EurOpen

News Letter

Number 3

Autumn 1991

Volume 11

EurOpen Forum For Open Systems

Volume 11, Number 3, Autumn 1991

ISSN 1011-4211

Editorial Team

Frances Brazier
Publication Executive

Department of Computer Science
Vrije Universitei
de Boelelaan 1081
1081 HV Amsterdam
The Netherlands

Telephone +31 20 5485588
Facsimile +31 20 6427705

Email frances@cs.vu.nl

Alain Williams
Editor

Parliament Hill Computers Ltd
7 Prospect Street
Caversham
Berkshire RG4 8JB
United Kingdom

Telephone +44 734 461232
Facsimile +44 734 474194

Email addw@phcomp.co.uk

Gina Baikenycz
Typesetter

The Instruction Set
City House
190 City Road
London EC1V 2QH
United Kingdom

Telephone +44 71 253 5121

Email gina@inset.co.uk

Joe Di Mascio
Printer

Rank Xerox Business Services Ltd
89-91 Hampstead Road
London
NW1 3EL
United Kingdom

Telephone +44 71 388 5607
Facsimile +44 71 388 8841

Publisher

EurOpen
Owles Hall
Buntingford
Hertfordshire SG9 9PL
United Kingdom

Telephone +44 763 73039
Facsimile +44 763 73255

Email europen@EU.net

Contents

Editorial

Alain D D Williams
Parliament Hill Computers Ltd
Caversham
United Kingdom

Email I addw@phcomp.co.uk

Technology Reviews

As announced in the summer there is a new section in this issue. You will find starting on page 71 a collection of write ups on backup products.

J M van der Waaij introduces the topic and provides are few criteria to help you choose between what is covered on the following pages.

It is hoped that the Technology Reviews will prove useful to you - perhaps something to keep for future reference. Please contact the vendors directly for more information.

Hot Topics

John Crowcroft describes for you work that he has been doing on Multi-media communication. He demonstrated it at a recent LUUG (London UNIX Users' Group) meeting where half of the evening's lecture was delivered in live real time colour, from New York onto the monitor of a Sun workstation. Very impressive. Turn to page 2.

Client-Server applications seem to be the way of the future, but there is a long way to go before it becomes an every-day technology. Daniel Bernstein describes the UCSPI interface and toolkit which tries to solve some of the problems. I know that the use of such toolkits helps to make my life easier, how about you? See page 7.

Conferences

Dieter Glauss attended the Tromsø conference and gives his enthusiastic impressions on page 11.

The next EurOpen conference is in Budapest in September. A full programme of tutorials and technical sessions can be found starting on page 16. As you can see this is another event that must not be missed.

When you leave Budapest you will naturally want to go to the SUUG conference in Vladimir the week after - find out about it on page 31.

Another event to write into your diary is UNIX '92 (page 38).

Desk Top Publishing

This will be continued in the next newsletter where the topic will be desk top publishing.

The latest generation of tools are very different to the traditional UNIX troff. They can revolutionise your presentation quality and give great productivity improvements at the same time. Look at what happened to this newsletter when we made the change!

Vendors of such products are invited to write about their offering. The article should contain the following information:

- What the product is called.
- Major features.
- What are the Unique Selling Points.
- Why it is better than the competition in general.
- How long has the product been available, how well has it sold?
- What support is available?
- What documentation is available?
- Performance, how fast is it?
- Are different (human) languages supported in the user interface, are configurations for European languages available 'off the peg'?
- Availability - ie what systems, what counties? Up to 2 names and addresses may be given.
- Pricing information. If it is modular give module prices (eg the network may be extra). For base price comparisons please quote the price for an 8 user 386 box.

Products that are not available in Europe or only available in one European country will be rejected.

Advertising space is available and may be booked in the same issue.

All those interested in participating should contact me by 23 September, final copy will be expected by 14 October.

The next topic (for the Spring 1992 issue) will be Spreadsheets.

Multimedia Packet TeleConferencing, Systems and Workstations

J. Crowcroft & I. Wakeman
UCL Department of Computer Science
University College London
United Kingdom

Email j.crowcroft@cs.ucl.ac.uk

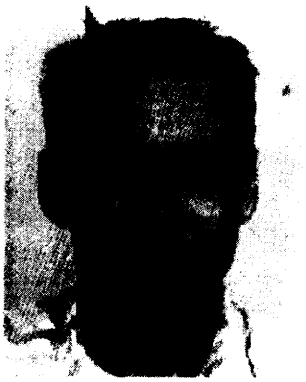

John was educated at Trinity College, Cambridge and University College London and has an MSc with distinction, Computer Science and is currently working for a PhD in "Lightweight Protocols for Distributed Systems"

He is a Member of British Computer Society, a Chartered Engineer and a member of Internet Advisory Board, End to End Research Group since its inception.

He is on the Technical Editorial Board, IEEE Networks Journal and is a member of ACM SIGCOMM Program Committee, 1990

His interests are: Data Communications including Network and Transport Protocols, Routing and Congestion Control, Distributed Systems, Multimedia Systems.

Abstract

The Department of Computer Science, University College London like many other research organisations recently, has started working on Multi-media communication. There are two main areas of activity:

Firstly we are building wide area communications support for real time traffic such as video and voice, and synchronisation between this and other traffic. Since we already have two teleconferencing rooms at UCL, a further area of work is in integrating these services into workstations.[1]

Introduction

Last year, the Department of Computer Science, University College London, commissioned what we believe to be the world's first Multimedia TeleConferencing system over International Packet Switched Networks.

The UCL Conferencing Facility is described in a review article in [TRICOMM]. In the rest of this article, we describe ongoing research on video conferencing under the following three headings:

- The Communications Architecture & Protocols:
  - i. Workstations commonly use the TCP/IP protocols [TCP]. The ST protocol [ST] is an extension of these to support packet video and voice. Some research is going on in how to make conventional Internet systems also support real time traffic.
  - ii. Future networks will be based on Asynchronous Transfer Mode(ATM) - [ATM] this has impact on workstation network interface and protocol support designs.
  - iii. Video Coder/Decoders (CODECs) not only provide Analog to Digital conversion, but also highly efficient image compression, both for compact storage, and timely transmission. They will soon be incorporated in workstation hardware.
- The impact of real time protocols on the Workstation Architecture and the Operating System.
  - i. If video and voice is carried in ST we must support this in the workstation. We also need support for video & audio device I/O.
  - ii. If we are running our teleconferencing over an ATM network, we must support ATM network interfaces.
  - iii. If we need image compression due to lack of bandwidth, this needs to be integrated with image capture, display and transmission.
- Integration of video and voice with text and graphics on the workstation screen involves adding functionality to Windowing Systems.

Communications Architecture

Protocols in End Systems (host computers/workstations) and Intermediate Systems (Routers/Bridges) both need to be modified to support video and voice traffic as well as ordinary data:

- They must support low variation of bandwidth, delay and error

1. This work was supported in part by DARPA under Contract Number N00014-86-0092. Any views expressed in this paper are those of the authors alone.

for video and voice traffic[2]

- Multicasting packets (and multicast forwarding in routers) increases efficiency of use when supporting more than two way conferencing over LANs and WANs, by up to n times for n participants. Since video requires high bandwidth, this is important.
- Traffic from different streams may belong to the same conference and thus needs synchronising (not just lip-synch but mouse pointer movement and other data must be synchronised with video/voice).

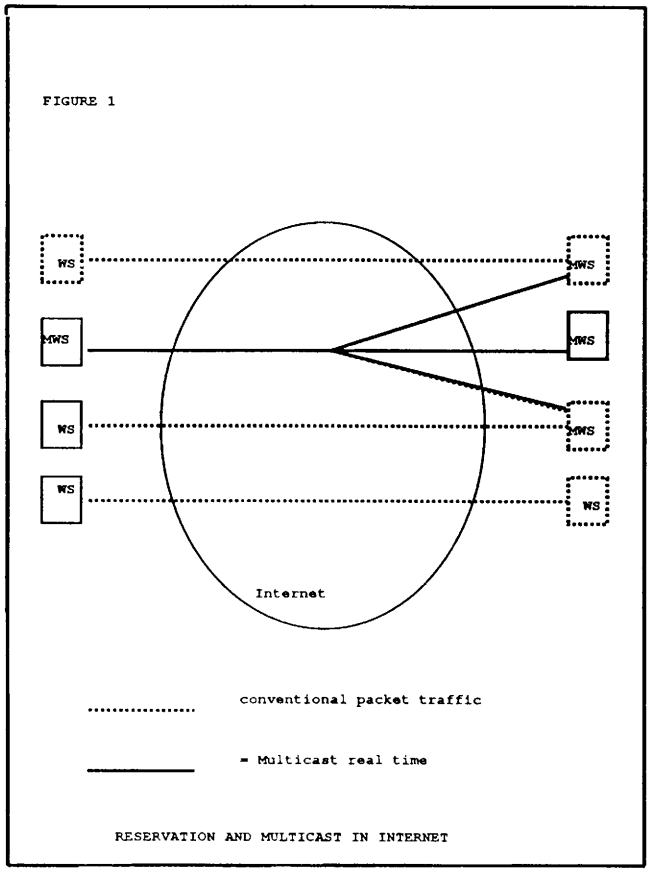

Figure 1: b/w guarantees without reservation and multicast for packet efficiency

ST versus IP or CLNP with multicast

The communications architecture UCL is currently using is based around the DARPA Internet Protocol suite, but with certain enhancements to support real time traffic.

These are provided by the ST protocol, running in Butterfly gateways for end and intermediate systems, or in Sun Sparcstations for end systems.

To provide video compression so that we can run our conference over typical wide area networks, analogue video is converted into compressed digital video by a CODEC.

The ST gateway then runs two further protocols to carry the video and voice:

- Packet Video Protocol/PVP
- Voice Traffic protocol/NVP

Each of these is tuned to the special requirements of video and voice.

ST is a variant of the Internet Datagram Protocol (IP). The main difference is that it provides resource reservation facilities for PVP and NVP to declare requirements in advance. A control message is sent using the Stream Control Message Protocol to setup paths through the network from the source of a conference site to all the destinations ("Multicast"). This message carries a "FlowSpec" to all the ST/IP [Multicast] gateways which says what the delay/ bandwidth/error requirements of the requester are.

In fact, the ST protocol specification has little to say about how the "FlowSpec"is enforced or how multi- destination delivery is achieved- it is largely left to underlying mechanisms.

Because of this, we are investigating whether the mechanisms for loose resource reservation and multicast can be implemented for conventional IP hosts and routers, so that we can run on almost un-modified parts of the Internet.[3]

Clearly a Connection Oriented Network Service would be an alternative choice, so long as the CONS suppressed retransmissions which would cause undue (and unnecessary) delay variations over long haul lines. It would also be necessary to have multicast facilities and synchronisation points in such a CONS stream.

Asynchronous Transfer Mode

Some underlying technologies provide multicast and resource reservation already. One of the emerging standards for future wide area networks is the Asynchronous Transfer Mode of the Broadband ISDN service. This is highly suited to video and voice traffic (not surprisingly, as it is developed by Telecom Companies as a future replacement for the worldwide phone system). This is achieved mainly through using small, fixed size packets.

If we can attach our workstation to such a network, it will have a lot less work to do to maintain the characteristics needed for video and voice.

On the other hand, it may have to do a lot more to run ordinary data traffic as we shall see below.

CODECs, Varying the Bit Rate

Recently there has been a lot of work on standard image compression techniques. We now have JPEG for static pictures and MPEG (from ISO) and H.261 (from CCITT) for moving picture storage and transmission. [MPEG]

Chipsets for these are being developed rapidly, so that the cost of an Analog to Digital (A2D) converter which can do image compression may be very low within a year or two.

Moving image compression has two main components of interest:

- Intra-frame compression, where each frame/image is separately compressed (using something like DCT).
- Inter-frame compression, where some kind of image differencing is done (this is highly efficient for video without scene changes).

There are interesting interactions between packet loss and image de-compression algorithms when inter-frame compression is in use. Some inter-frame compression techniques still send an independent "re-synchronising" frame (intra-frame compressed) to anticipate problems like this.[4]

---

2. Delay, in principle, must vary by less than 1/2 a frame time for video (e.g. 1/60 or 1/50 a second) and 1-200 msecs for voice. Bandwidth can vary if it results in loss, but only up to a few 10s of percent.
MPEG proposes both forward and backward prediction with reference points at least 5 frames apart - this gives 6 frames of delay in both the encoding and decoding processes.

3. IP could easily be replaced with the ISO CLNP, if there was a standard for multicast.

Multi-media Workstation Architecture

Once the network is in place to carry traffic around with real time characteristics, we need to get this traffic in and out of the workstation. As well as keyboard and mouse, we need camera and microphone (and possibly scanners). There are then various choices as to paths for analog and digital data through the workstation, and onto the network, together with appropriate places to perform A2D conversions, and possible picture (and audio) compression/decompression:

- Video display on a window on the workstation screen. (Cards are available to do this now on many popular workstations).
- There are two further levels of integration. One is to put the CODEC on aboard in the workstation, and have it talk to the analogue video controller, with some other channel to interact with the windowing system.
- The second is to have the digital video stream be converted straight into the frame buffer (in the correct place), and avoid the use of a CODEC for conversion between the network and on-screen. The CODEC is still needed for camera to the network. Indeed, we would need the CODEC to talk through a high(-ish) speed serial link to the workstation, which would then packetise and synchronise this with audio and other signals.

TCP/IP and ST support

ST/IP from a set of remote sites and reverse path traffic are processed as follows:

- ST from remote sites arrives at the Ethernet interface (or high speed serial if that connects to remote site), and is de-multiplexed from IP up to ST to the PVP and NVP protocol modules in the workstation.
- The ST stream then needs to be decoded by a packetiser and the CODEC. This can be done by running the packetiser code as a kernel process, feeding the output(s) to the high speed serial link out to the CODEC which then runs the analogue signal to the video card input (which then talks (out of band) to the windowing system, and puts the video in the frame buffer in the appropriate area, and thus on the screen.
- The reverse path involves analogue signal from the camera (or a switch) to the CODEC.
- Then the compressed coded traffic goes through the high speed serial input to the workstation, and is packetised.
- Then its put through PVP, and ST/IP and out onto the network.
- Audio is similar, but makes use of the audio hardware on the workstation to do the coding/decoding, and then maps this to a standard audio encoding through the NVP process, and thence to/from the net.

The current video protocol stack is like this:

Analogue Video Camera	Video Card
Serial Line	Serial Line
PVP	PVP
IP	IP

The audio protocol stack is like this:

Analogue Audio Camera	Speaker
NVP	NVP
IP	IP

When using ordinary packet switched protocols to carry the video/voice traffic, we need to be able to schedule the processing of arriving packets so that they can be delivered to output devices smoothly (we cannot overrun or starve a loudspeaker or framebuffer).

This Real Time support in operating system, is rather an anathema to many versions of UNIX: Either all the work must be in the kernel or we must change the scheduler.

Workstation ATM support

When the network is delivering many small fixed size cells, video can be digitised (uncompressed) directly from a camera into the network, with the workstation simply providing call setup and teardown. Inbound video could be received direct into the workstation framebuffer.

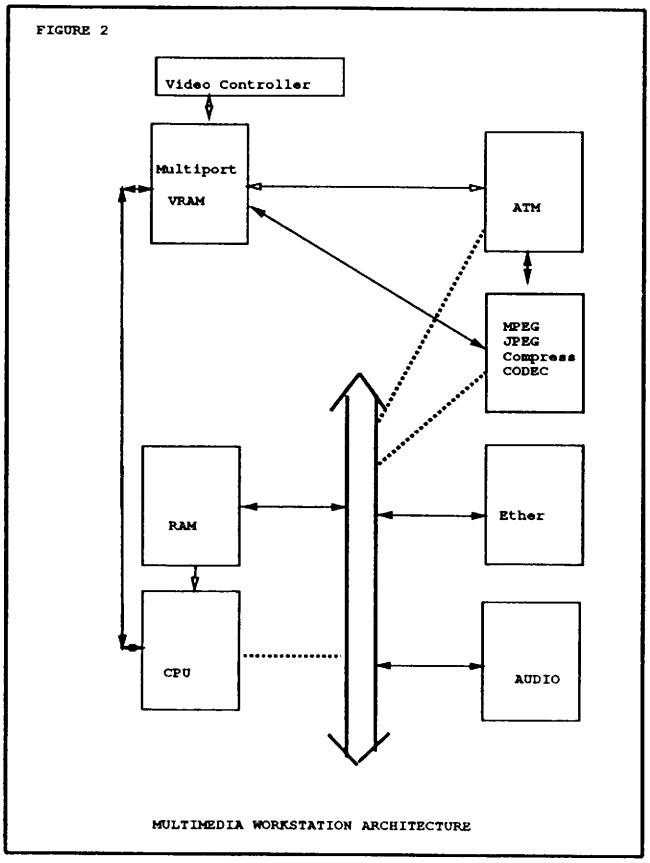

Figure 2: Analog and Digital Data paths

In this case, we no longer need the full real time support for the video and voice streams. However, we do need means to synchronise events between the video stream and data (mouse movements).

Compression- MPEG on Workstation

We envisage MPEG chips or the like being available in the same way as audio, very soon (sooner than ATM networks).

The minimum effect of having a low cost compressed digital video I/O device will be to make wide area networked multi-media more accessible, and local area more popular. Until we have ATM, however, the workstations and routers will more work to do to

4. The present UCL system uses a proprietary CODEC. A CCITT version will shortly become available. We believe it will be feasible to place different CODECs on a single analogue switch and relay video this way.

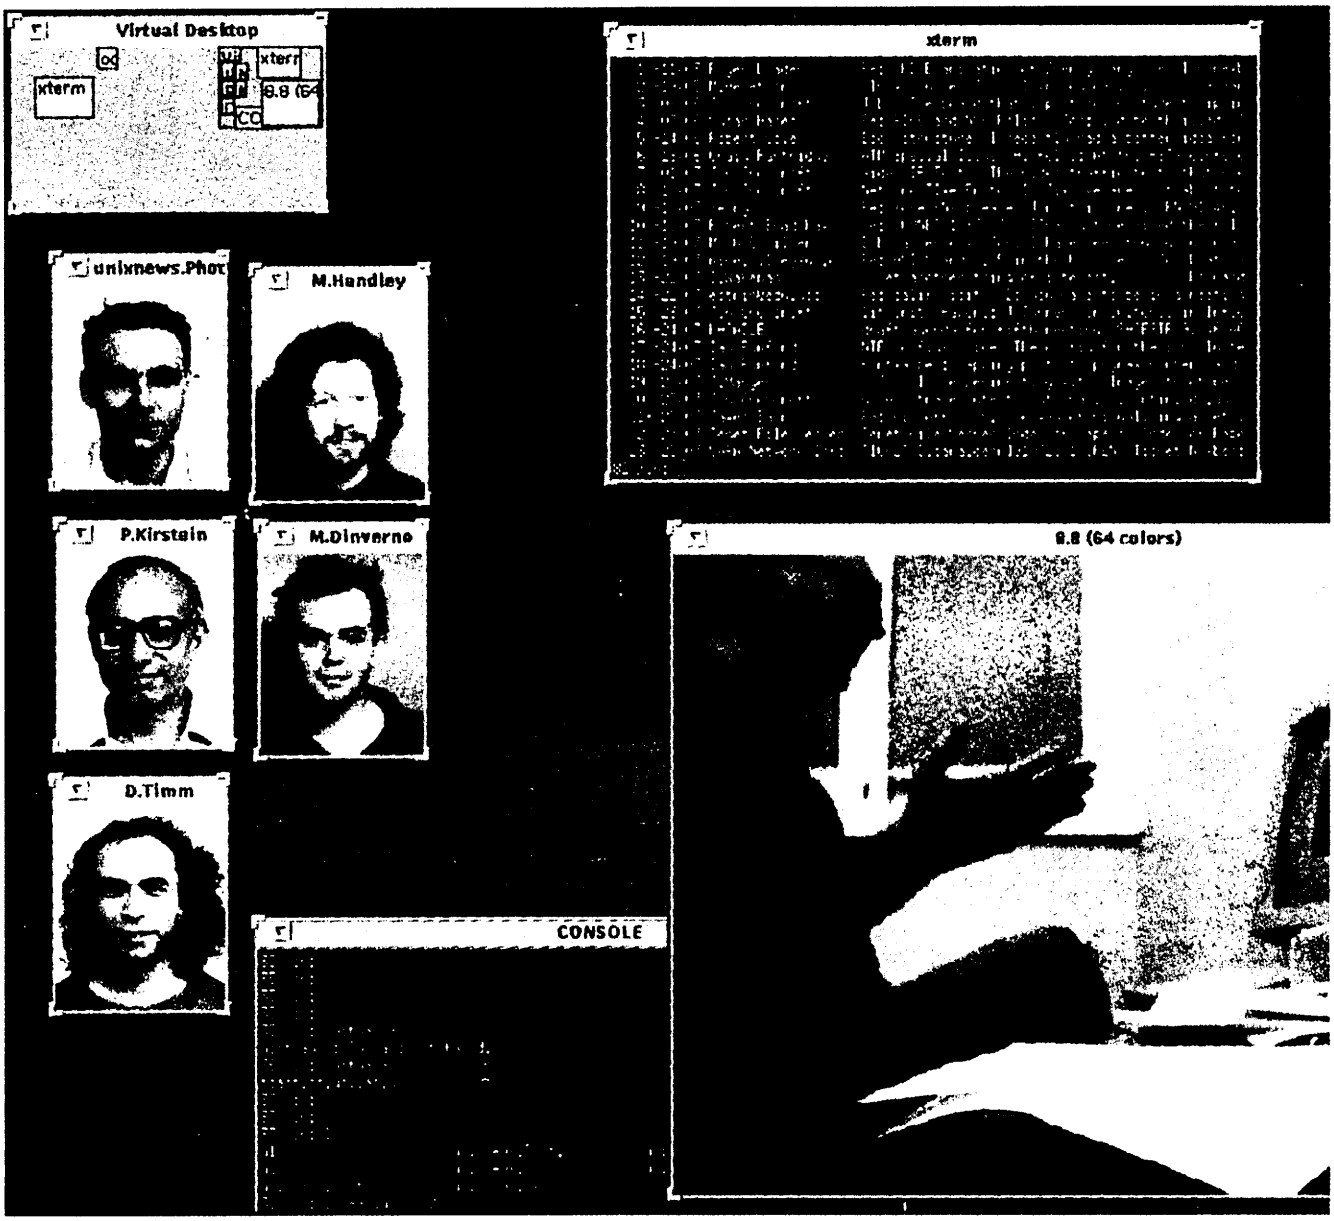

Figure 3: Cyber Space Window Manager

maintain the delay/bandwidth/error characteristics. Video hardware will not obviate the need for clever routers and scheduling.

Shared Windows and Shared Views

Many people have started to add the facilities to networked windowing systems [X] to share views of normal data (text/graphics) on screen [Shared X]. Some work has been done on adding in video windows [VEX]. Some window managers support extended views of a workspace [TVTWM].

To share the output window from an X Client to many workstations is reasonable easily achieved, and with a reliable multicast transport protocol, can be done efficiently[5]. Essentially, the XClient makes contact with an X redistribution server which in turn contacts all the participants servers. Some out of band management (part of floor control) mechanism chooses who has input (which server keyboard/mouse events reach the client) at any one time (or there may be a free-for-all).

5. An application could be shared by taking the command language or interface to the application, and distributing that rather than the window's view on the application. If the application does file or database access, one could even share the access at that level, and replicate the application itself. The disadvantage of this approach is that one then needs to implement some type of transaction processing to avoid consistency problems on access to shared data. The advantage is increased concurrency. Sharing a terminal emulator is so straightforward and allows all programs with a text based command language to be trivially shared, that the more application specific approach seems unnecessary.

Choosing where to add digital video to a networked window system involves several choices:

- Is the video sent by an ordinary client as a sequence of window bitmaps of some kind to the display server, using the ordinary protocol?
- Is there an extension to the windowing protocol to carry video (possibly encoded using MPEG)?
- Is the video client/server a separate (but integrally controlled) system, due to its continuous/real time nature?

We are interested in how we can combine these facilities with the scheduling necessary in UNIX and with window management, and more importantly, with Floor Control and Navigation.

- Floor control is essentially concerned with who has input control to the shared workspace.
- Navigation is mainly concerned with finding who and what are available to share workspaces.

We are devising a combination of these facilities to form what might be termed a Cyber Space windowing system [Gibson].

The current prototype can be seen in figure 3.

Conclusions

Re-designing networks, workstations, window systems and operating systems to support real time multi-media conferencing is not easy. None of these components was designed to meet the requirements in the first place. We believe that re-engineering in situ is hard, but still possible.

Acknowledgements

The work described here involved contributions from many disparate groups. In itself it is a small extension of previous work done elsewhere. In particular most of the technical background work for a packet switched conferencing system was designed under the DARPA Wideband Packet Satellite project, by Bolt Beranak and Newman Inc, in the US. Steve Casnerat ISI/USC is responsible for the packetizing and multi-media switch/camera control programmes. The LIVENet project at UCL, in particular Richard Beckworth, the LIVENet director, and Gordon Jameson, the director of the UCL Audio-Visual Department at UCL and Gerry Jolly have helped. Finally, many at UCL, ULCC and in the US have contributed to the smooth installation and operation of the UK-US link and packet switching equipment.

References

[TRICOMM] Multimedia Teleconferencing over International Packet Switch Network, Crowcroft, Kirstein Timm, Proc of IEEE TriComm '91, Chapel Hill, North Carolina, USA, April 1991.

[TCP] The DARPA Internet Protocol Suite, B.M. Leiner, R.H. Cole, J.B. Postel, D. Mills, Proceedings INFOCOM85, IEEE, March 1985.

[ST] RFC1190 Topolcic, C.,ed. Experimental Internet Stream Protocol: Version 2 (ST-II). 1990 October

[ATM] Asynchronous Transfer Mode- T1.M1Report

[Multicast] Crowcroft, J., Paliwoda, K. "A Multicast Transport Protocol" Proceedings of the SIGCOMM 88 Symposium, 1988.

[MPEG] Standards and the Emergence of Digital Multimedia Systems, Communications of the ACM, Special Issue, April '91, Vol34, No 4.

[X] Scheifler, Robert and Jim Gettys, "The X Window System," ACM Transactions on Graphics, vol. 5, no. 2, pp. 79-109, April, 1986.

[Shared X] Copyright (C)1990 by Digital Equipment Corporation. Author: Michael P. Altenhofen, CEC Karlsruhe.

[VEX] Brunhoff, Todd, "Pleasing The Eye," UNIX Review, vol. 7, no. 10, pp. 64-72, November, 1989.

[TVTWM] Tom La Strange, Solbourne Computer; Jim Fulton, MIT X Consortium; Steve Pitschke, Stardent Computer; Keith Packard, MIT X Consortium; Dave Payne, Apple Computer.

[Gibson] Neuromancer, Count Zero, Mona Lisa Overdrive, 1985-1990.

UNIX Client-Server Program Interface, UCSPI-1991

Daniel J. Bernstein
UCSPI Administrator

Email brnstnd@nyu.edu

Introduction

This document describes the UNIX Client-Server Program Interface, UCSPI (ooks-pie), a standard program-based interface to client-server communications tools under UNIX and its variants.

UCSPI is designed to solve several problems. First, although most networked applications today are essentially independent of the underlying communications protocol, each program must provide direct support for each protocol. To run a TCP/IP program over DECnet or through other IPC mechanisms requires rewriting and recompiling the program. In contrast, an application written for UCSPI will work without change over any communications medium.

Second, no standard tools are provided for shell scripts to take advantage of networking. A UCSPI-compliant communications tool is automatically usable within a shell script or from other languages.

Third, current applications have no mechanism for passing information about a communications link. A typical TCP/IP connection to a BSD machine may go through inetd, an access controller, and one or two applications. Each level may do several system calls, disk reads, and network requests for a name lookup and possibly authentication, instead of getting the information once and passing it down to the next level. In the case of logins, a user shell cannot even determine reliably where the connection is from. In contrast, a UCSPI-compliant tool will collect all necessary information and pass it through environment variables to the lower layers.

A program is not compliant with UCSPI-1991 if it does not satisfy any rule marked MUST. It is compliant if it satisfies all the rules marked MUST. In that case, it is unconditionally compliant if it satisfies all the rules marked SHOULD, and conditionally compliant otherwise.

The General UCSPI Interface

UCSPI clients and servers are executable programs. They MUST accept a command line in the following general format:

```
<program> <options> <address> <userprogram>
<arg...>
```

Here <program> is the name of the client or server, <options> are zero or more option arguments, <address> is a protocol-specific address, <userprogram> is a user-specified program to run upon any connection, and <arg...> are zero or more arguments to pass through to <userprogram> without interpretation.

UCSPI killers are executable programs. They MUST accept a command line in the following general format:

```
<program> <options> <address>
```

Each client, server, and killer MUST support some protocol, as defined below. <address> SHOULD take up a constant number of arguments, as defined by the protocol. <options> SHOULD be processed by the getopt standard; in particular, an argument of -- SHOULD terminate <options>.

If <program> is a client or server, it MUST change descriptors 6 and 7 and certain environment variables as described below before executing <userprogram>. <program> MUST NOT interpret, quote, or unquote <arg...> in any way. <program> may or may not fork before executing <userprogram>. <program> SHOULD NOT change its process state in any other way before executing <userprogram>. <program> SHOULD close descriptors 6 and 7 immediately upon entry.

<program> MUST NOT assume that any other particular descriptors are open or closed. <program> may assume that a certain number of descriptors are available. <program> SHOULD NOT assume anything else about its process state; in particular, <program> SHOULD NOT attempt to find its controlling terminal device, and <program> SHOULD NOT use *getlogin()*.

Servers

If <program> is a UCSPI server, it MUST wait for a client to connect to <address>. It SHOULD NOT accept connections from clients that do not support the same protocol. Upon

accepting a connection, it MUST spawn <userprogram> with the given <arg...>, descriptor 6 reading from the connection, and descriptor 7 writing to the connection. It MUST then continue accepting connections to the same <address>.

<program> MUST set the following environment variables for <userprogram>: PROTO=<PROTO>, where <PROTO> is the communications protocol used by the connection; <PROTO>LOCAL=<serveraddress>, where <serveraddress> is <address> in a server format defined by the protocol; <PROTO>REMOTE=<clientaddress>, where <clientaddress> is some information about the client in a client format defined by the protocol. <PROTO> MUST consist of characters selected from ABCDEFGHIJKLMNOPQRSTUVWXYZ, abcdefghijklmnopqrstuvwxyz, and 0123456789; it MUST NOT be longer than 100 characters. <serveraddress> and <clientaddress> SHOULD use solely printable ASCII characters, SHOULD have a predefined maximum length, and MUST NOT contain newline or null.

A server SHOULD accept the following options: -q to turn off all error messages; -Q to turn on error messages, negating any previous -q; -4 to immediately write the value of <PROTO>LOCAL, followed by a newline, to descriptor 4, and then close descriptor 4.

If the server receives signal SIGTERM, it MUST exit with exit code 0; this SHOULD NOT terminate any current connections. If the server cannot perform its functions, it MUST exit with a nonzero exit code. The server SHOULD interpret SIGINT the same way as SIGTERM.

Clients

If <program> is a UCSPI client, it MUST connect to a server at <address>. It SHOULD NOT connect to a server that does not support the same protocol. Upon connecting, it MUST spawn <userprogram> with the given <arg...>, descriptor 6 reading from the connection, and descriptor 7 writing to the connection. It SHOULD NOT make further connections.

<program> MUST set the following environment variables for <userprogram>: PROTO=<PROTO>, <PROTO>LOCAL=<clientaddress>, and <PROTO>REMOTE=<serveraddress>. Here <PROTO>, <clientaddress>, and <remoteaddress> are under the same rules as for servers, but note that the roles of <PROTO>LOCAL and <PROTO>REMOTE are reversed. The protocol may allow <clientaddress> and <serveraddress> to be only partially defined for a given connection, so the client's <PROTO>LOCAL may not agree with the server's <PROTO>REMOTE, and vice versa.

A client SHOULD accept the following options: -q to turn off all error messages; -Q to turn on error messages, negating any previous -q.

If the client cannot perform its functions, it MUST exit with a nonzero exit code.

Killers

If <program> is a UCSPI killer, it MUST tell a server at <address> to immediately stop accepting connections, just as if the server process had received SIGTERM. It SHOULD NOT have this effect upon a server that does not support the same protocol. The server may or may not accept the kill request.

A killer SHOULD accept the following options: -q to turn off all error messages; -Q to turn on error messages, negating any previous -q.

If the killer cannot perform its functions, it MUST exit with a nonzero exit code.

Protocols

A protocol MUST define at least the following information: Its name, <PROTO>. The format and meaning of <address>. The number of arguments <address> takes, or a notice that it does not take a constant number of arguments. Client and server address format. The maximum lengths of <serveraddress> and <clientaddress>, or a notice that they may grow to any length. The address character set, if it is not merely alphanumeric.

A protocol may define the following information, and more: Supported options. The forking behaviour of clients and servers. Any changes to the process state not required by UCSPI. The implementation of the underlying protocol, and how it interacts with other protocols. The nature of the descriptors passed to <userprogram>. The preferred names of the client, server, and killer programs. The circumstances under which a server will ignore a killer.

It is expected that all protocols will provide reliable two-way full-duplex strictly ordered not necessarily timed stream communication, though some protocols may also provide other types of communication (e.g., expedited ["out-of-band"] transmission, in which the stream is no longer ordered).

All clients and servers must obey any rules imposed by a protocol in order to support that protocol. If a program supporting a protocol cannot be unconditionally compliant with UCSPI-1991, the protocol is at most conditionally compliant.

Protocol Management

Public protocol definitions may be registered with the UCSPI Administrator. The UCSPI Administrator will normally refuse a registration request only on the grounds of a namespace problem, and in that case he will suggest an acceptable name.

All protocol names beginning with X are reserved for experimental or nonstandard use. All protocol names beginning with x are reserved for extended namespaces whose management is allocated to other authorities.

Further requests for interpretation or revision of UCSPI should be sent to the UCSPI Administrator. The UCSPI Administrator may designate one or more reference implementations of a registered protocol, and where the reference implementation establishes an interpretation not required by this document but not in conflict with this document, that interpretation should be considered binding until the UCSPI standard is revised.

The following protocol names have been registered as of 16 June 1991:

- TCP
- UNDOM

Availability

The clientserver package, including UCSPI-91, was posted to alt.sources at the end of June. You can also get it from stealth.acf.nyu.edu in pub/hier/clientserver/*. Please send any comments (or implementations on top of different protocols) to me at brnstnd@nyu.edu.

Why a C++ User Group?

Mike Banahan
Chairman European C++ User Group
Principal Consultant
The Instruction Set
London
United Kingdom

Email mikeb@inset.co.uk

Mike Banahan is Chairman of the ECUG and a Principal Consultant with The Instruction Set, the market leader in the provision of Open Systems training.

Mike is well-known for his leading role in UNIX, C, C++ and Object Oriented technology and for his popular books, 'The C Book' and 'UNIX The Book'. As Chairman of the European C++ User Group and together with his teaching and consultancy roles at The Instruction Set (for example, technical consultant to X/Open over much of the last 5 years), Mike is well equipped for his role on The Instruction Set's courseware design team, with special responsibility for the Programming Languages and Object Oriented Technologies curricula.

The first and most obvious question to ask is why bother to form a separate user group for C++ - indeed, does it need one at all? We think that it does, so I'll go on to outline the reasons below. I'd also like to explain why we felt that the needs of the C++ community don't necessarily align directly with the EurOpen constituency and hence why the group is not currently affiliated with EurOpen.

Let's get one of the most popular misconceptions out of the way to begin with. Many people who haven't used C++ in depth take the view that C++ is just a small extension to C; they assume that if you can program in C, then there are a few additional bits of syntax to learn and bingo! you are programming in C++. Undoubtedly one of the greatest marketing advantages of C++ is that it does appear to offer a simple migration path - if an organisation is already using C and has a substantial investment in C code, then the fact that C++ is almost an exact superset of ANSI C is very attractive. Choosing to move to C++ is seen as being simple: just run your existing code through the compiler, fix the very small number of places where the incompatibilities have appeared and then start introducing the C++ features where you need them. You don't have to use all of the Object-Oriented features of C++ if you don't need them, so there's a case to be made for simply using C++ as a "better" C.

Unfortunately, this simple view of things tends to come unstuck before long. The reasons are twofold: C++ is not just a small extension to C; what's more, if you use it for Object-Oriented programming, you end up writing very different programs than you would in C. I've tried several times to convince experienced C programmers that they are going to find C++ a much bigger mouthful than they thought but the result is nearly always the same - "I learned C in weeks - C++ can't even be as hard as that" - and it's usually a couple of months before they come back and say "You were right". To quote George Colouris of QMC: "C++ is a 200 pound gorilla of a language". Although the syntactic differences between C++ and C are undoubtedly limited, this belies the radical change of thinking practices that are needed to make use of it. Learning C is a relatively linear process (especially if you already know Pascal or FORTRAN); each new construct adds to what you know already without requiring much back-tracking to understand it. You pick up the syntactic elements, struggle a bit with pointers and file I/O, then off you go.

C++ differs in that each major new concept seems to go back and affect a lot of what you have already learned. The nearest analogy that I can find is a mathematical one: instead of being a linear problem (like C), C++ seems to have introduced the extra term into the equation to make it non-linear. It doesn't have smooth boundaries, taking up only a limited amount of mental space, but it has become fractal instead! The upshot is that instead of taking a few weeks to learn adequately, it is a several month long process to become thoroughly familiar with the core language features.

The second issue is the use of C++ for Object Oriented programming. For the majority of C programmers this is a very, very large paradigm shift indeed. Though you may find this hard to believe, there is plenty of both anecdotal and written evidence that it takes between a year to eighteen months to convert someone from the C/Pascal/FORTRAN styles of programming to one which genuinely exploits the power of the Object Oriented model. Once again, the usual reaction is to deny it ("I'm much cleverer than the rest of them"), then once the year is up to realise just how much work remains to be done.

Overall, combining the learning of a large new language (I've used both Ada and C++ and they feel about as big as each other), and a very different style of programming and design is not something to be undertaken lightly. Fortunately there are clear benefits very early on. Beginners discover that C++ programs can be made much more robust than C programs very quickly, with core dumps gradually becoming a fading memory, a moderate degree of code re-use is achieved quite soon, programs rapidly develop a better-designed quality and are usually more understandable. The long-term benefits of high code re-use, much higher productivity, better design and better maintainability are compelling and now demonstrated as being no pipe-dream. What happens is that people typically radically underestimate how much work they have to put in to start using C++, but just as the scale of the effort starts to become obvious, the paybacks start to become evident too.

So, it's hard to learn, there is a big paradigm shift, but it's worthwhile in the end. So what? Well, we believe that there is not only a lot to learn, but a lot to be learned too, even by those with a lot of experience already. There are substantial intellectual challenges at the design and implementation levels, together with organisational and managerial questions to be resolved. This subject goes much deeper than is implied by just another Algol-like language. As just one example, consider the requirement for persistent data, such as a database. Abstract data types and strong encapsulation are the direct antithesis of the model used by relational databases; although at present there are commercial database products which aim to provide data storage and access for C++ users, they are all relatively young. Query languages and retrieval techniques are in their infancy.

Overall, there is a whole new community of users with different problems, different ideas and different perspectives. That's why we felt that there was a real need for a user group with its own newsletter and conferences, so that this community can meet and exchange ideas.

Many EurOpen members will still be of the view that, given the tight links of C with UNIX, and the fact that C++ comes from the same stable, it would be natural for a C++ user group to operate in conjunction with EurOpen. We agree that there is some justification for this view, but we think that at the moment it would be premature to form links which were too close. At present the majority of C++ compilers in use are being used not on UNIX or similar systems, but on PCs! There is an incredibly good fit between the Object Oriented model of software design and the implementation of Window-based User Interfaces like the ones used by Apple and the Microsoft Windows product. The DOS C++ compiler vendors are not shipping hundreds of copies a month, but thousands. As a died-in-the-wool fan of UNIX for years, I have switched to using DOS as my primary development environment because of the quality of the tools that are available and their astonishingly low cost.

We felt that by starting the C++ User Group independently from EurOpen we were following the most appropriate policy. Many members of EurOpen will be uninterested in the activities of a C++ User Group and many C++ User Group members will be uninterested in the activities of EurOpen. Personally, I believe that we may well develop very close links in the future and that at some point affiliation with EurOpen is an option. For the moment though we intend to make progress on our own, but to ensure that there is a close dialogue with every interested party.

We have now scheduled our first conference, to be held in London at Imperial College over the weekend of 30 November - 1 December and warmly invite EurOpen members to attend. There will be talks on a range of technical and organisational topics together with an exhibition of products of interest to the C++ community. For further information, please contact me at the address below:

Mike Banahan, Chairman
European C++ User Group
City House
190 City Road
London EC1V 2QH
United Kingdom

Email ecug@inset.co.uk

Telephone +44 71 253 5121
Facsimile +44 71 251 2853

Are You Serious About C++?

Join the ECUG!

The European C++ User Group (ECUG) represents the interests of the professional C++ community. It provides a forum for the exchange of ideas and an opportunity to meet people with similar problems to. We haven't got all the answers! But you can share experiences with people who are asking some of the same questions as you are.

As an independent, non-profit making organisation, the ECUG offers a valuable support network.

Membership benefits include a subscription to the ECUG quarterly newsletter and reduced fees at conferences. The inaugural Conference is being held over the week-end of 30 November - 1 December 1991.

Membership Application Form

Name _____ Position _____

Company _____

Address _____

_____ Country _____

Tel _____ Fax _____ Email _____

Individual Membership £50

☐ Please invoice my company (Purchase Order No_____) ☐ Cheque enclosed

Please send to: Mike Banahan at the above address. Photocopy this form if additional copies are required.

Report from the EurOpen Conference in Tromsø

Dieter Glauss
Siemens Nixdorf
Ulm
Germany

Dieter Glauss was born 19.07.1963 in Friedrichshafen. In 1991 he finished his studies in mathematics and economics at the university of Ulm. He has been working with UNIX for 7 years. Special interests in databases, software ergonomics and software engeneering. Now a freelancer with Siemens Nixdorf in Ulm.

First there was only the wish going to the Conference in Tromsø, because the announcement sounded very interesting. I had just passed my last exams to complete my studies at the university and therefore I was a student with plenty of time. This would be the best and the last chance to visit a conference just for fun. I have heard a lot of the quality of EurOpen Conferences, but I had no exact information about the Spring Conference 1991 at all.

The main topic was about Distributed Systems with all the problems and some solutions. I think that this area of computer science will become more and more important and therefore it couldn't be wrong to learn something about it. But how should I pay the conference fee, the flight and the accommodation? Ernst Janich told me to have a close look to the announcement, especially to the section about student grants. The only thing I had to do, was to fill up the forms and send them to EurOpen. Because of my being very late, I had no hope to get a grant. The closing date for the conference in Tromsø was 1 March 1991. I sent a fax with my application on 28 February to EurOpen and two weeks later I received the positive answer.

Looking at a map of Europe you will realised that Tromsø is north of the arctic circle. From the end of May until the end of July this area presents a special spectacle; the midnight sun. Realising that this would be a opportunity to see a part of Europe that never comes back, I decided to make some holiday there. My girlfriend and I went there four days in advance. We took the plane from Frankfurt to Oslo and from Oslo to Tromsø. I was very excited because this was my first flight. It's a wonderful feeling sitting in a plane, watching a part of the earth from high above and seeing how small the world is.

My first impression of Tromsø was coldness. It was very windy and at about zero degrees. The weather was fantastic. It changed nearly every hour. We had snow, rain, storm, sun and fog sometimes all in one day. The temperature was between four degrees Celcius below zero up to eighteen degrees.

Tromsø is the largest city-municipality in Norway (2258 square kilometres) with 50,000 inhabitants. The central part of the *Gateway to the Arctic* - a name that was given because of the numerous polar expeditions started here - is on an island surrounded by snow-covered mountains and clear water. The scenery in this area is particularly delightful. We had made some long walks and were very impressed. It is indescribable. Tromsø is the world's northernmost university city. Nearly everything here is the world's northernmost: the brewery, cathedrals, bishop, planetarium and so on. Tromsø has been the meeting place for people from different cultures - that's why it is also called the Paris of the North. There are some nice cafes and pubs. But it is very expensive, a half litre beer cost 35 nok.

But let's now talk about the Conference itself. There were to main parts. First the tutorials and then the real conference. The tutorials took place in a hotel, the technical sessions in the *Kulturhuset*, where there is plenty of room for exhibitions, meetings and other events.

Monday

The tutorials started. First I had to go to the conference desk in the Kulturhuset. At the registration I got two files with the papers for the tutorials, the proceedings, some informations about Tromsø, the tickets for the social event and a coffee cup all in a blue backpack. With this backpack you were able to recognise everyone who belonged to the meeting. Carrying all this on my back I went to the Conference Hotel near the Kulturhuset.

The topic of the first tutorial was "Distributed Systems with UNIX" held by Dr. Peter Meinen. First he explained some basic architectures for system integration, load balancing and resource sharing and told us what the main design goals are: Separation of components for truly parallel program execution, isolation of single component faults and failure recovery without interrupting the complete system. Transparency for hiding the separation to the users and the application programmers. This means transparency in access, location, currency replication, failure, migration, dynamic reconfigurability and scaling. Complete transparency in practice is almost impossible to fulfill. The real world is heterogeneous and therefore you need clearly defined paradigms and standards. The chosen paradigm should be logically consistent and easy to understand. The main paradigm is the client-server paradigm with the remote-procedure-calls. He explained very clearly how RPC works. Starting with a short introduction to networks and protocols he told a lot about XDR, sockets, streams and TLI. The last part of the session was the most interesting part. It concerned methods of handling client or server crashes and authentication with public keys. The tutorial ended with some practical lessons in building distributed systems.

A very nice example for a client server architecture was our daily lunch, that was offered in the conference hotel. The buffets were the servers and we were the clients. First all clients tried to get a warm meal at the server for warm things, the queue was completely full and the queue manager told the waiting clients, that there are two other servers free. Very fine I thought so, let's go to the free servers. First I felt unhappy with my decision, because there was only cold food offered. But from a closer look I saw that there were all the things that the arctic kitchen is famous for: scrims, salm, sea wolf jam ... and seagull eggs. I tasted all the goodies and then there was no need for a warm meal. My table neighbour came from Tromsø, so I had the chance to get firsthand information about the Norwegian life. He told me, that it is not a problem to live for two months without any sunshine, because they have the polar light. I cannot believe that a light similar to the moonlight can be a substitute for the sun. But on the other hand the midnight sun compensates the darkness in winter.

Tuesday

The topic of the tutorial on Tuesday - the main reason I went to Tromsø - was *Introduction to Mach* held by Nafar Bitar, a software engeneer in the operating systems group of the Apollo Systems Division of Hewlett Packard Company. We studied many details of the microkernel. The name MACH (Multiple Asynchronously Communicating Hosts) was chosen because it sounds pretty fast. First there were some other names discussed for example SOS (Supercomputing Operation System) but this name is used. Another suggestion - MOOSE (Multiprocessor-Oriented OS and Environment) - first was favourite but a professor meant that it will sound very silly talking in computer science about moose. Why is there a need to develop a new operating system like MACH? UNIX was good because of the simple interface, concept, portability of a large class of uniprocessors, the extensive utility program library and the easy way to combine utilities. But the actual kernel degenerated to a kitchen sink and it is very difficult to integrate primitives for multiprocessors, distributed systems and shared program library.

The main goals in the MACH design are exploit parallelism in the OS and in applications, the support of large, sparse address space with flexible sharing, transparent network resource access, compatibility with existing environments and a portable implementation.

Bitar had shown us all the basic functions of the microkernel. We learned how the virtual memory management, the task/thread management external memory management and the interprocess communication work. A basic idea of the implementation of the kernel is the strict separation between machine-independent and machine-dependent components. I think that the study of the microkernel will be very useful to understand how operating systems works.

At lunch most of the clients learned their lesson from yesterday - the cold buffet was full, therefore I had a warm meal but it wasn't as good as the things I had the day before and so I also went to the server with the cold food.

Wednesday

The technical session took place in the Kulturhuset. 335 delegates have found their way to Tromsø. The first lecturer was Michael D. Schroeder. He delivered a very good overview what distributed systems are and what they have to be. First he took a look at the good old days of timesharing with no naming problem (uniform names for users, files, functions etc.), only one management domain, coherent filesharing, simple, reliable mail systems and uniform authentication. But the new distributed systems also have some good points. Information and resource sharing in wide geographical and organisational spread, use of cost-effective work stations, growth in small increments and large range of size and autonomy from decisions, vendors, software versions, and crashs (not all crashs at one time). The ideal distributed system may be a system which meets the requirements from both centralised and decentralised systems. Schroeder argued that interfaces are the key to open systems. Each service must be defined by an interface (meaning by a contract). In the second part of his talk he focused on naming, security and availability issues. We need systems be able for fault recovering in less than a second. This lecture was a very good introduction into the field of distributed systems.

With Shape Muellender's look to the good and bad things about Amoebea the presentation of the main streams in operating systems started. He listed also the things he would have done differently, if he had to do it all over again. Some aspects of Chrous, distributed an UNIX compatible, presented by Michel Gien followed afterwards. After lunch Simon Patience gave a even more commercial lecture about OSF/1. Dave Presotto presented Plan 9 with a two hour talk in 40 minutes. Though he did it in a very fast manner, people were quite happy with the presentation. The basic idea is to have only files to manage.

The ways these new systems work are very interesting. My favourite is Plan 9, because it based on such a simple idea and it looks very easy to work with Plan 9. But also OSF/1 with its MACH Kernel is a fine thing and as a look at the exhibition showed this seems to be the way industry is moving.

This noon the lunch client-server architecture had changed. We now had to be the servers. Our task was only to eat all things brought from the Clients They asked with very good meal for our mouth and stomach work.

After the coffee break, Andrew Schuelke tried to explain some efforts of UNIX international, but I didn't see the goal.

A forum closed technical session at the first day. All speakers of the day chaired by Micheal D. Schroeder first talked about some questions about security and real time support. Then the question concerned standards. A debate between OSF and UNIX International was attempted, but these two major organisations surprisingly agreed with each other. But the audience was very enthusiastic. I think a typical discussion with the result:

We need standards but we don't want accept a standard defined by anyone else.

Sape Muellender suggested that if using a standard to use only one standard. Dave Presotto saw the main problem of standards in a loss of performance because of a lot of layers.

This evening the first social event was scheduled. A trip on a ferry trough the fjord with free shrimps and beer. Although the weather

wasn't good, we had rain and snow, we were in a tremendous mood.

The competition from HP with a calculator as a first prize was one challenge on the ferry. The question was:

How many MIPS has the University of Tromsø bought from Hewlett Packard?

The suggestions ranged from *not enough* to *1 and 27,000 MIPS*. Some of these very good answers was rewarded with some wine, the first prize went to Mary Bitar of Boston.

My main challenge was to get the shrimps out of their skin and to discuss the last days and experience with other delegates. I think when we came back to the harbour most people were full of new impressions of the fantastic scenery and arctic beer. It was wonderful.

Thursday

The technical session started with the keynote speaker Bruce D. Shriver, who had changed the topic of his talk the night before, animated by the first day lectures. He talked about microprocessor systems. We must distinguish between distributed systems, for them the applications may not be rewritten and parallel systems which need new applications because automatic parallelism isn't possible.

To give a characterisation of ideal microprocessor systems, Shriver cited Schroeder that such systems should be coherent, manageable and secure - a blend of the best from centralised and decentralised systems. Such a design cannot only be realised by chip-developer. They must get help from compiler-builders, because this is the only way to get processor systems with a developing environment that supports the system.

Then Shriver gave a short overview of 1995's shared memory MP characteristic's:

- 8 - 16 processors
- 50 - 250 MIPS per processor
- 75 - 300 peak MFLOPs per processor
- 100 MHz (CISC) or 100 - 200 MHz (RISC)
- 256 - 512 Mbytes storage per processor
- 256 KB - 2 MB cache
- 64 bit address, and
- 1 - 5 Gbytes/sec I/O transfer rate

Another very important remark was on standardisation - he hoped that the work that is done within standardisation wouldn't reduce the work on distributed operating systems as Amoeba, Chorus, Plan 9, MACH etc.

Then after a very important coffee break - I had no breakfast, but that wasn't the problem, I also hadn't any coffee until this break and that was the point. A coffeeholic can get really weak without his brown warm drug - the conference went on with Simon Patience from OSF instead of Brad Johnson - he had to work - about DCE, nothing new for me because of the Monday tutorial and then Dario Avallone presented AxIS: a configurable software engeneering environment. There are some nice ideas how to integrate user-defined production tools.

Gordon Blair from the Lancaster University showed how to use distributed systems for a multimedia environment.

Lunch was the same procedure as the previous day and was also very good.

After lunch the conference went on about scheduling. The next two lectures were concerned with load balancing strategies. First Wouter Joosen described the design and implementation of a load balancing environment. They have implemented a testbed to evaluate several load balancing strategies.

The decentralised algorithm for load sharing in networks presented from Guy Bernard in the next talk seemed to be a nice candidate for the testbed. The advantage of this algorithm is that it is decentralised but needs expensive broadcast messages only occasionally.

After the coffee break Rob Pike started his show: He didn't talk about *Process Sleep and Wakeup on a Shared-memory Multiprocessor* because these could be read about in the proceedings. He wanted to talk about his window system 8.5 under Plan 9. First he gave a short impression of his main goals he wanted to reach. 8.5 had to be and is very small. Source is 1/20 of X and documentation only a 1/100. He said it isn't true that window systems must be big and complicated. The existing systems are only so big because their designers are big and complicated. His system based on the simple idea of plan 9 -- all you have to manage are files. There are the three basic files /dev/cons (Console) /dev/mouse and /dev/bitblt. The system reads from /dev/cons to get character input, from /dev/mouse to get information about mouse-position and status and write to /dev/bitblt for graphics. These devices a process or process group could bind and use to their dialog. Each client using the window works in the same namespace. The short introduction Rob Pike gave was very impressive and I think it would be very interesting to work with the system. The main question is how to get it.

The Conference Dinner this evening was a great success. We had some nice talks at a really international table - American, English, Dutch, French and German people. This evening we did not only talk about computers but also about politics and what life is like in the different countries.

For diner we had yellow caviar to start then reindeer with potatoes and beans and at last a strawberry ice cake. It was delicious.

Friday

The last day of the conference was very hard for me, because my head was overfilled with the information I had heard the last four days and it was no place for any new talks in my brain. Nevertheless I walked to the Kulturhuset absolutely determinated to pick up as much as possible.

The keynote speaker A.J Herbert talked about distributed objects. The problem in client server architecture is that the access is procedural, therefore the procedures must have remote data. He supposed objects as a solution for the main problems in distributed computing. To manage a separated heterogeneous system the need of a new engeneering model is essential. On base of abstract data types it is possible to define distributed objects.

The next issues about languages were first presented by Henri E. Bal from the Vrije University of Netherlands in a comparative study of five parallel programming languages. There are far more than 100 parallel programming languages. To compare these languages he took for every important paradigm one language:

- *SR* (message passing)
- *Emerald* (concurrent object based)
- *Parlog* (concurrent logic language)
- *Linda* (tuple space) and his own language
- *Orca* (distributed shared memory)

He compared these languages by programming with them. He implemented in every language some typical problems, like matrix multiplication, the All Pair Shortest Path problem, the Travelling Salesman Problem, alpha-beta search, and successive over relaxation.

For the next talks my concentration was so bad, that I could not follow them. I used this time to have a closer look at the exhibition. In the forum of the Kulturhuset HP and IBM showed their latest work stations. Addision Wesley and Prentice Hall presented a lot of books about Unix, C etc. This was a very good opportunity to get a little impression of some books that I had heard about. Digital presented their new OSF/1 system and so I could see how OSF/1 based on an MACH Kernel works.

After lunch I was fit again and so I followed the lecture from V. Tschammer who described a domain-based support of service administration and server selection. The last talk of the conference was held by John T. Kohl. He spoke about the experience they had with the authentication service of Kerberos and explained how it works. This was very hard stuff at the end of a long week and I decided to read this later in the proceedings.

A forum discussion chair by Strack-Zimmermann concerned the question whether we need distributed systems and how we can trust them finished the Conference. Some people criticised that no distributed applications were presented and reported from their own experience with distributed applications; the problems and also advantages they had.

Summary

The EurOpen Conference had fulfilled what I had expected; Many new ideas, interesting people and a fantastic landscape. The organisation of the Conference was perfect.

The proceedings and papers of the tutorials are absolutely necessary because to hear about so many new ideas in such a short time and compressed form has the effect that only a little part can be kept in mind. After the conference I had read a lot in the papers to freshen up the talks and to learn more about the things I have heard.

The social events were really great. It is very important to have one evening (or two) where all delegates come together.

If time allows I will visit the Autumn Conference in Budapest.

Call for Tutorial Proposals

The tutorial programme is one of the major elements of a EurOpen conference. In recent years classes on a wide range of subjects have been taught by recognised authorities in their field. Past classes have included System V and Berkeley Kernel internals, System Administration, UNIX networking, X Windows Concepts, UNIX standards to name just a few.

It has been policy to extend the range and variety of tutorials on offer. In order to do this new tutorials are constantly required.

EurOpen is now seeking proposals for classes for our 1992 conferences to be held in Jersey (April) and Utrecht (November). Proposals from tutors are invited on any topic, although past experience has shown that advanced tutorials are better attended by delegates than introductory ones.

All EurOpen tutorials are taught in English and usually last a full day, although half day classes are accepted from time to time. A full day class will normally be supported by between 80-200 slides, together with up to 200 pages of additional material. Previous experience has shown that classes can be successfully taught by either one or two tutors.

As well as seeking new tutors I am also seeking new tutorial subjects, so would like to see a class on a particular topic at a future conference please let me know, but remember the subject must not be so specialised that only a very small number of experts would have an interest in it.

Please send your proposals and suggestions to me at the address below. I will also be pleased to discuss possible proposals with you.

Neil Todd
EurOpen Tutorials Officer
Owles Hall
Buntingford
Herts SG9 9PL
UK

Email neil@pio.gid.co.uk

Facsimile +44 763 73255

EurOpen Autumn Conference 1991

16 – 20 September 1991

EurOpen (formerly EUUG) — The European Forum for Open Systems is a non-profit making organisation. It is an international federation of UNIX and Open Systems User Groups.

The EurOpen membership is comprised primarily of institutional and individual members of affiliated National Groups. There are at present 22 National Groups with a total membership of approximately 6,000 coming from such diverse backgrounds as software engineering, information technology, computer manufacturers, end-users, software houses, universities and research centres. EurOpen runs two major conferences each year on a European level.

The Autumn '91 Conference and Tutorials will be held in Budapest, Hungary in conjunction with EurOpen Hungary at the Conference and Education Centre, Budapest on 16th-20th September. The event will be accompanied by an exhibition, 'UNIX goes East'.

Hungary is the first milestone for UNIX and EurOpen to go East. Hungary made the first pioneering move to join EUUG (now EurOpen) in 1988, and was accepted as a member at the Lisbon meeting of the EUUG Governing Board. At that time no one imagined the dramatic changes which would take place in the whole Eastern block, changing its entire shape during the past year and a half.

UNIX has long traditions in Hungary. Although UNIX and all types of workstations were kept on the COCOM list until recently many Hungarians worked in UNIX environments in Western countries, and became enthusiastic about the ideas UNIX and open systems represent. Hungary had made several unsuccessful attempts to buy UNIX licences during the early eighties. After the failures they developed a Version-7 compatible native version from scratch called HUNIX. Nowadays the growing UNIX community in Hungary has access to all UNIX licences, and the whole philosophy of open systems and EurOpen has received broad acceptance in Hungary.

The capital of Hungary is one of the most scenic capitals in Europe lying on both sides of the river Danube. The population of Budapest is 2 million (one fifth of the whole country). It is the main industrial and cultural centre of the country having seven universities and many research and development institutions in information technology.

UNIX Goes East Open Systems Show

16 – 20 September 1991

Early Autumn is the best and most beautiful season in Hungary (average temperature is about 18 Celsius). EurOpen Hungary, the hosting National Group of the European Unix Community, will do its best not only to ensure a smooth running Conference but will also organise technical visits to Universities and Hungarian Companies by request and offers opportunities for participants to get acquainted with one of the most attractive European Capitals, Budapest, and with other beautiful spots in Hungary.

We encourage you to come and spend in Budapest not only the Conference days but also the weekend before or after it.

Conference and Tutorial Enquiries

Mrs Helen Gibbons
EurOpen
Owles Hall
Buntingford
Herts. SG9 9PL
United Kingdom

Telephone +44 763 73039
Facsimile +44 763 73255

Email europen @EU.net

Exhibition Enquiries

Mrs Maria Toth
John von Neumann
Society for Computing Sciences
Budapest
Báthori utca 16
H-1054
Hungary

Telephone +36 11 329 349
Facsimile +36 11 318 140

Programme of Events
Tutorials

Monday
16 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial M1 – Posix for Programmers

Tutor Donald Lewine

POSIX is an international standard that defines the interface between a portable application and the operating system. POSIX is based on historical UNIX system models. This tutorial is aimed at experienced programmers, pointing out what the potential pitfalls of converting to POSIX are and how they can be avoided.

The tutorial does not cover ANSI C. The student is assumed to know about portable C software. Instead the tutorial spends the time covering interfaces new to POSIX:

- The POSIX file system.
- The POSIX process model.
- POSIX signals.
- POSIX terminal I/O (What, no ioctl()!).
- Obtaining information at run-time.
- Internationalization and national profiles.

Each topic is illustrated with detailed examples. There are also detailed examples of converting BSD and System V applications to be POSIX conforming. The tutorial covers how to get around things that are in BSD and/or System V but missing from POSIX.

Biography

Donald Lewine has been writing computer programs for fun and profit since 1960. He has been teaching Computer Science in the State-of-the-Art (evening) program at Northeastern University for the past nine years, including courses on Assembler, VAX/VMS, PASCAL, C and UNIX.

Mr. Lewine spent 13 years with the Digital Equipment Corporation developing operating systems and central processing units. He was technical director for the MicroVAX program when he left.

For the past seven years, Mr. Lewine has been with Data General Corporation, and is currently Director of Engineering. In this role he has been developing the AViiON family of open systems. He is a founder and member of the board of 88open, a member of the Board of Directors of Unix International, and Data General's representative to the Open Software Foundation.

Tutorial M2 – System V Release 4 internals

Tutor Bill Rieken

Objectives

Upon successful completion of this tutorial students will be able to describe the major components of the SVR4 kernel and the new features that have been added to the SVR3 release. This is a brief overview course covering concepts and facilities. It is not an intensive hands-one course with access to source code.

Audience

This one-day tutorial is intended for those who want to know the basic structure and flow inside the SVR4 kernel, and what makes it different from previous releases. Attendees must have an understanding of operating system internals in general, and some knowledge of an other UNIX system kernel would be very helpful. C programming experience is not required, as only a few, very short code segments are presented and a few C data structures are shown to illustrate some key kernel control blocks.

Outline

Topics that will be covered include the virtual file system (VFS) architecture, process scheduling and real time enhancements, new exec switch to load a variety of object module formats, new virtual memory architecture with a variety of "object managers" (vm segment drivers), POSIX/BSD job control and signals, kernel memory allocator, STREAMS enhancements, fifo/pipes and mounted STREAMS, STREAMS tty subsystem with line discipline/X.25 pad modules, secure Service Access Facility and ttymon replacement for getty, transport-independent network selection and name to address mapping service, and, if time permits, Remote File Sharing (RFS) and Network File System (NFS).

Programme of Events

Tutorials

Monday
16 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial M2 (Continued)

Instructor

Bill Rieken has over twenty years experience as a computer programmer and educator. He taught at the University of Wisconsin and Southern Illinois University, where he began using UNIX version 6 in 1976. Bill provided UNIX System Support and Training at AT&T Bell Laboratories in New Jersey for five years from 1983 through 1988. He provided AIX System Support at the IBM Palo Alto Scientific Center for two years from 1988 through 1990.

He currently teaches TCP/IP/STREAMS networking, X-Windows/Motif programming, and OSF/MACH/UNIX system administration and internals courses. His undergraduate work in Mathematics and Computer Science was done at U.C. Berkeley, and his graduate work in Computer Science was done at Stanford University. He also has an M.S. in Mathematics from the University of Nevada in Reno. Bill is an International ACM Lecturer and a principal founder of ***.sh consulting*** based in Santa Clara, California.

He has written fourteen UNIX training books for the courses he teaches, and co-authored a new book on **"Adventures in UNIX Network Applications Programming"** published by John Wiley.

Tutorial M3 – Decomposition/Generalisation Methodology for C++ programs

Tutor Vaclav Rajlich

Instructor

Vaclav Rajlich, Professor, Department of Computer Science, Wayne State University, Detroit, MI 48202.

Intended audience

Programmers, system analysts, and managers with a reading knowledge of C. Knowledge of C++ and object oriented programming is not required.

Course description

The tutorial presents a methodology for top-down design of object oriented systems in C++. The methodology is a variant of stepwise refinement, extended to the object oriented systems. It alleviates the difficulty of finding the right objects. Steps of decomposition and generalization are explored in detail. Decomposition/Generalization methodology is compared to bottom-up program development.

Object oriented constructs of C++ are explained in the tutorial.

Biography of the author

Vaclav Rajlich published numerous papers on software methodologies and tools. He has been particularly interested in methodologies for object oriented languages (C++) and object based languages (Ada).

He is a professor and former chair of the Department of Computer Science at Wayne State University. Before, he was with the University of Michigan in Ann Arbor. He was born and raised in Czechoslovakia.

Programme of Events

Tutorials

Monday
16 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial M4 – System Administration of a Local Area Network

Tutor Evi Nemeth and Bob Coggeshall

This tutorial covers solutions to problems most often encountered by administrators of a small, rapidly growing facility. Emphasis will be on networking and resource management issues.

We assume a thorough knowledge of Unix from a user's point of view and basic familiarity with the problems confronted by a system administrator.

Topics covered include

Management Issues: spanning administrative domains, distributing information, documentation, problems of scale.

Networking Software:
Routing, static vs. dynamic routing, using gated.
NFS/YP, the Network File System and NIS (Yellow Pages) system for distributing files on the network.
SLIP, serial line IP for network connectivity using high speed modems over telephone lines.
BIND, the Berkeley Internet Name Domain system for dynamic lookup of hostnames and network addresses.
uucp, the unix to unix copy system for store and forward networking over telephone lines.

Communications:
Mail, the sendmail program and its configuration file, with a focus on the things that make sendmail.cf confusing.

Appendices: extensive appendices are included:
References
Vendor list
Public domain software
LAN articles
COPS articles
Internet resources and tools
Sendmail configuration checker
Forms for obtaining IP address and domain name

Software:
Some of the tools we describe will be available on the EurOpen Conference tape.

Evi Nemeth, biographical information

Evi Nemeth, a faculty member in Computer Science at the University of Colorado, has managed Unix systems for the past 15 years. She is co-author of the best-selling "Unix System Administration Handbook" (Prentice-Hall, 1989) and has taught numerous system administration tutorials. She is a member of the Board of Directors of USENIX.

Bob Coggeshall, biographical information

Bob Coggeshall has been administering Unix systems for over 10 years; most recently he managed the Unix system administration group for the University of Colorado from 1985-90. He has taught tutorials on system administration and network design and management throughout the US. Bob is now president of Cogwheel, Inc., a small company in Boulder, Colorado doing network hardware/software design and consulting.

Programme of Events

Tutorials

Tuesday
17 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial T5 – New Kernel facilities in 4.3 BSD-Reno

Tutors Mike Karels and Kirk McKusick

This tutorial is directed to systems programmers that have taken a course on 4.3BSD internals, or who have had at least a year of experience working on the 4.3BSD kernel, or who have worked on 4.3BSD derived kernels such as SunOS, System V Release 4, or Ultrix.

No licence is required for this tutorial.

Description

This tutorial presents a detailed discussion of the new facilities that have been released in the 4.3BSD-Reno distribution. The tutorial will cover the four main areas outlined below. Presentations will emphasize system organization, data structure navigation, and algorithms. This tutorial retains about 50% of the material from our previous tutorial, "Beyond 4.3BSD: Advanced Kernel Topics".

1 Networking
- OSI protocols
- routing and routing sockets
- new networking buffering scheme
- TCP improvements (header prediction, compressed SLIP)

2 Filesystems
- virtual filesystem interface
- changes to the local filesystem
- new quota implementation
- NFS implementation
- memory-based filesystem
- dead filesystem

3 POSIX
- termios
- new process structure
- process groups and controlling terminal
- new system call interface

4 System Interface Changes
- virtual memory
- ktrace facility
- get kernel info
- changes in system call interfaces
- authentication
- bstreams

About the speakers

Dr. McKusick got his undergraduate degree in Electrical Engineering from Cornell University. His graduate work was done at the University of California, where he received Masters degrees in Computer Science and Business Administration, and a Ph.D. in the area of programming languages. While at Berkeley he implemented the 4.2BSD fast file system and was involved in implementing the Berkeley Pascal system. He is currently the Research Computer Scientist at the Berkeley Computer Systems Research Group, continuing the development of future versions of Berkeley UNIX. He is president of the Usenix Association and a member of ACM and IEEE.

Michael J. Karels is the Principal Programmer of the Computer Systems Research Group at the University of California, Berkeley. Since the release of 4.2BSD, he has been the system architect for Berkeley UNIX, continuing the development of new versions of BSD. Michael Karels received his B.S. in Microbiology at the University of Notre Dame. While a graduate student in Molecular Biology at the University of California, he was the principal developer of teh 2.9BSD UNIX release of the Berkeley Software Distribution for the PDP-11. He is a member of the ACM, the IEEE, teh Internet Engineering Task Force and several POSIX working groups.

Both speakers are co-authors of the book "The Design and Implementation of the 4.3BSD UNIX Operating System".

Programme of Events

Tutorials

Tuesday
17 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial T6 – Introduction to X-Windows User Interface (½ day tutorial)

Tutor Bill Rieken

Objectives

Upon successful completion of this tutorial students will be able to use and customize their X-Windows environment using **X commands** and **X resource** (default) **files.**

Audience

This half-day tutorial is intended for experienced computer end-users (with little or no programming experience) who want to learn **the "point and click" X-Windows user interface.** Experience with UNIX commands and shell programming is very helpful, although not mandatory prerequisites. Examples use OSF/**Motif**, AT&T/Sun **Open Look** and **SCO Open Desktop.**

Outline

Attendees are introduced to X-Windows, a portable, non-proprietary, graphic windowing system for distributed applications executing over a network of computers. Topics include how to start an X Server and a Window Manager, opening a window, moving and resizing a window, hiding a window in an icon, "cut and paste" from one window to another, and use of "pop-up" menus, as well as remote execution and how to change window managers, fonts and/or colours. X commands such as **xterm, uwm, mwm, olwm, xrdb, xhost, xclock, xwd, xinit, xl lstart and startx** are covered, as well as customization keywords used in **.Xdefaults, .Xresources, .xl lstart, startxrc, .olinitre, .mwmre and .uwmre** files.

Tutorial T7 – Unix Programming Environments (½ day tutorial)

Tutor Bill Rieken

Objectives

Upon successful completion of this tutorial students will be able to use many of the UNIX programming environment commands to expedite program development and facilitate system builds and distribution.

Audience

This half-day tutorial is intended for experienced **UNIX programmers** who want to learn the **"UNIX Programming Environment".** Experience with UNIX commands, shell programming, and one or more programming languages is required to benefit from this session.

Outline

Attendees are introduced to the **sdb** symbolic debugger, the **SCCS** source code control system, and the **make** system build utility. Other programming utilities such as **lint** "strong typing" checker, **ar** program library archiver, **time** and **prof** performance measuring tools, **nm** symbol table lister, **cc, ld,** and **as** compiler and linker options, and **diff, cmp,** and **dircmp** file comparators are also described.

Programme of Events

Tutorials

Tuesday
17 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Tutorials start at 09.30

Tutorial T8 – An Introduction to UNIX System Security
Tutor Matt Bishop

Intended Audience

This tutorial is intended for UNIX users and system administrators who are concerned about system security; no prior knowledge of UNIX security features is presumed.

What This Tutorial Covers

From a discussion of security policies, requirements, and general mechanisms, we proceed to the basic UNIX security mechanisms including file permissions, the setuid and segid bits, passwords, and other tools. We next examine the password mechanism in considerable depth, ranging from generic "how to pick a good password" to details of how the password hashing scheme works and how passwords are stored. Then comes file system auditing and precautions one can take when running as the superuser, or as an ordinary user. We discuss network security for the internet and uucp-based networks. We then consider the art of writing programs for UNIX systems; which library calls to avoid, dangerous programming practises, safe and unsafe environments, and so on. Finally, we consider several miscellaneous topics such as restricting user environments and user and staff awareness.

About the Instructor

Matt Bishop has been looking at UNIX security since 1980. After graduating from Purdue University, he worked at the Research Institute for Advanced Computer Science at the NASA Ames Research Center, where he specialized in computer security. Currently he is an assistant professor at Dartmouth College, where he teaches operating systems, computer security and cryptography, and software engineering. He is a member of the Privacy and Security Research Group, which advises the Internet Activities Board on issues relating to security on the internet, and has chaired the last two UNIX Security Workshops.

Only EurOpen National Group or Direct Members are permitted to attend Tutorials. Except for T6 and T7, each Tutorial lasts for one whole day and will start at 09.30 and finish approximately at 17.30

Programme of Events

Provisional Technical Programme

Wednesday
18 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Technical Sessions start at 09.30

Keynote: The State of UNIX: Dennis Ritchie (AT&T Bell Laboratories) (USA)

Coffee

UNIX and Virtual Reality: Mike Griffin & Richard Mitchell (Reading University) (UK)

UNIX learn and intelligent lesson selection: Peter Brusilovsky (Int. Centre for Scientific and Technical Information) (SU)

The QEF/QEI Model for Software Component Consistency Dependency Determination, and Construction Recipe Ordering: David Tilbrook (Sietec Open Systems Division) (CA)

Lunch

Distributed Systems 1
Give a process to your drivers!: François Armand (Chorus systèmes) (FR)

Multimedia Synchronisation and UNIX: Dick Bulterman, Guido van Rossum, and Dik Winter (CWI)(NL)

Using a Wafer-Scale Component to Create an Efficient Distributed Shared Memory: Aarron Gull (City University) (UK)

Tea

Performance Measurement

Benchmaking at AFUU, the SSBAs (Christophe Binot and Nhuan Doduc) (FR)

Near real-time measures of Unix-like operating systems: Mario Cambiaso, Susanna Delfino and Giancarlo Succi (Universita di Genova) (IT)

Steppingstones – some remarks on measuring X11 performance: W Kriechbaum & F Pestenhofer (IBM AIX-FSC) (DE)

BOFS

Performance measurement

Thursday
19 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Technical sessions start at 09.30

Security Workshop

Keynote: Security and Open Working in the Networked Academic Community: Denis Russell (University of Newcastle upon Tyne) (UK)

PhLOGIN: Alain Williams (Parliament Hill Computers) (UK)

Coffee

Distributed Systems 2

MANIFOLD A Language for Specification of Inter-Process Communication: Farhad Arbab & Ivan Herman (CWI) (NL)

A Distributed Concurrent Implementation of Standard ML: David Matthews (University of Edinburgh) (UK)

Load Balancing Survey: Dejan Milojicic, Milan Pjevac & Dusan Velasevic (Institute "Mihajlo Pupin") (YU)

Lunch

Networking

Unix in a Novell Environment: Gabriella Ivanka & Gyorgy Leporisz (HUUG) (HU)

Managing the International X.500 Director Pilot: Colin Robbins (X-Tel Services Ltd) (UK)

A design overview of XLookup: D S Mahl (Brunel University) (UK)

Tea

Parallel sessions: (a) Industry Directions

OSF / UI / ACE

Parallel sessions: (b) Work in Progress:

Short presentations. Titles and speakers to be announced at the Conference.

BOFS

EUnet

General Information

Conference & Exhibition Venue

The Tutorials, Technical Sessions and the Exhibition will be held at:

The Conference & Education Centre
Villányi út 11-13
Budapest
Hungary

Telephone + 36 11 666 399
Facsimile + 36 11 667 410

The Conference Centre is situated in the 11th district of Budapest, on the Buda side, not far from the Danube, between the beautiful Gellért Hill and the romantic "Bottomless Lake", but at the same time has very good traffic connections with all parts of the city.

From all the hotels reserved for delegates, the Conference Centre can easily be reached by one tram or bus line.

See map on page 12.

Conference Hotel

Hotel Korona
Kecskeméti u.14
Budapest 1053
Hungary

Telephone +36 1 1174 111
Facsimile +36 1 1183 867

This new four star hotel is situated in the very centre of the Pest side having excellent public transport connections with all sites associated with the EurOpen Autumn '91 event. The Hotel has its own guarded car park (for an extra charge). The Conference & Exhibition Centre can be reached from the Hotel Korona by tram 47 or 49 and Bus 1. Journey time is approx. 10 minutes.

If you wish to book your accommodation in the Korona Hotel please indicate this on the booking form on Page 17.

Flights and Transport

Budapest is situated in the heart of Europe and can easily be reached from all major European cities. Budapest airport has direct flight connections with 42 European cities. MALEV, Hungarian Airlines – the official carrier of the Conference will be glad to make your flight reservation. Delegates using MALEV flights will receive special tickets to a raffle – first prize is a full refund of the Conference fee.

Budapest Airport terminals are located 15 and 20 km from the city. There is a regular bus service every 30 minutes from the terminals to the city terminal. The bus ride costs approx. 60p and takes approx. 30 minutes. A taxi will cost approx. £6.00 from the airport to any of the city hotels.

The public transport in Budapest is very cheap compared to other western cities.

Visa's

Some delegates **may** have to apply for a VISA in order to travel to Hungary.

The citizens of the following countries, listed below, do **not** require a visa in order to visit Hungary:

Argentina, Austria, Belgium, Bulgaria, Canada, People's Republic of China, Cuba, Cyprus, Czechoslovakia, Denmark, Finland, France, Germany, Great Britain, Greece, Iceland, Ireland, Italy, Republic of Korea (South-Korea), Liechtenstein, Luxembourg, Malta, Monaco, Mongolia, Netherlands, Nicaragua, Norway, Poland, Romania, San Marino, Soviet Union, Spain, Sweden, Switzerland, USA and Yugoslavia.

If your country is not listed above please contact your local travel agent or the Hungarian Embassy in your own country.

The Exhibition

The Exhibition "UNIX Goes East" will be held at the Conference & Education Centre from 16th-20th September and will be open to delegates throughout the Conference.

"UNIX Goes East" will attract companies and specialists from Middle and Eastern Europe. Outstanding companies including DEC, IBM, UNISYS have already indicated their intention to exhibit.

All enquiries to:

Ms. Gabriella Aranyos / Mr. Sandor Bartok / Mrs. Maria Toth
John von Neumann Society for Computing Sciences
Budapest
Bàthori utca 16
H-1054
Hungary

Telephone +36 11 329 349
Facsimile +36 11 318 140

Programme of Events

Provisional Technical Programme

Friday
20 September 1991, The Conference & Education Centre

Registration from 08.00 in the main Entrance Hall
Technical sessions start at 09.30

Systems

Keynote: The 4.4BSD Architecture: Mike Karels (University of California) (USA)

Virtual Swap Space in SunOs.: Peter Snyder, Howard Chartock (Sun Microsystems) (USA)

Coffee

Systems Administration

File System Organization: The Art of Automounting: Martien van Steenbergen (Sun Microsystems) (NL)

Effectively Monitoring the Performance of a Heterogeneous Departmental Network: Martin Beer & Shaun Hovers (University of Liverpool) (UK)

Lunch

Distributed Systems 3

StormCast – A Distributed Application: Dag Johansen, Gunnar Hartvigsen (University of Tromsø) (NO)

Location-Independent Object Invocation in Open Distributed Systems: Herman Moons (Leuven University) (BE)

Communicating Database Objects: Agnes Hernadi, Ferenc Jamrik, Gabor Janek, Elod Knuth
(Hungarian Academy of Sciences) (HU)

Tea

Interactive User Interface Design – the Teleuse Approach: Achim Brede (Bredex GmbH) (DE)

An International Hotel Reservations System Using Loosely Coupled UNIX Systems: Gary Bilkus (BLIX Ltd) (UK)

Close

Additional Papers

An Implementation of a Process Migration Mechanism using Minix: Sylvain Louboutin (University College Dublin) (IE)

A Public Access Interface to the OSI Directory: Paul Barker (University College London) (UK)

HAWKS – A Toolkit for Interpreted Telematic Applications: Carl Verhoest (Télésystèmes Innovation) (FR)

How to Pay & Book

Electronic Mail at the Conference

Delegates can be reached during the Conference by EUnet mail. Messages will be printed, sealed and posted on the message board. To reach people at the Conference use the following address:

firstname_lastname_organisation@europen_conf.EU.net
or
firstname-lastname-organisation@europen-conf.uucp

Delegates should limit the amount of messages forwarded to this address.
Distribution lists should **not** be forwarded.
Delegates will also be able to send mail.
Worldwide TELNET access via InterEUnet is also planned.

Language

The official language of the Conference will be English. No translation will be provided.

Liability

EurOpen will not accept any responsibility for damage to property or injury to persons during the entire event. Participants are recommended to arrange for their own personal travel and health insurances.

Payments

EurOpen is a European Federation of National Groups and as such prefers to use the European Currency – ECU's – for payments. To pay in ECU's please note the two methods of payment.

1 By Direct Payment to EurOpen's Bank Account

The Bank of Scotland — Account No. 41791 ECU 01
International Division — Bank Sorting Code: 80-20-13
Operations Department
PO Box 86, 120 St. Vincent Street
Glasgow G2 5DZ, Scotland

Please tell your bank that you will pay all charges so that EurOpen will receive the full amount.

2 By UK Cheque or Bankers' Draft, made payable to EurOpen and drawn on a UK bank. Eurocheques are acceptable, but each cheque must be 170 ECU's or less.

EurOpen is setting up facilities in Europe for delegates paying by Credit Card in ECU's, but unfortunately this is a very new facility and not fully functional at the time of printing this booklet.

If you have real difficulties in paying in ECU's you may pay in £ Sterling using the following prices **only.**

N.B. Please note that these are fixed prices and may not be converted in accordance with any exchange rate at any time. If you choose the £ option it must be at the prices stated.

Costs

Tutorials		ECU
Tutorial per person before 1st August	Members only	420
Tutorial per person after 31st July	Members only	480
Student Price – 50% reduction*		
Tutorial and Payment per person on the door*	Members only	580
½ day Tutorial before 1st August	Members only	210
½ day Tutorial after 31st July	Members only	240
½ day Tutorial and Payment per person on the door*	Members only	290
Conference		
3 Day Conference before 1st August	Members	400
	Non-Members	500
3 Day Conference after 31st July	Members	460
	Non-Members	560
Student price – 50% reduction*		
3 Day Conference Payment on the door*	Members	560
	Non-Members	660
**These can only be accepted if space allows.*		
Conference Dinner – extra ticket		50
Conference Tape		100

Note: Closing date for all bookings is 11th September, 1991.

Costs

Tutorials		£
Tutorial per person before 1st August	Members only	300
Tutorial per person after 31st July	Members only	345
Student Price – 50% reduction*		
Tutorial and Payment per person on the door*	Members only	410
½ day Tutorial before 1st August	Members only	150
½ day Tutorial after 31st July	Members only	173
½ day Tutorial and Payment per person on the door*	Members only	205
Conference		
3 Day Conference before 1st August	Members	285
	Non-Members	360
3 Day Conference after 31st July	Members	330
	Non-Members	400
Student price – 50% reduction*		
3 Day Conference Payment on the door*	Members	400
	Non-Members	465
**These can only be accepted if space allows.*		
Conference Dinner – extra ticket		35
Conference Tape		70

Note: Closing date for all bookings is 11th September, 1991.

How to Pay & Book

To pay in £ Sterling please note the two methods of payment:

1 By direct Payment to EurOpen's bank, which is:

The Bank of Scotland
61 Grassmarket
Edinburgh EH1 2JF
Scotland

Account No. 00613997
Bank Sorting Code: 80-31-50

Please tell your bank that you will pay all charges so that EurOpen will receive the full amount.

2 By UK Cheque or Bankers' Draft, made payable to EurOpen and drawn on a UK bank. Eurocheques are acceptable, but each cheque must be £100 or less.

If you wish to pay by Credit Card (Visa, Access, Mastercard OR Eurocard) please use the £ Sterling prices only. Card details appear on the booking form.

Please indicate clearly on the booking form whether you are using ECU or £ Sterling for payment. Your invoice will then be raised in this currency.

To book a place at the Tutorials and/or Conference, complete one booking form for each person and return it with full remittance or with evidence of payment to:

EurOpen
Owles Hall
Buntingford
Herts SG9 9PL
United Kingdom

Telephone +44 763 73039
Facsimile +44 763 73255
Email europen@EU.net

Use a photocopy of the booking form for each person. Please note that bookings can only be accepted when accompanied by payment or evidence of payment.

Telephone bookings can be accepted when paying by Credit Card. Fax bookings can be accepted when paying by Credit Card or direct to the bank (proof of payment to be faxed also).

The EurOpen Secretariat will acknowledge your bookings by sending a receipted invoice together with full details for registration.

Cancellations

It is regretted that no refund of fees will be possible in cases of cancellation, unless the cancellation is made more than one month before the start of the Conference. No cancellation will be accepted unless it is sent to the EurOpen Secretariat in writing.

Booking Form for Conference and Tutorials

Please complete this form and send it, with cheque or evidence of payment, to *EurOpen Secretariat, Owles Hall, Buntingford, Herts SG9 9PL, UK* (Block capitals please). Please note that forms sent without cheque or evidence of payment will be returned to you unregistered.

Surname _____ Usual first name _____

Company / Organisation _____

Address _____

Country _____ Post / Zip Code _____

Telephone / Fax / Telex / Email _____

EurOpen member? (National Group or direct) Yes ☐ No ☐ Student? Yes ☐ No ☐

Please read the sections on "Costs" and remember that pre-booking saves money.

Please indicate the currency to be used for payment: ECU ☐ £ ☐
and complete this form using the relevant prices stated on page 13.

Conference

Please reserve me a 3 day place for the Technical Sessions. _____

Tutorials *(members only)*

Please reserve me a place for Tutorial No: _____ on Monday 16th September _____

Please reserve me a place for Tutorial No: _____ on Tuesday 17th September _____

Do you require Vegetarian meals? Yes ☐ No ☐

Extra ticket for Conference Dinner ☐
(One is included in the Conference price)

EurOpen

Please enrol me as an institutional member of EurOpen via the appropriate national group.

Yes ☐ No ☐

Tape

Please reserve me a copy of the Conference tape.

☐ 1600 bpi, 1/2" reel _____

☐ 1/4" QIC-24, cartridge _____

Total amount _____

Please read the section **"How to Pay and Book"**

Payment Method

☐ UK Cheque, Banker's Draft or Eurocheque. The Cheque must be enclosed.

☐ Direct Payment. The bank advice note showing details and date of payment must be enclosed.

☐ by Visa

☐ by Access / Eurocard / Mastercard

*All bank charges must be borne by you and not EurOpen – please tell the bank this. EurOpen **must** receive the actual amount due.*

Name as it appears on the card (block capitals) _____

Address of cardholder _____

Card Account Number _____ Date of Expiry _____

Signed _____ Date _____

Executive Report

Helen M W Gibbons
European Forum for Open Systems (EurOpen)
Buntingford
Hertfordshire
United Kingdom

Email europen@EU.net

Helen Gibbons is the business manager of EurOpen and is contactable at the EurOpen secretariat.

Governing Board

The EurOpen Governing Board held its Spring meeting in Tromsø over two days, 18 and 19 May. Representatives were present from the National Groups of France, Germany, Italy, Norway, Great Britain, Spain, Iceland, The Netherlands, Finland, Ireland, Hungary, Austria, Portugal, Yugoslavia, Denmark and Belgium.

During the meeting applications for affiliation were accepted from Algeria and Tunisia, making 22 National Groups in all. Further contacts are being established with representatives of Groups being formed in Greece, Bulgaria and Poland, and improved contacts are sought with Turkey.

The most important items under discussion at the meeting were future conferences, publications, EUnet, the financial situation and future subscription algorithms, new by-laws, budget priorities, Working Groups and participation and co-operation with other bodies.

Conferences

Conferences are a large part of EurOpen activities and the Autumn 91 Conference is already organised for Budapest on 16-20 September - full details appear in this Newsletter. The next important event will be in Jersey on 6-9 April 1992 and will take the form of a workshop in co-operation with USENIX. There will be two days of Tutorials and two days of Technical papers. Joint chairs for this event will be Frances Brazier and David Tilbrook. Neil Todd, who has been officially appointed EurOpen Tutorial Officer will be organising the Tutorials.

In Autumn 1992 a major event organised jointly by EurOpen and UniForum and encompassing a 3 day Conference, 2 days of Tutorials and a 3 day Trade Show will be held in Utrecht, The Netherlands. It is intended that this event, entitled "OpenForum" will be organised annually in the Autumn from 1992 onwards, but every care is being taken to ensure that this does not have a detrimental effect on other National Group Conferences close to it, such as organised by GUUG.

It is hoped that Spain might be the venue for Spring 1993. The Executive Committee has been tasked with preparing a full future programme while awaiting the appointment of a new Events Manager.

Publications

EurOpen has several publications available for sale, and a full list is to be found elsewhere in this newsletter.

All National Groups have been asked to submit copies of their own newsletters to the Publications Officer, Frances Brazier, marked up with abstracts in English of the most interesting papers so that these could be picked out and translated for use in the EurOpen newsletter. It is also thought to be a helpful service to publish lists of abstracts so that members could know what was available from each National Group.

The EurOpen newsletter is a prime visible service of EurOpen and constant efforts are being made to improve the quality. It is sent free to every member of EurOpen, but the price to non-member subscribers has just gone up by $^1/_2$ ECU per issue to 30 ECUs per year.

It is intended to continue to publish an EMail Directory, but this will in future incorporate the Membership Directory as well. Apologies have been expressed to all concerned for the late appearance of the last Email Directory which had been printed in Yugoslavia as an experiment, but had encounted difficulties.

Thanks are however expressed to YUUG for the tremendous efforts that they made.

As present EUnet is considering if in the future Directories can be made available on central site.

AFUU have recently published a Portability Guide, based on the results and findings of their Working Group on Migration and Portability. Efforts are at present being made to see if this can be translated into English for publication within Europe.

CommUNIXations

UniForum's international magazine has been distributed twice now to the EurOpen membership. The first time it was considered as a service to the membership but as it did not contain the hoped for international information, a second issue has been sent out free of charge.

Reduced subscription rates are offered to members and details can be obtained from Owles Hall.

EurOpen Working Groups

AFUU, The Executive Committee and Jean-Michel Cornu in particular, have been making strong efforts to launch EurOpen Working Group activities, but these have been slowed down by lack of response from National Groups.

Having completed the initial organisation, AFUU has passed over the co-ordinating function to the EurOpen Secretariat. Initially the work is directed towards finding out which countries have the same interest in a common topic. The Working Groups are intended to small groups of between 5-15 people who meet regularly to exchange information and discuss a specific topic and then produce their findings for the benefit of all.

If anyone is interested in having a presentation on the subject of Working Groups made to their Group they should contact the EurOpen Secretariat.

The benefits of Working Groups are that they bring together people interested in a specific topic; they give the opportunity to share "Working Outputs" between National Groups; they open the door to funding from external organisations; and they help to position EurOpen as a major player in the Open Systems market.

Italy and France have already set up Working Groups and it is understood that Hungary, Ireland, Algeria, Belgium and Germany are planning to do so shortly.

Jean-Michel Cornu writes about recent Working Group progress else where in this newsletter.

Changes

Major changes are being discussed both to the constitution and by-laws and to the method of calculating contribution payments. The Governing Board by-law sub committee, under Georges Schild, has put together a complete new draft set of by-laws and these are being refined for consideration at the next Governing Board meeting. The Executive Committee is currently looking at ways in which the change over can be most efficiently effected.

On contributions, it was agreed at the last Governing Board meeting that the method of contribution calculation would be changed so that the contributions from the National Groups would be based on the following principles:

- The annual budget of EurOpen's Income and Expenses
- The gross National Product for each country
- The number of labels submitted by National Groups for service distribution
- Corrective factors

The Governing Board sub committee dealing with this subject is in the process of finalising the scheme.

Finances

In the past few years EurOpen has succeeded in ending the year with a small profit, which is a satisfactory situation for a non-profit making body, but unfortunately the investment income, which should have been used for reserves in order to compensate for the effects of inflation, had to be used for the general operations, which are expanding rapidly. Tight control on the financial situation will continue to be maintained.

AFUU, which is the largest group within EurOpen, and therefore, by far, the main contributor, financially, has had some problems in paying its current contribution under the present scheme. The Governing Board therefore agreed that other National Groups should increase their contribution on the basis of 2 ECUs per member more this year to EurOpen to compensate a 10% reduction for economy of scale granted to AFUU in 1991, while waiting for the results of the fees sub committee. EurOpen already operates a grant scheme from reserves for small new groups and those with hard currency problems.

Relations with other Groups

Relations with other Groups continues to be actively pursued through the joint events with UniForum and Usenix already mentioned, and through co-operation on Publications and Standards.

EurOpen was also represented at the meeting organised by X/Open and UniForum in Dallas, USAearlier this year along with representatives from Open Systems Users Organisations from all over the world. A common statement of issues of concern from all the organisations represented has been prepared and published. Further meetings have been planned.

The Executive Committee

Johan Helsingius has resigned from the Executive Committee, and this leaves a present Executive of Michel Gien, Frances Brazier, Norman Hull, Kim Biel-Nielsen, Nigel Martin and Ernst Janich (with Teus Hagen and Neil Tood as active helpers).

The Executive has met twice since the Governing Board meeting in Tromsø and will meet again on 2 August in Paris. It undertakes a vast amount of work and all concerned give freely of their time in the service of EurOpen. This load is gradually being eased by the introduction of paid managers. Glenn Kowack is already on board as EUnet Manager and it is hoped that a General Manager will soon be appointed.

Soviet UNIX System Users' Group

Peter L.Brusilovsky
International Centre for Scientific and Technical Information
USSR

Email plb@plb.icsti.su

The First Conference of the Soviet UNIX System Users' Group (SUUG) was held in Moscow during In 29 October - 2 November 1990.

The purposes of the Group are:

- promote the exchange of ideas and regularly issue of bulletins in the area of common interest between users of the UNIX operating system, arrange an annual conference and support the information network, etc;
- cooperation between users of UNIX-like systems in the USSR and drawing up users' recommendations;
- communication with foreign and international organisations; promote the exchange and distribution of information between the group's members.

However, the Group does not intend to collaborate only with the users of the concrete operating system. The global purpose of the Group is the infrastructure for program research and development in all software branches.

The SUUG was officially registered on 7 September 1990.

The Conference

The Conference was the first serious action of the Group. More than 250 representatives from Moscow, Leningrad, Novosibirsk, Taganrog, Vilnus and other cities and nearly 20 foreign guests took part in the Conference.

The reports and information presented there reflected the recent state of the Soviet software development in UNIX environment and its perspectives in the near future. The main interest was caused by the reports of the foreign guests. It is necessary to mention the extremely high level of the invited guests. Among them were UNIX veterans such as Edward Gould (MT Xinu), Michael O'Dell (Bellcore), Bill Joy (Sun Microsystems) and others.

Beside the reports devoted to the development UNIX system as whole (Michael O'Dell and Ray Davis from Convex spoke about installation of UNIX on supercomputers, Daniel Klein from Carnegie-Mellon University - about the safety of passwords, Judith Grass from AT&T - about C++ tools) one can notice two groups of reports. The first concerns the problem of supporting the process of software elaboration and management. It was mentioned in the reports of Ed Gould concerning the administration of Xinu project and David Tilbrook (Nixdorf, Canada) - about the technology and further development of big program complexes. The second group of reports was devoted to the principles of open systems creation. Among them was Bill Joy's lecture about the open systems philosophy, professed by Sun Microsystems; Johan and Jaana Helsingius delivered a report concerning the methodology and technology of the creation of open systems on the basis of the existing and elaborated standards; the report by Jaap Akkerhuis was devoted to the systems created to use the integrated kinds of information (multi-media design). Nandor Horvath made a report concerning the E-mail network in Hungary.

Besides the informal presentation of Sun Microsystems in Bill Joy's report, two well known American firms were explicitly presented at the conference: Intel (Dmitry Rotow lecture) and Santa Cruz Operation (Roger Chalke lecture).

The importance of the Conference can be compared with such biggest actions as Moscow International Computer Club forum, Moscow PC World Forum, "Informatika-90" exhibition.

The Joint Meeting

As part of the Conference program the joint SUUG members' meeting was held. It fulfilled the process of the SUUG arrangement. The new members of SUUG board and preliminary 1991 budget project were approved at this meeting.

The SUUG decided to hold the 2nd Conference in September 1991, and a workshop in November 1991. For more information see below.

In addition to arranging the 2nd Conference, the SUUG plans for 1991 to activate its efforts aimed on the creation of SUUG net and to issue bulletins and other publications.

An Invitation To SUUG '91

The 2nd SUUG National Annual Conference will be held in Vladimir, from 23 to 28 September 1991. The main items for discussion are the current state and perspectives of development of the UNIX operating system in the Soviet Union and abroad, development of application software tools and programs in UNIX environment.

There will be papers from invited foreign speakers as well as native Russians. Panel discussions on various problems will be organised. The abstracts of the papers will be published before the beginning of the conference.

Conference will be held in the Kljazma hotel in Vladimir. Vladimir is an old Russian town, about 300 km from Moscow. We are planning a compact meeting (of about 120 attendees).

Please contact:

Alexander Fridman
IPIAN, 117900 Moscow, Vavilova st. 30/6
Telephone +7 095 930-0516
e-mail: <af@alf.public.su>
Peter Brusilovsky,
ICSTI, 125252, Moscow, Kuusinena st., 21b

Telephone +7 095 198-7055
Facsimile +7 095 943 0089

Email plb@plb.icsti.su

The SUUG Workshop On UNIX And Applications

Moscow, 25-28 November 1991

The organisers: Soviet UNIX systems Users Group (SUUG), International Centre for Scientific and Technical Information (ICSTI)

The workshop will discuss the problems of the development of UNIX application systems, the current state and developing perspectives of UNIX-based application systems market. The workshop is oriented on users and authors of the following systems:

- Data Base Management Systems
- spreadsheets and integrated systems
- business and industrial graphics
- CAD/CAM
- E-nets and E-mail
- command languages and shells
- tools for interface design
- AI and knowledge based systems
- tutoring and help systems

During the workshop reviews, demonstrations and discussions will be presented. The exhibition of hardware and UNIX application systems will be presented.

The workshop will be hold in Moscow, 25-28 November 1991 in the ICSTI building.

If you would like to take part in the workshop or in the exhibition, please contact the Executive or Program committees for more detailed information.

Workshop Executive Committee

Elmira Muhutdinova
INEUM, Moscow, GSP-1, Vavilova 24,
Soviet UNIX System Users' Group

Telephone +7 095 939-0056

Email elmira@jvd.msk.su

Program Committee

Peter Brusilovsky, Alexander Fridman - addresses as above.

DKUUG Report July 1991

Mogens Buhelt
DKUUG secretariat
Brønshøj (Copenhagen)
Denmark

Email mogens@dkuug.dk

DKUUG (Danish UNIX systems User Group) was established in November 1983 and has grown steadily, measured by the number of members as well as in activities. The membership count on the 20 July 1991, and the membership fees for 1991 are:

	Number	Fee/year
Grand members	18	560 ECU
Organisational members	317	280 ECU
Individual members	56	76 ECU
Total membership number:	391	

Grand members are organisational members, that pay double fee and have more services available. They are typically larger vendors etc.

We have no figures for the distribution of members by classes like hardware vendors, software vendors, end users, universities etc., but the share of vendors is quite large.

Manpower

The board from November 1989 was re-elected at the Annual General Meeting in November 1990. After Tonny Andersen's recent resignation the board consists of:

Keld Jørn Simonsen	University of Copenhagen, chairman, chairman of administration, external, and standardisation sub-committees, representative in EurOpen.
Peter L. Petersen	Aalborg University Center, vice chairman.
Kim Biel-Nielsen	UNIWARE Denmark a/s, chairman of member meetings subcommittee, representative in UniForum.
Kim F. Storm	Texas Instruments, chairman of DKnet subcommittee.
Kim G. Frei	Dansk Data Service.
Mary-Ann Frydendahl	
Hans Kierulff	Merkur Data A/S, chairman of PR and marketing subcommittee.
Anne Hertz Kühnell	Gartner Group Scandinavia.
Gitte Rasmussen	Dansk Data Elektronik A/S, chairman of newsletter subcommittee.

In addition to the board members approximately 12 people have joined one or more of the subcommittees on a voluntary basis, and six people are "employed" part time as paid assistants: two "backbone watchers", two newsletter editors, and two people at the secretariat.

Events

We have about 10 member meetings (conferences) each year. For 1991 and the start of 1992 they were/will be:

- 1991-02-13--14: UNIX & Communication (Workshop)
- 1991-03-21: Management systems.
- 1991-04-18: Graphical user interfaces and multimedia.
- 1991-06-03: Communication and network.
- 1991-06-04: System development and CASE.
- 1991-06-20: Standardisation
- 1991-08-29: The UNIX market - a snapshot
- 1991-09-26: Publishing and image processing
- 1991-10-31: Client-server solutions
- 1991-11-28: "Open" systems - how open?
- 1992-02-06: Executive information systems

The member meetings are held during working hours, either for the whole day or only for the afternoon. The participation fee for the whole day meetings are typically 150-200 ECU, and the meetings typically attract 60-80 attendees. The half day meetings are free for members, and they typically attract 40-60 participants. Non- members are also invited to the meetings, and they pay about 50 ECU more than the members.

In addition to the daytime member meetings we have a workshop (club evening) one night per month (except holiday periods) in

Copenhagen and in Aalborg. These meetings are for members only, and they are free.

The annual income for member meetings was in 1990 about 60,000 ECU, and the expenses were about 65,000 ECU.

We expect to participate this year (as usual) with a booth at the exhibition "Mikrodata 91", 23-27 October 1991.

Network

DKUUG's "subsidiary company" DKnet runs the Danish backbone. 90 sites are connected, including 29 on news. Network managers: Jørgen Jensen and Kim Chr. Madsen. Email netpasser@dkuug.dk; Telephone +45 31 39 73 22:

The annual income from network fees were about 87,000 ECU in 1990, and the expenses were about 73,000 ECU. The aim is balance the two.

Budget

The budget for member meetings and DKnet is stated above. Besides that, our income from member fees were about 100,000 ECU in 1990, and the most important expenses per year were:

- EurOpen fee: 19,000 ECU
- Secretariat (fee, stationary, postage, phone etc.): 23,000 ECU
- DKUUG Newsletter: 11,000 ECU
- Travels, Standardisation participation: 12,000 ECU
- Exhibitions, PR 6,000 ECU

Member services

Own publications:

- DKUUG Nyt (Group newsletter) - 10 issues per year.
- Dansk UNIX markedsoversigt (UNIX products directory): 1 issue per year.

Own, informal publications:

- Member list
- Information for new members
- Info on tape distribution
- Info on DKnet (mail and news) (new issues as required)
- Sale to members at special prices:
- UNIX-bogen (UNIX - the book, Danish version)
- EurOpen Newsletter (additional subscriptions)
- UniForum UNIX Products Directory
- CommUNIXations (subscription to International Edition of UniForum Newsletter)

Tape distribution:

- 10-12 different tapes on stock. Tape formats: QIC-24 cartridge and 0.5", 1600'

NEWS from the IUUG

Brian O'Donovan
(Secretary of the IUUG)
Digital
Galway
Ireland

Email odonovan@cs.tcd.ie

Things seem to be looking much better for the IUUG now than they were last year. Our monthly meetings have been continuing and seem to be well attended. Our last three monthly meetings featured Annrai O'Toole talking about COMMANDOS, Wayne Sulivan talking about TeX and Robert O'Dea talking about OSF/Motif. After a short summer break these talks will be restarted in September with a talk by Donagh Mc Cabe from Digital about Network Applications Support (NAS) which is Digital's architecture for how distributed applications should be built.

However, the main news from the IUUG is that we intend to hold the first ever IUUG (annual ??) conference on 23 September. The conference will almost certainly be held in Trinity College, Dublin. The details of this conference had not been finalised at the time I wrote this article; although they will probably be finalised by the time you read the article... However, I can tell you that we will be bringing in several internationally renowned UNIX experts to speak at the conference. The keynote speaker will be someone you all know very well.

By the time you read this article all the arrangements will have been made for the conference. You can find out more by contacting the conference organiser Annrai O'Toole from Trinity college (Telephone +353 1 772941 or email aotoole@cs.tcd.ie). This event will provide you with an opportunity to hear many well known speakers who would not normally visit Ireland. I would advise you to take up this opportunity.

The next piece of news is that the IUUG is now a company named IUUG Ltd. This change will have very little effect upon the running of the IUUG, but it will give added protection to the members in case the group becomes insolvent. Under the old scheme it would have been possible for the members of IUUG to be held personally liable for any debts than might be incurred.

In keeping with this principle we have also split off the networking activity into a company named IEunet Ltd. This allows a certain amount of operating independence to the network while at the same time keeping it under the control of the IUUG itself. You should not notice any immediate change in the running of the service because the same people who have been running your network are also the directors of the new company.

Unfortunately, this new company will have to charge VAT on all network bills. However, many of the network subscribers will be able to claim this VAT back.

The network infrastructure has recently been significantly upgraded. A leased line has been installed to UKC. Hence the mail routing machine has direct access to the US Internet. Some new 9600 baud modems are on order so the incoming lines should not be busy for as long. Negotiations are at an advanced stage to get a local computer manufacturer to donate a very powerful machine to the IUUG for use as a mail router.

Tim Murphy has started a network information server which can be used to retrieve various sources as well as IUUG related documents. The principle of this network information server is that you send an email message to it and depending upon the content of the email it decides what to send you. To find out more about this service you should send an email containing the single word "help" to info-server@maths.tcd.ie. The information server will reply with a message giving details of how to access it.

The IUUG newsletter has already appeared once (possibly a second will appear before this article is published). Tim Murphy (tim@maths.tcd.ie) is acting as editor for this newsletter and he would be very grateful to receive any articles that you might write.

Report on i2u Annual Convention

Valeria Schiavi
University of Milan
Italy

Email valeria@siemil.uucp

Computer Center of the University of Milan, Italy, Ms. Valeria Schiavi works now as freelance. Her task usually consists in assisting the technical staff of EDP sectors involved with UNIX administration and data communication. Inside i2u she works on behalf of the Executive Board to support i2u activities, acting as technical interface between i2u members and "others".

i2u annual convention

The i2u annual convention took place on 15-17 May 1991.

The Convention was made up of the three "classical" components

- Tutorials
- Conference
- Exhibition

as most national and international UNIX Conventions are.

The Conference itself began on 16th May and the first session focused on a very appealing theme: "UNIX evolution". Invited speakers were P. Weinberger for Plan 9, M. Gien for Chorus, A. Forin for Mach, and K. Bostic for BSD4.4. The style of presentation was much more technical than commercial - a style that UNIX users appreciate very much.

All recent projects in the operating system area adhere to a modular architecture, where a set of system servers are layered upon a micro-kernel. A UNIX or UNIX-like program interface is preserved.

P. Weinberger opened the session by reviewing the evolution of computing environments, from mainframe to work stations and distributed systems, and the consequent necessity for new operating system models. He stressed that the goal for Plan 9 is ease of use and easy administration in distributed environment, characterised by terminals equivalent to diskless work stations; CPU multiprocessor servers from which big computations can be exported; file servers; fast network connections. This goal is achieved by means of an object oriented architecture, where objects are mainly files.

He admitted that Plan 9 hasn't yet the industrial strength necessary to make it a product.

M. Gien introduced Chorus-v3 as micro-kernel solution to the UNIX model. He stressed that the new architecture addresses system builders rather than application writers, providing a structured and modular way to cope with the increased complexity of operating systems development. In contrast to the scheme of Plan 9, specialised object servers are provided for file/process/memory/device management. They rely on the tools provided by the "Nucleus", a minimal real-time kernel that integrates distributed processing and communication.

A. Forin described the key components of MACH, the basic features of micro-kernel, and how UNIX (or other operating systems) are implemented on a MACH kernel as an application program. Hardware management (virtual memory, scheduling and device management) are provided by the kernel, while UNIX services are provided by the UNIX server, a MACH task with multiple threads of control.

Current advanced researches deal with real-time and enhanced security.

The presentation of K. Bostic was very pragmatic, as he stressed the role of (existing) UNIX as available product on which users have invested a lot of resources and which they can really work with.

At the end of this session a very pointed question was put to the speakers: *"How do you plan to distribute these new systems?"*.

Answers:

- BSD 4.4: When ready, non-AT&T components will be distributed as public domain software, maybe by ftp.

- Chorus: They sell the system.
- Plan 9: Tapes are already available for non-profit users at minimal cost and by licence agreement with AT&T.
- MACH: Requires AT&T licence for AT&T software. Apart from that, it's free... but be aware:
  << You get what you pay for... >>

A special session was dedicated to the evolution of UNIX standards.

J. Andrews outlined the international scheme operated by X/Open in order to provide a mark of conformance to vendors. He gave useful information for European and Italian operators and introduced the new XPG4 program.

OSF's technologies and directories and the UNIX SVR4 Roadmap were illustrated by P. Wahl and M. Rizzi respectively.

As a result of supporting standards, OSF can now offer DCE and Motif as products which can be layered on OSF/1 as well as on other operating systems, while UNIX SVR4 tends to enlarge the capabilities of UNIX as an open system by offering the support for multiprocessor technologies, OSI platforms, enhanced security.

In this same session the theme of security was covered by the presentation of two Siemens-Nixdorf implementations of UNIX, which are classified in class C2 and class B2, respectively, of the Orange Book.

Another general session concerned market trends.

All the speakers outlined the increasing penetration of UNIX systems into the marketplace, not only in the UNIX traditional engineering market but also in the commercial one.

The most leading factors are the advantages in price/performance ratio compared to proprietary solutions, especially due to powerful Intel and Risc platforms, and to the emerging interest for innovative solutions as network computing, UNIX desk-top work stations, Oltp and fault-tolerant architectures.

A significant step ahead will be the binary portability, that is now a critical issue.

Other overlapped sessions dealt both with user experiences and with new products and technologies.

Topics covered were networking and distributed applications; performance evaluation and capacity planning; multimedia; OLTP and real-time; software development tools and case; parallel and multiprocessor architectures; PC integration; porting; user interfaces.

User experiences confirm that UNIX in Italy is now installed in large as well as in small companies and is used for wide and critical applications. Just a couple of examples: the traffic control system, and the supervisor system of thermo-electric national power plants are based on UNIX.

The i2u organising board remarked that this year the attendance was characterised by a number of end users and software houses much greater than the number of suppliers. Therefore they feel with a certain satisfaction that they addressed the right audience.

Announcement And Call For Papers Convention UNIX '92

5th Conference and Exhibition

23 - 27 March 1992

CNIT, Paris - La Défense

Open Systems: Inter-operability

UNIX is a registered trademark of UNIX System Laboratories in the United States and other countries.

The Association Française des Utilisateurs d'UNIX et des Systémes Ouverts (AFUU) is organising, in collaboration with the Bureau International de Relations Publiques (BIRP), CONVENTION UNIX '92 from 23 to 27 March 1992, at CNIT, Paris-La Défense.

This event is centered on tutorials and conferences dedicated to UNIX and Open Systems. A display of hardware and software organised by more than 150 exhibitors, will take place at the same time.

CONVENTION UNIX '92 is an even more important event than usual. 1992 will provide the occasion for us to celebrate the tenth anniversary of AFUU. Consequently, extra care will be taken in the organisation of the CONVENTION, both from the point of view of the quality of the conferences, tutorials and the exhibition, as of certain events whereby we shall celebrate this 10th anniversary together. The organisers are open to all suggestions in line with this aim.

The Theme

Users have more and more sophisticated operational requirements and the facts prove that these needs can no longer be satisfied in the sole context of proprietary products. Users want open and evolutionary systems which make the best use of the variety of solutions offered to them.

Software and inter-operability is one solution. It may be defined as the ability which applications and systems of different origins to cooperate to produce an integrated service for the user. In the same way as we have been able to come to terms with the hardware heterogeneity, inter-operability enables bridges to be built between different types of software components in order to develop fully open systems.

The purpose of CONVENTION UNIX '92 is to clarify and illustrate this notion of inter-operability, both from the technical and economic points of view. What does it provide to the users and to the companies? What are the benefits for the users? What are the economic benefits? How is it done? What are the concrete and demonstrable examples of it?

The Programme Committee would like to receive submissions in the following three categories :

- Man/machine interfaces:

  They provide the user with an ultimate view of the realisation of his needs. Hence, the man/machine interface is one of the keys to inter-operability. By facilitating communication with the machine (integration of voice, data and picture), they should lead to better productivity and to "inter-operability" between men and machines. Although the ergonomics and operator comfort of existing tools are now an established fact, the variety of interfaces still remains a barrier to be overcome.

- Applications:

  Applications are the essential components in the creation of services (data management, transactional processing/OLTP). How do they cooperate? Are the emergence and the adherence to standards, and API definitions, sufficient to guarantee the interaction between components? How can compatibility be defined? Does compatibility between interfaces guarantee the portability of applications? How can we re-use these components to improve return on software investment?

- The tools of inter-operability:

  What are the bases on which inter-operability between men, interfaces and applications may be founded? What are the basic services necessary to the applications and to the man/machine interface to permit cooperation (distributed systems, communications protocols, RPC)? What are the languages and methods which facilitate the implementation of inter-operability (programming and object-oriented languages, interfaces generators)? What are the consequences on software engineering? What are the development costs?

The purpose of this Conference is to create a vast forum in which information can be exchanged about experiences with inter-

operability, the technology, the risks and benefits and the applications in everyday use. The papers may cover users experiences, industrialists' and manufacturers' strategies, technological innovation in the research world, or even the broad principles in a particular field.

Tutorials

The tutorials will provide the listeners with lectures on clearly defined subjects. Their purpose will be to present the state-of-the-art of the most important points in the implementation of inter-operability. As it was the case last year, these tutorials will be led by experts of national and international fame. The persons interested in lecturing at a tutorial are invited to contact the chairman of the Programme Committee as soon as possible.

Dates to Remember

- 27 September 1991:

  deadline for the receipt of full papers or extended abstracts by the Convention Secretariat

- 25 October 1991:

  notification to authors of the Programme Committee's decision

- 13 December 1991:

  deadline for the receipt of the final texts by the Convention Secretariat

Submission of Papers

Paper proposals should have a title, the name of the author and the organisation to which he belongs. These proposals should include the complete text (5 to 10 pages) or at least an extended abstract (2 pages). They will be examined on the basis of their quality, originality, and the adherence to the general theme of the Convention (the inter-operability of Open Systems) and to the directives given :

- Man/machine interfaces:

  graphics, multi-media, ergonomics, learning

- Applications:

  DBMS, transactional processing, standards implementation

- Tools:

  distributed processing, RPC, multi-personality systems, communication between UNIX and other systems, object-oriented languages, software engineering.

The Programme Committee wishes to include both technical papers and syntheses of different approaches. The quality of the lectures depends to a large extent on the facilities which are provided to the Programme Committee to enable it to make its decision on the acceptance of a paper. Consequently, candidates are strongly encouraged to supply a complete version of their text as soon as possible.

The official languages are French and English. A simultaneous translation service will be available for conferences and tutorials.

Proposals should be sent to:

A.F.U.U.
Secrétariat de la CONVENTION UNIX '92
11, rue Carnot
94270 Le Kremlin-Bicetre - France

Telephone +33 1 46 70 95 90
Facsimile +33 1 45 89 94 20
Telex 263887 F

Email afuuconf@inria.inria.fr

Recommendations to authors:

For the final version of your paper, if it is accepted, it is essential that we receive, in good time :

- a camera-ready copy (that is on paper) of your article
- the source text which will be processed during the preparation of the proceedings

Please note that if your text includes drawings or illustrations, it is essential that we have a printed version of them.

Extended abstracts are not subject to the recommendations below. The source of your document should reach us on one of the following media :

- electronic mail addressed to Philippe Dax (dax@enst.fr) ;
- QIC-24 (150 Mb) cartridge ;
- 3 1/2" or 5 1/4" floppy disk (MSDOS format) ;
- Mac floppy disk.

The format of this source may be one of the following :

- "nroff" or "troff" using macros "mm", "ms", "tbl" and "eqn" ;
- "TeX" or "LaTeX" (if you use non standard macros, please attach them to the source) ;
- ASCII mentioning the characters which are accentuated ;
- WORD (DOS or Mac) ;
- figures should be supplied in PostScript or "pic" whenever possible.

The composition of the proceedings should be in "troff" ("mm" macros) and "TeX", producing the PostScript output which will be sent directly to our printer.

Useful Addresses

Convention Chairman

Sylvain Langlois
Electricité de France - Direction des Etudes et Recherches

Telephone +33 1 47 65 44 02
Facsimile +33 1 47 65 30 01

Email sylvain@cli53an.edf.fr

Programme Committee Chairman

Francois Armandd
Chorus systémes

Telephone +33 1 30 64 82 36
Facsimile +33 1 30 57 00 66

Email francois@chorus.fr

UKUUG Report

Mick Farmer
Department of Computer Science
Birkbeck College
London, UK

Email mick@cs.bbk.ac.uk

Mick is the Secretary of the UKUUG. His primary interests are Ornithology (restricted to the Western Paleartic at the moment because of cost) and Oenophilism (especially pre-1962 Bordeaux, pre-1980 Burgundy, and 1945 Port). His secondary interests include Software Consultancy (to pay for the above primary interests) and Distance Learning Methods (especially interactive video and hypertext). When not pursuing these and other interests he can be found at Birkbeck College (London) where he teaches in the Department of Computer Science.

He lives in Lewisham (South East London) with his wife Sue and a TV called Sonya. His neighbours have two children and a dog.

Start Bit

Is there some aspect of UNIX, or open systems, that you wished we spent more time and space on? Perhaps you feel we should run more workshops and/or tutorials on specific subjects. If so, why not suggest it as a Special Interest Group (SIG) within the UKUUG? We'll help with the infrastructure and starting up, so send e-mail to ukuug-sig-info containing your views. Don't be shy, it's you we answer to!

Liverpool Conference, 15-17 July 1991

I've just returned from a successful UKUUG conference in Liverpool. It followed the usual pattern of starting at lunch time on Monday going through to lunch time on Wednesday. I won't describe the talks as the abstracts can be found elsewhere in this issue.

Bill Barrett and I arrived on Sunday evening. After a quick chinese meal near the station we woke up Hugh Conner, who'd also arrived early, to discuss his arrangements for our next conference in Edinburgh this winter (this demonstrates we're really working hard :-).

It's raining hard on Monday morning as delegates arrive for registration. There's a panic when we find that the pads for the clip boards haven't been delivered; despatch student helper round to Smiths for replacements! The conference finally kicks off after lunch with David Korn talking about his Korn shell. Other speakers argue the merits of different windowing environments as we battle with the catering department about the need for two coffee points during the breaks. We go into town for a meal and return to the bar to find the cogniscenti in full flow, about HCI mainly. It's somebody's birthday so there's a round of "Happy Birthday" accompanied by one of David's sons (just one Kornetto?) on the piano.

Tuesday's a full day and everything's going well except the FaceSaver that refuses to talk to a Liverpudlian Apple LaserWriter. Eventually, I take off to buy a few "thank you" presents for the conference organisers. Martin Beer is into model railways so I go to the mecca of toy train enthusiasts, Hattons, and purchase an N gauge, BR Black 5. I hope he's going to like it! The good news when I return is that the FaceSaver is now running (I knew we should have brought our own breakout box). As we board the coach outside the halls of residence a group of American tourists are getting off their coach to take pictures of the "Penny Lane" road sign. For some reason they seem more interested in our coach which has "Magical Mystery Tour" in large letters down the side and pictures of four young musicians on the front. The conference dinner is in the National Maritime Museum, part of the redeveloped Albert Docks complex. Unfortunately, they don't let us look round before we sit down so we pay extortionate prices at the bar! Afterwards, Dave Edmondson suggests a quiet pub he found the previous night. It turns out to be Karaoke evening and the bar is full of young Liverpudlians singing familiar songs!

The final morning sessions go well and we begin to wind down. There's a sudden rush to catch the FaceSaver before it closes and we even sell a few tee-shirts.

The conference competition was to write a verse or two of a Beatles lyric incorporating a UNIX flavour. The winner was Steve Platt and his entry is given below.

Our thanks to Martin Beer and his team at Liverpool University for organising a successful conference. I hope they enjoyed it as much as we did!

Edinburgh Conference, 16-18 December 1991

Our winter conference is being hosted by Hugh Conner at Herriott-Watt University. Although the winter meeting is traditionally about networking, Hugh tells me that this event will have a marked scottish flavour as well. For further details contact our secretariat at Owles Hall or send e-mail to ukuug-conf-edinburgh@ukc.ac.uk.

All our events are open to EurOpen members at the same price as UKUUG members.

LISA Workshop, February 1992

We intend to hold a workshop in February 1992 concentrating on aspects of Large Installation System Administration (LISA), connected with being part of a networked environment. It is hoped that topics such as configuring and running a Domain Name Server (DNS) and multi-protocol mail routers will be discussed. For further information contact our secretariat at Owles Hall or send e-mail to ukuug-wshop-lisa@ukc.ac.uk. All our events are open to EurOpen members at the same price as UKUUG members.

Belfast Conference, Summer 1992

For further information contact our secretariat at Owles Hall or send e-mail to ukuug-conf-belfast@ukc.ac.uk. All our events are open to EurOpen members at the same price as UKUUG members.

Manchester Conference, Winter 1992

For further information contact our secretariat at Owles Hall or send e-mail to ukuug-conf-manchester@ukc.ac.uk. All our events are open to EurOpen members at the same price as UKUUG members.

Local User Groups

There is currently an increasing interest in creating and/or joining a local user group from our members. This activity is strongly encouraged, so here is a list of existing or embryonic groups. These groups need your support for them to prosper; please contact the local organiser.

A Cambridge local user group is about to be formed. Your local organiser is Jane Hills from Logica (e-mail to ukuug-lug-cambridge@ukc.ac.uk) and she tells me that an initial meeting will be held at Logica, followed by some liquid refreshment at the Alma public house! Contact Jane to give her your support.

The London local user group continues to meet informally at the Fritzroy Tavern (Charlotte Street) on the last Thursday of each month. The local organiser, Andrew Findlay (e-mail to ukuug-lug-london@ukc.ac.uk), should be contacted regarding the exciting winter programme he has devised. As an example of the high standard, our last formal talk was given by Jon Crowthorne of University College and you can find his paper elsewhere in this issue.

The Midlands local user group has held an inaugural meeting and would like to hear from people willing to give support. Contact Kevin Hopkins (e-mail to ukuug-lug-midlands@ukc.ac.uk) for details of the next meeting.

An Oxford local user group is also coming into being. Although nothing concrete has happened so far, contact the local organiser, Andrew Newman (e-mail to ukuug-lug-oxford@ukc.ac.uk), for more details.

FaceSaver

By the time you read this, you will be able to retrieve faces from the UKUUG Software Archive at Imperial College. For details send e-mail to ukuug-faces-info@ukc.ac.uk.

Stop Bit

We're still looking for UNIX nostalgia from the old days. Any lyrics that you can recall gratefully received. To show what we're looking for, here's the winning entry from our Liverpool conference. The manual page is a nice touch!

```
5L      Liverpool Competition Entry       5L

TITLE
"A hard read-write" (or "The dump tape blues")
SYNOPSIS
to the tune of "A hard day's night" by the Beatles
AUTHOR
Steve Platt
LYRICS
You've had a "hard read-write",
       while dump(8) was working like a dog.
Your options must be right,
       when you're appending to a log.
'cause if youre fstab(5) is true(1),
       all those files that pipe(2) through,
will make(1) you feel all write(2).
You need it, "hard read-write",
       from your server, with all the mail(1).
You have to mount(2) it right,
       else your spooling just might fail.
For when your cache is write-through,
       you'll lose those header lines too,
without that "hard read-write".
REFRAIN
In your $HOME,
       everything seems to be gone ...
In your $HOME,
       "File_System_Full" equals one (1)!
SEE ALSO
"I shot the serif", etc.
17 July, 1991
```

EurOpen Publications

EurOpen publications may be ordered from the Secretariat at Owles Hall.

The publications available are listed on the next page, it is planned that EurOpen members will be able to also order National Groups publications from one single point: the secretariat.

The EurOpen publications available are:

EurOpen Newsletter
Your EurOpen National Group membership gives you one free copy of this newsletter, you may order extra copies of the newsletter for distribution within your organisation. Please contact your national group who will arrange this for you.

Proceedings
Proceedings from some past conferences are still available. The list of those that you can order is on the next page.

European E-Mail directory
This is a must if you are a serious e-mail user. It provides you with a way of referencing e-mail sites thoughout Europe by means of several different indexes.

The second edition is now available.

USENIX Publications

We receive requests for USENIX publications and, as a result of our close cooperation with that organisation, we are pleased to announce a service by which a range of publications can be ordered through EurOpen. The range includes:

;Login:
This is the USENIX newsletter which is published every two months, and carries articles on a variety of topics of interest to UNIX users everywhere and usually runs to around 40 pages.

Computing Systems
This is the journal of the USENIX Association, a quarterly publication which is devoted to the analysis and understanding of advanced computing systems. It is perfect bound with a printed spine for ease of reference and usually runs to around 100 pages.

Proceedings
Some proceedings from past USENIX conferences and workshops are available from stock - the current list is enclosed with the attached order form.

4.3 BSD Manuals
These manuals have proved popular with EurOpen members. We now new have stocks at the EurOpen Secretariat.

Ordering procedure. An order form is enclosed with details of how payment can be made by credit card, by direct payment to the bank or by certain types of cheque or banker's draft.

Publications available though EurOpen

			£	ECU
;Login:	The USENIX newsletter	6 issues/year	20.00	30
Computing Systems	The USENIX Journal	4 issues/year	30.00	44
USENIX Proceedings	Anaheim Conference	June '90	25.00	37
	C++ Conference	1991	25.00	37
	USENIX Conference	Summer '91	30.00	44
	C++ Conference	Apr '90	30.00	44
	Washington DC Conference	Jan '90	27.50	40
	Graphics Workshop V	Nov '89	19.00	28
	Dist & Multiproc Workshop	Oct '89	33.50	49
	Large Inst Sys Admin III Workshop	Sept '89	15.00	22
	Baltimore Conference	June '89	24.00	35
	UNIX Trans Proc Workshop	May '89	13.50	20
	Software Management Workshop	Apr '89	24.00	35
	San Diego Conference	Feb '89	33.50	49
	C++ Conference	Oct '88	33.50	49
	C++ Workshop	Nov '87	33.50	49
	Graphics Workshop IV	Oct '87	17.00	25
	Washington DC Conf	Jan '87	20.00	30
	Graphics Workshop III	Dec '86	17.00	25
4.3 BSD Manuals (EurOpen Members Only)	User's Manual Set Programmer's Set System Manager's Manual		60.00	88

			£	ECU
EurOpen Proceedings	Dublin	Autumn '83	2.00	3
	Nijmegen	Spring '84	5.00	7
	Cambridge	Autumn '84	5.00	7
	Paris	Spring '85	5.00	7
	Copenhagen	Autumn '85	10.00	15
	Finland/Sweden	Spring '87	20.00	30
	Dublin	Autumn '87	20.00	30
	Munich	Spring '90	20.00	30
	Nice	Autumn '90	25.00	37
	Tromsø	Spring '91	25.00	37
Directories	European E-Mail directory, 2nd edition		20.00	30

All prices include post and packing.

The price for ;Login: and Computing Systems is for a one year subscription.

AFUU Publications available fromEurOpen

The following publications are written in the French language

		£	ECU
Dossier Benchmarks AFUU 1989 Results - 132 pages	Edition : Mars 1990 Summary send upon request	24.00	35
	Edition: Mars 1991	34.00	50
Réussir avec UNIX	Edition : 1988 - 15 pages	4.00	5
Petit guide destiné aux directeurs informatiques d'organismes ayant fait récemment le choix d'UNIX.			
Il traite de : L'environnement de développement sous UNIX L'environnement d'exploitation sous UNIX La formation La portabilité Les outils d'UNIX Les outils du marché			
Vivre avec UNIX 2 - Administration du Systéme	Edition : 1988 - 16 pages	4.00	5
Petit guide destiné aux techniciens des organismes qui ont fait récemment le choix d'UNIX			
Sommaire : Introduction Systéme de fichiers configuration de l'espace disque Création du systéme de fichiers Gestion des périphériques Gestion des impressions Administration quotidienne Assistance et environnement de l'utilisateur Sécurité Performances			
(Vivre avec UNIX 1 - Utilisation du systéme (Epuisé))	Edition 1988 - 16 pages	34.00	50
CONVENTION UNIX 90 - Conferences proceedings		14.00	20
End-users and solutions conferences proceeding	Edition Mars 1990 - 165 pages	14.00	20
Technical conferences proceeding	Edition Mars 1990 - 198 pages	14.00	20
CONVENTION UNIX 90 - Tutorials			
Are available the following tutorials (the others are no more available)			
Sendmail, Annexe 1 et 2	Edition Mars 1990 - 112 et 205 pages Par Yves Devillers (Inria)	11.00	15
Postcript	Edition Mars 1990 - 119 pages Par Gilles Dauphin (Telecom Paris)	11.00	15
Langage C++	Edition Mars 1990 - 77 pages Par Frédéric Lung (Consultant)	11.00	15
UNIX Systéme V administration	Edition Mars 1990 - 160 pages Par Michel Wurtz (Institut Géographique National)	11.00	15

UKUUG Publications available from EurOpen

	£	ECU
UNIX — The Legend Evolves 1990 Summer Proceedings, Royal Lancaster Hotel, London	30.00	44

UniForum Publications available from EurOpen

	£	ECU
Your Guide to POSIX	7.00	10
POSIX Explored: System Interface	7.00	10
POSIX Update: Shell and Utilities	7.00	10
Network Substrata	7.00	10
Network Applications	7.00	10
1991 UNIX Products Directory	40.00	60
International Editions of CommUNIXations	Price on application from Owles Hall	

NLUUG Publications available from EurOpen

		£	ECU
UNIX en Standaardisatie	November 1988	17.00	25
UNIX & Mens-Machine Interaktie	May 1989	17.00	25
UNIX & Connectivity	November 1989	17.00	25
Systeem Beheer "Gaan Open Systemen enVeiligheid samen?"	May 1991	17.00	25

EurOpen Tape Distribution

	£	ECU		£	ECU
EurOpenD1 R6	42.00	60	EurOpenD13	42.00/60.00	60/86
EurOpenD2	42.00	60	EurOpenD14	42.00/60.00	60/86
EurOpenD3	42.00/60.00	60/86	EurOpenD15	82.00/63.00	118/90
EurOpenD4	50.00/60.00	72/60	EurOpenD16	42.00/60.00	60/86
EurOpenD5	42.00	60	EurOpenD17	42.00/60.00	60/86
EurOpenD6	78.00/96.00	112/114	EurOpenD18	42.00/60.00	60/86
EurOpenD7	24.00/60.00	35/86	EurOpenD19	42.00/60.00	60/86
EurOpenD8	42.00/60.00	60/86	EurOpenD20	24.00/60.00	35/86
EurOpenD9	52.00/60.00	75/86	EurOpenD21	42.00/60.00	60/86
EurOpenD10	33.00/60.00	48/86	EurOpenD22	42.00/60.00	60/86
EurOpenD11	42.00/60.00	60/86	EurOpenD23	42.00/60.00	60/86
EurOpenD12	42.00/60.00	60/86			

The first price listed is for $^1/_2$" 9-track reel tapes, the second one is for $^1/_4$" QIC-24 Sun format cartridge.

ORDER FORM FOR PUBLICATIONS AND SOFTWARE DISTRIBUTIONS

This page may be photocopied for use. Please print!

Name: _____ _____

Company name: _____

Address: _____

_____

_____

E-mail address: _____

I would like to order the following publications: _____

_____

_____

I would like to order the following tapes: _____ _____

_____

_____

Please note that for distributions D1, D2 and D4 a copy of your source licence agreement with AT&T for at least UNIX version 7 should be enclosed. Note also that you have to be an EurOpen member (or a member of a national UUG) to obtain tapes at list prices. Non-members will have to pay DFl 300,- per tape extra as handling fee. Please enclose a copy of your membership or contribution payment form when ordering. Do not send any money or cheques, you will be invoiced.

All 1/2", 9-track, reel tapes come in tar format, 1600 bpi. 800 bpi is possible on request. Cartridge tapes come in tar format, written with dd, with a blocking of 126b. This is a so-called QIC-24 format, written on a Sun. QIC-11 is available on request.

Tape format, either 1/2" 9-track, or 1/4" cartridge: _____

Copy of EurOpen (or national UUG) membership (or payment) form enclosed? Yes / No

Copy of AT&T source licence enclosed? (For D1, D2, D4.) Yes / No

"I declare to indemnify the European Forum for Open Systems for any liability concerning the rights to this software, and I accept that EurOpen takes no responsibilities concerning the contents and proper function of the software."

Date: _____ Signature: _____

Please see payment details on next page.

I enclose my payment in the sum of £_____ /I have completed my credit card details on next page

---

* delete as applicable

PAYMENTS

EurOpen is a European Federation of National Groups and as such prefers to use the European Currency - ECUs - for payments. To pay in ECUs please note the two methods of payment.

1. By Direct Payment to EurOpen's ECU Bank Account

The Bank of Scotland
International Division
Operations Dept
PO Box 86 120
St. Vincent St.
Glasgow G2 5DZ
Scotland

Account Number: 41791 ECU 01
Bank Sort Code: 80-20-13

Please tell your bank that you will pay all charges so that EurOpen will receive the full amount.

2. By UK Cheque or Bankers' Draft, made payable to EurOpen and drawn on a UK bank. Eurocheques are acceptable, but each cheque must be for 170 ECUs or less.

EurOpen is setting up facilities in Europe for delegates paying by Credit Card in ECUs, but unfortunately this is a vey new facility and not fully functional at the time of printing this newsletter. If you have real difficulties in paying in ECUs you may pay in £ sterling using the following methods of payment:

1. By Direct Payment to EurOpen's bank. Please tell your bank that you will pay all charges so that EurOpen will receive the full ammount due. The Bank of Scotland Account Number: 00613997 61 Grassmarket Bank Sort Code: 80-31-50 Edinburgh Scotland EH1 2JF

2. ☐ By VISA ☐ by ACCESS/EUROCARD/MASTERCARD

Name as it appears on the card (block capitals) _____

Address of card holder _____

_____

Card Account No: _____ Date of Expiry: _____

Signed: _____ Date: _____

Calendar of UNIX Events

This is a combined calendar of planned conferences, workshops, or standards meetings related to the UNIX operating system. The information here is collected by those listed below after an idea by John S. Quarterman of Texas Internet Consulting. The information comes from the various conference organizers, ;login:, Communications of the ACM, CommUNIXations, and many others. We encourage others to reuse this information, but we ask for proper acknowledgment, for example by including this statement.

If you have a UNIX related event that you wish to publicise then contact either John S Quarterman at *jsq@tic.com*, Alain Williams at *addw@phcomp.co.uk*, Susanne W Smith at *sws@calvin.wa.com*, or Carolyn Carr at *carolyn@usenix.org* giving brief details in the style that follows.

Abbreviations:

APP	Application Portability Profile
C	Conference or Center
CC	Computer Communication
G, MD	Gaithersburg, Maryland
GM	Generam Meeting
LISA	Large Installation System Administration
MHS	Message Handling Systems & Application Layer CommunicationProtocols
OSE	Open Systems Environment
S	Symposium
SEDMS	Symposium on Experiences with Distributed and MultiprocessorSystems
T	Tradeshow
U	UNIX
UG	User Group
W	Workshop

1991

September 3-6	SICON 91', *Raffles City CC, Singapore*
September 9-13	OIW, *NIST, G, Maryland, USA*
September 10-12	European Sun User Group CT, *NEC, Birmingham, UK*
September 11	IUUG C, Network Applications Support (NAS), *Trinity College, Dublin, Eire*
September 16-20	EurOpen, *Budapest, Hungary*
September 23	IUUG, Annual Conference, *Trinity College, Dublin, Eire*
September 24-27	AUUG CT, *Darling Harbour, Sydney, Australia*
September 25	RevCom, NesCom, *IEEE HQ, New York, New York, USA*
September 26	IEEE Standards Board, *IEEE HQ, New York, New York, USA*
September 26	DKUUG C, Publishing and image processing, *Denmark*
September 26-28	GUUG C & T, *Wiesbaden, Germany*
September 27	Sinix, GM, *Hilton Hotel, Singapore*
September 30-October 3	LISA, *USENIX, San Diego, California, USA*
October	RIPE, *Geneva area, Switzerland*
October	GUUG W, EUnet, *Dortmund, Germany*
October 1-2	EXUG C, *Cambridge, UK*
October 7-11	Interop, *CC, San Jose, California, USA*
October 10-11	Multi-User C Show, *UniForum Canada, Montreal, Quebec*
October 16	IUUG C, C++, *Dublin, Eire*
October 21-25	IEEE 1003, *Parsippany, New Jersey, USA*
October 25	Sinix, GM, *Hilton Hotel, Singapore*
October 30	IEEE CS SCC/SAB, *Nashville, Tennessee, USA*

October 31	DKUUG C, Client-server solutions, *Denmark*
October 31	Sun UG-NL, *Amsterdam, The Netherlands*
November 6-8	ISO/IEC JTC1 SC22 WG15, *Stockholm, Sweden*
November 7	User Friendliness C,T NLUUG, *Ede, The Netherlands*
November 13	IUUG C, Postscript Extensions,
November 14	APP/OSE Users Forum, *NIST, G, Maryland, USA*
November 14-15	JUS Symposium, *Osaka, Japan*
November 27-29	AFUU C, Journees UNIX De Grenoble, *Grenoble, France*
November 28	DKUUG C, "Open" systems - how open?, *Denmark*
November 29	Sinix, GM, *Hilton Hotel, Singapore*
December 2-6	IETF, *LANL, Santa Fe, New Mexico, USA*
December 3-4	JUS UNIX Fair, *Tokyo, Japan*
December 4	RevCom, NesCom, *IEEE HQ, New York, New York, USA*
December 5	IEEE Standards Board, *IEEE HQ, New York, New York, USA*
December 8-11	Sun User Group C, *San Jose, California, USA*
December 9-13	DECUS S, *Anaheim, California, USA*
December 11	IUUG C, Large Systems Administration, *Dublin, Eire*
December 16-18	UKUUG C, *Edinburgh, Scotland*
December 9-11	Sun User Group, *San Jose, California, USA*
December 9-13	DECUS S, *Anaheim, California, USA*
December 27	Sinix, GM, *Hilton Hotel, Singapore*

1992

January 13-17	IEEE 1003, *Irvine, CA (location tentative)*
January 20-24	USENIX, *Hilton Square, San Francisco, California, USA*
January 20-24	UniForum, *Moscone Center, San Francisco, California, USA*
February 6	DKUUG C, Executive information systems, *Denmark*
February 19-21	PCI (Publication & Communication), *Chicago, USA*
March	INDC 92', *IFIP TC6, Finland*
March 4-7	Computers in Libraries, *Meckler, Westport, Connecticut , USA*
March 11-18	CeBIT 92, *Hannover, Germany*
May 24-27	AFUU, Convention UNIX 92, *CNIT, Paris-La Defense, France*
April 6-10	IEEE 1003, *Atlanta, Georgia, USA (location tentative)*
April 13-17	EurOpen, *Jersey, UK*
May 4-8	DECUS S, *Atlanta, Georgia, USA*
May 18-22	ISO/IEC JTC1 SC22 WG15, *New Zealand (tentative)*
Summer	UKUUG C, *QueensUniversity,Belfast,NorthernIreland*
June 2-4	UNIX EXPO West, *Anaheim CC, Anaheim, California, USA*
June 8-12	USENIX, *Marriott, San Antonio, Texas, USA*
June 21-24	Sun Users Group, *Washington DC, USA*
July	PCI, *Boston, USA*
July 13-17	IEEE 1003, *Chicago, Illinois, USA (location tentative)*
September 8-11	AUUG C T, *World Congress Centre, Melbourne, Australia*
Autumn	ISO/IEC JTC1 SC22 WG15, *Denmark*
October 6	WG15, *Denmark*
October 19-23	IEEE 1003, *Montreux (location tentative)*
October 26-30	Interop *Moscone C, S.F., CA, USA,*
November 25-29	EurOpen/UniForum, *Amsterdam, Netherlands*
December	UKUUG/UKnet, *Manchester, UK*

1993

January 11-15	TCOS WG, *New Orleans, Louisiana, USA (location tentative)*
January 25-29	USENIX, *Town & Country, San Diego, California, USA*
March 15-18	UniForum, *Moscone Center, San Francisco, California, USA*
March 24-31	CeBIT 93, *Hannover, Germany*
April 5-19	TCOS WG, *Boston, Massachusetts, USA (location tentative)*
April 26-30	EurOpen, *Southern Spain (tentative)*
June 21-25	USENIX, *Cincinnati, Ohio, USA*
July 12-16	TCOS WG, *Hawaii (location tentative)*
Autumn	Europen/UniForum, *Utrecht, The Netherlands*
October 18-22	TCOS WG, *Atlanta, USA (location tentative)*
October 25-29	Interop, *Moscone C, S.F., California, USA*

1994

January 17-21	USENIX, *Hilton, San Francisco, California, USA*
February 14-17	UniForum, *Dallas CC, Dallas, Texas, USA*
March 16-23	CeBIT 94, *Hannover, Germany*
April 18-22	EurOpen, *Switzerland (tentative)*
June 6-10	USENIX, *Boston, Massachusetts, USA*
September 12-16	Interop, *Moscone C, S.F., California, USA*
Autumn	Europen/UniForum, *Utrecht, The Netherlands*

1995

January 16-20	USENIX, *Marriott, New Orleans, Louisiana, USA*
February 20-24	UniForum, *Dallas CC, Dallas, Texas, USA*
May 1-5	EurOpen, *Scotland (tentative)*
June 19-22	USENIX, *Hilton, San Franciso, California, USA*

1996

March 11-14 — UniForum, *Moscone Center, San Francisco, California, USA*

1997

March 10-13 — UniForum, *Moscone Center, San Francisco, California, USA*

Organising Bodies

NIST/NBS/POSIX
Roger Martin
National Institute of Standards and Technology
Technology Building, Room B266
Gaithersburg, MD 20899, USA

Telephone +1-301-975-3295

Email rmartin@swe.icst.nbs.gov

IEEE Computer Society
P.O. Box 80452
Worldway Postal Center
Los Angeles, CA 90080, USA

Telephone +1-202-371-0101

UniForum (was /usr/group)
2901 Tasman Drive
Suite 201
Santa Clara CA 95054, USA

Telephone +1 408 986 8840
Facsimile +1 408 986 1645

./usr/group/cdn
241 Gamma St.
Etobicoke, Ontario M8W 4G7
Canada

Telephone +1-416-259-8122

Tracy MacIntyre
Exhibition Manager
EMAP International Exhibitions Ltd.
Abbot's Cour-t
34 Farringdon Lane
London EC1R 3AU
United Kingdom

Telephone +44-1-404-4844

AUUG Inc
P.O. Box 366
Kensington
N.S.W.2033
Australia

Email uunet!munnari!auug
auug@munnari.oz.au

Telephone +61 2 361 5994
Facsimile +61 2 332 4066

AMIX, c/o IPA
P.O. Box 919
Ramat-Gan
Israel, 52109

Telephone +972-3-715770
Facsimile +972-3-715772

Email amix@bimacs.bitnet
amix@bimacs.biu.ac.il

Japan UNIX Society (JUS)
#505 Towa-Hanzomon Corp. Bldg.
2-12 Hayabusa-cho
Chiyoda-ku, Tokyo 102
Japan

Email bod%jus.junet@uunet.uu.net

Telephone +81-3-234-5058

UNIX Fair '88 Association
1-1-1 Hirakawa-chu,
Chiyoda-ku, Tokyo 102
Japan

Singapore Unix Association - Sinix
20 Bideford Road #11-05
Wellington Building
Singapore 0922

Telephone +65 734 3256

DECUS U.S. Chapter
219 Boston Post Road, BP02
Marlboro, Massachusetts 01752-1850
USA

Telephone +1-617-480-3418

DECUS Europe
1-3, chemin Anneville, Box 176
CH-1213 Petit-Lancy 1
Switzerland

Telephone +41 - 22 - 709 42 64
Facsimile +41 - 22 - 792 25 03

DECUS Munich (for Germany, Austria, Hungary):
DECUS Muenchen e.V.
Freischuetzstr. 91
D-8000 Muenchen 81
Germany

Telephone +49 - 89 - 95 91 - 44 30

The European C++ User Group
c/o The Instruction Set
City House
190 City Road
London, EC1V 2QH
United Kingdon

Telephone +44 71 253 5121
Facsimile +44 71 251 2853

Email ecug@inset.co.uk

EXUG
The EXUG Secretary
145 St Neots Road
Hardwick
Cambridge, CB3 7QJ
United Kingdom

Telephone +44 954 211860
Facsimile +44 954 211244

Email exug@unipalm.co.uk

USENIX Association Office
2560 Ninth St., Suite 215
Berkeley, CA 94710
USA

Telephone +1 415 528 8649

Email office@usenix.uucp

National Expositions Co., Inc. (UNIX EXPO)
15 West 39th Street
New York, NY 10018
USA

Telephone +1-212-391-9111
Facsimile +1-212-819-0755

Sun UK User Group
Sue Crozier
Sun Microsystems, UK

Telephone +44 276 20980

Sun User Group the Netherlands (SUG-NL)
Clementine Voest
Sun Microsystems Nederland

Telephone +31 33 501234

Sun User Group, Inc.
Peter H. Salus
1330 Beacon Street
Brookline, MA 02146
USA

Telephone +1 617 232 0514
Facsimile +1 617 232 1347

Email peter@sug.org

Sun Expo/PCI

Telephone +1 512 250 9023

USING
P.O. Box 1077
Lisle,
Illinois 60532, USA

UniForum NZ Secretariat
PO Box 585
Hamilton
New Zealand

TRUUG
Esref ADALI,
Professor of Control and Computer Engineering,
TRUUG UNIX '90 Chairman,
Istanbul Technical University
Ayazaga
Istanbul
Turkey

Telephone +90 1 176 3586

EurOpen National group addresses can be found at the back of this newsletter.

Here is a list of acronyms that you might find useful:

ACE	Advanced Computing Environments
ACM	Association for Computing Machinery
AFUU	The Association Française des Utilisateurs d'UNIX
AUUG	The Australian UNIX systems Users Group
DECUS	The Digital Equipment Computer Users Society
ECUG	The European C++ User Group
EurOpen	The European Forum for Open Systems, (was EUUG)
EXUG	European X User Group
FNUG	Federation of NCR User Groups
GUUG	The German UNIX Systems User Group
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
Interex	The International Association of Hewlett-Packard Computer Users
JUS	Japan UNIX Society
MCNTI	Moscow International Center of Science and Technical Information
NCR UUG	NCR UNIX User Group, Inc.
NIST	The National Institute of Standards and Technology
NLUUG	The Netherlands UNIX Users Group
NSF	National Science Foundation
SAB	Standards Activities Board
SERC	NSF/Purdue/Forida Software Engineering Research Center
SUUG	Soviet UNIX Users' Group
Sinix	The Singapore UNIX Association
UKUUG	The United Kingdom Unix systems Users' Group
USENIX	The Professional and Technical UNIX Association
UniForum	The International Association of UNIX Systems Users

Europen Software Distribution

Frank Kuiper
Centre for Mathematics and Computer Science
Amsterdam, The Netherlands

Email europen-tapes@EU.net

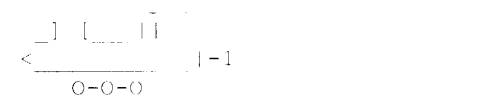

Finally! It's here. I mean of course the summer we have all been waiting for for so long. And now of course people start to complain about it being to hot :-) Oh well, nothing's perfect.

Summertime also means for me trying to get some things done I normally don't have or don't take the time for. Including this column. Last time I missed the deadline for publication because of to much other work to do. This means I now have to update you on all that has happened for the last half year, on software distribution matters that is. Here goes.

First of course (what else :-), the announcement of a new tape. The latest example of getting more "real nice stuff" to EurOpen members is called EurOpenD23 "Tromsø" conference tape. It consists of network environment software, among which you can find:

- amd - automount daemon
- athena - software from the MIT Athena project, including: hesiod, kerberos, palladium & zephyr
- routed - for making sure every packet finds its way.
- python - an interpreter language, C, shell and awk like, and very powerful
- ftp - a version without most of the (security) bugs found in proprietary systems
- telnet - reach another host on a network
- ncsatelnet - the same, but different; version 2.2tn
- osi - some gossip
- nfswatch - keep an eye on your ethernet usage
- tcpdump - see what actually comes over your ethernet wire
- slip - both version 3.x and 4.0
- tip - a little help in using modems, with SLIP support
- cisco - some help in using this kind of box
- ka9q8 - version 8.0
- modems - a little help in setting up your modems
- snmp - goodies from different places
- tunnel - a driver for routing IP packets over other networks
- uupc - wanted over and over again
- kermit - some people just can't do without
- pcbridge, pcip, pcrout - for those connection PC-LAN's to the real world and some other little gadgets.

A lot of people collected their pre-order tape during the conference. We find this to be a quick way of getting software to people. Of course we had some trouble getting the tapes to Tromsø on time, and making sure enough copies were available. I want to thank Magnar Antonsen from the University of Tromsø for logistic help and specially Tore Solvar Karlsen from the University of Oslo for actually making sure enough tapes were available in Tromsø. With the conference less than a week ahead, he managed to ftp the software from Amsterdam, getting hold of enough empty tapes and making sure that the tapes actually got copied and labelled. Thank you Tore.

Unfortunately, something went wrong handing out the 1/2" reel versions of the conference tape. People were not aware that there actually were two tapes! It didn't quite fit on one. Those who have collected a 1/2" 9-track, reel tape will receive the other half in due course. Those who picked up a 1/4" cartridge have nothing to worry about. They have everything that was supposed to be on the tape.

And all this is not enough. For the next conference, the one in Budapest, 16-20 September, we again plan to have a conference tape issued. You can book this tape one the conference booking form (which will come out shortly after you received this newsletter). On this tape will be most of what's available on the "Start Kit" tape. We think that many people in eastern Europe will now be better able to join their national and the international networks, and therefore could very well do with a little help from our end.

People interested in ISODE should note this: version 7.0 of the ISODE software (EurOpenD14) is due to come out. By the time you read this it is probably available. If not and you do want the

new version, just order EurOpenD14 "ISODE" and specify version 7.0. As soon as I have it available I will send it off.

One last thing. EurOpen is working on centralising communications with its members. For that reason, all requests for EurOpen software distributions should, as of now, be directed to Owles Hall. Possibly in the future Owles Hall will also start to ship tapes and take care of invoicing, but until that time, tapes will be send from me straight to you.

That's it. If you haven't had your vacation yet, I hope for you it will be sunny. For those who had their vacation: remember, it can't be pouring *every* summer! Good luck next year.

Below you'll find the usual list of available distributions. As always, anyone is invited to make their own tools, games, etc. available for publication on an EurOpen tape. Please contact me for more details. Don't hesitate, just put the results of many nights of serious programming and hacking in the public domain, and you might even become famous!

This is a list of all the current (July 1991) EurOpen Software Distributions. It is a short description of the available tapes. Any changes to the contents of the tapes, as well as announcements of new tapes will be placed in the EurOpen Newsletter. Those of you who have ftp access to the outside world, be aware that on mcsun.eu.net (192.16.202.1), you can find in the ~ftp/EurOpenD-index directory files with contents lists of the available EurOpen distributions. These files are compressed tar tv listings. I am working on a method so you can automagically, by e-mail, easily find out which program is on which distribution. For the moment you will have to e-mail, call, or write me to find out.

The first price listed with each distribution, is for 1/2", 9-track, reel tapes in tar 1600 bpi format, the second one is for distributions on 1/4" cartridge tapes in (Sun) QIC-24 format.

Note that you have to be an EurOpen member (or a member of a local UUG) to obtain tapes at list prices. Non-members will have to pay an extra handling fee per tape.

EurOpenD1 R6:
UNIX V7 system, specially made for small DEC PDPs (11/23, 11/34, etc.). The Kernel supports the UK terminal driver. V7 source licence minimum.
Price: £42 60 ECU - (No cartridge version available.)

EurOpenD3:
Starter Kit The tape contains many public domain mail, news, networking utilities, tools, attributes and information. It is THE tape for all who want to share in the joy of really feeling connected to the world in general, and the Unix TM world in particular! Programs you will find on this distribution are:

mail:	mh, elm2, sendmail (5.61 currently), smail, ida, mailway
network:	ka9q, snmp, uupc, uucp over x25
news:	bnews, cnews, nn, nntp, rn, vms, tmnn
gnu-tools:	gawk, grep, make, mh, tar
docmentation:	iso3166, ethernet-info, several RFC's
misc:	kermit, tn3270, vacation

In regard to the documentation a note. It is our intention to have National Groups supply information for the tape, in the local language, about how their network is set up, how to connect, what is and what is not (yet) possible, and some further guidelines. As we still have to build this database of local information, it will (hopefully frequently) happen that this Starter Kit distribution will be kept up to date and incorporate new software and documentation over time. This of course depends largely on the amount of time I can make available for this and how input the local groups give. I will do my best.
Price: £42/£60 60/86 ECU

EurOpenD4:
Software tools, sampled (in 1982) by the Software Tools Users Group. Most of the software is written in Ratfor, for which a Fortran support tool is included. This tape is available in different formats: DEC RSX, DEC VMS, UNIVAC, IBM MVS, UNIX tar, MIT line feed format, and MIT card format (80 columns).
Price: £50/£60 72/60 ECU

EurOpenD5:
Currently not available. See tape EurOpenD20 for new benchmark software.

EurOpenD6:
(USENIX 83.1) USENIX tape, containing contributions from various UNIX System Group Members. Created in 1983. This is a licence dependent distribution: V7, V32, SIII, V6 or no licence disclosure available.
Price: £78/£96 112/140 ECU

EurOpenD7:
UNIX|STAT Version 5.2. A collection of about 25 data manipulation and analysis programs written in C by Gery Perlman (1985).
Price: £24/£60 35/86 ECU

EurOpenD8:
A collection of useful software, based on the so called Copenhagen tape (EurOpen UNIX conference autumn 1985).
Price: £42/£60 60/86

EurOpenD9:
A collection of useful software, based on the so called Florence tape (EurOpen UNIX conference Spring 1986).
Price: £52/£60 75/86 ECU

EurOpenD10:
MMDFIIb. Multichannel Memo Distribution Facility (version IIb). This is a powerful, domain oriented mail system with access control and the ability to communicate over a variety of network systems including TCP/IP, JANET, UUCP, PHONENET, etc. It has been ported to a variety of UNIX's including but not limited to 4.[123] BSD, 2.9 BSD, System III/V on a variety of different hardware. You should first obtain a licence agreement by sending a message to euug-tapes@EU.net. Return the signed licence with your order.
Price: £33/£60 48/86 ECU

EurOpenD11:
This is the 'Boat' tape; the Helsinki EurOpen 1987 spring conference. It contains about 25 Megabytes of programs, games, etc. Including: jove, less, nag, news, rn, uEmacs, uuencode and larn.
Price: £42/£60 60/86

EurOpenD12:
This is the Dublin EurOpen 1987 autumn conference tape. It contains about 26 Megabytes of programs, games, etc. Including: copytape, crc_plot, fastgrep, jove, kermit, notes, uupc, nethack, cron, sendmail, mh, Recipes, brl-gw, isode, pcip, pctelnet.
Price: £42/£60 60/86 ECU

EurOpenD13:
The conference tape for the London EurOpen 1988 spring conference. It contains, amongst others, the following items: cake, chat, config, copytape, graphedit, kermit, little- st, mcc, mstools, news, pd-diff, pdtar, perl, postscript, psfig, pshalf, shar, rpc, moria4.85, omega, arc, backup, smail, sush, watcher, and much, much more.
Price: £42/£60 60/86 ECU

EurOpenD14:
This is version 6.0 of this non-proprietary implementation of some of the OSI parallel protocols suites as defined by the International Organisation for Standardisation (ISO), the International Telegraph and Telephone Consultative Committee (CCITT), and the European Computer Manufacturer's Association (ECMA). [Version 7.0 coming soon on a computer near you.]

This release is coded entirely in C, and is known to run under the following operating system without kernel modifications:

BSD 4.2 and 4.3, Ultrix, AT&T UNIX SVR2 and SVR3, AIX, HP-UX, ROS, Pyramid OsX

Since a Berkeley UNIX system is the primary development platform for ISODE, the documentation and source are somewhat slanted towards that environment. The tape contains some 12Mb of both tools and documentation in machine readable form. The EurOpen will send you a tape only.
Price: £42/£60 60/86 ECU

EurOpenD15:
The complete X11 Windowing system, as distributed by MIT, release 4: X11R4. Do to the vast growth in user contributed software, this distributions now totals 60 Mb in compressed form. This results in two 1/2", 9-track tapes, and one 600 ft, Qic-24, 1/4" cartridge. This includes the core system, as well as much, very much user contributed software. No patches, however.
Price: £82/£63 118/90 ECU

EurOpenD16:
This is the Brussels EurOpen 1989 spring conference tape, and consist entirely of software from the GNU project from the Free Software Foundation. Last update: autumn 1990. On this tape you will find: ispell, g++, awk, gcc, gdb, Cscheme, emacs, lisp-manual, libg++, binutils, bison, ghostscript, gas- dist, gawk, gnews, gnuchess, make, oops, pace, ps-emacs, scheme, sed1, tar and torture.
Price: £42/£60 60/86 ECU

EurOpenD17:
This tape contains the software for ET++. From the abstract of the "Autumn 1988 EurOpen Conference Proceedings": "ET++ is an object-oriented application framework implemented in C++ for a UNIX environment and conventional window system. The architecture of ET++ is based on MacAPP and integrates a rich collection of user interface building blocks as well as basic data structures to form a homogeneous and extensible system." It totals some 18Mb of software that the people of the Institut fuer Informatik of the University of Zurich were so kind to let us, mere mortal souls, play with. Have fun.
Price: £42/£60 60/86 ECU

EurOpenD18:
This is the "Vienna EurOpen 1989 autumn conference tape", and consists entirely of games! There is a SUN specific set, a set for the X Windowing System environment, and a general useable set. All the games supplied are working, and have been tested at CWI by our "Games Keeper <play@cwi.nl>". For many games he added additional features, not found in the originals. Some of the games included are: for SUN: Asteroids, Mahjongg, Othello, Qix, Sdi, Tetris. For the X environment: Xtrek, Xgo, Xwanderer, Xrobots. General games: Nethack, Adventure, Atc, Empire, Reversi, Yahtzee, Trek73, Backgammon, Corewars, MazewarsV, Vtrek, and lots, lots more. If this doesn't bring some fun back into using computers, I don't know what else can:-)
Price: £42/£60 60/86 ECU

EurOpenD19:
This is the "Munich EurOpen 1990 spring conference tape", and consist entirely of graphics material. Conversion programs, display tools, toolkits to build you own display program, and off course images, lots of pictures to play around with.
Price: £42/£60 60/86 ECU

EurOpenD20:
This tape contains benchmarking software and is named "AFUU/SSBA 1.2, benchmarks". The French group have done a good job creating a tape with all the necessary tools, so you can finally bring your machine down to it's knees, and see what it is really worth.
Price: £24/£60 35/86 ECU

EurOpenD21:
This is the "Nice EurOpen 1990 autumn conference tape", and consist of a number of different kinds of software, like: dtree, abc, new versions of various mail and news utilities, and PP5.0.
Price: £42/£60 60/86 ECU

EurOpenD22:
This is the "Postman Pat PP5.0" distribution. PP is a Message Transfer Agent, intended for high volume message switching, protocol conversion, and format conversion. It is targeted for use in an operational environment, but may also be useful for investigating Message related applications. Good management features are a major aspect of this system. PP supports the 1984 and 1988 versions of the CCITT X.400 / ISO 10021 services and protocols. Many existing RFC 822 based protocols are supported, along with RFC 1148 conversion to X.400. PP is an appropriate replacement for MMDF or Sendmail.
Price: £42/£60 60/86 ECU

EurOpen D23:
The "Tromsø EurOpen 1991 spring conference tape". Also as the "Network" tape. On this distribution you will find a large number of programmes having to do with networked environments, like" amd, athena, routed, python, ftp, telnet, ncsatelnet, osi, pcbridge, pcip, pcrout and some other little gadgets.
Price: £42/£60 60/86 ECU

USENIX Association News for EurOpen Members

USENIX Association
2560 Ninth Street, Suite 215
Berkeley, CA 94710, USA

Email office@usenix.org

Summer 1991 Conference and Exhibition

The Summer USENIX Conference and Technical Exhibition was held 10 - 14 June 1991 at the Opryland Hotel in Nashville, Tennessee.

The conference's focus was "Multimedia: For Now and the Future". Voice, video, animated graphics, touch and music add to the versatility and capability of the modern computer in the expression of ideas. But each new interface multiplies complexity on the system level and challenges our ability to obtain rapid and easy access to resources. The papers, panels and demos presented during the week highlighted the challenges of incorporating these into a single environment. The technical sessions covered the following topics:

- File Systems
- Hypermedia
- Multimedia Data Rates and Synchronisation
- Multimedia Publishing
- String and Things User Interface
- System Implications of Compression
- Audio Conferencing
- The multimedia demonstrations included:
- Spatio-Temporal Editing Using
- Multi-Layered Image Synthesis (HDTV)
- DIDDLY: Digital's Integrated Distributed Database LaboratorY
- The Architecture of the IRCAM Musical Workstation
- Software Technology at NeXT
- The MIT Media Laboratory
- Integrating Real-Time Video with Sun Workstations
- A Workstation- based Multimedia Environment for Broadcast Television

The Invited Talks included:

- Overview of Hypertext
- UNIX and MIDI for the Masses
- From Blazon to Postscript
- The Korn Shell Past, Present and Future
- Overview of Motif
- Scaling Up - Automating System Administration
- Networks: Friend or Foe?
- C Programming Style

The two panels covered "Window Pains" and "Software and Intellectual Property - Who Owns Your Work?"

Keynote Address

Paul Lansky, noted composer and musicologist, and chair of Princeton University's Music Department gave the keynote address. Peter Salus wrote the following review:

"Lansky's own music attempts to integrate the everyday and the commonplace into the artistic repertoire, and examples of compositions involving traffic noise, chatter, shopping-mall ambience, and household sounds provided the counterpoint to Lansky's themes.

"One of his main themes was the variance among the tempi of technology (fast), science (allegro ma non troppo), and the adagio of the arts. "Art moves at a glacial pace," he stated.

"Moreover, art is constantly responding to the environment. Thus, the invention of the synthesizer 25 years ago influenced both popular and formal composers. In fact, Lansky sees rap music as growing directly from the mechanical ambience of early synthesized music.

"Moving from Mozart, through Nietzsche's remarks on Wagner, to Babbitt, Boulez, and Cage, Lansky discussed the fact that "the multiplication of media carries the danger of desensitisation. There is little joy left for the commonplace experience."

"Furthermore, movie music and the ubiquitous Walkman make it increasingly difficult "to distinguish what we take in through the media and directly." Lansky said that he saw software design "as a part of artistic creation and self expression. The distinction between the tools and the things we do with tools is vanishing.""

Best Paper Award

The Program Committee named Matthew Blaze and Rafael Alonso of Princeton University as the authors of the best student and overall paper. Their paper was entitled "Long-Term Caching Strategies for Very Large Distributed File Systems" and is available in the conference proceedings.

Forthcoming Conferences and Symposia

- Fifth Large Installation Systems Administration Conference, *30 September - 3 October 1991, San Diego, California*
- Second Mach Symposium, *20 - 22 November 1991, Monterey, California*

  20 November - Tutorial

  Richard Draves, Writing a Multi-Threaded Mach 3.0 Server
  David Black, Writing an External Memory Manager

Richard Draves will lead a tutorial analysing the process of writing a multi-threaded server, with particular attention paid to the complexities of using Mach IPC. David Black will demonstrate how to create an external memory manager; discussion will center on the intricacies of developing an efficient (and well-behaved!) external manager.

These tutorials will explore concepts and rationale as well as real examples. They are oriented towards programmers who already have some familiarity with using Mach IPC and VM. Each tutorial is a half-day, so conference attendees may take part in both. The tutorials will be priced separately from the conference registration fee.

USENIX Winter 1992 Conference

20 - 24 January 1992, San Francisco Hilton, California

The Association's two annual technical conferences are well recognised as the leading forums for the communication of new research and investigation of important developments in UNIX, UNIX-related and advanced computing systems as well as for the tutorial program offerings. Software professionals and technical managers will tackle questions of immediate importance to advance computing systems development and management.

Schedule Of Events

Tutorial Program

Monday and Tuesday, 20-21 January

Introductory as well as advanced, intensive yet practical, example-filled tutorials by leading experts will focus on topics essential to successful technical management of UNIX, UNIX-like advanced computing systems, X windows, Mach, the C and C++ programming languages and related areas of interest.

The tutorial program at San Francisco will include topics such as:

- Programming in Perl
- Introduction to the TCP/IP Suite
- C++
- System VR4 Internals
- Network and System Security
- UNIX Programming Tools
- Advanced System Administration
- Network Programming
- Programming with OSF/Motif
- Programming X Windows
- OSF/1 Internals
- Introduction to X Toolkit Intrinsics

Refereed Papers and Invited Talks

Wednesday through Friday, 22-24 January

Some believe that UNIX standardisation efforts have killed innovation. And yet, we need innovation, and opportunity for it abounds. Large write-once disks make the current filesystem untenable. Even the 2 gigabyte file limit built in all through the system breaks. Gigabit networking clogs an I/O model designed to push hundreds of kilobytes per second, not hundreds of Megabytes. System administration for thousands of machines? Programming tools for distributed workgroups? Object-oriented and visual programming? Microkernels with client/server architectures? RAID disk arrays? Transcontinental file servers? What's a programmer to do?

Papers formally reviewed and accepted by the Winter 1992 Conference Program Committee will attempt to answer these questions.

Invited Talks

A full series of invited talks will teach the tricks of using standard UNIX tools and tackle the difficulties of system administration and integration. Suggestions for new and interesting topics as well as submissions proposing a particular session are welcome. Proposals should include a brief outline. Be sure to emphasise why your topic is of general interest to our community and what your main focus would be. Send these submissions to:

Andrew Hume
Bell Labs 2C-515
Murray Hill
NJ 07974
USA

Email ITusenix@usenix.org

or

Email uunet!usenix!ITusenix

or

Facsimile +1 908 582-5857

Birds-of-a-Feather Sessions

Tuesday-Thursday evenings, 21-23 January

Conference Headquarters: San Francisco Hilton

The Hilton is located in the heart of downtown San Francisco, two blocks south of Union Square. Discount air fares and special, affordable room rates at the Hilton and nearby hotels will be available to USENIX Conference attendees.

UniForum, scheduled concurrently with the USENIX Winter Conference, takes place 22-24 January at San Francisco's Moscone Convention Center.

Materials containing all details of the technical and tutorial program, conference registration, hotel and airline reservation information will be mailed in October 1991. If you wish to receive the pre-registration materials, please contact the USENIX Conference Office (see below).

- Third Symposium on Experiences with Distributed and Multiprocessor Systems - III, *26-27 March 1992, Newport Beach, California*

  (See Call for Papers in this issue.)

Information on Conferences and Workshops

Contact the USENIX Conference Office at:

22672 Lambert Street, Suite 613
El Toro
CA 92630
USA

Email judy@usenix.org or {uunet,ucbvax}!usenix!judy

Telephone +1 714 588 8649
Facsimile +1 714 588 9706

The USENIX Association

USENIX, the UNIX and Advanced Computing Systems professional and technical organisation, is a not-for-profit association dedicated to

- fostering innovation and communicating research and technological developments,
- sharing ideas and experience, relevant to UNIX, UNIX-related and advanced computing systems,
- providing a forum for the exercise of critical thought and airing of technical issues.

Founded in 1975, the Association sponsors two annual technical conferences, and frequent symposia and workshops addressing special interest topics, such as C++, Mach, systems administration, and distributed/multiprocessor systems. USENIX publishes proceedings of its meetings, a bi-monthly newsletter ;login:, a refereed technical quarterly, Computing Systems, and is expanding its publishing role with a book series on advanced computing systems. The Association also actively participates in and reports on the activities of various ANSI, IEEE and ISO standards efforts.

If you would like information on membership or on ordering USENIX publications (proceedings, manuals, Computing Systems, or the Association's newsletter, ;login:), please contact the USENIX Executive Office at:

2560 Ninth Street
Suite 215
Berkeley
CA 94710
USA

Email office@usenix.org

Telephone +1 415 528 8649
Facsimile +1 415 548 5738

Call for Participation

Symposium on Experiences with

Distributed and Multiprocessor Systems (SEDMS III)

26-27 March 1992, Newport Beach, California, USA


IEEE COMPUTER SOCIETY

In association with:

The Software Engineering Research Center (SERC)

In cooperation with:

ACM SIGARCH, SIGCOMM, SIGOPS and SIGSOFT (Pending), IEEE-CS Technical Committees on Distributed Processing, Operating Systems, Software Engineering, and Design Automation.

The goal of this symposium is to bring together individuals who have built, are building, or will soon build distributed and multiprocessor systems. SEDMS III will provide a forum for individuals to exchange information on their experiences, both good and bad, including experiences with coding aids, languages, debugging and testing technology, reuse of existing software, and performance analysis. The presentations should emphasise the lessons learned from use of such systems and tools.

Extra-long breaks between sessions and work-in-progress presentations will be provided to facilitate a workshop-like atmosphere during parts of the symposium. We will also have discussion panels on submitted themes.

Six copies of each submission or panel proposal should be sent to the programme committee chair (spaf@cs.purdue.edu) to arrive no later than 1 November 1991. Submissions of full papers are invited on any topics related to the theme of the symposium. The committee will give preferential consideration to submissions describing experiences with actual systems - papers describing purely theoretical work will not be accepted. Panel proposals should include a description of the relevance to the goals of the SEDMS, and the qualifications of the participants suggested.

Important Dates

- Submissions due: 1 November 1991
- Notifications mailed: 20 December 1991
- Camera ready copy due: 24 January 1992

For further information, contact:

General Chair

George Leach
AT&T Paradyne
MS LG-129
PO Box 2826
Largo,
FL 34649-2826
USA

+1 813 530 2376
reggie@pdn.paradyne.com

Programme Chair

Gene Spafford
Software Eng. Research Center
Dept. of Computer Sciences
Purdue University
W. Lafayette,
IN 47907-1398,
USA

+1 317 494 7825
spaf@cs.purdue.edu

Programme Committee

- John Barr, Motorola
- David Cohn, Univ. of Notre Dame
- Yousef Khalidi, Sun
- Andrzej Goscinski, Univ. of New S. Wales
- Mike O'Dell, Bellcore
- Satish Tripathi, Univ. of Maryland
- Howard Katseff, AT&T
- Thomas Wilkes, Univ. of Lowell
- Susan LoVerso, Encore
- Dan Duchamp, Columbia Univ.
- John Nicol, GTE Labs
- Bharat Bhargava, Purdue Univ.
- Marc Pucci, Bellcore
- Elie Milgrom, Louvain Univ., Belgium
- William Bain, Intel
- Karsten Schwan, Georgia Tech

X11 Release 5 Now Available

Gary Henderson
IXI Limited
Cambridge
United Kingdom

Gary is an experienced software consultant with over 4 years knowledge of working with the X Window System. Gary has hands-on experience of porting X and Motif to a variety of platforms and is also a seasoned toolkit programmer. His role at IXI involves lecturing in X and Motif programming as well as heading up the Motif Development Team with responsibility for IXI's Motif Development Kit for Sun product.

On 9 August, the MIT X Consortium issued to its members a new release of the X Window System, X11 Release 5. As these releases are only made every 18 months or so and contain significant new technologies, this is an important event for anyone working with X.

X11 Release 5 includes scalable fonts technology, support of internationalisation and the PHIGS international 3-D graphics standard. The latter is particularly important as the inclusion of the PHIGS Extension to X or PEX provides a standardised way of supporting 3-dimensional graphics. It has been possible to generate 3-D graphics on X as vendors have added their own extensions to the X server but this has been in a stand-alone form. This release marks the first time that there is networked support for 3-D graphics.

PHIGS or Programmer's Hierarchical Interactive Graphics Systems has been criticised in the past for its complexity and the demands on processing power to draw the 3-D graphics. However, as an established international standard, it was natural for the X Consortium to select PHIGS for the basis of the 3-D version. There are proprietary graphics libraries such as Silicon Graphics's GL which has achieve wide recognition to the extent that Network Computing Devices (NCD), the leading X terminal supplier, may be supporting GL alongside PEX.

The new PEX files take 10.5Mb memory and take the form of a server extension and client side library. The client uses two processes, one to handle X events and the other the user application.

The new font server technology widens the font repertoire for X users. A single font server sitting on a network node can provide a large selection of fonts in a variety of styles in all point sizes. Each point size of the different fonts are not stored in a separate file as in earlier X releases but as a single copy which can be scaled up or down as required. This considerably reduces font storage space and increases the range of point sizes available.

There are two sample font scaling implementations; one using bitmap fonts which allows existing X fonts to be scaled and the other using font outlines contributed by Bitstream Inc of Cambridge, Mass. The latter is said to produce better quality images but needs special outline font files. Bitstream supply a Charter outline font in normal, bold, italic and bold-italic.

Colour has now been standardised so that you will no longer get any colour variation when you move from one terminal to another. Originally, colour was represented as RGB values but in X11 Release 5, this system has been improved by a device-independent colour system at the Xlib level. Developed by a European standards body, the Commission Internationale De L'Eclairage, the CIE technology provides a standardised representation of colours from one display to another. Several methods of describing colours are implemented including CIE u'v'Y, CIE XYZ and Tektronix's TeKHVC systems.

A good deal of work has been invested into providing adequate support for Internationalisation. This is obviously key to the success of X worldwide and will no doubt lead to the Open Software Foundation and Unix International providing further internationalisation support for their GUI products - Motif and Open Look. In fact, the next release of Motif, v1.2, will be based on X11 Release 5.

The work takes into account the ANSI and ISO standards for internationalisation. The previous releases of X11 could only effectively deal with 8-bit character sets such as error messages

and the resource database and had no way of inputing the larger character sets. X11 Release 5 now supports characters used in Japan, China and other parts of the Far East.

Apart from these major new technologies, Xlib has been optimised - routines that the toolkit calls frequently have been improved some by a factor of 300. The server now supports true Save Unders, rather than saving the contents of the entire window, only the obscured section is now remembered. There is now support for hardware cursors for systems that have this feature.

The initial release of X11 Release 5 only includes the core technology (about 74Mb in size). However, the full release on 10 October will also contain user contributed software but be warned it will be rather large!

If you would like to receive your copy of X11 Release 5, then tapes are available from the X Consortium in Cambridge, Mass (price had not been decided as we went to press). Alternatively, IXI are supplying the MIT tape here in Europe at 295 pounds, for further information call IXI on +44 223 462131.

EWG News

Jean-Michel Cornu Consultant
69 rue de Seine 91130 Ris Orangis
France
Telephone +33(1)69.43.48.47
Facsimile +33(1)69.06.83.16

Email Jean-Michel.Cornu@enst.fr
Email Jean-Michel.Cornu@afuu.fr

Jean-Michel Cornu is a French consultant involved each time there is a group or a committee on the planet. That means that he participates to several users organisations and standard bodies, but has also created research Working Groups on fundamental physics and a worldwide amateur TV network. In EurOpen he proposed to coordinate the support to the creation of Working Groups.

The presentations of the Working Groups are continuing. I expect at least 13 countries to have arranged a WG presentation during their conferences (and started WGs?) Please, contact me if you have at least one person interested in one topic (It seems to be not too difficult, but is sufficient to initiate some thing). The "secret" of the creation of a Working group fits in 5 questions:

- Which topic are you interested in?
- Who will follow the issue? (normally, at least 4 or five participants of the first meetings are known by the instigator even if a national mailing may help to find some other people interested).
- What is the scope of the group?
- What kind of output do you expect? (producing a guideline, articles in the newsletter or more ambitious output, ensure that people will have the motivation to continue after the first meeting)

and ...

- What is the date of the first meeting?

If you can answer to the two first ones, you are on the right way, and the momentum that we have produced may be a good tool for you.

Budapest week will see many WGs events. Here are some of them:

- EWG meeting on Tuesday 17 September 1991; 09h00-12h30

  For those who have created (or will create) WGs, it is a good opportunity to meet each other and get contact with people interested in the same topic than you. Luigi Bertuzzi, chairman of the EWOS EG-CAE, will probably come to see how we can collaborate to the European standardisation work

- EurOpen/Benchmark WG on Tuesday 17 September 1991; 14h00-18h00

  This will be the first meeting of the EurOpen Benchmark working group. See the article bellow for more information.

- BOF on EWG presentation on Thursday 19 September 1991; 17h00-19h00

  For those who want to know what is EWG, and get a first contact with people involved in it.

Some other events will occurs such as a presentation to the governing board and a meeting of the coordination team.

We have prepared a database of all the WGs at national and EurOpen level, with information such as the chairman, contact, scope, output... I hope it will be and interesting tool for you to get contact and boost your own working group. I expect to publish it in the next issue of "EWG news".

We also get many contacts with the European and international community. I have already told your about EWOS, but we got also contacts with the CEC (we have to meet after Budapest), Allen Hankinson, chairman of the steering committee of UniForum WGs We also had many discussion with Andrew Cash and Paul Tanner of X/Open to know how we should be involved in X-TRA, the user enquiry of X/Open. Henk Hesselink is our representative to IEEE, he can bring you many interesting working documents about all the topics raised in IEEE TCOS. You can contact him at henk@ace.nl.

The following is a short presentation of what is going on in the benchmark area within EurOpen.

Benchmark Column

Linpack Results and WXorking Groups

Nhuan DODUC
Framentec - Cognitech
Paris - La defense

AIUUG, I mean our Algerian colleagues, created a benchmark WG on the 8 July 1991 with the help of the French Benchmark group. On that day, a tutorial on Classical and Standard benchmarks was presented to about 35 very interested people in the conference room of the PTT ministry. The chair of this group is

Mustapha Ben Bouzid
SONELGAZ, Unite de recherche
2 bd Salah Bouhakouir
Algeria

Telephone +213 2 612304

An announcement paper is available, which presents the scope of this group: Benchmark analysis and existing results of the SSBA 1.21 study.

GUUG met the same week, on 11 July 1991, in Dortmund, to create its own benchmark group. This was a formal creation, because this group has already existed for a long time, and have done a lot of work, as shown by all the SSBA results published in iX. The main members of this group are: Torsten Beyer, Michael Kusche, Michael Kiener and Ralph Hulsenbusch. The temporary contact is:

Ralph Huelsenbusch
iX - Multiuser Multitasking Magazin
Heinz Heise Verlag Gmbh & Co KG
Helstorferstr, 7
D-3000 Hannover 61
Germany

Telephone +49 511 547 47 24

This kernel group has made contact with other people interested in the same topic (e.g. in UK to coordinate and share work). They will probably concentrate on Graphical and network benchmarks.

See you in Budapest on Tuesday 17 September AM, where the three groups (Algerian, German and French) will meet with other individuals from other National Users groups. They will present their goals and working plan.

(This paper is extracted and modified from the regular benchmark column in Tribunix number 38, 9/91).

EUnet Report

Glenn Kowack
EUnet Chief Executive

Email glenn@mcsun.eu.net

Glenn has spent most of the last 15 years in the computer industry, primarily managing software development for operating systems and networks. For the last few years he has been a consultant for a variety of companies including UNIX International, Harris Semiconductors and various small startup companies.

He has degrees in mathematics and psychology from the university of Illinois. During the 80's he helped to found a community radio station.

The major changes in EUnet during the first six months of my tenure have been focused on "organisation building". That is, putting in place many of the procedures, agreements, and relationships required for EUnet to evolve further. Here's an overview of progress since January of this year.

Administrative Changes

Office at NIKHEF in Amsterdam

EUnet has rented office office space at NIKHEF, the National Institute for Nuclear and High-Energy Physics in Amsterdam. Our office is conveniently part of the same campus as our technical operations center at CWI. In fact, the CWI and NIKHEF buildings are connected by a common tunnel. EUnet now has a place for meetings, day-to-day management, and for miscellaneous operational support. Our postal address is:

EUnet
c/o NIKHEF
Postbus 41882
1009 DB Amsterdam
The Netherlands

Telephone +31 20 5925124
Facimile +31 20 5925155

Geography buffs will be impressed to note that although our office is on the second floor, it is only about 1 meter above sea level.

Budget Reorganised

The EUnet budget has been completely reorganised: the number of line items have been roughly tripled, and the overall budget now runs to approximately 7 pages. Strictly speaking, volume is not necessarily valuable. However, in this case, the additional information is required to completely describe what's actually going on and will permit us to exercise greater control over how we spend money. There is greater detail regarding income from News and IP subscriptions and Backbone contributions for the trans-Atlantic line, and regarding expenditures such as line leases, meetings, capital equipment, and management overhead.

The newly reorganized budget was reviewed and approved at the EurOpen meeting in Tromsø, Norway in May, and at the Breukelen backbone meeting in July.

The format of this year's budget is also somewhat "overbuilt". That is, it is generalized and extensible so that it will serve as a template for next year's budget. I expect that future budgeting will be quicker and easier. The 1992 budget, which is being developed as we go to press, will be presented for approval at the EUnet and EurOpen meetings in Budapest in September.

Backbone Customer Survey

We have completed an internal survey of expected customer growth at each of our backbone sites. This survey lets us look at how many customers we have, for each of our different types of services, at each of our backbones, for each quarter in 1991 and 1992. This is an important tool in helping us decide when and where to add more resources (for instance, 'phone lines).

EUnet has seen substantial growth already this year, from 1660 sites on the first of January, to over 2030 sites on 1 July. This equals 22% growth for the first half of the year.

Pricing Reorganisation

At this summer's Backbone meeting in Breukelen, The Netherlands, we decided to analyze and recommend changes to our pricing structure, for further discussion at the September Backbone meeting in Budapest. The reorganised budget and backbone customer survey are key data in the creation of these recommendations.

EUnet/CWI Technical Management Agreement

As many of you will recall, EUnet's European backbone in Amsterdam has long been operated at the Center for Mathematics and Informatics (CWI) in Amsterdam. In the past, both technical and administrative work for EUnet were done at CWI on the basis of informal agreements. Our new agreement, which runs from 1 July 1991 through 31 December 1992, provides for some significant changes: All administrative work, including billing and accounting, will be done by the EurOpen Secretariat, Owles Hall. This has two significant consequences. The first is that CWI no longer bears any of the financial risk of running the net. This was mutually agreeable since EUnet did not think it fair for CWI to be at risk for decisions over which they had limited influence. The second is that EUnet will now retain 100% of any surplus left after network operations are paid for. These additional funds will permit EUnet to grow more quickly.

Archive Consultancy Position Opened

EUnet is interested in developing and offering advanced services. We have decided to hire an archiving specialist to analyze future archiving services, improve the existing operation of the archive, and to propose new archive services and operations.

This position will be a 6-month consultancy, to culminate in the proposal mentioned above.

Changes Among the National Nets

Backbone Change in Switzerland

EUnet's backbone in Switzerland was formerly operated by the SWITCH organization in Zuerich. In recent months, SWITCH management has decided that they wish to operate outside of the typical backbone operational requirements. At the same time, the Swiss UNIX User Group, the CHUUG, also decided that they wanted changes in the way SWITCH was operating the backbone. When the two groups could not come to an agreement, the CHUUG and EUnet decided to conclude the relationship with SWITCH and to start a new backbone under the wing of the CHUUG.

The new backbone, in full operation since June, is now growing quickly and providing network services across Switzerland.

New Backbones

Tunisia's membership in EurOpen was officially ratified at the Tromsø meeting. Since that time a new backbone was formed under the helm of Mondher Makni at IRSIT - the Regional Institute for Informatics and Telecommunications in Tunisia.

Poland

EurOpen is in mid-process of certifying a new user group in Poland. In the meantime, we have agreed to informally allow a group at Poltra, a company in Warsaw, to temporarily act as a "cooperating site" with EUnet. This means that EUnet will accept dialup (UUCP) mail from Poltra while we work to certify an official backbone. This will permit us to start connecting to sites in Poland as quickly as possible.

Bulgaria

In recent weeks, we have been getting a lot of mail from Daniel Kalchev of Digital Systems, a small company in Varna, Bulgaria. With Daniel's help, Norman Hull (EurOpen's Executive Board member in charge of new user group development) arranged a meeting in Bulgaria in July to begin to establish a User Group there. As in Poland, we have asked Daniel's group to temporarily act as a cooperating site.

Luxembourg

The last hold-out in the EC has finally come forward. A group of researchers in Luxembourg are coming together to form a user group. Since Email connections to Luxembourg are so good already, we do not need to set up a cooperating site. Instead, we are going to quickly bring the Luxembourg group up to user group status, and select a new backbone there in the near future.

The only remaining countries in Europe without an operating backbone, or one on the way, are Romania and Albania.

NORDUnet Agreement

EUnet has joined a consortium with NORDUnet, which has acquired a leased line from Amsterdam to KTH (the Royal Technical Institute) in Stockholm. Our participation in the consortium also permits our Nordic backbones to share infrastructure with other networking organizations in the Nordic countries. EUnet collects funds from each of the participating backbones (Iceland, Norway, Sweden, Denmark, and Finland) and pays for our share of the overall resources.

CERN-Amsterdam Link Consortium

At the end of last year, a consortium was formed to acquire and operate a 256 Kilobaud line between Amsterdam and CERN in Switzerland. The participants include EUnet, CERN, IBM, SURFnet (of The Netherlands), NORDUnet, and NIKHEF. After several months of technical problems at the PTT level, the new line was finally accepted and put into operation in May of this year.

EUnet is looking for additional opportunities to share resources with other networking organizations. These can result in very significant cost savings since the price-per-bit of a telephone line decreases quickly as the bandwidth rises.

Technical Issues

Telebit Modem Donations

Telebit Corporation has generously donated five "Trailblazer Plus" modems to EUnet. These modems will be distributed to new backbones and sites in Eastern Europe. In exchange for the modems, EUnet sites will provide Telebit with feedback on the performance of their modems in the very difficult telecommunications environments of these countries. This is the first step in a continuing donation/feedback plan being developed by EUnet and Telebit. Everyone at EUnet is very grateful that the people at Telebit have been so generous.

Trans-Atlantic Link Upgraded from 56k to 128k

EUnet has operated a 56 kilobyte digital link to UUnet in Falls Church, Virginia, USA for several years. That link provides us with connectivity to UUnet, NSFnet (and other nets in the United States), and is our source for Network News groups which do not originate in Europe. Continuing growth in EUnet sites and traffic now require that we increase our bandwidth across the Atlantic. The new link will replace the old link and transmit at 128 kilobytes. This link will prepare us for significant growth over the next year.

Statistics and Operational Standards Working Group

Under Joy Marino, of the Italian Backbone, we have established a Statistics and Operational Standards Working Group. The mission of this group is to develop and propose standards for collecting statistics about EUnet performance and behaviour. This data will be very important in deciding what sizes of resources (particularly lines and routers) will be needed in the future. The next step will be to define a "ladder" of operational standards for the backbones. This ladder of standards will be used by each backbone to point the way to future operational improvements. It is hoped that each backbone will steadily advance up the ladder. We hope that this will be in important vehicle for improving EUnet services and reliability.

So, there we are. Many modest changes to form the base for additional growth, development, and change in 1991 and 1992.

Puzzle Corner

Mick Farmer
Department of Computer Science
Birkbeck College
London, UK

Email mick@cs.bbk.ac.uk

Mick is the Secretary of the UKUUG. His primary interests are Ornithology (restricted to the Western Paleartic at the moment because of cost) and Oenophilism (especially pre-1962 Bordeaux, pre-1980 Burgundy, and 1945 Port). His secondary interests include Software Consultancy (to pay for the above primary interests) and Distance Learning Methods (especially interactive video and hypertext). When not pursuing these and other interests he can be found at Birkbeck College (London) where he teaches in the Department of Computer Science.

He lives in Lewisham (South East London) with his wife Sue and a TV called Sonya. His neighbours have two children and a dog.

Hello peeps,

Solution to Puzzle Number 19

The simplest solutions are $\frac{77}{0.77}$ and $\frac{7}{0.7} \times \frac{7}{0.7}$. Obviously any other digit can be substituted for 7. My thanks to Michel Fingerhut for providing the simplest solution (least number of additional symbols).

Solution to Puzzle Number 20

The simplest solutions are $\frac{4! + 4.4}{0.4}$, $\lceil 4 \times 4 \times 4.4 \rceil$ and

$$\frac{4! \sqrt{4 - \sqrt{0.4}}}{\sqrt{0.4}}$$

Puzzle Number 21

Here's a nice problem involving probability.

Alfred, Brian, and Peter agree to fight a pistol duel under the following conditions. After drawing lots to determine who fires first, second, and third, they take their places at the corners of an equilateral triangle. It is agreed that they will fire single shots in turn and continue in the same cyclic order until two of them are dead. At each turn the man who is firing may aim wherever he pleases. All three duelists know that Alfred always hits his target, Brian is 80% accurate and Peter is 50% accurate.

Assuming that all three adopt the best strategy, and that no one is killed by a wild shot, who has the best chance of surviving? Harder, what are the exact survival probabilities of the three men?

Puzzle Number 22

This puzzle can involve Diophantine equations, but there is another solution involving the concept of *negative* coconuts.

Five men and a monkey were shipwrecked on a desert island, and they spent the first day gathering coconuts for food. Piling them all up they went to sleep for the night.

When they were all asleep one man woke up, and he thought there might be a row about dividing the coconuts in the morning, so he decided to take his share. He divided the coconuts into five piles. He had one coconut left over, which he gave to the monkey, and he hid his pile and put the rest back together.

The next man woke up and did the same thing. He had one left over, which he gave to the monkey. All five men did the same thing one after another; each taking a fifth of the coconuts in the pile, with one left over for the monkey.

In the morning they divided what coconuts were left, and they came out in five equal shares. How many coconuts were there in the beginning?

Puzzle Number 23

Our third puzzle in this issue comes from the realm of recreational logic.

3. In 1918, on the day that the armistice of World War I was signed, three married couples celebrated by having dinner together.
4. Each husband is the brother of one of the wives, and each wife

is the sister of one of the husbands, i.e. there are three brother-sister pairs in the group.

5. Helen is exactly 26 weeks older than her husband, who was born in August.
6. Mr White's sister is married to Helen's brother's brother-in-law. She (Mr White's sister) married him on her birthday, which is in January.
7. Marguerite White is not as tall as William Black.
8. Arthur's sister is prettier than Beatrice.
9. John is fifty years old.

What is Mrs Brown's first name?

Question

What sort of puzzles do you like to see in this column? Let me know by e-mail preferably (mick@cs.bbk.cs.uk).

Loads-a-puzzles,

Mick

UKUUG Seeks Editor!

The UKUUG seeks an editor for **news@uk**, our bright, wonderful, and witty newsletter that covers all that's interesting, and fit to print, on the UK scene. This bright, wonderful, and witty editor will collect material, produce camera-ready artwork, and coordinate printing and distribution.

It will be a Unigram-style publication appearing, initially, six times per year.

Candidates should:

- Live in the United Kingdom.
- Have reasonable net connectivity.
- Have available equipment and experience to allow them to typeset a newsletter.

The UKUUG will pay all costs and a fee (to be agreed) upon successful production of each newsletter. Help will be available in the form of advice, contact with a printer, and generally setting up.

For more information please contact:

UKUUG
Owles Hall
Buntingford
Herts SG9 9PL
England

Telephone +44 763 73475
Facsimile+44 763 73255

Email ukuug@ukc.ac.uk

USLE Column

Dr. Alan Brown
Principal Consultant, specialising in TUXEDO
USLE
London
United Kingdom

There has been much discussion about USL's Transaction Processing System, TUXEDO. Alan Brown joins us in this edition to talk about the product.

Alan Brown is Principal consultant at UNIX System Laboratories Europe based in London. He has responsibility for TUXEDO in Europe - this includes the provision of pre- and post-sales technical support, training and consultancy.

For further information on this column, please contact Gill Mogg on gill@eul.uucp. Gill is Marketing Manager at USLE.

The TUXEDO Transaction Processing System Release 4.1

Introduction

The trend today is towards using networks of high powered mini/ workstation style machines as part of the Open Systems movement. To explain today's definition of distributed, open online transaction processing (OLTP) computing, I will begin with a short look at the major evolutionary steps that preceded it. (See Figure 1)

Starting with a batch transaction processing environment, demands for data integrity and availability became more and more pressing, leading to the need for online transaction processing. Businesses in which data that once was sufficient on a weekly basis now need it available every day; international companies need consolidated reports for world-wide regions to be drawn up from reports for separate offices; many businesses need to know instantaneously about their transactions. A transaction processing system can provide the framework for meeting these demands as it provides realtime data access and updates while combining multiple business systems into one coherent application.

The last few years have seen an ever intensifying move towards Open OLTP. The benefits of Open Systems in general apply as well to a transaction processing environment. Open OLTP has the benefit of being portable and interoperable. Constrained MIS budgets leave many MIS executives seeking to level current investments with increasing demands for IT solutions and processing power. Looking toward OLTP, they might realise that their preferred solution consists of a mixture of hardware platforms, databases, networks, LANs, presentation managers and the like. Using open OLTP not only allows the interconnection of current products but also protects the investment through compatibility with future developments

This article will discuss the different components of the TUXEDO Transaction Processing System Release 4.1 developed by UNIX System Laboratories. The TUXEDO System is a mature, available OLTP product supported on more than twenty hardware platforms, by nine operating systems and currently deployed in over fifty applications. As the roots of TUXEDO lie in the UNIX operating system, it is an open OLTP system, supporting unlimited front-end interfaces, network protocols and resource managers. TUXEDO has been evolving since 1978 and is fast becoming recognised a the standard for open OLTP solutions. In particular, it offers the following key capabilities:

- TUXEDO provides application designers and programmers with a state-of- the-art framework for building OLTP applications.
- Openness at all levels of the TUXEDO System architecture with true heterogeneity across all layers of an OLTP application in combination with product components that adhere to standards giving the customer upward compatibility with future hardware and software investments, to take advantage of price/performance and yet still protecting past investment.
- Control over distributed data and functionality gives the user a unified view of a distributed application.

TUXEDO is an internationalised product that allows the user to be presented with diagnostic and system messages in their language of choice, reflecting national date, time and currency

conventions. This is as specified by the X/Open Portability Guide Issue 3.

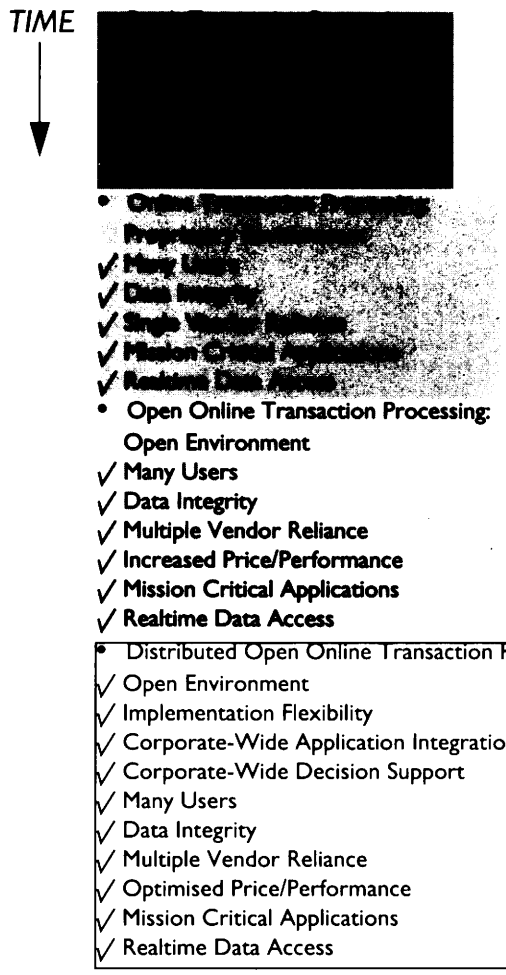

Figure 1: The move towards Open Online Transaction Processing

TUXEDO also gives an application writer a common software platform across a multitude of systems both UNIX-based and proprietary systems (e.g. VMS(TM), AIX., HP-UX, Sun Os(TM) and ULTRIX(TM)). This is achieved by using highly portable code based on industry standards including SVID, POSIX and XPG.

A modular approach allows the product components of TUXEDO to be networked and integrated into other OLTP products. The System comprises two basic parts, which can be licensed and deployed separately: the transaction manager TUXEDO System/T, and a high performance database management system, TUXEDO System/D. Two additional components will shortly be available: TUXEDO /WS, workstation extensions to System/T, and TUXEDO/HOST, enabling System/T to use mainframe services.

The following discussion will now look at the four TUXEDO components.

TUXEDO System/T

The Client/Server Model

The client/server model is the basic structure of the TUXEDO System, providing location transparency by mapping the logical name of a service to the physical address of the server that can perform that service. (See Figure 2)

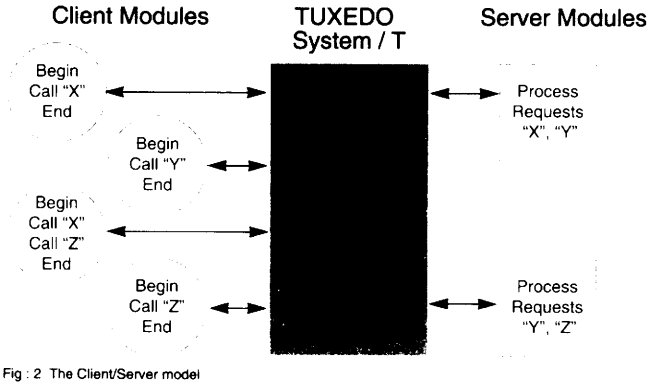

Fig : 2 The Client/Server model

Figure 2: The Client/Server Model

A client process collects input, makes service requests, receives replies to requests and puts out results. A server process accepts a service request, performs the work (which may include additional service requests for the original request), and, when it has finished, either returns results to the requester or forwards the results obtained so far to a new service. Server processes are stateless: several servers may offer the same service and repeated requests for the same service may go to different servers. The request/response model supports asynchronous and synchronous communications.

System/T provides a number of features that enhance the client/server model. These are described below. An application client will use a well defined interface called the ATMI (Application to Transaction Management Interface) to pass requests to a server process. Similarly the server will respond using the same ATMI. This ATMI heightens the underlying client/server configuration and transparently controls transactions on behalf of the client.

The Name Server: Bulletin Board

The Bulletin Board (BB) is the heart of System/T. One of its main functions is to map the service name a client uses to an internally maintained physical address to which a request can be sent. In that way, the BB allows clients to request services by name rather than sending requests to a specific address. Clients do not have to know which server handles their request.

The BB also keeps statistics to aid in load balancing when deciding where a client's request should be sent. This enables System/T to keep track of how many outstanding requests exist and to which servers they are destined. It then routes the request to the server most likely to process it first. The BB is implemented as a piece of shared memory to which all client and server processes attach. Consistency is guaranteed by a combination of user- and system-level locking.

Operations, Administration and Maintenance (OA&M)

System /T allows an application designer to centrally define the hardware, software and networking resources that make up an OLTP application. It can be stated where servers and services are supposed to run and where they should be migrated to in the event of a processor failure. In addition, various characteristics can be assigned to the application's software resources, including processor placement and scheduling information, process recovery criteria and time-out periods. System /T also provides

for central configuration management and dynamic tools for starting, stopping or administrate a distributed OLTP application.

Servers can be dynamically started or stopped making only selected services available. Various parameters such as time-out interval, priorities and load factors may also be changed dynamically. If a processor fails or needs maintenance, the server and services on the out-of-service processor can be migrated to another processor without interruption to the running application.

To enhance application availability, robustness features are built into System /T, including process viability checks, time-out checks, automatic server re-start and recovery procedures. In distributed and multiprocessing environments, System/T can increase the availability of an application by replicating servers and services across several processors. Application data can be partitioned across processors participating in the application and accessed by data-dependent routing of service requests. This enhances the distribution of services and the application's resiliency to select processor failures.

To ensure maximum throughput, System /T automatically performs load balancing and scheduling throughout the system. It uses per-service load factors and keeps totals on outstanding work to deliver a particular request to the server that can process it most quickly.

System/T applications can naturally extend over a set of machines on a LAN without special attention in the application. Communications are handled by System/T using the network independent library (TLI) or the sockets interface.

Transaction Control

Ideally, an open distributed transaction processing system imposes no restrictions on the application?s choice of RMs. Employing the XA interface defined by X/Open, TUXEDO System/T communicates with all resource managers that are XA compliant.

The distributed transaction comprises two general elements: the transaction manager handles the global part, it keeps track of local transactions participating in the global transaction and handles all commit and recovery decisions. The resource manager/server process controls the individual pieces, associates them with the identifier for the global part and carries out the decisions of the TM as they affect the local transaction. To participate in a two-phase commit, the resource manager (RM) must be able to start, precommit, commit and abort a local transaction.

The RMs that participate in a distributed transaction provide support for a two-phase commit presumed-abort protocol:

- They have to offer subroutines that begin, precommit, commit and abort a local transaction.
- They must not make independent commit or failure handling decisions as part of recovery for precommitted local transactions associated with a GTRID without informing the TM.

Each global transaction must have a Global TRansaction IDentifier (GTRID) which is unique across the OLTP system. The use of GTRID is required for recovery as well as a unique identification of distributed units of work. In a two phase commit protocol, a commit coordinator requests participants to precommit. When all participants have reported successful precommit, the GTRID, the coordinator id and a list of ids for RM participants must be logged. The GTRID includes the coordinator id and the RM must put it into stable storage. This allows the transaction manager to obtain a list of GTRIDS in case of a failure and request the status of the global transaction for the GTRID.

The function of the commit coordinator is the management of a two-phase commit protocol. TUXEDO System/T maintains internal information about the participants in a server process, and when an application requests the transaction to be committed, the coordinator of the transaction uses the internal information to begin the commit process by sending precommit requests to participants. The TUXEDO System/T code for commit services is written using common transaction manager/resource manager interface subroutines; different implementations of these subroutines, included in the respective RM libraries, are used to build commit servers.

TUXEDO System/D

TUXEDO System/D is a high performance database management system specifically designed to support the needs of an OLTP application running under UNIX System V. (See Figure 3, TUXEDO System/D architecture)It provides a set of tools to build and administer such applications by implementing an independent user level file system which stores database and transaction log information on a privately managed set of raw disk partitions. Since all the I/O is synchronous and bypasses the UNIX buffer pool, data from committed transactions is written to stable storage providing database consistency and recoverability.

Fig : 3 Tuxedo System / D Architecture

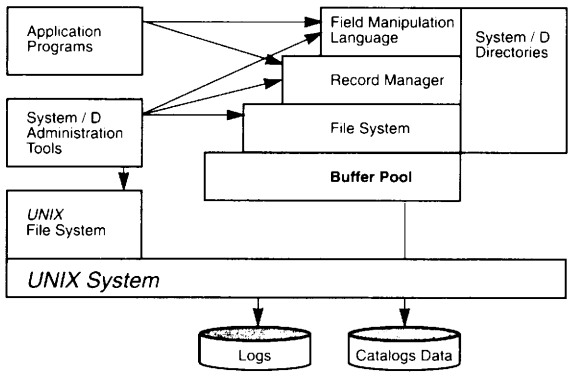

Figure 3: Tuxedo System/D Architecture

TUXEDO System/D is compliant to X/Open's XA interface definition. System/D uses a redo/no-undo strategy with updates logged to separate devices before being applied to the database. TUXEDO System/D also supports the use of embedded ANSI SQL. In case of a system failure, the system only has to redo the updates of committed transactions that are not in the database. High performance database access in System/D has been achieved through the extensive use of caching (using shared memory) optimised for DBMS access, extent-based disk allocation, support of multiple transaction consistency levels and by allowing programmers to navigate through the database with the use of a record at a time interface.

Enterprise Transaction Processing (ETP)

The product direction for the TUXEDO Transaction Manager beyond Release 4.1[1] will expand System/T to embrace proprietary, non-UNIX System V systems to provide for a truly heterogeneous environment. Enhancements will extend System/T client support to include the interaction of workstations and PCs

1. TUXEDO /WS and TUXEDO /HOST will be generally available from USL fourth quarter 1991.

with System/T servers, and the addition of gateways to embrace servers within external OLTP systems. (See Figure 4)

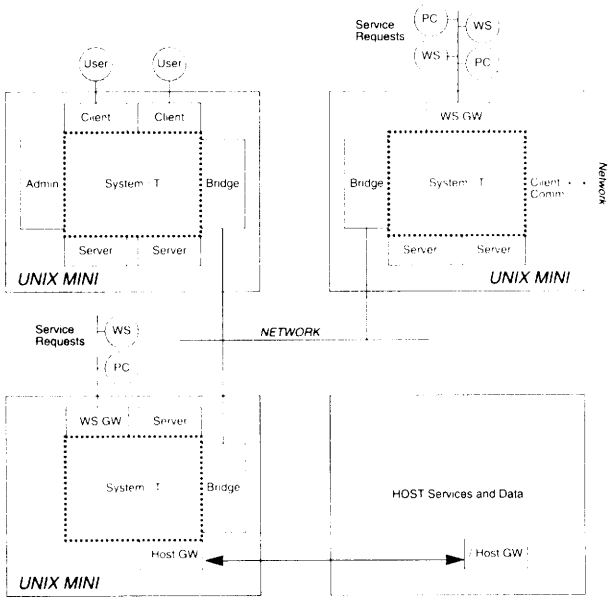

Fig : 4 TUXEDO Release 4.1 Enterpnse Transaction model

Figure 4: TUXEDO Release 4.1 Enterprise Transaction Model

Workstation Support - /WS

The workstation interface (/WS) will allow an application to use its existing base of workstations and PCs to develop and run System/T client programs. These client programs have access to System/T application servers throughout the network. This facility will provide the advantage of removing terminal character processing overhead from server processors and places the overhead on intelligent processors. It also allows application developers to use the presentation facilities and screen managers available for workstations and PCs. In addition to a full client programming interface, further security measures will be implemented to control access to application servers. Gateway software will provide all necessary functionality to establish the LAN connection between the client and server, manage communications and perform the necessary encoding and decoding of messages.

/WS can be utilised on both disk-based and diskless desktop computers that run either the UNIX System or MS-DOS. Full access from diskless systems will be dependent on properties of the network such as remote file access. The MS-DOS version of /WS will support the NETBios interface to the available network provider. The UNIX System version of /WS will support both TLI (Transport Layer Interface) and BSD (Berkeley Software Distribution) Sockets networking interfaces.

/WS will provide maintenance of investment in current systems and will allow the user to take better advantage of desktop power and hardware trends.

Host Support

Major corporate computing resources such as services and data reside on machines that currently run proprietary OLTP software. TUXEDO System/T's new /HOST feature extends the client/server model into the surrounding host environment and allows a System/T application to provide transparent access to remote, external services and data. System/T resident gateway servers represent host resident servers and services on the local System/T node and manage communications with the host computer. Through the gateway servers, host resident servers and services are viewed as being part of a single System/T application. Consequently, the availability of host resident services can be accessed through the standard System/T ATMI programming interface.

With /HOST, USL will provide a System/T gateway for the IBM MVS/CICS environment. The LU 6.2 protocol will be used to perform peer-to-peer communications between the UNIX-resident and the CICS-resident gateways. Software will be provided that will allow a UNIX System/T node to be linked via an LU 6.2 communications library to a CICS application programming interface and CICS administrative utilities. Data conversion will be performed by System/T for the ATMI supported data structures; user defined encode/decode functions will also be supported.

Summary

TUXEDO is a mature, open, on-line transaction processing system, developed to meet widespread market needs. Where applications writers are concerned, client applications can be written not only in C or Cobol but also using CASE tools and some 4GLs (e.g. Ally), resulting in increasing ease of application development.

The TUXEDO system is continually evolving as is demonstrated by the forthcoming release of /WS and /HOST, and with the increasing demand for OLTP in all sectors, TUXEDO will underly most of the TP systems in the Open Systems arena.

The system gives applications developers a great deal of freedom as it allows client applications to be written not only in C or Cobol but also in CASE and some 4GLs (e.g. Ally).

Introduction to back-up and restore

J M van der Waaij
PTT Telecom
The Netherlands

Why would you go through all the trouble of making back-ups? After all, all the data is stored already so why copy it to put it somewhere on a shelf? Ideally, 'the computer' should take care of that. Put this way the lack of quality of both hardware and operating systems becomes apparent. As hardware and software are, quite understandably, not as good as they ideally should be, we (should) all make back-ups as an insurance against troubles.

The back-up process, we'll go into the restore problem later, has several aspects of interest. The choice of the options is dependent on the local situation. How critical is the data? How fast does it change? How many Mega/Giga-bytes? How much money is available? How much time can you spend doing back-ups? All these questions influence the basic three aspects of back-ups:

- back-up media
- back-up format
- back-up strategy

Back-up Media

The choice of back-up media in the UNIX environment is basically one of the following:

Type	*Size*
QIC tape (Quarter Inch Cartridge)	150 Mb
9" tape reels	120 Mb
8mm DAT tape (VHS system)	1-5 Gb
4mm DAT tape (either DATA/DAT or DDS)	1-2 Gb

Depending on the amount of data, the amount of available money and the amount of available time (for changing the tapes) one these media types will be used. Most machines have a QIC tape unit nowadays. The DAT units are becoming more popular. For some situations optical (WORM) or magnetical-optical (MO) disks could be a good solution.

Back-up Format

Since UNIX escaped from Bell Labs there has been a discussion going on. Should 'tar' or 'cpio' be used. The controversy has reached its peak now that we have 'pax' that does both. While the UNIX world was busy struggling with this problem, the rest of the world, meaning the big guys out there, all use an ANSI standard format: ANSI X3.27. It has all the advantages of the UNIX tools and more. However, somehow it seems to have escaped our notice. Maybe EurOpen should do something about it.

Back-up Strategy

If the data on your system is crucial, or rapidly changing, more back-ups should be made than if they are not. If you look closer into the problem there appear to be parts of the disk that are more crucial (your application data) than other parts (the operating system part). To save time you want to have a monthly or so back-up of the UNIX part of the filesystem and maybe 2 backups a day of the precious application data. A back-up utility will have to be able to handle this. In the end the back-up strategy comes down to a matter of efficiency and cost calculation.

Other Aspects of Back-up

How many time have you written the tapes you are using now? When you make the next back-up, are you sure that the tape does not contain valuable data? To be sure of the answers to these questions you need a tape administration mechanism. By labelling the tapes and having an administration tool that automatically checks which tape is used, if it can be used and how many times it has been used. Many time consuming, costly problems and accidents can be prevented if this is done. Alas these mechanisms are nowhere to be found in "standard" UNIX.

Restore

If, in five minutes time, a user calls: "I have accidentally deleted such and such a file and I need it now!" How much time will it take you to find or bold the correct tape and reload the file? Even if the user is *not* your boss? Maybe the tape administration utility should also keep track of tape-indexes. The utility should have a retrieve utility to aid your search. Imagine that your building catches fire tonight. How much valuable data will be irretrievably lost? And if that one tape that you have stored at home is unreadable, what will be the costs?

Conclusion

It is possible to do the back-up/restore process on UNIX systems with the standard tools. However, more than basic back-up/restore cannot be done with them without a lot of work. As soon as you have crucial data, many machines or little time, buying a profession back-up/restore tool should be easily justified. Since UNIX is a professional operating system based on open standards it is high time a X/Open or POSIX committee does some work on this.

However, even with the best tool, you yourself are still responsible for using it.

On the following pages you will find outlines of several products. Naturally every supplier claims that his product is the best. Above we have discussed back-ups in general. But first a few guidelines to help in selecting the best tool for your situation.

- The tool should be able to use both "tar" and "cpio" format and maybe even ANSI X3.27.
- It should be possible to specify what to back-up and what not.
- a good index of saved files should automatically be created.
- It should be possible to restrict usage of the restore function.
- The tool should work with your back-up device.
- It should be possible to specify when the back-up should be performed (timed back-ups).

Optionally tape administration should be automatically done.

ARCHIVE*SQL

Computertime Network Corporation

Brief Product Description

ARCHIVE*SQL is a systems software product that provides a robust, integrated and automated backup and archiving facility. ARCHIVE*SQL is an entire storage management system governed by a powerful on-line catalog. ARCHIVE*SQL is the only product which handles relational databases *and* all you computer system files.

ARCHIVE*SQL operates in stand-alone, client/server or fully distributed multivendor VAX/VMS and UNIX computer environments. ARCHIVE*SQL will also support any other vendor's proprietary computer system within the network.

With ARCHIVE*SQL, you can backup, archive, retrieve and move your data to off-line, near-line or on-line destinations, *anywhere* in the network...onto any storage device, or directly into another database.

Product Description

ARCHIVE*SQL contains the foundation technology for a total service management system that runs independently of operating systems, databases, file systems and networks. It provides automated facilities for the management, operation and distribution of software and data throughout a distributed processing environment.

ARCHIVE*SQL solves these problems by allowing users to:

- Backup, archive, retrieve and move data to off-line, near-line and on-line locations (storage device or within database tables) anywhere in the network.
- Backup, archive and retrieve any file structure supported by a host operating system.
- Archive and retrieve relational database tables or portions of these tables directly and without the need for any intermediate files or operations.
- Use standard SQL and other database operation commands including referential integrity, to select database information for archiving or retrieval.
- Use a "script" facility to specify rules associated with archiving, backup, retrieving and moving data including pre and post operation logic.
- Use a relational database to store comprehensive information about archived data and operation requests, including backup, archiving and retrieval specifications.
- Use standard encryption algorithms that ensure data storage and network data transfer security; or integrate your own encryption algorithm.
- Use a standard data compression algorithm that delivers storage space savings and network traffic savings in the order of 30%; or integrate your own compression algorithms.
- Organise, manage and allocate storage devices and media, manually or automatically, using techniques commonly found only in "mainframe" environments.
- Use on-line transit and data staging areas to support unattended storage/retrieval operations, optimise network data transfers and implement automated storage migration.
- Use a "link" facility to configure and specify rules associated with network data transfers including schedules, file transfer protocols and multi-node routing algorithms.
- Use menus and fill-in-the blank forms for ease of use, or a direct command-line interface. Context-sensitive help is available at all times.
- Use an application programming interface to provide direct access to all ARCHIVE*SQL functions from within application software.
- Use automated operations facilities to schedule and automatically execute the above tasks.

Product Architecture

The heart of the system includes the Kernel, Network Manager and Cataloguer, which uses the power and flexibility of relational database system.

ARCHIVE*SQL is unique in its support of both non-database files and relational database structures. The Kernel provides all common services such as device management, scheduling, data formatting, compression, encryption, exception handling, error recovery and event management and notification. The Kernel contains both a host operating system file system interface and an SQL interface that allows the use of standard SQL scripts to select rows from any specified table or group of tables as well as perform any other permitted database operation. The resulting data can be subsequently inserted into the same or any other table, int he same or different database, located on the same or remote computer. The SQL interface is designed to support multiple database engines.

The Cataloguer provides interactive on-line access to all the information necessary to handle all ARCHIVE*SQL operations. The Catalog contains information about; all storage media utilisation and location; the device configuration of each remote site; all archive, retrieval and network transfer requests and operations; all schedules; all archived data and its characteristics; and the logical grouping of all storage media. Since Catalog information is itself stored in the host relational database, the relational database's standard report and query facilities can be used to generate ad hoc or customised reports additional to those supplied with the product.

The Distributed Network Service Manager provides the product's network protocol independent interface and contains the product's internal network management and recovery services. The Network Service Manager allows users to specify their own

data transfer utilities as an alternative to those supplied by the native network protocol.

The Automatic Virtual Operation (AUTOVOP) automatically performs all user and system scheduled operations including any related network data transfer steps. Schedules can be based on date/time as an absolute date in the future or relative to the current date/time, either on a single shot or repetitive basis. Schedules can also be based on the execution of another request. This allows both independent and sequenced operations to be scheduled and executed automatically.

ARCHIVE*SQL provides multiple end-user interfaces (with multilingual support) directed at users of varying sophistication and subject to different levels of security authorisation. All users can access an easy-to-use menu system. Context sensitive help is available at all times. The direct command-line syntax interface also serves as an application programming interface providing direct use and access by on-line application software.

Price

ARCHIVE*SQL cost $3,000-$36,000 (USD) per CPU, depending on size. (8 user 386 box - $7,500 USD)

Support and Maintenance

Annual hot-line support and updates - 15%.

Availability

Since 1990 in North America. Distributed in England, France, Belgium and Italy through exclusive and non-exclusive distributors.

Unique Selling Points

Only product to support UNXI and VMS.

Only product to support heterogenous networks.

Only product to support relational databased directly in addition to regular files.

Product is also a complete storage management system (media and device librarians, unattended operations etc.) supported by robust on-line catalogue.

Sales Contacts

North America
Gail Binmore
VP, Sales Archive*SQL
Computertime Network Corporation
10340 Cote de Liesse
Montreal, Quebec, Canada H8T 1A3
Telephone +1 514 633 9900
Facsimile +1 514 633 9886

France & Belgium
Mr Oliver Barnabot
Directeur
Echo-Soft Technologies
Tour de Bureaux Rosny 2
Av. Charles de Gaulle
93118 Rosny Sous Bois
Paris, France
Telephone +33 1 48 94 75 75
Facsimile +33 1 48 94 26 80

England
Mr Rod Fenwick
Director
SQL Systems International PLC
Almners Priory
Almners Road
Lyne, Chertsey KT16 0BH
United Kingdom
Telephone +44 932 567 677
Facsimile +44 932 568 168

Italy
Dr. Arsenio Pica
Senior Partner
Confor S.r.l.
Via E Gianturco, 1
00196 Roma, Italy
Telephone +39 6 361 4267
Facsimile +39 6 323 0230

MIRROR and EZWORM

Sysnet

Probably the most secure of all backup solutions is the Write Once Read Many, or WORM, disk drive. Once data has been written onto a WORM cartridge, the only way to remove it is to destroy the media. In addition the WORM cartridge is removeable, giving unlimited offline storage like magnetic tape, but with longer shelf life and faster access to stored data.

Unfortunately, most current UNIX implementations, including the "pure" SV 3.2 and SVR4, do not directly support the attachment of WORM devices, and it is for this reason Sysnet has developed two software packages, one called "MIRROR" and the other called "EZWORM".

MIRROR is a set of utilities and a SCSI device driver that allow the partitioning of the WORM media, and creation of a mountable UNIX filesystem in each of the up to 15 partitions. Each partition can also be used for multiple archival backups, using the standard UNIX tar and cpio commands. The special SCSI device driver is needed for most UNIX implementations, because even if they come with a SCSI driver, it normally does not allow full control and access to WORM drives.

For those who do not require the WORM filesystem capability, and only wish to use the WORM drives for backup, using tar and cpio, Sysnet has a product called "EZWORM". This packages the device driver with a (text) menu system for ease of use, and allows backups using directory, date and wildcard parameters, or by means of a LIST, created with the EZWORM list editor.

MIRROR is available for the Intel/386 versions of UNIX, including SCO UNIX V and Xenix, Interactive 386/ix and UNIX, and "pure" AT&T System V 3.2 and SVR4. The device driver supplied is for the ADAPTEC SCSI host adapters, 1540 for PC-ISA bus, 1640 for Micro-channel and 1740 for EISA bus. Up to seven drives may be installed on one SCSI host adapter (the eight SCSI device is the host) and they can all be WORM drives, or a mixture of devices. In the case of Interactive 386/ix we also provide extensions to the Interactive HPDD (device driver), allowing that to be used with the other types of SCSI controllers it supports. EXWORM is currently available for Interactive 386/ix and SCO UNIX.

The approach taken by Sysnet to create filesystems on the WORM media differs from that taken by other products and proponents, in that we make no attempt to make the WORM drive look like a "normal" read/write drive. The result is that we create a true UNIX filesystem on the WORM media, which is mounted as a read only filesystem using the standard mount command. This means that a WORM filesystem written by MIRROR can (at least theoretically) be mounted on any UNIX V system that has the same type of WORM drive. One caveat to this is that the system must use or recognise the same partitioning scheme, which is definitely the case if it uses the Sysnet device driver.

Other methods generally make use of the virtual filesystem capability provided by many UNIX systems, such as the SUN VFS. This has the drawback of at least one additional layer, with its associated overhead of time and complexity. Another approach is to use a hardware controller that maps sectors on the WORM media, making it look to the software like a read/write drive. Both these methods have the additional disadvantage of consuming a lot of disk space rewriting sectors, a reasonable estimate is probably up to 40% sector overhead.

In general, we believe that both these approaches are somewhat ridiculous, and users requiring that type of capability would be much better served by using an erasable optical drive instead. If you need the special security and audit features provided by a WORM drive, then it should be used as a non-erasable device, and not disguised to look like a read/write device!

As well as supporting WORM drives, the Sysnet driver also supports erasable optical drives, and the latest multifunction drives, which can accept both write once (WORM) and erasable cartridges. Panasonic, Pioneer and Hewlett Packard all offer multifunction drives. The erasable optical drives are not quite functionally equivalent to a normal hard disk drive, however. In addition to being a lot slower in access time, they also (except Panasonic) require a separate erase pass over a sector before it can be rewritten.

This obviously dramatically slows throughput, especially with UNIX, unless the driver software can pre-erase sectors and maintain its own free list. The Panasonic drive does allow direct over-write, and also stores 1 gigabyte per platter, compared to the usual 650 megabytes for "standard" WORM or erasable cartridges, but it does this by use of a non-standard cartridge and format. Note, however, that even "standard" cartridges do not offer proven interchangeability between drives made by different manufacturers.

To increase online storage capacity even further, optical disk autochanges, or "jukeboxes", are now available from a number of sources, including Hewlett Packard and Panasonic. These typically contain two multifunction drives, and up to 50 cartridges that can be loaded and unloaded by commands via the SCSI bus, or a serial port. In the USA they cost in the region of $30k, versus $3-$4k for single multifunction drives, and as little as $2-$3k for WORM drives, but they do offer virtually unlimited storage.

Sysnet products are available in all European countries by direct order from our California office. They are also available through the Programmer's Shop UNIX Catalogue. The documentation provided with MIRROR includes a comprehensive manual, which has proved adequate for reasonably competent UNIX users. EZWORM comes with a brief introduction, and is quite intuitive. Both programs include installation scripts which install the driver and utilities, and rebuild the kernel. Technical support is available by phone or telefax from Sysnet in California, we have not yet been reduced to "limiting" the availability of support, most problems and questions occur at installation, and are usually easily dealt with.

MIRROR is prices at $450.00 per system for quantity one, regardless of the number of users. EZWORM, which does not have filesystem capability, is $99.00. There are various quantity

discount and royalty plans available to Resellers, integrators and VARs. Most copies of MIRROR have been purchased by integrators for resale to various US Government agencies, who seem to have quite stringent backup and audit trail requirements.

For more information contact:

Doug Vincent
Sysnet
125 Blueberry Drive
Scotts Valley
California 95066
USA

Telephone +1 408 438 5120
Facsimile +1 408 438 6676

REELBackup™

The Product and Function

REELbackup™ provides disk to tape backup by handling the five tasks associated with backup operations:

- Routine backups
- Crash recovery
- Individual file recovery in responsed to users requests for an off-line copy of a file
- Scheduling and control
- Tape management - where the tapes used for archiving are monitored and maintained.

REELbackup's Major Features

- Full screen monitor
- On-line index
- Full screen interface for all activities
- Capabilities to:
  - — store multiple files on one tape
  - — automate file archiving
- Flexible backup schedules based on weekly, monthly, quarterly and yearly cycles. A backup schedule is a sequence of backups described on a day-by-day basis. Each day can have full or incremental backup performed.
- Security features include:
  - — tape drive locking
  - — file locking
  - — user restrictions on files they may locate, reconfigure or otherwise alter
- Tape management and control features include: support for standard cpio and tar formats; electronic labels that identify each tape under REELbackup's control; label checking to prevent incorrect tape mounts; support for IBM or ANSI standard labels; off-site rotation controls; and periodic cleaning and replacement schedule capabilities to minimise tape failure.
- Three modes of operation: Self-service, Background and Unattended Backups.
- Simultaneous use of up to 8 tapes for backups and recoveries.
- Administrative controls and reports: Ten administrative reports detail system activity and status. Several can be automatically generated and electronically mailed to appropriate personnel. Includes Missed Backup reports.
- Tape Mount verification by electronic cables on each tape.
- Tend different backup levels.
- Full history of all backup activities.
- NFS network support.

REELbackup's Advantages over UNIX

- Flexibility
- Administrative controls
- Ease of use
- Mainframe-like features
- Speed
- Tape management
- On-line managment
- Ability to locate a particular copy of a lost file
- Reliable way to control access to files on backup tape
- Security features

REELbackup's Unique Selling Points

- Client-server architecture minimises network bottleneck by making it possible to perform all of the selection of files for backup and any data compression on each local disk node before transmitting the data across the network to the tape node.
- Dynamic routing and load balancing of backup disk data to tape based on system configuration and drive availability.
- Sophisticated checkpoint/restart to minimise delay caused by faulty tapes or drives.
- Controls the backup process on multiple devices anywhere on the network.
- Designed for a heterogenous network environment; all of backup data is written in machine-independent format making it very portable.
- Management of the entire backup process to ensure control, security and reliability.

REELbackup's Advantages over Competitive Products

REEL backup typically offers the following advantages to other backup/archive products:

- Superior lead through
- Better file archiving
- Support of standard cpio tar
- Easier to configure
- Support of nearly every tape drive including; 3480 cartridge, 9 track (6250, 1600, 850BPI), QIC cartridge, DAT and 8mm
- Friendly, functional interface
- Comprehensive restore functions
- Excellent documentation

Documentation

REELbackup Master Guide includes sections for system administrators, operators and users. Included with product.

Performance

REELbackup is designed so that there are no inherent performance limitations. Performance is baseed on the underlying hardware.

Price Information

REELbackup is not priced per number of users, but rather by number of nodes.

Therefore, REELbackup is priced at $1500 for one 386 machine, and $2000 for two 386 machines (any number of users). Please note that this price is also applicable to 486 machines and many other systems such as DECstations 2100, 3100 and 5XXX.

Also, SCH provides multiple site discounts for customers who purchase more than one copy of the product.

Availability

REELbackup is available world-wide from:

SCH
Three Centennial Plaza
Cincinnati
Ohio
USA

Telephone +1 513 579 0455
Facsimile +1 513 579 1064
Telex 24 1665 (SCH CIN)

SM-*arch* Version 2.0

Network Backup and Extraction System

Major Features

SM-*arch* is a comprehensive backup, extraction and media library management system for complex networks, supporting remote file access and remote backup devices. Features like multiple simultaneous backups, overlapped rewind, parallel data verification, incremental backup and data compression allow large network-wide backups to be performed efficiently. Files can span multiple volumes. The backup format is tar compatible.

Backup classes control the scope of a backup and can be defined to include or exclude specific files or directories. Backup schedules can be predefined, allowing automatic unattended backups. E-mail is sent to designated users indicating success or failure of the backup. Retention stretch option permits, for example, the weekly backup at month end to have a longer retention, effectively making it the monthly backup.

An on-line catalog of files backed up is maintained. This allows transparent recovery of lost files with the user specifying only the name of the lost file or directory; the user need not know the backup media, number of volumes or their mounting order for a successful restore.

SM-*arch* has its own security features, offering an additional level of password protection. In addition, authorised users can be given varying levels of access. Security of data backed up can be ensured by encryption. to extract from an encoded file, the user has to provide the extraction key.

SM-*arch* is easy to use with on-line help and an identical functionality across platforms with Motif/Open Look and character based interfaces. SM-*arch* users the client-server model backup. Thus, while it is available on various UNIX platforms like the IBM RISC System/6000, Sun 3/Sun 4, SPARC stations and Intel 386/486, it offers network extensions to support PCs running under MS-DOS and OS/2, Macintoshs and DEC VAX computers running VMS.

SM-*arch* can be installed in a single workstation/disk environment and upgraded gradually to include other nodes on the network.

SM-*arch* Ver. 2.0 was announced in May 1991 and is getting good attention. It will be delivered to a number of client sites in June/July 1991. Its predecessor Version 1.0 (single user, with character interface) has sold well in the United States.

Improvement Over Straight UNIX

- Multi-level incremental unattended backups
- Multiple simultaneous backups; parallel rewind and write verification
- Backed up files can span over multiple volumes
- ANSI tape label support
- Tape volume management
- Backup duplication avoided by treating link files separately
- Extraction process transparent to the user
- Motif, OpenLook and character based interfaces
- Full network support allowing backup of other UNIX and non-UNIX nodes. Remote backup devices supported
- Additional security features
- E-mail support
- On-line help

Unique Selling Points

- Unattended backup
- Transparent extraction
- Motif, OpenLook, Character mode interface on various UNIX platforms
- Compress/crypt for space saving and security
- Network backup permitting access to drives connected to different machines
- DOS, OS/2, Macintosh and VAX/VMS platform backup permitted
- tar compatible while supporting ANSI label support

Why Better Than Competition?

- In general, more functions and features as listed above
- Motif/OpenLook interfaces
- Client-server network model allowing UNIX and non-UNIX nodes to be backed up on local or remote devices

Documentation/Support

User Manual describing installation, setup, regular use and trouble shooting. Index provides cross-referencing.

90 day phone/fax support included in the price. Annual support/maintenance contract available.

Performance

Backup speed 10 MB per minute without compression and encryption using Exabyte 2.2 GB 8 mm tape drive on IBM Risc System/6000, model 320 running under AIX 3.0. Other systems comparable.

National Language Support

Not available currently. Will be provided in the third quarter of 1991.

Availability

Server platforms supported: AIBM Risc System/6000, Sun 3, Sun 4, SPARC stations, Intel 386/486 based machines.

Available all over Europe. Distributor/OEM arrangements are currently being worked out. Contact:

Vinod Gupta
Software Moguls Inc
One Southwest Crossing
Suite 510
Eden Prairie
Minnesota 55344
USA

Telephone +1 612 942 7824
Facsimile +1 612 942 7699

Pricing

SM-*arch* Server Modules	
for RS/6000, Sun, SPARC Stations	$4,8000
for Intel 386/486 based PCs	$1,2000
SM-*arch* Client Module	
for Workstations & VAX/VMS	$600
for PCs, Apple Macintosh	$150

All prices are CIF in US dollars.

SysAdmin

Manix Systems

SysAdmin provides comprehensive UNIX system management and reporting utilities, including:

- account maintenance,
- multilevel file backup and restore,
- filesystem maintenance,
- network (LAN, WAN and UUCP) maintenance
- file archiving,
- printer management,
- security monitoring and auditing,
- terminal, process, and, cron maintenance,
- tape library management, and,
- resource accounting and chargeback services

either for a single machine or for a network of heterogeneous systems.

SysAdmin provides a consistent, user-friendly graphical user-interface for various types of output displays and UNIX variants regardless of the underlying operating system. Both character cell terminals are supported through the use of curses and bitmapped displays are supported using X/Windows and Motif.

Although SysAdmin is a complete system administration tools, we concentrate on the Backup and Restore functionality for the purposes of this article.

Major features

The following types of backups are supported:

Physical:	Users may make physical copies of volumes direct to tape or other filesystems.
Logical:	Full and incremental backups are supported using any standard UNIX tool including cpio, tar, dump or wbak. Users may also specify file types or directories to be excluded from backup.
Unattended:	Backups may be scheduled to run from cron or from the command line or user interface of SysAdmin.
Networked:	SysAdmin provides the functionality to backup any filesystem or directory to any tape device with a local area network.

In addition, an intelligent front end is provided for backups to span multiple tapes and tape labelling allows multiple backups to be done to a single high density tape. An optional tape library facility is also provided so that backups automatically get written into the tape library database.

All the standard UNIX tools remain usable, unless tape labeling and volume maintenance are used, in which case tools are provided for reading the tapes transparently.

A large array of media is supported including: floppies, 9-track tape, streamer tape, 8mm tape, DAT, and auto-change devices such at the DATstacker.

Unique Selling Points

- Backup/Restore of any filesystem to any tape drive with a network.
- Tape labelling, volume identification
- Consistent user interface regardless of platform
- Fully networked,
- Comprehensive, not a bunch of disjoint packages
- User API so customers can modify/add features at will
- Can specify what users have privilege for any operation or resource

Product Maturity

SysAdmin was the first commercially available system administration tool and was delivered in 1985. There are over 700 installations worldwide with large UK/US/Canadian installations with smaller sites throughout Europe including France, Germany, Netherlands, Belgium, Austria, and Switzerland.

Support & Training Services

30 days start-up support is provided with any product, and annual maintenance is available at extra cost. This support includes hot-line support and UUCP mail or bug reporting.

Training courses are available on site or at our premises in Dublin. Customisation of SysAdmin is also available to meet specific user requirements.

Documentation

A comprehensive user manual is delivered with SysAdmin and includes:

- Installation and Configuration Manual (with worked examples)
- User Interface Guide
- Full User Manual
- GUI/API Programmer's Guide
- User Level Command Reference (and On-line documentation)
- Context sensitive help is also provided.

Observed Performance

Backup performance is closed related to the hardware used, however an average of 600 Mb per hour for networked backup and 2.2 Gb per hour for local backup is typical.

Native Language Support

All message text is distributed in a single source file, so can be translated by any end user. All help text is also provided is source format with a standard Binary release.

Supported Platforms

The following hardware platforms and associated UNIX variants are currently supported:

Altos, Apollo, Arix, Arete, AT&T, Alliant, Amdahl, CCI, Compaq, Concurrent/Masscomp, Convex, Cray, Data General/Aviion, Encore/Gould, FPS/Celerity, DEC, Fujitsi, Harris, HP, IBC, IBM, Integration Solutions, Intergraph, MIPS, Motorola, Multiflow, NCR, Pyramid, Sequent, Siemens Nixdorf, Silicon Graphics, Sun, Tandem, Unisys, Wang.

Distributors

Europe: Manix Systems
1 Haigh Terrace
Dun Laoghaire,
Co. Dublin, Ireland
Telephone +353 1 280 3128
Facsimile +353 1 284 4987
Contact: Sean Byrne

Benelux: Pink Elephant Management Services
Zwartelaan 24, Postbus 106
2270 AC Voorburg
The Netherlands
Telephone +31 70 369 4231
Facsimile +31 70 387 4294
Contact: Martin Siebert

We are at present actively searching for distributors for France, Germany, and Italy.

Pricing Information

Licencing is on a per node basis. Quantity discounts, site licences, source licences and corporate agreements are available. Price for a single end use 'client' SysAdmin is IR £1,850 for a SCO/386 platform ranging up to IR £4,600 (GBP £9,200) for mainframes such as Amdahl, Convex, and Cray. SysAdmin includes two other products, UniMenu - a Unix menu interface, and SysAudit - a system security monitoring package.

SysAdmin itself has the additional modules at extra cost: Network Management, NIS & NFS Management, Quota Management, Remote software maintenance and Distribution, Tape Library Management, and Performance Monitoring and Resource Accounting.

A new product, SysMaint, shortly to be announced, covers only the areas of Backup, Archiving, Restore, and User and Group Management providing the same features of SysAdmin at lower cost. SysMaint provides a smooth upgrade path for customers wishing to tackle the area of backup/restore now and to use the same product for all other areas of system administration at a later time with a low upgrade cost and a consistent set of tools.

Other products offered by Manix Systems include UniMenu, SysAudit, and JobAcct, a resource chargeback system.

Xi-Bar

Xi Software Ltd

Xi-Bar has been available for two years and has been adopted by some of the major corporations in the UK. Telephone hotline support for 90 days is provided free of charge. Xi-Bar is also available on a free of charge timed trial period. Once the rial period is completed, the user who has decided to purchase Xi-Bar only has to call for a reverification code and pay the licence fee to have full use of the product on his chosen machine.

Xi-Bar only costs 250 pounds on an 8 user 386 based SCO or Interactive machine, a very small price to pay for the peace of mind that routine backup is occuring on a continuous basis and that files are recoverable in an orderly manner. Xi-Bar is available throughout the world, and with the opportunity for the user to modify the help and comment files on a global or per user basis, it is clear that the user can have these functions on screen in their native tongue. Distribution agreements through Europe are expected to be signed before the end of 1991.

Xi-Bar is a software package which allows the system manager to historically record the whereabouts of each file and every modification to that file, dating it for future reference, giving the current information on the screen can only be a bonus. This package catalogues the tapes and floppy disks used so that all files can be quickly and easily located, the system manager is more likely to ensure that the backup of the system is a priority which occurs regularly and painlessly. It could be vital to the user, after a system failure, to be able to recover the lost data quickly and efficiently. Unfortunately it is only after a disaster that one discovers the weakness of the backup system available. Very few users are prepared to tolerate a simulated crash to test the recovery procedures or the capability of the provided software for this purpose.

This utility written in the UK by Xi Software Ltd, called Xi-Bar, provides a complete environment for Backup, Archive and Retrieval under UNIX. This program gives the user full control of the backup environment via the computer screen in a user friendly package. The facilities include :-

- Full Backup of nominated file systems.
- Continually running incremental backup.
- Off-line archive of little used files.
- Intelligent Media Management of tapes and all other removable media.

With screen-based facilities letting the user define the details of the host system, examine the catalogue of saved files and media, and to initiate backup or retrieval operations, Xi-Bar is a powerful asset to the system manager. The backup records are held on the computer so that the necessity for keeping records on paper are eliminated. Individual files, or groups of files can be recovered as required with optimisation as required by the user.

Xi-Bar was written with the user in mind. It has been modified as users come on line to incorporate modifications which are requested by them which will provide a more flexible product. It is after all the user who determines the usefulness of a utility and thus the amount of usage that utility will achieve. As we saw before the backup function is all to often ignored by the UNIX system manager because it has become an onerous task. Xi-Bar is usually set to run automatically in the background making the user completely unaware of its existence until something is lost or needs to be recovered.

Files become changed as the user works, some get deleted some get added and some just get modified. A good backup program will operate behind the scenes, periodically noting the changes, selecting tapes or other media and writing them out. A well administered system will have someone responsible for looking after these tapes or discs. A good backup program will prompt the system manager to load the backup media periodically and hence if a file becomes lost at some time it can easily be restored. To insure that this restoration process is regularly undertaken it must firstly be made easy for the system manager to remember that the process needs to be undertaken, and secondly the process itself must not be complicated. On both of these areas UNIX, in any of its standard guises, falls short of the requirement. In Xi-Bar there is a full screen representation of the records in the current directory which allows the manager to view the status of all files listed.

The user can recover either a single file, a group of files, or a subdirectory with all its contents - which could include further subdirectories - by making a mark against the files requiring recovery and quitting the directory. When the user arrives at this point he is prompted to continue, or not, with the recovery process. If he decides that the process should continue the appropriate recovery jobs are submitted. A record of the tapes or other media is then submitted to him. The operator then loads the media in any order until the recovery process is complete. Files then become reinstated in the appropriate directory.

At the setup stage the operator only needs to define the devices to be used for the backup and recovery, allocate and label the sets of media for full and incremental backup, define the times for automatic backup, create the script files to specify the backup areas, and add the appropriate cron entries. Once the system is running the operator only needs to review the media records to add new tapes, remove old tapes, remove old ones and expunge any out of date records. It is also necessary to note the media used in the last full backup in case there is a catastrophic failure in the system which will necessitate a full recovery.

There are four types of backup, all of which are available to the user of Xi-Bar. The most common type of the backup operation is the one where files created or modified since the last backup are written out. This is called incremental backup. A full backup will copy out all files in the specified parts of the file system regardless of any changes which may or may not have been made since the last backup. A root backup will usually be made to give exhaustive cover to the full file system. In the event of a catastrophic loss only the most recent tape containing this saved copy, and subsequent incremental backups, made after the latest root backup, will be necessary to give a close approximation to

the state of the system before the loss. The fourth method of backup is archiving, where it is necessary to keep files for long periods of time whilst not filling the available memory space on the computer. Xi-Bar provides media indexing facilities which are necessary to enable fast accessing of this information when it is eventually required.

We have seen previously it is necessary to have good facilities for backup and archiving of the files on the computer system. It is of little comfort to the user to have the best facilities on his system fr backup and archive if there is a difficulty in recovering the information to the system.

When the operator decides to recover files, he may load as many as are required by as many users who need them recovered. The system will recover as many of the files as is possible in one pass. If an error is detected on the media a message is signalled to the operator who can then recover them from a previous backup which has not been affected in the same way. A full recovery which will be necessary after a catastrophic failure will be necessary. This is a simple process to the user who has Xi-Bar installed on his system. Unfortunately the user realises when it is too late that the support he has for file restoration is inadequate. It becomes a matter of BUYER BEWARE.

Contact: Peter Theobald
Xi Software Ltd
6 West Burrowfield
Welwyn Garden City
Herts AL7 4TW
United Kingdom

Telephone +44 707 371405
Facsimile +44 707 335801
Email pbt@xisl.co.uk

Book Reviews

Reviewed by Simon Kenyon of ICL - Information Technology Centre, Dublin, IRELAND <simon@itc.icl.ie>

X Toolkit Intrinsics Programming Manual OSF/Motif Edition

Adrian Nye and Tim O'Reilly
O'Reilly and Associates Inc., December 1990
ISBN 0-937175-62-5.
(US) Price $30, Soft Back, 632 pp,

This is volume 4 in the X Window System series published by O'Reilly and Associates.

The blurb for this book says that it is a "complete programmer's guide to the X Toolkit (Xt)" and that the book "uses the Motif widgets to demonstrate how to use existing widgets, but is equally applicable to and provides a good introduction to programming with any other widget set based on Xt, such as the MIT Athena widgets, or AT&T's OPEN LOOK widget set".

This aside, what is this book? It is a rehash of an earlier edition which described the Intrinsics using the Athena Widgets as examples. It starts off with a chapter describing what the X Window System is all about and then swiftly moves on to discuss the X Toolkit (Xt), which is described as "providing a simplified approach to graphical user-interface programming".

The book is structured as a tutorial. An example application is developed, becominging increasingly more complex as new features are added to it. The culmination of this is the recoding of the example using a new widget, the BitmapEdit widget.

The intention of the book is that the reader should become a proficient Xt programmer, and for my money this objective is achieved. The book arrived just as we in the ITC were about to port a large application from XView to Xt. We already had the previous edition of this book; but it was my copy that kept disappearing off my desk. This is perhaps more relevant than my subjective opinion.

I do have a number of gripes about this book however. The first is that the typography is awful. It really gets in the way of the learning process. For a book about constructing graphical user-interfaces, faking the screen images makes the images worse than useless. They give the reader no clue as to what the screen should actually look like.

The second point is that the running example leaves a few "exercises for the reader". It is like a lecture in mathematics, where some intermediate steps are left to the student. This caused me some grief.

The description of OPEN LOOK and Athena are sketchy and in the case of OPEN LOOK, out of date.

The final point is that why are all the interesting bits "beyond the scope of this book".

There are no glaringly obvious errors that I could see, apart from some problems with the resource entries accompanying the examples.

This is a book for the programmer who wishes to program with widgets, and Motif in particular. It is for someone who needs to know, as there is effort required. This effort is a lot less than wading through the MIT supplied documentation.

Recommended.

X Window System User's Guide OSF/Motif Edition

Valerie Quercia and Tim O'Reilly
O'Reilly and AssociatesInc.
December 1990,
ISBN 0-937175-61-7.
(US) Price $30, Soft Back, 709 pp

This is volume 3 in the X Window System series published by O'Reilly and Associates.

This book is an introduction to the X Window System. It describes the basic concepts of X and gives details of all the MIT supplied X clients. It describes the mwm window manager and brief overviews of the various features of a Motif application and is only superficially different from the previous edition. As such does not deserve the title of an OSF/Motif Edition.

Unlike Volume 4 in this series, there is no incentive to overcome the poor typography, as the information that this book contains is not too hard to find elsewhere. This means that this book is not going to serve its intended audience, which is the novice X user.

The other volumes in this series are indispensable to the X Window System user. A much sharper presentation and a much more thorough treatment of Motif is required, if this book is to achieve the same status.

Not recommended.

UNIX Tutorial Review

UNIX Tutorial

K&A Software Engineering

Reviewed by Peter Theobald of Xi Software Ltd

Email pbt@xisl.co.uk.

This product is intended to assist the DOS user to migrate to the UNIX environment.

The general impression given by the packaging and documentation is probably unfortunate because the software does give a good basic grounding to the uninitiated into the terms used in UNIX. The presentation of the material on the screen is good, clear and easy to follow. The user is able to run through the entire programme without having to refer to the documentation, as all of the necessary keystrokes and prompts are given on the screen.

There are however some underlining points which might be considered for future releases. It would be nice to be reminded, when returning to the tutorial, after a spell away, exactly to which point one had arrived in the previous session. It might also be a reasonable request to have the program tell one that they appeared to be having some difficulties on a particular aspect and that it might be in their interest to repeat a particular section rather than the correct answer being given, keystroke by keystroke, to allow the student to complete the exercise without understanding how or why they had arrived there only to be met with congratulations at having reached the goal. It was particularly noticeable that after deliberately typing some code incorrectly, and being corrected on every character, that I was told that I was well on the way to becoming an excellent "C" programmer.

A few more examples in the earlier exercises which allowed the student to work things out for himself rather than being given the answer would probably help them to absorb the information more quickly.

The cost of £165 plus VAT would seem to be a little excessive for this product in its present form. If it were made more interactive with regard to the error messages, some of the comments made above rectified, and the documentation were produced in a more permanent form, the price would become more acceptable.

UNIX Tutorial is however a useful guide to the first steps in learning UNIX to those who have been used to DOS.

Contact

Corporate Hi-Tech Services

Telephone +44 793 420069
Facsimile +44 793 611437

Book review

Turbo C++

Ira Pohl
Benjamin Cummings
ISBN 0-8053-6017-4
(UK) Price £22.45, Paperback, 279pp

Reviewed by Lindsay F. Marshall, Computing Laboratory, University of Newcastle upon Tyne

Email Lindsay.Marshall@newcastle.ac.uk

Ira Pohl has written several books on programming and so has considerable experience of the genre. It also means that he can practice software engineering - not only does he reuse material from his other books but replicates parts of this one as well. There is of course nothing intrinsically wrong with this, but some of the examples look decidedly careworn and should be reworked to make them more object-oriented.

The author has taken the opportunity provided by this new book to reorder his material. This has lead to a slightly better lead in to the maze that is C++ but as usual the reader soon gets lost in a jumble of features. For example, why the author feels it necessary to introduce operator overloading on page 18 is a mystery, particularly as it is not used again for 100 pages.

This kind of muddled thinking is typical of C++ texts, as is over emphasis of the language's relationship with C. Pohl is less guilty of this latter fault than most authors, but only because he assumes his readers have a working knowledge of C. However he still feels it necessary to relate new ideas back to the earlier language rather than let them stand by themselves. That being said, the material is in a sensible order and emphasises the importance of learning object-oriented programming, though without much real help to understanding it. It does however, fail my two standard tests of a C++ book by introducing both friend functions and inlining far too early.

The Turbo C++ aspect of the book is puzzling. Turbo C++ is one of the more standard of the commercially available language processors and does provide an elegant environment in which to work. But the book has minimal coverage of the environment and even though each chapter has a section of "Turbo C++ Considerations", the information contained in them is germane to all recent versions of C++. Another way of marketing this would have been to call the book simply "C++" and say "featuring release 2.1" on the cover. The only real contribution to the book seems to be that there is an effusive forward written by the president of Borland International, Inc.

When all is said and done, this is not a bad book and is an improvement on the author's last book on the language. Perhaps in an iteration or two he will produce the perfect text.

Book Review

The C Programmer's Companion: ANSI C Library Functions

R S Jones
Silicon Press, 1991
ISBN 0-929306-09-0

Reviewed by Mick Farmer of Birkbeck College, University of London

Email mick@cs.bbk.ac.uk

As the title implies, this book describes the C library functions in detail. It also includes details of the types and macros defined in the standard headers. A reasonable number of examples are also included.

As the author points out in the preface, most books on C concentrate on the language, discussing facilities such as functions, macros, pointers, statements, and types. The C library is often only mentioned in passing, usually when a particular function or macro from the library is being used in an example. This is because, historically, the library functions were considered part of the C programming environment provided by the underlying operating system. Even Kernighan & Ritchie's The C Programming Language didn't specifically discuss library functions. ANSI C became a standard in 1989/1990. Recognising that library functions are really part of C, and to ensure compatibility between compilers, ANSI C has formally made the library part of the language specification (in my copy of the ISO draft standard, section 4 - concerned with the library - is nearly half the complete publication).

The overall organisation of this book is good. The introduction first demonstrates the need for headers in C programs before going on to introduce the fifteen standard headers that ANSI C compilers must provide. Each chapter then discusses the contents of one or more standard headers and shows how the more interesting macros and/or functions are used, via simple recipes or more complicated examples. Limits and restrictions, such as user-defined identifiers clashing with library identifiers, are pointed out in the appropriate context. There is an extremely short bibliography (six publications), but the index is acceptable - every function, macro, and type from the standard headers are included.

I came to this book as a competent C programmer with many years experience and as an author (including a book and a video on C :-). The book is very easy to read and-doesn't assume any great experience on the part of the reader. Still, there were many facets of the library that I hadn't come across, or thought about before, that the author talks about in a straightforward, uncomplicated style. Even high-boredom items such as `<limits.h>`, `<float.h>` and the modes and format specifiers in `<stdio.h>` are dealt with in a refreshing way.

In spite of this, there are places where the author has been sloppy. The book is full of references to header files, whereas the ANSI standard, via footnote 86, explains that a header is not necessarily a source file. The notion is, of course, that the contents of a header are available by some means through the `#include` preprocessor directive. The author states in the preface that the examples have been tested thoroughly. However, when discussing searching and sorting utilities, the author consistently gets wrong the relation between the comparison function and the key and the array element being compared. Also, in some examples, there is a difference between the declaration of members of a structure (fname, lname) and their use (first, last). Perhaps it's a case of last-minute editing! There are a few typographical errors, mostly of the missing word variety.

In conclusion, this is a good book for both the beginning C programmer, who wants to know about the library, and the experienced C programmer, who's interested in the finer semantic details. Like many of the books from Silicon Press, this is a useful addition to your bookshelf, especially the cheaper, paperback, edition that I have in front of me.

Book Review

Open Systems for Europe
Elliman and Sanger
Chapman & Hall, April 1991
ISBN 0-412-378507
(UK) Price £50, Hard Back, 186 pp

Reviewed by Peter Theobald of Xi Software Ltd

Email pbt@xisl.co.uk

Open Systems for Europe is a collection of papers given in December 1988 aimed at informing the computer user of the present and future. The book is aimed at "IT managers, purchasing officers or procurement officers and all those within user organisations who seek to justify the business case for an Open Systems IT strategy".

As a historical record of events that occurred 18 months prior to publication in an industry that produces solutions to obsolete the previous operating standards every two to five years may not appeal to the targeted audience. It is possible however that there are some people who still need to be persuaded that there is a need to consider the advantages of open systems over bespoke systems running proprietary operating systems. These people will find that some of the case studies from major corporations will encourage them to consider the implications in detail before embarking on a half hearted trip into a maze which has been created by marketing hype.

All of the articles address their respective topics in a workman like manner albeit in rather dated terms. One may ask whether some of the contributors might now have differing views from those given in 1988 particularly when one considers the way that the power of computers has increased and the costs to the end user has decreased.

The purchasing officer may find some difficulty in finding his way around the vast array of abbreviations that are endemic in the vocabulary of those involved within the computer industry and in this book. Every effort appears to have been made to assist the reader on this point but to the non computer buff the three and four letter abbreviations can be a cause of distraction if not frustration.

This book therefore when seen in context presents the potential user of computer systems with a record of the anticipated situation at this time from a historical point of view without the benefit of the understanding that the marketing forces would have caused the case for open systems to become overwhelming. The question then becomes centred on how much further the standardisation will be taken. The later articles in the book deal with the applications seen as necessities by the larger users in 1988 which will become essential to the small to medium sized users of today.

Book Review

MH & xmh, E-mail for Users and Programmers

Jerry Peek
O'Reilly & Associates, Jan 1991, ISBN 0-937175-63-3. (UK) Price Not yet Available, Soft Back, 555 pp,

"The Octopus Book"

Reviewed by Andrew Macpherson.

Email A.Macpherson@sakura.uucp

This is a book about electronic mail, or rather about MH which is an extremely rich an interface to electronic mail, and its partner XMH which gives a graphic front end to that functionality. How does one separate the book about the program from the program itself? When I started with the Octopus Book I had not used MH at all before, and the distinction was blurred, I could claim to be a serious user of electronic mail systems, but that was about it. I even had to compile up the latest release of MH to be in sync with the text.

Some months on (yes this review is late) the split is much easier to make. MH is a complex system, and the book sets out to present two ways of using it, the style is a friendly, informal tutorial and very readable. The sheer volume of the book is daunting but one swiftly realises that there are really two books and a fairly extensive manual set between the covers. Chapters 4 through 7 describe the command line interface, while 13 to 15 have XMH covered in 88 pages. Much relieved one can move on to read the half of immediate interest.

The book's practical approach gets one up out of the armchair immediately. It may be possible to read, and believe 'this is what happens' in the text, it is more satisfying by far to work through the examples on ones' own system.

Starting with the basic interface one reads, creates, sends and files messages. Then comes the fun of customisation \(em how do you like to list the contents of a mailbox? Do you want to include the original message automaticly in a reply? All the features needed to function as a mail user are well covered. In XMH adding accelerator buttons to print the displayed message is easily done.

And yet... at the end of the section one has had a good tutorial on using mail but one has read a lot of pages without ever touching on auto-filing, "I'm on holiday" messages, bulletin boards or anything else of what one would term the 'power-features' of MH. That first time through one was disappointed by the carefully crafted focus.

A week later I was rude enough to send an automatic "Thanks I'll get round to your message" note to Byron Rakitzis when he replied to a question I had asked about his 'rc' shell. Properly taken to task I turned back to the Octopus Book to see how to do something about it. The information was there, and that is a mistake I will not be making again. My view was changed, and I now appreciate that I had built up the wrong expectations while installing the programs. Read the book first, and borrow someone else's system to work through the examples.

So who should buy this book? A pre-requisite is of course that one will be using or maintaining MH. That in turn begs the question of whether one should be.

Since ucbmail / mailx are rather fundamentally flawed in address handling, something else is needed, and MH is one of the few mail systems that really does obey the standards in every detail. For the mailtool user XMH is a wonderful alternative, and the various emacs modes are great. On the command line though MH is ultimately so rich and flexible that anyone trying to cope with it probably needs this book to get going. I have now learned MH, and I shall continue using it, even running it on my Xenix PC, but did I really want to go that far?

Abstracts

From NLUUG conference at Ede on 8 May 1991.

Many thanks are due to Emile van Dantzig <emile@tmu.nl> who translated many of these abstracts from Dutch.

Securing Open Systems in Open Networks

David Dack
HP-labs, Bristol, United Kingdom

Mr Dack will give an overview of the conference theme, entitled "Securing Open Systems in Open Networks." No paper was provided in time for the proceedings.

Perl for management

Johan Vrmans
Multihouse, Gouda, The Netherlands

Email jv@mh.nl

The base of UNIX system management is traditionally formed by shell scripts, exploiting UNIX tools like "sed", "grep" and "awk". This conforms to the UNIX principle "Keep It Simple". Usually, these scripts are rather optimistic. They stem from a sequence of interactive commands, and as long as everything goes well they behave as they should. Exceptional situations and unusual error conditions are seldom handled properly, if at all. Moreover, since the individual tools have been developed independently, their use and integration is not always a trivial task.

Using one consistent programming language, like C, is better, but most of the handy features of the UNIX tools are not easily accessible from plain C programs.

The tool "Perl" combines the properties of a decent programming language with the handy features of most UNIX tools. Perl offers

- data extraction and manipulation with extended regular expression patterns,
- access to almost every system call and C library routine,
- the possibility to create and make up reports easily,
- no hidden limits on string lengths, buffer sizes etc.

All in the form of a programming language with built-in symbolic debugger (ever tried to debug an awk program?) and much more.

Nowadays, system management has to conform to high quality, security and reliability standards. Therefore, better tools are needed than simple shell scripts using out-of-the-box tools that are sometimes designed for this purpose, but most of the time they are not.

A still growing user population with very active USEnet support may lead to the conclusion that Perl could be the answer to this challenge.

Administering Technical Workstations

Howard Farmer
Shell UK Information and Computing Systems, London, England

Within the Oil industry, there is a trend towards the use of high performance UNIX workstations, for interactive geological, petro-physical and reservoir engineering applications. These workstations are purchased from a variety of vendors, selected on grounds of application availability and price/performance. Two major problems arise, that of distributed system management and distributed systems security. These have become multi-vendor problems requiring multi-vendor solutions.

In the past, UNIX system management has been performed on an ad-hoc basis, using an arcane collection of single vendor tools and generic UNIX utilities. As the management environment becomes more complex, these methods become untenable, being too skills intensive and expensive to apply.

Formal (vendor-neutral) standards for system management are at an early stage of development and may not be ready for commercial application for some time. To wait for International standards before tackling the management problem, would place unacceptable limitations on business development. Shell have therefore taken a pragmatic approach of commissioning the development a multi-vendor Facilities Management (FM) scheme using currently available technology.

The FM scheme has been designed to handle the projected growth in distributed computing over a (minimum) 2 to 3 years period. We would then envisage a migration to standards based management techniques.

The submission provided a case study, discussing the successes, failures and frustrations involved in implementing a multi-vendor distributed management and support environment for Open systems, within a large commercial organisation.

NFS Client-Server Performance

Jos van der Meer, Frans Wessels, Maarten Westenberg
Sun, Nederland

Email maarten@idefix.sun.nl

NFS Client-server performance *) is a very important issue. Often we find ourselves in a situation where a client-server configuration is not performing optimally. Fortunately, there are several ways to improve the performance in such a situation, but in general not all of them will be equally successful.

The document describes the way in which we feel that machines in a client-server configuration, using NFS, should be configured.

In its simplest form, a client-server configuration consists of three independent performance determining factors: The NFS server, the NFS client and the Network. Determination of bottle-necks in such a configuration is difficult, as we deal with a complex set of interacting systems which behaviour is non-deterministic.

Often, we tend to look for improvement by reconfiguring the (Ethernet) network, or by adding disks or memory, thereby looking at one of the bad performing elements in the configuration in isolation.

In our document a different strategy is outlined, based upon the way clients and servers interact. Therefore, the total configuration should be considered in order to perform best, and we will see that this can result in reallocation of resources instead of adding new ones.

The approach will be pragmatic, we will abstract from technical details as much as possible. However, for a better understanding of the NFS client-server mechanism, we will also provide some technical background in this area.

*) The document deals with NFS servers only, database servers are therefore not discussed.

System Management on large scale

UNIX at the PTT Telecom department

Michiel van de Waaij
PTT Telecom, Groningen

PTT telecom is a large UNIX account. About a hundred systems for diverse purposes are distributed all over the Netherland, while a strong growth of the number of systems is to be expected. I & AT (the central department supplying information facilities and automation for the PTT), is on the road to setup a central support and management system for all these UNIX systems.

The scale, the decentralisation and the geographical spread involves a number of specific problems. Moreover often many sites lack well trained UNIX system managers. This paper will focus on the current solutions worked out at PTT I & AT for management of users, backup and media and remote performance guarding.

UNIX System management!

J.C.H. Aalders
Faculty for Technical Mathematics and Informatic's,
Technical University Delft

Computer protection does not just incorporate measures against penetrability of computer systems, but measures that aim to guarantee continuity and correctness of the data handling.

Protection against penetrability was not a design criterion of UNIX. Programmed protection measures form an addition. This addition offers in the "standard UNIX versions" such protection provided that the users are willing to cooperate. UNIX is penetrable for each user that is willing to appropriate super user privileges. Efforts to raise more thresholds within the current UNIX system lead to loss of functionality.

Data communication did imply a new chance to realize protection against penetrability as intrinsic aspect of the (distributed) system. However practice is different.

Protection in a UNIX Network Environment

Frans Schippers
Xirion B.V., De Meern

Email frans@xirion.nl

With the current process of development of local area networks it has become more and more difficult for an organisation to establish the network and connected systems reliability demands.

The management model of the current network software allows transparent data exchange only between systems which trust each other. For university development centers ease of use and performance were more important issues than well anchored permission control. This is due to the fact that the first LANs were formed by a limited number of well managed reliable large computers connected by means of expensive hardware and simple protocols implemented on kernel level. The financial and technical barriers were assumed to be sufficient to prevent cracks.

Personal work stations brought a fundamental change in this situation. These workstations work mostly in a client server model, to receive, process and submit information. The required data transport can be handled in a transparent way by the network software because the work stations are to be expected reliable. However, these work stations are often freely accessible in a room and the actual management will be done mostly by the users themselves. That is why the organisation is not able to locate reliable functioning with enough certainty. They will be classified as unreliable within the management model. Forthcoming consequences are that serious limitations should be imposed on communications. We have to accept the fact that the current infra structure and the management model are not valid for a secure and reliable exchange of data between central servers and personal work stations.

In this presentation security aspects of data exchange in networks will be highlighted points wise. Authentication of systems and persons, securing data storage and transport and demands in relation to new developments for UNIX in a networked environment.

Data and System organisation, a first move: The Art of Automounting

Martien van Steenbergen
SUN Nederland

Email mvs@sun.nl

This paper will be presented at EurOpen Budapest conference. Although NLUUG awarded this paper with the best presentation of the conference, Martien asked me to wait until a new and improved abstract and paper will become available.

SNMP Simple Network Management Protocol

Daniel Karrenberg
Centre for Mathematics and Informatics, CWI Amsterdam

Email Daniel.Karrenberg@cwi.nl

SNMP, the Simple Network Management Protocol, is a protocol enabling management of networks including hosts, routers, bridges, modems etc. It is possible to retrieve information from network components and to control them.

Topics to be covered will be:

- SNMP overview
- The Management Information Base (MIB)
- Applications
- Public Domain SNMP Implementations

Product Presentations

Network Management Tools

Huub van Baal
Hewlett-Packard, Amstelveen

To be able to manage and control todays UNIX networks in a proper way, Hewlett Packard (a company with a large experience in workstation based networks) has designed a number of interesting tools for the management of complex UNIX networks. Many of these networking tools are part of (or proposed to be part of) OSF developments like OSF/DCE and OSF/DME. Almost all facilities are based on NCS (the Network Computing System). NCS is an integral part of OSF/DCE, the Open Systems Federation Distributed Computing Environment.

This paper will cover a short overview of available tools and a more in-depth presentation of:

- Omniback, a network wide backup tool
- OpenView, a network management help

Omniback is a distributed backup tool based on NCS in order to use resources (tapes and disks) in the network to realize a fast and adequate backup procedure. Features are: parallel backup, diverse schedule mechanisms, unattended operation, automatic registration on several levels, etc. With good reason can be spoken about a Total Backup Management System.

OpenView encloses a complete family of network tools and services to manage and control local and wide multi vendor area networks. OpenView is the result of Hewlett Packard's broad experience with multi vendor network solutions, test and measuring tools.

OpenView does not just incorporate hardware based tools, but also a software management tool integrated with the hardware, allowing the user to control, manage and monitor the network both locally and remotely.

Integration and use of tools in Sun environment.

Frans Wessels
Sun Microsystems Nederland B.V., Amersfoort

<wallstreet%frans@sun.nl> System management, including network and system management, is going beyond just installing and maintaining software on a number of computer systems.

SunNet Manager is a product offering a basis for management of complete configurations. However more additional products will be necessary to meet the current needs. An integration of these products is required.

This presentation will deal with two cases of such an integration. A number of hints and tips concerning performance measurements of a client server model will be part of this presentation.

ACS Access Control System for UNIX systems

Marten van Gelderen
Xirion bv, De Meern

Email marten@xirion.nl

The "problem" of cooperating colleges in business, a department or a project will be solved in one of the two classical ways: . one will log on with an own user account and will subsequently "cd" (change directory) to the directory where the common activities will take place. One will log on with a common user account (i.e. account) that is connected to the desired working directory and no "cd" is necessary.

The disadvantage of the first solution will be that the fundamental ownership of the files are hard coupled to the owner of the files. The owner can change the protection of files in such a way that other project participants will be unable to enter his files.

The disadvantage of the second solution will be the uncertainty who is doing what under the same user account. The user account root is a good example of the above painted situation. What's more a serious side affect in the second solution is the disadvantage of the need to learn new passwords when a project participant leaves the company.

ACS, the Access Control System for UNIX based systems, designed by Xirion, offers a golden mean between these two extreme ways of working, actually offering no solution at all: no permissions and everybody can log on to his own user account, or every one logs on to a common user account and loses his personal identity "for ever" in the system. ACS uses techniques well known in a UNIX environment. For example a local password will permit access to a certain sub domain for a user or a group of users under certain conditions.

phLOGIN Product Overview

Alain Williams
Parliament Hill Computers Ltd, Berkshire, U.K.

Email addw@phcomp.co.uk

phLOGIN is a security enhanced, forms based replacement for the standard UNIX getty, login, rlogind and telnetd programs. phLOGIN is intended to be used in turnkey environments and/or where good access control and audit trails are important.

phLOGIN allows the systems manager to restrict login by user, group, time and day of week. He may also provide his own script for further tests (eg. that a particular user is not on holiday). These restrictions may be specified terminal by terminal. All login attempts, successful or otherwise, are logged.

Terminals may be locked globally (eg. for backups) or individually (eg. if a break-in attempt is detected on a particular line).

When told the terminal type phLOGIN acts in full screen mode. The date and time may be displayed and periodically updated as may the MOTD file contents.

The screen layout and user messages are completely customisable which gives good native language support and allows terminal specific escape sequences to be used to provide appealing visual effects.

The amount of information given to the user explaining why a login has failed may be closely tailored. This allows the system manager to choose the trade-off between security and "helpfulness".

The user's environment, home directory, terminal modes may be set up and the application program run directly --- avoiding a login shell and the inherent security flaws with that.

phLOGIN may be used to set the user's maximum file size (ulimit). This is particularly useful where database applications are being run.

Options permit immediate login to a specified user and application when the terminal is switched on or, alternatively, after the user has hit the return key. This autologin feature provides easy startup where a terminal has a dedicated use; if the terminal performs a display only function it does not even need a keyboard.

A failed login is followed by progressively longer pauses to discourage brute force break-in attempts.

phLOGIN is available on all versions of Xenix and UNIX, for direct terminals and networks: rlogin & telnet protocols.

UNIX management tools

Arnold Landzaat
C/S/E/ Consultants, Maarsen

Procurement of computer systems will be the first step for an organisation in the land of automation. The supplier will be often very helpful with the installation and after sales support. After a certain period the company will have other concerns. The automated vein of life can crash and cause other creepy cases, like undesirable visit of other people's data. Protection is playing a big role in every computing centre as well in an UNIX environment. The UNIX image of "everything must be possible" does not contribute to maturity. However UNIX can be secured perfectly by means of helpful tools, as long as we consider just average securing. This type of securing varies from login procedures, password check's up to backup systems and security auditing. This broad view on the company/computer security implies a well contemplation of directions of management/comply with and introduction of standards and procedures.

Of course backup must be based on a well considered philosophy. Essential will be the frequency of backups and administration of the tape library in a professional manner. In a network with many active UNIX based systems, it would not just be nice if all diverse UNIX systems (read versions) were backed up in the same manner, but save a lot of time as well.

Backup tools exist that allow intelligent tape-unit management. Automatic selection of the UNIX system with the largest tape capacity for the backup of the workstations, is an example of an application of such a system. Portability of backup tools regarding the absence of tar or cpio on for instance Digital VAX/VMS based systems can play an important role in a selection procedure of backup system.

An other aspect of the security issue will be misuse of data. Who will use what files and what commands and what data and commands may be accessed or not? Does a database need extra protection measurements and are all modes set well? Is it possible to trace commands that can cause serious damage to the continuity of the computer system? Do users need protection against their own creativity? Actually these are matters to be decided by the upper management. However one fact is established:

The system manager will be regarded to be responsible for the management of data, the ascription of privileges the continuity and protection of company confidential data to avoid misuse. With regard to the growth in complexity of networks and stressful circumstances because of lack of personnel the system manager can not live without helpful tools. One time investment of these tools offer a day to day assistance to the UNIX system manager.

The design of these tools enable computer management based on structured procedures where colleagues can assist in emergency cases rather then protected personal computer management.

CenDar

a CENtralized Data ARchiving system

Gertjan van Oosten
West Consulting BV, Delft

CenDar is a reliable data archiving system for use in a heterogeneous network with UNIX and VMS based computer systems, interconnected by an Ethernet local area network using the TCP/IP protocol suite.

The aim of CenDar is to provide the user with a reliable tool that enables him to archive sets of files (projects) to an archive device attached to one of the machines in the network. CenDar makes it also possible for the user to retrieve his files again in a user-friendly and simple way. Extensive report facilities are provided to ease the retrieval process. Both these functions can be performed without any knowledge about the archiving configuration.

The architecture of CenDar is based on the concept of the server-client model. The server has an archive device attached and is equipped with the database that is used to store the meta-data of the archive files. The majority of the CenDar activities are executed on this central server. It is called the CenDar node. The client computers attached to this server are called Archive nodes.

The current implementation of CenDar assumes the CenDar node to be a UNIX computer utilising a SCSI-based subsystem as the archive device. Both UNIX and VMS computers can be set-up as archive nodes.

Tutorials

Performance Analyses and UNIX

Ger Austen
ATComputing, Nijmegen

Email ger@atcmp.nl

AUDIENCE TARGET GROUP

This tutorial is intended for those who wants to know about the different methodologies to measure, predict and improve performance of UNIX based systems. The primary target group are experienced system managers.

PERFORMANCE ANALYSES AND TUNING

Performance analyses is the speciality where performance of a system (1) first will be measured, (2) then will be analyzed and understood, to lead (3) finally to improvement of the performance. For each of these subjects (1) to (3) do exist assisting tools and techniques. However these are not yet complete and standardised in UNIX.

TO MEASURE IS TO KNOW?

To measure performance is not easy. The subjective complaint "it is going so slowly" must be translated to an objective cause. UNIS SV and BSD offer a conglomeration of programs that are able to tell all a bit about performance in combination with application software. Many necessary measure tools are still missing.

TO KNOW IS TO MEASURE?

To know - after measuring - one should compose a picture of the functioning of the system and the application. A deeper understanding about how processes, kernel and peripherals function, is prerequisited supported by mathematical models that consolidate measurements and enable extrapolations to domains that are unable to measure now. The interpretation of all bits and pieces and the formation of mutual coherence approximates a classic piece of detective work.

TO KNOW IS TO IMPROVE!

After tracing the bottle neck it can be cleared away. At least if this is possible at all. Improvement of the capacity can be achieved by setting parameters in the kernel, by fastening the file system, by improving the applications, by controlling or moving processes and work load, or by extending or reconfiguring the hardware.

CONTENTS OF THE TUTORIAL.

This tutorial will cover mainly measurement and analyses techniques: measurements at system level, measurements at application level and measurements at user level. Further more the analyses of the measurements will be focussed: to learn to understand the measurements to enable future predicates of the performance of the system and to be able to improve the performance in a focussed way.

NIS and NFS

Sjoerd Mullender
CWI, Amsterdam

Email sjoerd@cwi.nl

The Network Information Service (NIS, formerly called Yellow Pages or YP) is a distributed lookup service, providing read access to a replicated database. The lookup service is provided by a set of NIS database servers, which communicate among themselves to keep their databases consistent. To setup the service, daemon processes have to be started and files created. Which processes must be started and which files must be created is described. The way in which new NIS servers can be added or new database maps can be created is also described.

The Network File System (NFS) provides nearly transparent access to file systems on other machines. In order to use NFS, several daemon processes must be started and files must be created. Which processes must be started and which files must be created is described. It is possible on many systems to use an automounter. How to setup the automounter configuration files is also described.

For both NIS and NFS, some of the details on how the services work are also given.

E-Mail for end-users and managers

Henk Hesselink
ACE Associated Computer experts B.V., Amsterdam

Email henk@ace.nl

This tutorial is for users and managers of electronic mail. Users are often confronted with the internal functioning of e-mail systems (strange addresses, undeliverable post, etc.). By means of handling a large number of aspects of e-mail superficially it will become clear to the attendees how to use e-mail in the best way, what can go wrong and what to do if it fails.

As long as time permits subjects to be covered will be:

- User agents (/bin/mail, Mail & mailx, MH, elm, mush etc.)
- Message Transfer Agents (sendmail, mmdf, pmdf, smail, etc.)
- the difference between a name, an address and a route
- domains and networks
- standards (RFSs, X.400)
- directory services (bind, X.500)
- address conventions
- address mapping, gateways
- mixed (non unique) addresses

Abstracts

From UKUUG conference at Liverpool on 15-17 July 1991.

Many thanks are due to Martin Beer <mdb@compsci.liverpool.ac.uk> who organised the conference.

The KornShell - Past, Present, and Future

David G. Korn
AT&T Bell Laboratories

Email dkg@ulysses.att.com

This talk will present a history of UNIX shells and the role that the shell has played in the UNIX environment. The KornShell, ksh, will be compared to its ancestors, the Bourne shell, the C-shell, as well as to some of its contemporaries such as tcsh, BASH, tcl and perl. A summary of the features of the current version of the KornShell will be summarised. Changes to the KornShell language that are required to allow the KornShell language to conform to the proposed IEEE POSIX 1003.2 shell and utilities standard will be presented.

For many users the command interpreter has been replaced by mouse driver graphical user interfaces which raises the question as to whether a shell is even necessary. The talk will focus on the role the shell is likely to play in the future and the direction that KornShell is likely to take in its next major revision.

Accessing Manual Pages more Easily: PELICAN, Natural Language Understanding and Information Retrieval.

Richard F.E. Sutcliffe
Department of Computer Science University of Exeter

Email richard@cs.exeter.ac.uk
Email richard@dcs.exeter.ac.uk

Unix is a very powerful operating system. However, its command interface is not as easy to use as one would like. One simple problem which assails beginners and experts alike is simply remembering which command is needed to perform some function. Of course one can use 'man -k' to scan the 'whatis' database but this often does not produce useful results for several reasons:

a) one must guess the right keywords, b) one is dependent on the accuracy and completeness of the descriptions which have been associated with the files, and c) one can often obtain vast numbers of commands in no particular order whose descriptions do not enable a suitable command to be selected.

What we present here is a system called PELICAN which aims to help in exactly this situation. PELICAN takes as input a query in English expressing an information need and produces as output a list of unix commands ordered by their relevance to the query. For example if you say 'I want to move a file from my directory' the system will output a list of commands with the command 'mv' at the top of the list. The user can then consult the 'man' entry for 'mv'.

Within Artificial Intelligence there has already been a lot of work which is directed towards helping people to use UNIX. For example the Unix Consultant (Wilensky et al 1988), OSCON (Guthrie, Mc Kevitt and Wilks), SINIX Consultant (Kempe 1987) and the work of Hong and McTear (1988) are all directed towards providing users with information about the operating system, including what commands to use in order to solve a specific task.

Graphical Interfaces in Education: A Case History

Tim Love
Dept of Engineering Cambridge University, Trumpington St, Cambridge CB2 1PZ

Email tpl@eng.cam.ac.uk

Cambridge University Engineering Department (CUED) has a yearly intake of about 350 undergraduates and 100 post graduates. Since 1984 we have been involved with Computer Based Teaching projects. Computers are widely used in the curriculum in teaching programming as well as in computer based practicals covering a wide range of engineering topics. User have varying degrees of computer literacy and interest in computing and may only use a package once or twice a year during laboratory exercises. In May 1989 we replaces an IBM mainframe and PCs with workstations. This paper reviews the transition from a command line interface to a WIMP-driven one under UNIX, and looks at how the Graphical Interface has affected working practises, teaching and program development. Though the conclusions are most

applicable to educational establishments, they could apply to any site undergoing the same transition.

CoOpLab: Computer Supported C-Operative Working by Evolution

T A Plant, M E Morris, P T Hughes, O R K Benson
BNR Europe Limited, London Road, Harlow, Essex, England

Email T.A.Plant@stl.stc.co.uk

The design and implementation of Computer Supported Co-Operative Working (CSCW) systems should consider and integrate both technological and user perspectives. CoOpLab is a prototype system developed to investigate such integration. This paper presents the aims, software architecture and initial implementation and results of CoOpLab. It also discusses the issues of co-operation and usability raised by the experimentation with CoOpLab.

MUCH: A UNIX-Based Hypertext System

Akmal Zeb, Antonios Michailidis Geeng-Neng You, Mahmoud Mhashi, and Roy Rada
Department of Computer Science University of Liverpool Liverpool L69 3BX, England

Email ajz@uk.ac.liv.cs

MUCH (Many Using and Creating Hypertext) is a collaborative hypertext system which is being developed on a network of graphical workstation in a UNIX environment. The first prototype of the MUCH system supported browsing hypertext documents like some other system (Emacs-Info, SuperBook etc.). The present prototype of the MUCH system consists of a front-end which consists of a number of Emacs Lisp programs, and a back-end (a process written in C, which communicates with Ingres database to retrieve and store information). The user interface is provided through an Emacs window, running under windowing system X. The Emacs window is further subdivided into multiple windows to display different information. Information in the MUCH system is represented as a network of concepts linked together, with the link pointing to the text block between two concepts. This multi-authoring system can be used to create and access hypertext. The MUCH system supports creation (of documents, annotation, notes), accessing (of annotations, text blocks associated with a document, notes, searching through documents via indices) and modification (of text blocks, annotations). Dynamic outline generation facility is provided which allows the user to view the outline in various ways. Documents in the MUCH system can be indexed manually and automatically. In each case, the user is provided with an alphabetically sorted list of terms which when selected gives the attached document(s) back to the user.

XTALK: An Experimental Talking Tool based on X Window DianGuo Zhao Information Technology Institute University of Salford Salford M5 4WT, England d.g.zhao@sysb.salford.ac.uk "Talking" in plain text with the unix "talk" command is no longer satisfactory for most users nowadays. One reason is that many applications are multi-media. In CAD/CAM, for example, a picture may be much more meaningful than a whole page of text. Another reason is that the talking environment has changed significantly. Window and icon base user interfaces, multi-media data, LANs, and high performance workstations have made plain text talking out of fashion. By talking here we mean two or more people(the talkers) on different workstations(or terminals) communicating both interactively and realtime. Interactively, because one talker must negotiate with other talkers from time to time. Realtime, because one talker's actions should be observed by others very quickly. A talking tool is a software that is used to start, control, and manage a talk. XTALK is an experimental talking tool based on the X window system. It provides an environment for two or more people to talk not only in text but anything that can appear in windows. It is intended to expose problems and explore ways for building a more powerful multi-media conferencing system. This paper presents the experimental work carried out during a XTALK prototype design and implementation. Different talking strategies are compared. Design and implementation issues of the XTALK are discussed. And finally exposed problems are investigated.

User Interfaces and Remote Working

N M Maclaren
University of Cambridge Computer Laboratory New Museums Site, Cambridge, UK

Email nmm1@phx.cam.ac.uk

Users working in one location often need to access systems at other locations, either nearby or far away; many experts believe that this use will increase over the next 10 years. Access to large shared systems can also be regarded as remote, because of the very long paths in the communications system. Modern user interfaces are often graphical or WIMP based (Window, Icon, Mouse and Pull-down menu), but these may be unavailable of unusable slow when using a remote system.

This paper describes the problems of using and designing modern interfaces (especially WIMP ones) for access to remote systems. It gives some examples of both good and bad design of user interfaces, and guidelines on how to design interfaces that are usable for remote working. It mentions how the popular windowing systems constrain good design, and indicate some ways that they could be improved. It explains why such redesign can also improve user-friendliness.

DOS & UNIX Interworking - The Challenge of the 90's

Steve Purdham
JSB Computer Systems Ltd Cheshire, UK

Email sp@jsb.co.uk

Which future computing strategy is the right strategy for the future?

If one was starting with a clean slate and possibly a blank cheque book the answer would be easy, or would it? For example of you were a UNIX guru or UNIX descendant the future would most likely be based upon a UNIX network utlising X window based applications and services, but if your background is DOS based then the solution is probably MS Windows 3 based PC networks with the LAN manager and suites of client/server applications utilising the latest in, say, SQL servers and/or RPCs (Remote Procedural Call) from Netwise or ultimately it could be even the NEXT Step.

No matter which is the ideal solution, once reality is brought into the equation there are a few restricting factors and hence the challenges ahead.

Historically computer users have as their desktop devices character terminals (eg DEC vt100 and Wyse50), during the 80s a migration to DOS based PCs has been in evidence and now the pressure is on the move towards Graphical Based desktop workstations utilising UNIX and X window. However the correct future computing strategy is one which can cope with the real world situation of Mixed environments and caters for interworking, integration and interoperability of the character terminals of yesteryear, the PCs of today and the workstations of tomorrow.

Wide Area IP in the UK Academic Community

Piet Brooks
University of Cambridge Laboratory, Cambridge, UK

Email pb@cl.cam.ac.uk

After many years of resisting IP on networks in the UK Academic Community, the Joint Network Team are running a pilot IP network (ShoeString) before a full service (Janet IP Service) is made available later this year. It is based on IP tunnels over X.25 calls to allow the existing JANET X.25 infrastructure to be used.

This paper is a brief introduction from the user's point of view to the progress being made, rather that a detailed technical description.

KBS Development using X windows: The MADE Development Methodology

Frans Coenen, Trevor Bench-Capon
Department of Computer Science, The University of Liverpool, England

Email frans@cs.liv.ac.uk

Mike Smith
International Computers Ltd Manchester, England

This paper describes the development of a Knowledge Based System (KBS) for British Coal using the Make Authoring and Development Environment (MADE). This environment is based on X windows and is designed to take the developers from a set of textual sources though to a final application. MADE is founded on a KBS development tool called KANT (Knowledge engineers assistANT). The advantages of MADE are firstly that it is self contained in that the development environment is also the final application environment. Secondly that the philosophy behind the methodology is such as to encourage the production of KBS which are readily maintainable through an isomorphic approach to knowledge analysis and KB development and the implementation of a suite of maintenance tools. The intention of isomorphism is that the Rule Base should faithfully reflect the sources from which it was derived. The development environment is designed to be implemented on SUN work stations operating under UNIX. The system consists of three Base Windows referred to as the KANT, MADE and MAPPE (Make APPlications Environment) windows. In the KANT Base Window the knowledge analysis part of the development takes place which results in an intermediate representation of the knowledge to be displayed by the system. This representation is then compile and will form the final executable application. The compiled MIR can be interacted with though the MADE Base Window, where various verification and validation support tools can implemented on the application. The MAPPE Base Windown provides the end user view.

You Can't Just Change the Toolkit (so try changing the tool)

Alistair George
Imperial Software Technology Ltd Reading, UK

Email apg@ist.co.uk

Graphical User Interfaces (GUI) for applications help users improve their productivity and reduce the learning time for new applications, but the complexity of the technology makes application development more expensive. Worse, it may be necessary to produce several implementations to support different look and feel standards.

A Look and Feel Independent Applications Programming Interface (LAFIAPI) is intended to address this problem, by providing a single toolkit with multiple implementations for different look and feel standards. An application written using the ASFIAPI can be converted to a different look and feel simply by compiling and linking it with the appropriate implementation of the LAFIAPI toolkit.

This paper discusses some of the problems of implementing a LAFIAPI, and the limitations of the approach. The conclusion is the developers of high quality product applications will gain only limited benefit from a LAFIAPI. As an alternative, we briefly describe a tool-based approach to supporting multiple look and feel standards, based on the Motif GUI builder tool X-Designer.

phLOGIN Product Overview

Alain Williams

Parliament Hill Computers Ltd, Berkshire, U.K.

Email addw@phcomp.co.uk

The paper will discuss the history, aims and design philosophy for the program. An overview of the major program features and how they are to be used will be given.

Porting issues on a utility like this are unusual and will also be covered. An indication of where further work is needed.

A Look at Open Look

David Keeffe
Sun Microsystems Ltd, Camberley, UK

Email David.Keefee@sun.co.uk

An overview of the Open Look Graphical User Interface is presented. Its philosophy, appearance to the user and specific components are described. The main Open Look toolkit products are described, along with short code examples.

Objects in the Office: The benefits for Computer Users

Chris Marshall
Hewlett Packard Ltd, Wokingham, UK

This paper describes the HP NewWave object oriented GUI.

The Object in HyperNeWS

Tim Niblett
The Turing Institute and University of Strathclyde, Scotland

Email tim@turing.ac.uk

HyperNeWS is an object-oriented User Interface Design Environment, which is designed to allow anyone, from novice to seasoned professional, to develop a very wide range of user interfaces. It uses a completely different model from the toolkit-based approaches of Xtk, Xview, Motif etc.

Central to the fulfillment of the design goals is the notion of object used in HyperNeWS. This notion is explained and justified in the context of the HyperNeWS2.0 system, which is under development.

The X11/NeWS Server

David Keeffe
Sun Microsystems Ltd, Camberley, UK

Email David.Keeffe@sun.co.uk

An overview of Sun's X11/NeWS server is presented. X11/NeWS is a single technology which merges the functionality of X11 and NeWS. Advantage is taken of shared attributes of X11 and NeWS windows, and a union of facilities is provided otherwise. Common usage of windows by NeWS and X11 protocols within a single program is discussed.

A Workstation-Based Environment for Digital Systems Design and Programming

C Charlton et al
Department of Computer Science, University of Liverpool, UK

The MATRESS project aims to provide a fully integrated software environment to support all aspects of the design, analysis and microprogramming of digital systems. Its implementation is based round an object-oriented model of design information, which is used to represent system design as it proceeds from inital specification throughout a process of hierarchical resolution. The system supports simulation of the design in progress, and interaction with simulation to verify behaviour and develop firmware. We describe here the current implementation of the MATTRESS system on UNIX workstations, using X windowns, and work in progress and planned towards its long-term aims.

Business Shell (BSH)

T Dawes-Gamble,
Information Technology Group
Texas Instruments Ltd. M/S TIL 47 BEDFORD MK41 7PA

Email tmdg@ti.co.uk

Business Shell is a screen orented shell designed to run on a standard character based screen.

The Business Shell (BSH) is compatible with CSH(1) but has two additional built in commands and a number of extra shell and environment variables that effect the behaviour of the shell.

The additional built in commands are menu and prompt. The additional shell variables are menu, cmd and noexclam. There is one additional environment variable MLC.

These additions allow the shell to present the user with a menu that provides a more friendly interface than the more traditional shells. Single key stokes can be used to get help about a menu or command presented by the menu. Programs can be started by a single key stroke also. The menu is displayed as three columns: name, type and brief description, There is a forth non displayed field that is the pathname of the help file. The type column indicates if the item is a command, sub menu, or additional help file.

Menus have two modes pick and point or command line. When in pick and point mode only those programs presented on the menu can be invoked. Command line mode allows any program in the path defined in the MLC environment variable to be invoked by typing the program name. The menu may be toggled between these two mode using the command key. If it is required to start the menu in command line mode then the cmd variable may be set.

Evolutionary Development of Commercial Applications

Brian Thompstone
BCT Business Systems Ltd, Lymm, Cheshire, UK

This paper describes a general approach to, and a particular method of, evolving products into new applications. It includes details of the use of the Cthru product for evolutionary user interface enhancement.

Incorporating Multimedia in the ANSA Distributed Architecture

Gordon Blair et al
Distributed Multimedia Research Group, Computing Dept, Lancaster University, UK

Email mpg@comp.lancs.ac.uk

No Abstract

Developing the User Interface for Headed Record Expertext

Martin Beer and Dan Diaper
Department of Computer Science, University of Liverpooll, United Kingdom

Email mdb@compsci.liverpool.ac.uk
Email diaper@compsci.liverpool.ac.uk

Headed Records Expertext (HRExpertext), is a new type of document management system that possess both a novel architecture and novel functionality. Architecturally, Headed Records Expertext combines some aspects of both expert systems and hypertext. The novel functionality arises because Headed Records Expertext is predicted to be capable of supporting desirable, large scale, multi-author document systems and will be able to provide, most importantly, a feasibly maintainable system. This paper briefly describes the architecture of a prototype system and a few of its applications. The paper then argues that the long term success of such a system will in part depend on the Human-Computer Interaction effort devoted to the system. Finally, some of the problems associated with designing a user interface, as an example of such Human-Computer Interaction issues, to a system with novel functionality are discussed.

Algeria - AIUUG
Nour Eddine Oulmi
4 Rue Beaufort
Alger 16000
ALGERIA
Telephone +213 2 744112
Facsimile +213 2 638775

Austria - UUGA
Friedrich Kofler
Schottenring 33/Hof
A-1010 Wien
AUSTRIA
Telephone +43 222 34 61 84
Facsimile +43 222 310 44 62
Email uuga@tuvie.at

Belgium - BUUG
Edgard Nyssen
BUUG Secretariat
VUB-FAC. T.W.ETRO/IRIS
Pleinlann, 2
B-1050 Brussels
BELGIUM
Telephone +32 2 641 2930
Facsimile +32 2 641 2883
Email buug@etro.vub.ac.be

Czechoslovakia - CSUUG
Sekretariat CSUUG
Výpočetni Centrum VSE
Nàm. A. Zápotockého 4
130 67 PRAHA 3
CZECHOSLOVAKIA
Email csuug@Czechoslovakia.EU.net

Denmark - DKUUG
Inge & Mogens Buhelt
Kabbelejevej 27 B
DK-2700 Brønshøj
DENMARK
Telephone +45 31 60 66 80
Facsimile +45 31 60 66 80
Email sek@dkuug.dk

Finland - FUUG
Outi Nyman
Finish UNIX Users' Group
Perustie 23 A 11
00330 Helsinki
FINLAND
Telephone +358 0 489 016
Email fuug@fuug.fi

France - AFUU
Miss Ann Garnery
AFUU
11 rue Carnot
94270 Le Kremlin Bicetre
Paris
FRANCE
Telephone +33 1 46 70 95 90
Facsimile +33 1 46 58 94 20
Email anne@afuu.fr

Germany - GUUG
Dr Anton Gerold
GUUG-Vorstand
Elsenheimerstr 43
D-8000 MÜNCHEN 21
WEST GERMANY
Telephone +49 89 570 7697
Facsimile +49 89 570 7607
Email gerold@lan.informatik.tu-muenchem.dbp.de

Hungary - HUUG
Dr Elôd Knuth
Computer and Automation Institute
Hungarian Academy of Sciences
H-1502 Budapest 112, P.O. Box 63
HUNGARY
Telephone +361 1 665 435
Facsimile +361 1 667 503

Iceland - ICEUUG
Magnús Gíslason
University Computer Center
Dunhaga 5
IS-107 Reykjavik
ICELAND
Telephone +354 1 694750
Facsimile +354 1 28801
iceuug-s@rhi.hi.is

Ireland - IUUG
Dr T G Murphy
School of Mathematics
Trinity college
Dublin 2
IRELAND
Telephone +353 1 772941
Facsimile +353 1 772694

Italy - i2u
Ing Carlo Mortarino
i2u
Viale Monza 347
20126 Milano
ITALY
Telephone +39 2 2520 2530

Netherlands - NLUUG
Sandra Scharff
p/a Transmediair BV
PO Box 297
3720 AG Bilthoven
THE NETHERLANDS
Telephone +31 30 291718/281820
Facsimile +31 30 292294
Email nluug-buro@nluug.nl

Norway - NUUG
c/o Balder programvare AS
Sentrumsbygget
PO Box 1344
N-1401 SKI
NORWAY
Telephone +47 9 870550
Facsimile +47 9 876766

Portugal - PUUG
Legatheaux Martins
Avenue 24 de Julho, n° 134, 7°
Lisboa
PORTUGAL
Telephone +351 1 395 06 42
Facsimile +351 1 67 18 76
Email sec@puug.pt

Spain - UUES
Jose A Manas
Usarios de UNIX de España (UUES)
Telefonica I&D
C/Emilio Vargas 6
28043 Madrid
SPAIN
Telephone +34 1 337 42 22
Email jmanas@dit.upm.es

Sweden - EUUG-S
Hans E. Johansson
NCR Svenska AB
Box 1206
S-164 28 KISTA
SWEDEN
Telephone +46 8 750 26 03
Email hans@ncr.se

Switzerland - CHUUG
Patrik Eschle
Viktoriastr 19
CH 8051 ZÜRICH
SWITZERLAND
Telephone +41 1 313 1326
Facsimile +41 1 312 4977
Email eschle@forty2.uucp

Tunisia - TNUUG
Mondher Makni
IRSIT BP 212
2 Rue Ibn Nadime
1082 Cite-Mahrajene
TUNISIA
Telephone +216 1 787 757
Fascimile +216 1 787 827

United Kingdom - UKUUG
Bill Barrett
Owles Hall
Buntingford
Hertfordshire SG9 9PL
UNITED KINGDOM
Telephone +44 763 73475
Facsimile +44 763 73255
Email ukuug@ukc.ac.uk

USSR - SUUG
Dr. Kuznetsov
Kosinskaja str, 16-3-65
Moscow 111538
USSR
Telephone +7 95 374 7049

Yugoslavia - YUUG
Milan Palian
Parex Institute
Kardeljeva No 8
YU-61000 Ljubljana
YUGOSLAVIA
Telephone +38 61 223463
Facsimile +38 61 214223
Email milan@parex.yu

EurOpen Secretariat

Owles Hall
Buntingford
Hertfordshire SG9 9PL
United Kingdom

Telephone +44 763 73039
Facsimile +44 763 73255
Email europen@EU.net