Extending the LAN via ISDN

Graham Knight[1]

Abstract

This paper describes work carried out in the Department of Computer Science, University College London to make use of Basic Rate ISDN for teleworking. The hardware, software and network architecture used are described, and the costs of using the systems and the performance they offer are discussed.

I.Introduction

This paper describes work carried out in the Computer Science Department at University College London (UCL) during the last 12 months to experiment with the use of ISDN. The work was done as part of the PROOF (Primary Rate ISDN OSI Office Systems) project which was sponsored by the European Commission as part of the ESPRIT programme. The objective of the PROOF project was to demonstrate the integration of ISDN with existing networks in the context of OSI office applications. In fact, most of the work done by UCL used Internet IP in the network layer and considerable experience was gained of using IP over ISDN for teleworking. The trend towards provision of ISDN interfaces as standard workstation features suggests that many others will be indulging in this sort of activity in the future. It is hoped that the UCL experience related here will be of value to such people.

At the start of the work we expected to use Primary Rate ISDN mainly as a way of providing flexible bandwidth to interconnect two LANs. Later, our interest shifted towards teleworking with the Primary Rate interface providing an access point for multiple Basic Rate interfaces in people's homes. To support both of these activities a Primary Rate interface was built and attached to a workstation so as to provide an IP router between Ethernet and ISDN. Sections II-IV below describe the basic mechanics of this set-up. Section V looks at the raw performance of the system, whilst Sections VI and VII takes a more subjective view using a home workstation under Unix and MS-DOS respectively. In section VII the economics of teleworking via ISDN are considered.

II.The ISDN in Britain

Primary Rate ISDN in Europe means 30B+D - 30 circuit-switched 64Kbps channels and one 64Kbps signalling channel. Although the relevant Physical Layer standards for this [1][2] are widely implemented, there are still significant national differences as far as D-Channel signalling protocols are concerned. This is especially so in the UK where a proprietary British Telecom signalling protocol called DASS2[3] is used; this pre-dates the CCITT Q.931[4] protocol. Primary Rate access is now generally available in the UK. Installation costs about $1500 for 30 channels plus the cost of installing the cable - which could be as much as $20,000. Rental is about $50 per month per channel.

Basic Rate ISDN is available anywhere in the UK served by a digital exchange (provided line quality is good enough). This means it is available to well over half the subscribers in the country. The service is 2B+D circuit-switched with Q.931 signalling. Installation is about $600 with rental being $60 per month per channel.

Call charges are exactly as for the PSTN. In BT's case this means a stepped tariff with time being purchased in "units" each of which costs about 6c. At peak times on a local call one unit buys 57.5 seconds of connect time. Off-peak this rises to 220 seconds.

III.Interworking with the ISDN

Early on in the work, and based on the analysis in [5] we considered two alternative approaches to

[1] Graham Knight is from the Department of Computer Science, University College London. knight@cd.ucl.ac.uk

interworking between the ISDN and our existing LANs. One was the "Telecommunications Oriented" (TCO) approach in which the objective was to make ISDN services available to all LAN workstations. Here, "ISDN services" includes bearer services (the. ability to make and receive ISDN calls) plus "Teleservices" (fax and so on). The other approach was the "Data Communications Oriented" (DCO) one. In this case the objective was to integrate the ISDN so that existing LAN hosts did not realise it was there - LAN applications would then work over the ISDN without modification.

Given the paucity of available ISDN services and end-user equipment the TCO approach did not look fruitful. Added to this we were already heavy users of LAN applications and the idea of having these available across the ISDN was attractive. Obviously there would be some performance penalty in going from a 10Mbps Ethernet to a 64Kbps ISDN B-channel. However, in contrast to the Ethernet, a B-channel is both dedicated and duplex so it was anticipated that the resulting performance would still be acceptable. Studies of traffic on our Ethernet [6] suggested that, in most cases, imposing a 64Kbps speed limit would not lead to unacceptable delays. For us, the DCO approach seemed to be the clear win

III.A. Protocol Architecture

Traffic on the UCL LANs is almost exclusively IP and virtually all the applications we use can work over IP. This includes the various OSI applications which were to be used in the PROOF project since these generally operate over the RFC1006 mapping[7]. To get the sort of transparent IP working over the ISDN that we wanted, the "gateway" between the ISDN and the LAN had to be low-level - either an IP router or a MAC bridge. We chose the router solution as this offered good control over which hosts and networks were allowed access to the ISDN.

At that time there were not even draft standards around for carrying IP over circuit-switched ISDN. We decided to adopt the simplest thing that would work and place IP datagrams directly into HDLC frames (Figure 1). Note that this means exactly what it says; we use only the core "framing" aspects of HDLC, we do not implement any protocol procedures. We have since discovered several variants on this framing used by various manufacturers and have added these to the UPI's repertoire. Currently we are in the process of adding the PPP protocol[8].

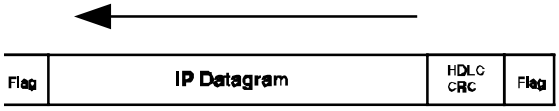

Figure 1. ISDN IP Frame

III.B. ISDN Channel Management

The transparency requirement meant that LAN hosts should play no part in the management of ISDN channels. Thus, one function of the ISDN router would be to open B-Channels as these were needed and close them when they were no longer required. Where the ISDN was providing a LAN-LAN interconnection the number of B-Channels in use could vary from 0 to 30, so quite fine-grain control of bandwidth would be possible. After simulation work [9] we adopted direct estimation of the offered load to a destination using an exponentially weighted moving average with the ***nth*** estimate of the bit rate R_n given by the formula:

$$R_n = aR_c + bR_{n-1}$$

where R_c is the current bit-rate estimate sampled over 200ms, $a + b = 1$. Currently $a = 0.2$ and $b = 0.8$.

IV.ISDN Hardware

The configuration is shown in Figure 2. This shows the Ethernet/ISDN router constructed from a SUN 3/120 workstation connected to the UCL Primary Rate ISDN Interface board (UPI) by a VME bus. This combination handles traffic from several workstations with Basic Rate ISDN interfaces (w/s).

IV.A. The UCL Primary Rate Interface

The UPI is constructed on a 6U extended double Eurocard which, when mounted on a special adaptor plate, fits into the enclosure of a VME-bus based SUN workstation.

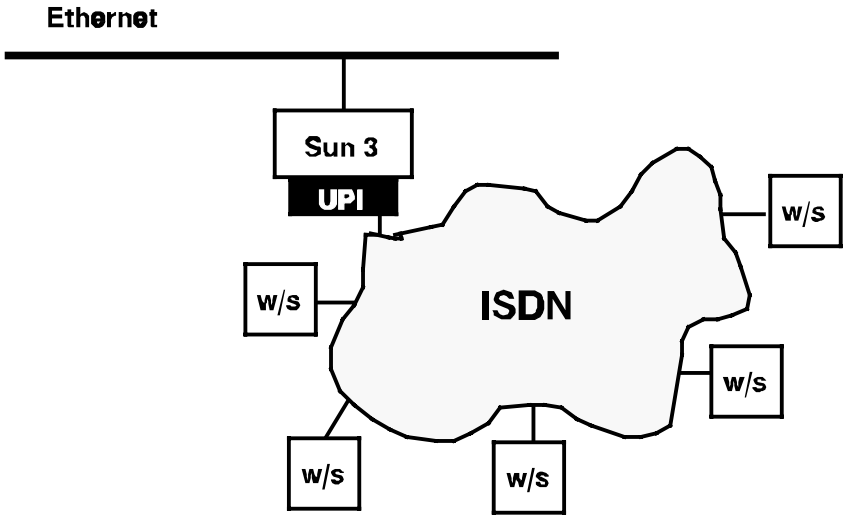

Figure 2 UCL ISDN configuration

The principal board components are as follows (see Figure 3):

- a VME bus interface,
- a Motorolla 68020 processor,
- an AT&T SPYDER synchronous protocol data formatter,
- 1 Mbyte of DRAM, shared between the VME bus interface and the 68020 processor,
- 256 Kbytes of SRAM, shared between the 68020 processor and the SPYDER. This is needed because the SPYDER requires fast access to memory for input.
- ISDN interface circuitry, consisting of Siemens ACFA (Advanced CMOS Frame Aligner) and IPAT (ISDN Primary Access Transceiver) chips.

The basic operation is as follows (again, see Figure 3):

- Data in the form of IP packets is written by the workstation to the shared DRAM across the VME-bus;
- the 68020 processor takes this data and places it into one of up to 30 buffer areas in the shared SRAM - one for each possible B-channel available for PRISDN access;
- the SPYDER takes the data and performs HDLC framing and time division multiplexing on it to produce a single serial data stream;
- this stream is then sent to the ACFA which performs the framing function on the serial data;
- the output from the ACFA is sent to the IPAT which performs the final analogue interface functions.

The reverse path involves a similar process with the ISDN circuitry performing the clock synchronisation and frame extraction functions.

IV.B. Basic Rate Access

Most UCL Computer Science staff use Unix workstations at work - and an increasing number of them use similar systems at home. However, the majority of home systems are MS-DOS PCs. We decided that both classes of home "workstation" should be supported. There has been quite a large number of PC-bus ISDN products on the market in Europe for several years. However, until 1992 few of

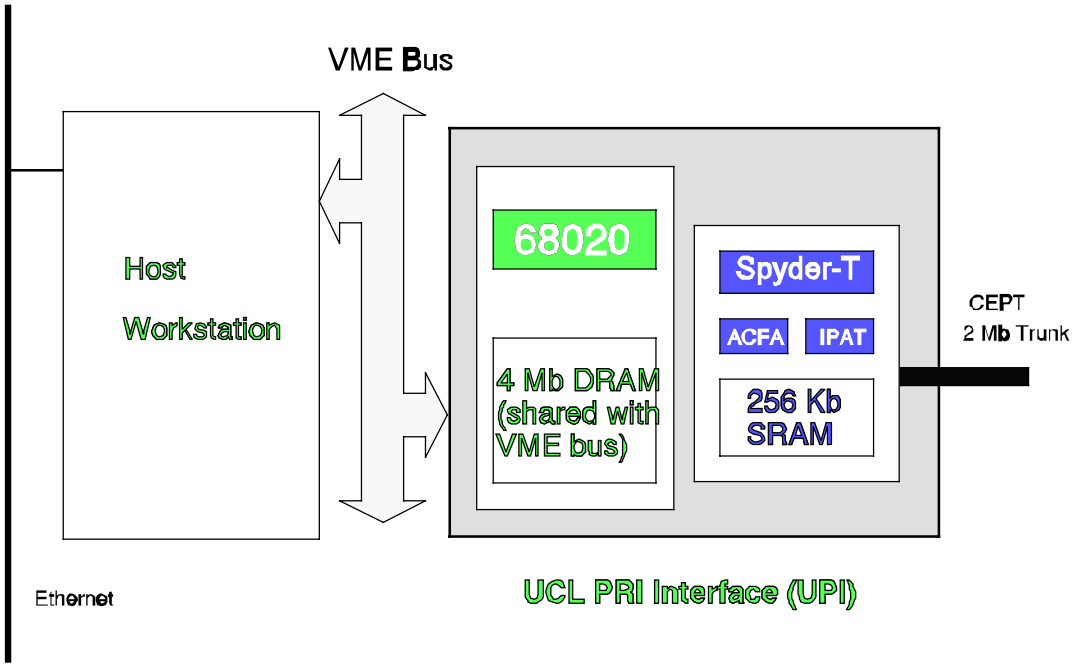

Figure 3. UCL Primary Rate Interface components

these were compatible with the UK Basic Rate ISDN (in spite, or perhaps because, of this conforming well to the European norms!), fewer still had any sort of IP support. We had more success with the Unix workstations as SUN S-bus cards were available from Diehl in Germany which worked with the UK ISDN. Furthermore, IP software was available for these cards from another German company (NetCS) and this used an IP encapsulation compatible with our own. These S-bus cards have been in operation for 18 months now and have proved very satisfactory.

We have now experimented with several PC-bus cards and have settled at present on some cards from kn-X in the UK. In fact these cards come with software to do remote bridging over ISDN and are able to emulate Ethernet interfaces in the PC. We use SUN PC-NFS or KA9Q in the PC to handle IP traffic which is then sent in MAC frames across the ISDN. The UPI strips off the MAC frames and passes on the IP part.

V.Performance

ISDN in the UK is a high-quality service, we see almost no line errors. Therefore "performance" relates mainly to the attached systems and their interfaces.

V.A. Throughput

All of the hardware involved is capable of transmitting packets at close to the theoretical maximum. For example, loop-back tests through the UPI go at around 63Kbps per B-Channel using 800 byte packets.

V.B. Delay

Although we have seen very long round-trip delays on calls from the UK to Japan (presumably due to satellite hops), for local calls round-trip delays are minimal. Of greater importance is the channel set-up delay. Generally this has been substantially less than one second. This comparatively fast set-up is the principle feature of the ISDN which allows

transparent working. Calls can be timed-out by system software and set up again as traffic demands without user intervention. The user sees a slight delay - perhaps in the echo of a typed character - when the channel is set up but this is not an unusual experience for a user used to working on a heavily loaded LAN and seems to be acceptable to our users.

VI.Using an ISDN workstation

The author has been fortunate enough to have had the use of a SUN IPX workstation, equipped with a Diehl card, in his home for the last 9 months. This has enabled some assessment of the effectiveness of ISDN-based teleworking to be made. In fact, "teleworking" is probably the wrong term to describe the work-pattern employed. The principle place of work remained UCL and the author travelled there on most work-days. The home system was used to support a more flexible style of working, with some full days spent working at home (one day a week was the target) and others split between home and office. "Home" in this case was just outside London about one hour's travelling time from UCL. "Flexible working" enabled the abandonment of the railway season ticket (current price about $1800 p.a.) in favour of a late start and off-peak travel. The author now enjoys an after-breakfast cup of tea and reads his email at home whilst his fellow commuters fight in the London crush. He then travels to work in civilised comfort at the reduced off-peak rate.

Had the work-place moved **entirely** to the home it would have made sense to provide large quantities of disc storage on the home system so that system was effectively independent of the office ones. Copies of all system and user files would be kept on the home system disc, access to the office systems would have been needed only for electronic mail and occasional file-transfers. Given that, in practice, the principle site was to remain the office one and that good communications were available, complete duplication of the office environment in the home seemed unnecessary. Instead ISDN would be used to provide transparent access to the office environment. As might be guessed, the arrangement which finally developed was something of a compromise with the home system having a 200MB disc containing the most commonly used system files with a little space left for the user's working files.

As time went by patterns of working were developed which sought to preserve as much as possible the illusion of transparent access to the office whilst at the same time being cost-effective. Three examples below illustrate some of the issues which arise.

VI.A. Processing Electronic Mail

It was considered essential that just one mail-box should be used and that this should be the existing one at UCL. The author uses an entirely text-based mail interface so one possible approach was to remote login to an office system and process messages in the normal way. This requires that an ISDN call should be open throughout the session and is especially inefficient when composing messages; it would clearly be more sensible to run the editor on the home system. However, if only a few messages are to be processed then this is probably the best approach.

Unfortunately, the author, like so many these days, receives a couple of hundred messages a day so a different approach was needed. One possibility was to run the mail interface on the home system accessing mail files remotely using SUN NFS. This certainly works and has the advantage that the composing of messages takes place almost entirely on the home system. If composition takes more than about a minute (during which time no ISDN traffic is generated as no file accesses are made) then the ISDN call times out and is disconnected. However, the particular structure of the mail system in use (the RAND MH Message Handling System) makes this rather inefficient and slow. The first thing that happens in a MH session is that the mailbox file containing the most recent incoming messages is "burst" so that each message is placed in a separate Unix file. From then on, most mail-handling operations involve file accesses - this results in a great deal of ISDN traffic. For example the *scan* command gives a summary of current messages. It works by separately opening and reading the header of every message file. Accessing a file for the first time is expensive in terms of NFS operations and this makes it slow across the ISDN. Running scan in this way results in about one message summary a second - not really acceptable. Furthermore, running the burst operation on the "home" workstation involves copying the entire mailbox file across the

ISDN and then back to the message files which can hardly be said to be efficient.

An attempt was made to identify the most file-intensive MH operations and to arrange for these to run remotely through the use of the Unix *rsh* command (*rsh* allows a single command to be run on a remote machine). Thus the *inc* command (which initiates the "burst" of the mailbox file) and the *scan* command were arranged to run on a UCL machine rather than on the "home" machine. *Show*, which displays individual messages, was run on the "home" machine. Unfortunately *rsh* itself is rather sluggish as it has to authenticate the user with the remote system on each command. User-perceived performance was little improved therefore.

The most successful system so far has been the use of a shell-script called "*shadow*" which is run at the start of a session. *Shadow* maintains duplicate copies of a user's mail messages on the home system. Each time it is run it tries to move the minimum quantity of data needed to achieve consistency between the "home" and the "UCL" copies. Once these copies have been made, the bulk of mail processing operations on the home system involve only local disc accesses. This is not quite the full story as MH allows related messages to be organised into multiple "folders" separate from the main "incoming mail" folder ("inbox"). Each folder is a separate Unix directory. To his shame the author discovered he has close to 10MB of messages in folders, most of which are rarely accessed. These folders are not duplicated by *shadow* as too much network traffic would be generated whilst verifying consistency. Instead, "soft links" are used to make these folders appear to be part of the home-system's file-space. Access to messages in these folders is therefore via NFS and is slow. This is seems acceptable given the low frequency of access to these folders.

The *shadow* scheme is a compromise. Clearly transparency is lost as there is no equivalent command on the office systems. However, once *shadow* has run mail processing proceeds exactly as at the office.

VI.B. Software Development

The author is responsible for the maintenance of a fairly large piece of software (the code for the UPI). The source for this software is stored on a file-server at UCL and is occasionally modified across the ISDN. Given that the file-store can be remote mounted across the ISDN from the home system, it would be possible to run the complete software development cycle (edit, compile, link, test) on the home system's processor and thus achieve complete transparency. In fact this does not work too well.

Editing is mainly fine. The whole file gets transferred across the ISDN into the editor's address space at the start of the edit. Once it is there are no more ISDN accesses until the edited version is written out. One difficulty arises from the modular structure of most systems which leads to software being managed in a multiplicity of small files. This leads to more file reads than would otherwise be the case and frequent recourse to programs like the Unix *grep* to help locate variables; obviously *grep* has to process every byte in the file and so is very inefficient when run across the ISDN.

Another major problem comes with program linking. This usually involves searching large, remote libraries. Worse still, access to these libraries is random not sequential. NFS is poorly optimised for this type of access and adopts a "read-ahead" strategy which causes it to read (usually) 8Kbytes even when only a handful are required. At 64Kbps 8Kbytes takes one second so things get very slow.

Once more we must make a compromise with transparency and be aware that, whilst editing should use the local processor, compiling and linking are best done via a remote login.

VI.C. Accessing on-line Manuals

In addition to the normal Unix manual pages Sun now provide a comprehensive set of guides and tutorials in Postscript format. A tool called "*answerbook*" is provided to browse this information. The volume of information provided is much too great to be stored on the comparatively small disc of the home workstation and so it must be accessed

Session	Elapsed Time (min)	Connect Time(min)	% Time Connected	Longest Connection (min)	Approx. cost ($US)
1	152	129	85	19	8.93
2	26	79	30	17	5.97
3	498	73	15	10	5.61
4	577	195	34	23	12.84
5	622	93	15	9	6.80
6	162	57	35	5	4.66
7	136	17	13	5	2.30

Table 1. Teleworking Session Statistics

remotely. One could run *answerbook* on the "home" machine accessing the remote information via NFS.

However, this would make searching the information extremely inefficient. Instead, answerbook can be run on a remote machine with the Graphical User Interface (GUI) being exported over the ISDN to the home workstation. Export of GUIs in this way is a natural characteristic of the X-Window system which is used to support *answerbook*. Obviously there is some performance loss when X is used in this way - display of a page of text takes from 2 to 3 seconds. However, the performance is quite acceptable given that use of the service is not all that frequent.

VI.D. Some pitfalls

Three examples show how apparently innocent activities can lead to excessive ISDN usage (and hence cost)

- The author uses a file-viewing program which occasionally polls the file to see if it has grown. These polls will keep an ISDN channel open when a remote file is being viewed.
- As noted in the *answerbook* example, it is convenient sometimes to run a program remotely but to have its windows displayed on the local console via X-windows. When so doing one should beware the type of program which asks to be informed every time the mouse cursor enters or leaves its window. This will cause ISDN calls to be set up merely to report spurious mouse movement when the application is, in fact, quiescent.
- Assuming that the local disc is not big enough to store all the system programs which might be required, some programs may be loaded from a remote file-server across the ISDN. If directories on the remote file-server are included in the user's search path they will automatically be searched for executable files. Mis-typing a command can cause futile and costly searches of such directories.

An essential aid to catching such problems is some sort of visual or audible indication that ISDN access is occurring.

VII.MS-DOS PCs

The experience described above was obtained using a SUN workstation and Unix. Several of our users are now using MS-DOS PCs with a variety of PC IP implementations - principally PC-NFS. The experience here has not been quite so satisfactory. In part this has been due to poor reliability of the networking software. We suspect this is a consequence of the generally poor support in MS-DOS for general purpose networking; with Windows, PC-NFS and ISDN drivers all competing for resources in a relatively unstructured way there are bound to be conflicts and it is not unusual for the

whole system to "hang". A more fundamental weakness comes from the general style of working from PCs to Unix systems. Although PC-NFS allows file sharing - and hence access to Unix files across the ISDN - it does not generally have the tools to process the files once it has accessed them. For example, we do not have any good PC tools for processing Unix mail files. This tends to lead to users using the ISDN mainly for telnet sessions which is very expensive in terms of ISDN connect time. In the future we will try and remedy this lack of PC tools.

VIII.Cost Effectiveness

Table 1 gives details of 7 recent "teleworking" sessions of varying lengths from a Unix workstation. There were a mixture of activities involved; mail processing at intervals, document preparation with references to on-line documents, software development etc.

Here, the "elapsed time" is the total length of the session for which the ISDN was potentially in use. The cost approximations include call charges and rental but make no allowance for installation. Calls were made at the local rate; about half at the "standard rate" (which buys 80 seconds for 6c) and about half at "peak rate" (57.5 secs for 6c). Evening work would be about one-third these costs.

Currently the return train fare from the home site to UCL out of peak hours is $7.28 (Normal fare $10.92). Thus, for the author, teleworking with an ISDN workstation is cost-effective if it results in an extra day at home. At present the UK ISDN tariff seems to be set fairly high in accordance with what is seen as a "premium" service. It is to be hoped that more reasonable tariffs will be available in the future which will swing the balance further in ISDN's favour.

IX.Conclusions and further work

ISDN can form the basis for efficient and cost-effective data communications. Its performance, especially with regard to call set-up delay, is good enough to allow transparent working with call control being in the system's rather than the user's hands.

Using ISDN as a basis for teleworking works well. However, the home workstation user needs to be aware of the location of data and the way it is processed by an application in order to avoid excessive transfers across the ISDN and hence increased costs and delays.

During the next year we expect to add several more users of Basic Rate workstations. We will also embark on a series of experiments involving packet video over ISDN. We anticipate having to tune the channel set-up and allocation strategies in order to support packet video adequately. At present, the management facilities on the gateway are crude. We expect to enhance these so that remote management, according to OSI principles, is possible.

X.References

[1] CCITT G.703, "Physical and Electrical Characteristics of hierarchical digital interfaces", CCITT Blue Book Fascicle III.4, 1988

[2] CCITT G.704, "Synchronous Frame Structures used at Primary and Secondary Hierarchical Levels", CCITT Blue Book Fascicle III.4, 1988

[3] British Telecom, "Digital Access Signalling System No. 2 (DASS2)", Volume 1 - Core Document - Issue 2, Draft 2, Sept 1988

[4] CCITT Q.931, "ISDN User-Network Interface Layer 3 - Specification of Basic Call Control", CCITT Blue Book, Fascicle Vi.8, 1988

[5] A. Cobb, "Interworking between IEEE 802 LANs and an ISDN. (Version 4.0)", York University TR CS-88-10, North York, Ontario, Jul 1988

[6] G. Knight, G.Pavlou, "Basic Rate ISDN Workstation Traffic Patterns" "Computer Communications" vol 13 no. 10 pp 587-594, Butterworth-Heinemann, December 1990.

[7] M.T Rose, D.E Case, "ISO transport services on top of the TCP: Version 3.", Internet RFC 1006, 1987 May

[8] W. Simpson, "The Point to Point Protocol (PPP) for the transmission of Multi-protocol Datagrams over Point-to-Point lines", RFC 1331, May 1992

[9] X Du, G. knight, "Further Study of Channel Management Issues in LAN-ISDN Gateway", "Computer Communications" vol 16, no 4, pp251-259, Butterworth-Heinemann, UK, April 1993

Author Information

Graham Knight has worked in the Data Communications and Distributed Systems group at UCL since 1980; initially as a Research Assistant, latterly as a lecturer. His interests include the use of ISDN for teleworking and network and system management. He has led UCL teams in three ESPRIT projects in these areas.