'It is good that, for the first time, the data is being published on the Internet. This will give wider and speedier access to the information, particularly for business customers.'

### Ionica Customers Retain Service

BT stepped in during December 1998 to provide some £3 million in funding so that Ionica, which went into administration on 29 October 1998,

could continue to operate for a period of up to three months. This prevented Ionica's 60 000 customers losing telephone service, including 999 cover, during the winter months.

Meanwhile, BT has attempted to contact every Ionica customer by telephone and mail to offer a smooth transition to BT or an alternative supplier.

By mid-February 1999, with engineers working flat out, BT had already connected more than 28 000 former Ionica customers and had a further 17 000 appointments booked.

Ionica's customers were offered BT's standard range of services, but they were free to switch to other service providers should they so wish.

### Free Internet Service From BTClickFree

---

BT has launched a free Internet service, BT ClickFree, the only free UK Internet service which requires no registration.

Customers can access the Internet simply and easily when they choose, paying only the cost of a local call or less if they belong to one of BT's call discount schemes, such as Friends and Family and PremierLine.

BT has also teamed up with Value Direct to give BT ClickFree users the ability to shop on-line for a range of products at lower prices than through any other access service or traditional high street retailer.

John Swingewood, BT's director of Internet and multimedia, said: 'BT has introduced ClickFree in response to customer demand for a free service combining reliability and the quality of service which they have come to expect from BT.'

'This is a positive and logical move for BT and its customers. We will continue to respond to our customers and to develop ClickFree to make it the most competitive service on the market. Our deal with Value Direct means that we have gone even further than just free Internet access. BT ClickFree users will be able to shop on-line for a range of electrical goods and obtain greater discounts which are not available through other Internet access services.'

'BT's existing pay-as-you-go service, BT Click+, has already been praised for its simplicity and it has continually performed well in benchmarking tests by industry analysts.'

BT ClickFree will include information and entertainment content from Excite. Also included is BT's free e-mail address service, talk21 which will be upgraded to POP3 in March 1999 allowing users to read and write e-mails off-line.

### BT, NEC and Nortel Networks in Third-Generation Mobile Technical Trials

---

BT has announced it has begun technical trials of third-generation mobile communications services, based on UMTS (universal mobile telecommunications system), in collaboration with NEC and Nortel Networks. These trials will strengthen BT's position as a leading authority on UMTS mobile multimedia communications.

Among several devices to be tested is a mobile telephone with built-in camera and video screen, capable of transmitting and receiving data and images at high speeds. The trials at BT's research centre at Martlesham, Suffolk, will help the companies involved assess the market for futuristic mobile multimedia, Internet protocol (IP), data and voice services using prototype UMTS networks and equipment.

Results will be fed into worldwide standards activities to ensure users' handsets work wherever they are, as well as being made available to BT's subsidiary companies, including Cellnet, and global joint ventures.

NEC and Nortel Networks are the first of several companies with which BT is working to strengthen its capabilities in this area. BT expects to conduct additional trials on other aspects of this service.

BT will test prototype terminal equipment which will be representative of the first generation of commercially available devices in the next millennium. It is envisaged that these devices will support a number of innovative mobile multimedia services, including:

- the automatic downloading of sports highlights into a user's handheld mobile videophone;
- the ability for field salespeople to view catalogue information in graphic formats via videophone;
- access to e-mail and the Internet via a mobile handset; and
- the ability to view the latest film trailers and purchase cinema tickets on-line.

### AT&T's UK Research and Development Laboratory

---

A building in England's Silicon Fen, set among 12th century architecture but housing a 21st century laboratory, is home to AT&T's first major research and development centre outside the United States.

The facility, the former Olivetti & Oracle Research Lab, was acquired in February and renamed AT&T Laboratories—Cambridge. It is co-located with one of the world's pre-eminent research universities. It houses 50 full-time researchers, working on networking, multimedia and mobile communications systems.

'There was a time when we believed that all good ideas originated within our New Jersey office complex,' said David Nagel, AT&T's chief technology officer and president of AT&T Labs. 'Today, AT&T Laboratories is branching out, first 2600 miles to the west to Silicon Valley, with facilities focused on IP development, then to a new Internet research lab at the University of California at Berkeley, and now 3600 miles north-east to the English countryside.'

'What better way to bring new products and solutions to market than by putting networking to work in our laboratories, linking together the people with the ideas and the expertise to make things happen, wherever they happen to be,' Nagel said.

### Global Forecasts for Local Loop Growth

---

Independent analyst group, Ovum (<http://www.ovum.com>), has released



research that charts the growth of the global local loop market. Ovum anticipates that the number of local lines installed each year will grow from around 72 million in 1998 to reach 195 million by 2005. Strong growth in data communications, stimulated by the Internet, is the most important driver. Other drivers include the development of alternative technologies, regulatory encouragement for access competition, and the increasing importance of mobility. These findings come from Ovum's new report *The Future of the Local Loop: Market Strategies*.

'The key competitive battleground in telecommunications in the foreseeable future will be the provision of broadband lines' said Adrian May, co-author of the report. 'Long gone are the days when operators could differentiate themselves on tariffs alone—nowadays the challenge is to be first to provide bandwidth at cost-effective prices. This is driving a huge demand for broadband technologies.'

### **BT and Butterworths Form Alliance**

---

Butterworths, the largest legal publisher in the UK, and BT have formed an alliance to develop a new multi-media service for the legal sector. The new package of tools for law professionals, to be delivered online to subscribing lawyers' desktops, will bring together Butterworths' unrivalled legal content with BT's expertise in advanced communication tools.

Nationwide market research, followed by discussions with focus groups, has established that there is an unmet demand among small-to medium-sized law firms for a bundle of services which would take them into the next generation in legal know-how and communication services—and help them keep pace with an increasingly important part of the service which cutting-edge City firms provide to their clients.

Top priority in the wish list is a secure document exchange service; next in line is a single point of access for a variety of legal information

(such as the commentary, statutory and case law resources represented, among others, by Halsburys Laws and the All England Law Reports), closely followed by standard forms and templates. Then come on-line conferencing tools for interest groups in the legal community, followed by electronic document archiving and storage.

The 'Legalconnect' package under development by Butterworths and BT is designed to meet these needs, while providing the confidentiality and reliability which is of such importance to the legal sector.

The key elements in the 'Legalconnect' package are being offered to selected law firms for a two month period starting in February, with a view to a launch of the full package nationwide later in the year. Initially, firms specialising in personal injury work in the West Yorkshire area will be participating in this exercise.

### **Internet Expected to Bring Fundamental Change**

---

A major survey by BT and *Management Today* reveals that 63 per cent of managers believe that the Internet will fundamentally alter the way they do business, or will open up new commercial opportunities for their businesses, within two years. Only 13 per cent thought the Internet would have little or no effect on their business.

Reflecting this view, 81 per cent of managers expected their company's investment in the Internet and other electronic business technology such as intranets and e-mail to increase over the next two years. Only five per cent respondents thought this would definitely not be the case.

Despite their apparent enthusiasm, however, many admitted to be learning on the job, with 68 per cent of managers saying that they were still finding out about e-business. Technofear still exists among some managers who are concerned about the effect of technology on their business. Of those managers interviewed, 45 per cent were concerned about the speed at which technology

becomes outdated and 22 per cent were worried about the cost of investment in new technology. Security is also a concern inhibiting the growth of e-commerce for some managers. A quarter had fears of potential access to confidential information and 19 per cent had worries about secure payment over the Internet.

The survey questioned more than 400 readers of *Management Today* from companies throughout the United Kingdom and found that electronic business is already an important aspect of most managers' working lives. Internet access is widespread (88 per cent) and 72 per cent are surfing at least once a week. 42 per cent had access to their company's own intranet and almost all, 98 per cent, were e-mail users.

Intranets emerged as an e-business technology earmarked for significant potential growth—26 per cent of those without access to an intranet expect their company to buy one within the next six months and three quarters of managers who already had access to an intranet but were not using it said they would start to do so within six months.

The survey found that managers view improved communications and access to information as the main benefits of the growth of electronic business. At the present time, most managers who surf the Web are using it to find out information (88 per cent) rather than for buying and selling products and services, but this is expected to change.

### **500 Million Internet Users Predicted by 2005**

---

Ovum, an independent analyst group, has carried out new research charting the continued development of the Internet. Ovum forecasts 206 million dial-up connections and 17.5 million permanent connections to the Internet by 2005, representing a four-fold increase on current figures. Permanent connections will become increasingly critical for large corporates and for organisations that depend on fast transmission of large data files. The findings are available

in a new report, *Internet Market Forecasts: Global Internet Growth 1998-2005*.

The key findings include:

- The United States market will become saturated after 2002 and most continuing growth will derive from developed Asia and Western Europe. Aggregate traffic will grow over 6 terabit/s by the year 2005.
- There will be a major capacity growth but no redistribution of traffic patterns. The most impor-

tant routes in 1998 will still be the most important in 2005.

- The international Internet backbone is on the point of collapse. If new fibre build across the Atlantic does not progress to plan there will be a chronic shortage of capacity to the United States. Ovum warns that this will result in a poor service that could threaten the development of the Internet on a global basis.
- The Japanese market will dominate the Internet in Asia, although

China will experience a significant growth of its Internet community but will continue to have very low Internet penetration rates.

- Value-added service provision is key to long term survival on-line. Internet access revenues from dial-up connections will decline in per-capita and real terms despite significant growth in subscriber numbers. Revenues from permanent connections will rise initially, but will also encounter a decline in per-capita revenues around 2004.

## forthcoming conferences

### April 1999

**AON99: All-Optical Networks ComForum**, 30 March-1 April 1999, Renaissance Orlando Resort, Orlando, Florida, Presented by the International Engineering Consortium. Registration/ further information: Tel: +1 312 559 4600, Fax: +1 312 559 4111, E-mail: [events@iec.org](mailto:events@iec.org), Web site: <http://www.iec.org>, Mail: 1999 All-Optical Networks ComForum, 549 West Randolph Street, Suite 600, Chicago, IL 60661-2208 USA.

**Internet and Digital Broadcasting Converged Services**, 12-13 April 1999, The Millennium Commodore, Paris. Post Conference Workshop: Regulating the Era of Digital Convergence, led by Andreas Bartosch, European Counsel, Gleiss Lutz Hootz Hirsch, 14 April, The Millennium Commodore, Paris, Further information from: Vision in Business Ltd, Tel. 0171 839 8391, Fax: 0171 839 3777, E-mail: [postmaster@visibis1.demon.co.uk](mailto:postmaster@visibis1.demon.co.uk)

**The NMT Congress, Revitalising Analogue Networks to Maximise Growth and Potential**, 13-14 April 1999, Hotel Intercontinental, Helsinki, Finland. Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com/nmt99>

**Mobile Internet '99, Developing the Potential of Mobile Internet/Intranet and Information Services**, 13-16 April 1999, London Marriott Hotel, London, Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com/mobint>

### SIM '99, Improving Operator Profitability by Increasing Differentiation

14-16 April 1999, Le Meridien Hotel, London, Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com/sim>

**LRIC and Cost Allocation for Interconnect Pricing**, 19-20 April 1999, Brussels, Further information from: Vision in Business Ltd., Tel: 0171 839 8391, Fax: 0171 839 3777, E-mail: [postmaster@visibis1.demon.co.uk](mailto:postmaster@visibis1.demon.co.uk), Web site: <http://www.visibis.com>

**European Telecommunications Law, Regulatory/Business Platform for the Next Millennium**, 19-21 April 1999, Conrad International Hotel, Brussels. One-Day Post-Conference Regulatory Workshop on 'Interconnection'. 21 April 1999, Conrad International Hotel, Brussels., Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com/eulaw>

**Developing Tariff Schemes for ATM Multiservice Networks**, 26-27 April 1999, Regents Park Hotel, London. Post Conference workshop led by Eutelis on 'The Marketing and Deployment of ATM Multiple Services', 28 April 1999. Further information from: Vision in Business Ltd., Tel: 0171 839 8391, Fax: 0171 839 3777, E-mail: [postmaster@visibis1.demon.co.uk](mailto:postmaster@visibis1.demon.co.uk), Web site: <http://www.visibis.com>

**The Internet Architecture and Internet Protocols**, 26-27 April 1999, The Fitzrovia Hotel, London, Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com>

**Transfer Pricing for Telecoms Operators**, 26-27 April 1999, Amsterdam. Separately bookable one-day conference: 'E-Commerce Taxation in Telecoms', 28 April 1999, Amsterdam., Further information from: Vision in Business Ltd., Tel: +44 171 839 8391, Fax: +44 171 839 3777, E-mail: [postmaster@visibis1.demon.co.uk](mailto:postmaster@visibis1.demon.co.uk), Web site: <http://www.visibis.com>

**Billing Systems 1999**, 26-29 April 1999, Olympia Conference Centre, London, Further information from: IIR Telecoms and Technology, Tel: +44 (0)171 915 5055, Fax: +44 (0)171 915 5056, E-mail: [billing@telecoms.iir.co.uk](mailto:billing@telecoms.iir.co.uk), Web site: <http://telecoms.iir.co.uk/billing/>

**Interconnection Latin America**, 26-29 April 1999, Loews Miami Beach Hotel, Miami, USA, Further information from, Scott Forbes, IBC Global Conferences, Tel: +44 171 453 5495, Fax: +44 171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com>

**Mobile Security and Fraud Prevention, How to Minimise the Risk and Maximise your Profit**, 27-29 April 1999, The Mount Royal Hotel, London, Further information from, Scott Forbes, IBC Global Conferences, Tel: 0171 453 5495, Fax: 0171 636 1976, E-mail: [cust.serv@ibcuk.co.uk](mailto:cust.serv@ibcuk.co.uk), Web site: <http://www.ibctelecoms.com/security>

# BRITISH TELECOMMUNICATIONS ENGINEERING

ISSN 0262-401X

Published in April, July, October and January by The Institution of British Telecommunications Engineers, Post Point 2D05, The Angel Centre, 403 St. John St., London, EC1V 4PL.

Formerly *The Post Office Electrical Engineers' Journal* Vols. 1-74: April 1908-January 1982.)

The Board of Editors is not responsible for any statements made nor the opinions expressed in any of the articles or correspondence in *British Telecommunications Engineering* or the *Structured Information Programme* unless any such statement is made specifically by the Board.

© 1999: The Institution of British Telecommunications Engineers.

*British Telecommunications Engineering* and the *Structured Information Programme* are printed in Great Britain by Headley Brothers Ltd.

## Subscriptions and Back Numbers

Annual subscriptions (including postage)—companies, universities, libraries and other bodies £60-00 UK, £65-00 overseas; private individuals £36-00 UK, £41-00 overseas.

Single copies and back numbers—companies etc. £15-00 including postage UK, £16-25 including postage overseas; private individuals £9-00 including postage UK, £10-25 including postage overseas.

Overseas customers paying by cheque should use a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London.

Orders can be paid by credit/debit card by calling +44 (0)171 843 7626.

Orders by post should be addressed to *British Telecommunications Engineering Journal* (Sales), PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL.

Cheques should be made payable to 'IBTE' and should be crossed '& Co'.

An index of recent back issues is on <http://www.ibte.org>

IBTE Members requiring back issues should contact the IBTE Office on +44 (0)171 843 7626.

## Advertisements

All enquiries relating to advertisement space reservations should be addressed to The Advertisement Manager, *British Telecommunications Engineering Journal*, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 (0)171 843 7626. Fax: +44 (0)171 843 7888.)

## Communications

All communications should be addressed to the Editorial Office, *British Telecommunications Engineering Journal*, PP: 2D05, The Angel Centre, 403 St. John St., London EC1V 4PL. (Telephone: +44 (0)171 356 8008. Fax: +44 (0)171 843 7888.)

## Copyright

The entire contents of *British Telecommunications Engineering* and the *Structured Information Programme* are covered by general copyright and the Board of Editors' permission is necessary for reprinting extracts.

Authorisation to photocopy items for internal or personal use, or the internal or personal use of specific clients, is granted by The Institution of British Telecommunications Engineers for users registered with the Copyright Clearance Centre's (CCC's) Transactional Reporting Service, provided that the base fee of \$5-00 per copy is paid directly to CCC, 27 Congress Street, Salem, MA 01970, USA. For those organisations that have been granted a photocopy license by CCC, a separate system of payment has been arranged. Payment is additionally required for copying of articles published prior to 1978.

0262-401X/99 \$5-00 + -00



# THE INSTITUTION OF BRITISH TELECOMMUNICATIONS ENGINEERS

## IBTE Local Contacts

The following is a list of local contacts to whom enquiries about IBTE membership and services can be directed.

Zone / Centre	Contact	Telephone Number
London	Carl Alderman	(0181) 666 5956
Martlesham Heath	Howard King	(01473) 642583
Midlands	Dave Plumb	(01203) 556034
North East	Geoff Jenkinson	(01226) 248141
North West	David White (in Manchester)	(0161) 600 2468
Northern Home Counties	Kevin Rees	0802 195219 (mobile)
Northern Ireland	David White (in Belfast)	(01232) 215354
Paris	IBTE Office	(0171) 843 7622
Scotland	Cameron Scott	0131 668 5804
South West	Tim Jeanes	(0117) 920 5481
Southern Home Counties	John Dymott	(01202) 554001
Wales	Mike Thomas	(01492) 538103

**IBTE Helpline: 0800 028 0209**

**IBTE On-Line (Internet): <http://www.ibte.org>**

**(BT intranet): [http://networks.intra.bt.com/ibte\\_nat/](http://networks.intra.bt.com/ibte_nat/)**

## BRITISH TELECOMMUNICATIONS ENGINEERING

### Journal Subscription Order Form

PLEASE COMPLETE THE FORM IN BLOCK CAPITALS AND RETURN IT, WITH THE APPROPRIATE SUBSCRIPTION PAYMENT, TO THE FOLLOWING ADDRESS:

The IBTE Administration Manager  
Post Point 2D05, The Angel Centre,  
403 St. John St.,  
London EC1V 4PL  
United Kingdom

Annual subscription rates (four issues, including postage and packaging) are as follows:

Companies, universities, libraries and other bodies: **£60-00 UK; £65-00 Overseas**  
Private individuals (to be paid for by personal cheque): **£36-00 UK; £41-00 Overseas**

Please supply four quarterly issues of *British Telecommunications Engineering*. I enclose a cheque/postal order to cover the yearly subscription as given above.

(Cheques and postal orders, payable to 'IBTE', should be enclosed with the order. Overseas customers should pay a sterling cheque drawn on a London bank, a sterling Eurocheque or a sterling travellers cheque payable in London. Customers may wish to use a local subscription agent. Customers wishing to pay by credit/debit card should call +44 (0)171 843 7626.)

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*

Voice-Data Convergence and the Corporate Voice-over-IP Trial	218
IPv6—The Next Generation Internet	225
A Platform for Electronic Commerce Service Trials	232
Agent-Based Solutions for E-Commerce	237
Whither Video?—Videotelephony in Perspective	245
A Review of Optical Wireless What is it and what does it offer?	251
Professional Mobile Radio—Service, Market Position and Technology	262
Developing Customer Service Skills	269
Structured Greenfield Planning	274
GenOSys—Automated Tool for Planning Copper Greenfield Networks	281
'The Eye of the Beholder'—Designing for Colour-Blind Users	291
Evolution of the BT UK Core Transport Network	296
Video-Based Virtual Clubs Using Multimedia to Re-present Technology to Women	301
Convergence Between Public Switching and the Internet	308
<i>Supplement to this issue—Telecommunications Engineering: A Structured Information Programme, Issue 30</i>	

