CSIRO Victoria Network Meeting 18-May-1989

AGENDA

| | | |
|---|---|---|
| 1. | Welcome | Ron Kerpen |
| 2. | Introduction/explanation. | Bob Smart |
| 3. | Background on various inter-related networking initiatives. | " |
| 4. | AARNet, Vic region | John Mann |
| 5. | Admin networking | David McCulloch |
| 6. | Library networking | Bob Smart (?) |
| 7. | Whip around | everybody |
| 8. | General discussion | " |

9. Should we arrange a bigger Vic CSIRO Network meeting for a wider audience with better publicity?

Robert Smart
CSIRO Div of Information Technology
55 Barry St
Carlton
Vic. 3083

May 16, 1989

CSIRO Management

Dear Sir/Madam,

Australian Universities and Colleges are in the process of establishing a Computer network: AARNet (Australian Academic Research Network). I believe an offer has or will soon be made to CSIRO to become a part owner and foundation member of the network. This letter is to present some arguments in favour of accepting that offer.

I am currently on the CSIRO Supercomputer Facilities Task Force. My role on that committee is to ensure that adequate networking infrastructure is provided to give CSIRO researchers convenient access to the supercomputers. It would be quite impossible for the allocated supercomputer money to provide links into every site or even to every state. The Task Force could only allocate enough money to provide a backbone network linking the major east coast capitals, and that was only possible because there is spare capacity available on CSIRO's CSIROTEL telephone network. The task force feels that membership of the Australian Academic Research Network is the ideal way for CSIRO to provide high speed and convenient supercomputer access to CSIRO researchers in other states, and to many of the researchers away from the capitals in all states. To provide our own 48Kb lines to Hobart, Adelaide and Perth would cost over $100,000 per year. However this is only a part of the extra costs which would be borne by individual sites if CSIRO is not a member of AARNet. I have already had queries from the CSIRO Marine Lab in Hobart about how they will obtain access to the supercomputers.

CSIRO membership of AARNet will cover the cost of CSIRO access to the link to the American research network, the Internet. The Supercomputer Facilities Task Force sees this link as a very important part of CSIRO's Supercomputer plans. It will allow exchange of information and programs with American computer centres having similar supercomputers. It will enable CSIRO Researchers to get convenient access to other types of supercomputers which may be more suitable for the solution of particular problems. The link will also have many important uses unrelated to Supercomputers. It will allow convenient exchange of information, programs, data files and preprints of papers with colleagues in America. The link will make life much easier for visiting scientists, letting them access their computers at home. This should make it easier to attract visiting scientists. It will also make life easier for CSIRO researchers who are visiting overseas. The link will also allow convenient communication with other Research networks attached to the Internet, and this includes New Zealand, Scandinavia and other parts of Europe already. If we don't join AARNet we will have to make our own arrangement to attach to the overseas link. I have a ball-park quote from the University of Melbourne of $40,000 per year to connect — this is based on current expenditure of $36,000 per year by CSIRO for overseas mail alone through the current overseas gateway run by the University.

CSIRO membership of AARNet will make it much easier and cheaper for CSIRO sites around Australia to communicate with each other. For example some leased lines, such as the one between Division of Manufacturing Technology sites in Melbourne and Adelaide, will become unnecessary. I know that the Division of Soil Technology is interested in links between their sites in Townsville, Canberra and Adelaide.

The CSIRO Library system in East Melbourne is acquiring new communications equipment that will enable it to be accessed very conveniently from the CSIRO TCP/IP network which will be established for Supercomputer access. The possibility of an improvement in network access to the CSIRO Library system has been greeted with considerable enthusiasm by Library people I have spoken to. This method of access will be much more widely available if CSIRO joins AARNet. This holds the possibility of considerable savings in the cost of CSIRONET X.25 micronodes.

Finally membership will promote cooperation between CSIRO and the Australian Academic world, and also with DSTO which will be connected to the AARNet even if it isn't a member.

Yours sincerely,

Robert Smart

DRAFT

# The CSIRO IP Network

Robert Smart

May 16, 1989

## 1 Introduction

Access to CSIRO Supercomputers will be based on TCP/IP. This decision was made by CSIRO Management and is a constraint on the Supercomputer Facilities Task Force. TCP/IP is the industry standard for scientific workstations and it is the current standard for connecting workstations to supercomputers, particularly in the USA where most relevant development is taking place. TCP/IP is also a good way to connect any classes of machines together, being implemented on everything from PCs, Macs and NGENs up to supercomputers. So I and the rest of the committee completely support the decision to base the networking on TCP/IP.

The objective of this document is to fill you in on details of the CSIRO-IP network, how it will work, how to connect to it, and what facilities it will provide in addition to access to Supercomputers. This document is a snapshot of current thinking. The details may vary slightly. While it is intended to encourage you to connect, this document will also present alternatives to direct connection.

## 2 What is TCP/IP

TCP/IP is a suite of network protocols developed initially by the American Department of Defence for their experimental research network. The first DoD networking was a single wide-area network (the Arpanet). In the late 70s local area networks started to spring up. The American researchers realized that wide area networking should link LANs together, not individual computers. The basis for this new networking software was the Internet Protocol (the IP in TCP/IP). This protocol defines how packets of data can be routed from network to network until they reach their destination. A network in this sense can consist of many entities linked together (e.g. the Arpanet, or any ethernet) or it can be just a single wire connecting two ends. The whole mess of interconnected networks is called The Internet.

The thing that makes IP work is that it is very easy to implement. Computers or specialized boxes can act as routers. They don't have to deliver packets in the order they receive them. They can break big packets up (in a carefully defined way) if they have to move the data over a network that can't handle big packets. They can throw packets away if they get overloaded. The ability to deliver packets out of order means that an IP network has no trouble using two communication lines joining the same points to double throughput.

Because IP does so little, the higher level protocols (normally TCP) have to handle the problems of retransmission of lost packets, reordering out of order packets, defragmenting fragmented packets. This turns out to be just the right division of labour between what is now called, in OSI parlance, the network and the transport layers. Note that TCP/IP is a bit of a misnomer: other protocols can work over IP, and in particular NFS uses UDP. TCP works hard because it guarantees a reliable data stream, but UDP doesn't do anything much, allowing the higher level to handle lost packets, etc. TCP is used in most IP applications, justifying the TCP/IP name, but the CSIRO (inter)network is just called CSIRO-IP combining brevity with accuracy.

The standard protocols that come with most TCP/IP software are FTP file transfer, TELNET remote login and SMTP for electronic mail. These are all designed to work between unlike computer systems. In a uniform BSD unix environment the r-series commands (rlogin, rcp, rsh, etc) are more convenient. Other major applications include NFS which allows you to access remote disks as if they were local and X-windows a distributed windowing system which separates the computational and display sides of windowed computer access. There are also well defined ways for running appletalk, OSI and even DECNET over an IP network, though the mechanisms are not sanctioned by the bodies defining those protocols. ACSnet connections can be made over an IP network. This is in wide use already and will be standardized in the next major release of the ACSnet software from The University of Sydney.

More details of what TCP/IP software is available is contained in Appendix A.

## 3 The Planned IP Infrastructure

The supercomputer funding can only provide a backbone for the TCP/IP network. Sites not fortuitously on this backbone will need to provide their own links to the backbone. The backbone will use the CSIROTEL telephone network to link Sydney, Melbourne and Canberra, plus a 48K DDS link to Brisbane. This connects the supercomputers themselves plus a link to the only significant remote group of current supercomputer users (in Brisbane). To provide access to other states, and to some remote places in the eastern states, this document assumes that CSIRO will be a part of the Australian Academic Research Network, AARNet. This seems likely, but if it fails this document will need a substantial rewrite. More information on AARNet is given later in section 5.3.

The following map assumes that the Melbourne supercomputer will be located in Clayton. If not an extra high speed link will be needed from Clayton to the ultimate location.

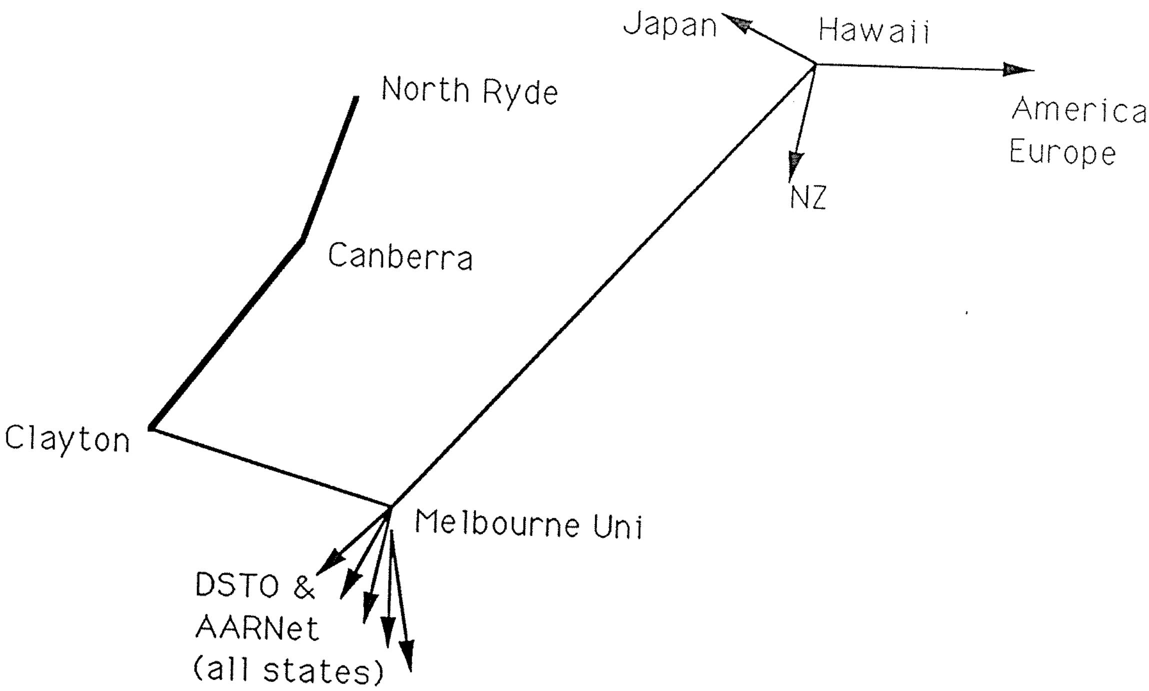


There will be a range of options for connecting to the CSIRO-IP network, from high speed links suitable for high speed graphics and remote file system access down to very cheap low speed IP links. It will be possible for users with no TCP/IP computers to gain asynchronous terminal access (transparent not line mode) to computers on the network. Options for connecting are described in more detail in a later section.

The cost of terminal servers and other equipment to make the IP network more useful must be borne by the Divisions using that equipment. Connecting sites need to pay for the leased lines and communications hardware and software they use. They also have to pay a small amount for connecting to cover the cost of supporting them at the backbone site. However the Supercomputer Task Force recognises that there is a chicken and egg problem, and so to discharge its duty to make allowance for adequate networking it will set aside from supercomputer funds some money to kick-start those networking activities which will ultimately be financed on a user-pays basis.

## 4 The Overseas Link: A Vital Component

A 56Kb TCP/IP link will be installed in late May to link Australian researchers to the American research TCP/IP network (the Internet). This link will go in to Melbourne University which has been the traditional gateway between ACSnet and the overseas networks. The Supercomputer Task Force agrees that the connection of CSIRO's IP network to Melbourne University and thus to that link is a high priority.

Connection to the Internet will make it possible for CSIRO supercomputer users and experts to collaborate closely with supercomputer users and experts in America (and probably beyond). They will be able to exchange code and information about bugs and techniques, and do so with a turn-around time in minutes.

Connection to the Internet will allow CSIRO researchers to access supercomputers in America that they are authorized to use (they may have to pay to do so). This will allow access to machines which are faster, or more specialized for the particular problem, or for which existing programs exist for the particular problem, or which have large current data files of relevant information not available here.

It will also allow all CSIRO (not just supercomputer users) access to many facilities: electronic mail (and other forms of information exchange) in minutes to or from overseas colleagues; access to many on-line databases without additional OTC charges; the ability to exchange files (programs, data, drafts of scientific papers, etc.) easily with colleagues overseas. Many Divisions already get substantial bills for overseas mail through the current gateway. These bills will disappear since access to the new international link will be paid for as part of CSIRO's AARNet membership.

## 5 Options for Connecting to The CSIRO-IP Network

There will be essentially two ways to connect to CSIRO-IP. The simplest will allow terminal access in a number of ways. This will allow file transfer with Kermit and similar mechanisms. This will be possible with leased lines or dial-up to terminal servers or computers on the network. It may also be possible via

X.29 PAD software from Austpac, from Csironet X.25 or from Admin's private X.25 network.

A much more desirable way to connect is to create a local IP network and connect that, directly or indirectly, to CSIRO-IP. Many sites already have an ethernet network which can carry TCP/IP traffic. If you have no local area network you can often connect an isolated computer to the wide area IP network without installing a local ethernet: but most sites will find a local ethernet a cheap and useful asset. You then need to set up an IP router from your ethernet to the outside world, and configure the software correctly on your local network. Taking the second first:

## 5.1 Setting up Your Local IP Network

IP networks have three distinct things to get right: names, addresses and electronic mail.

### 5.1.1 Names

For a brief discussion of the domain name system see Appendix C. Each machine on your network will have a full name which is globally unique. CSIRO sites have a choice of using a sub-domain of **oz** (the ACSnet domain) or a sub-domain of the **csiro** domain. Both **oz** and **csiro** are sub-domains of the Australian domain **au**. So typical computer names are:

**ateles.rp.csiro.au** The machine ateles at Radio Physics.

**mimir.dmt.oz.au** The machine mimir in Manufacturing Technology.

The argument for using **oz** is for compatability with current ACSnet software: SunIII (Sun = Sydney University Network). However to create a valid domain name acceptable in the Internet you have to add the **.au**. Since SunIII cannot cope with using the names in that form (with the **.au**), and since SunIV will cope just as well with the **csiro.au** domain as with **oz.au**, it seems sensible to choose to use the **csiro.au** domain for the full Internet host names of your machines.

To use the **csiro.au** domain you should register the subdomain you wish to create with the administrator of the **csiro.au** domain. This is currently Dr Trevor Hales of the Division of Information Technology, though he will probably be happy to pass this administrative chore to the network manager of the CSIRO-IP network when that position is filled. The subdomain should be a Division name which is recognizable to the world, and not too long. Current examples include **wool.csiro.au**, **rp.csiro.au** and **dit.csiro.au**. CSIRO components which are not Divisions will also fit into this scheme and we can hope to see names like **iict.csiro.au**, **bipc.csiro.au**, **library.csiro.au** and **hq.csiro.au**. I have suggested that you place computer host names immediately below (i.e. to the left of) your allocated subdomain giving a four component name. Normally that's enough: that's why I fought for a **csiro** domain within **.au** without which we would be in **csiro.gov.au**. However you can create more levels if you like: DIT decided to use site subdomains of **dit.csiro.au** with hosts within those. Multi-sited don't need to do this. You can avoid machine name clashes at distinct sites, and make it easy to work out where a machine is from its name, by using a different category of names at different sites. For example DIT uses star names at our Canberra site and fish names at our Melbourne site (though one day we might have trouble remembering why wanda is a fish name).

Of course you will define local nicknames for your machines so that, for example, the machines named above will be known at their local sites as just ateles and mimir. Remote users will have to give the full name to access those machines unless they have set up their own nicknames for those machines.

One of the tasks of looking after your local TCP/IP network is to maintain and distribute files mapping between host name and IP address. In BSD unix this file is /etc/hosts. When you join CSIRO-IP the management of that network will have the job of coordinating this information and distributing back combined comprehensive files like /etc/hosts. This situation will be almost instantly unwieldy and will soon be replaced by a distributed directory service whose components are name servers. At that time local sites will only have to maintain this information for their Division on one machine (maybe two for backup) and they will get access to name-address translation information from the whole Internet. Currently there is still a lot of TCP/IP software which doesn't understand name servers, but that is rapidly changing.

### 5.1.2 Addresses

More detail on how IP addresses are defined and used for routing is in Appendix D. Here are some points to watch.

When you get computers they typically come with example IP numbers. You can use these as long as you are not connected to anyone else. Before you can be connected to CSIRO-IP you must get official IP network numbers for the networks you wish to connect. The easy way to do this is to send electronic mail to hostmasterkl.sri.com. If you don't have electronic mail access then send an ordinary letter to: Hostmaster, SRI, ???, USA.

Remember that IP addresses are not really for the host, it is for the host-interface pair. So a machine with interfaces to multiple networks (usually acting as a router) will have separate IP addresses on each network.

Should you use several Class C network numbers for your networks? Or get a Class B network number and subnet it? The object of subnetting is to make your network seem simpler to the outside world. Subnetting lets a network composed of multiple little networks appear to be one network as far as the outside world is concerned. Routers outside your Class B net will not know anything about its internal structure. So, for example, if you have a Class B network number and put some subnets of it in Sydney and some in Melbourne then you will never be able to move packets between the two subnets over CSIRO-IP or over AARNet. A subnetted Class B network should be a physically compact group of networks, normally with one point of contact with the outside world. Subnetting is particularly convenient when you need to create little (sub)networks from time to time because of small networks like Appletalk segments linked to a Multigate.

Remember that IP network numbers belong to a network (e.g. an ethernet) and not to an organization. So if an ethernet has hosts from different Divisions attached, then they will have different domain names reflecting their different logical relationship, but they will all be within the same IP network (or subnet if you have a class B subnet). Some TCP/IP software will let you put different (sub)networks on the same ethernet but it is confusing, requires more management and is inconsistent with the way IP addressing is meant to work.

### 5.1.3 Electronic Mail

Configuring electronic mail systems is not always easy. This should become much easier on a IP network which is connected to the Internet. The management of the network, or of AARNnet, should provide standard configuration files which will work almost everywhere using modern mail software which queries name servers to find out how to get mail to other sites.

A word of advice: don't try to configure your mail software to give minimum mail addresses. Some people think that if you get mail from me on the same machine you should just see:

```
From:   smart
```

and if you are on the same network you should see:

```
From:   smart@wanda
```

and only remote users should see:

```
From:   smart@wanda.cng.dit.csiro.au
```

But you will find things work much better if the last of these is what everyone sees. It is the only format that actually complies with the RFC822 standard for mail messages.

## 5.2 Connecting up Your Local Network to the World

The first question is: who to connect to? There are three possibilities:

- Attach to the CSIRO-IP backbone. This is the most desirable, particularly for supercomputer access. Unfortunately apart from paying Telecom long distance charges this option is limited to Brisbane, Sydney, Canberra and Melbourne.
- Attach to the Australian Academic Research Network (AARNet). This will be connected to CSIRO-IP by a 48Kb line, giving good connectivity to all state capitals, and many remote centres which have significant Academic establishments.
- Attach to any other site that is already connected. This is undesirable because it introduces extra hops that packets have to make before reaching their destination. It *will* be a good idea for remote sites who are close to each other and wish to share the cost of Telecom leased lines. It is also sensible for those who already need a high speed link to a nearby site. Examples of this include those CSIRO Labs which are close to a University and have an ethernet connection to it: they will be automatically connected to AARNet.

Sites can connect to the CSIRO-IP backbone by buying a baby IP router box (about $14K) with a medium to high speed connection. Sites wishing to economize can use a host computer (e.g. Sun workstation) with suitable high speed serial board as a router and connect via that to an IP Router box on the backbone.

Using host computers as IP routers is not really recommended for high speed or important links. A host computer has other jobs which are likely to interfere with its role as a router (or vice versa). Also host computers have other things that can go wrong with them, require periods of stand alone use for software or hardware upgrades or fixes, etc. An independent network box is a better idea if you can afford. For a low speed link using a host as a router is a reasonable way to economize, and the amount of traffic means that it should not adversely affect (or be affected by) the hosts other duties.

A slower and cheaper connection can be done using the SLIP (Serial Line IP) protocol, or a variant. You can use a host computer or terminal server on your ethernet to connect via a serial line to a host computer or terminal server at the nearby backbone site. Currently Annex terminal servers and many sorts of computers (including Suns, PCs and NGENs) can be used for this. You can actually connect to any other nearby site on the network, but not connecting to the backbone will make packets take extra hops thus reducing performance.

## 5.3 Network Context: AARNet

The Australian Universities and Colleges are planning to set up a network around Australia based on 48Kb lines, and including support for TCP/IP.

The network is (perhaps tentatively) called AARNet for Australian Academic Research Network. It is intended to encompass Research organizations such as CSIRO. CSIRO has Representatives on the Management Committee (John O'Callaghan, Chief of DIT) and on the Technical Committee (Trevor Hales of DIT's Computing Network Group).

If CSIRO joins AARNet it will allow CSIRO to use the AARNet to connect it's research computers to each other, to CSIRO's and other supercomputers, and to promote cooperation between CSIRO researchers and colleagues overseas and in academia in Australia. In particular it will provide a convenient way for CSIRO researchers in Western Australia, South Australia and Tasmania to access the supercomputers and other facilities on the proposed CSIRO IP backbone. It will also provide convenient communications to remote CSIRO sites which are close to Academic sites, such as the Division of Soils in Townsville.

Even if CSIRO doesn't join AARNet as a whole, it will still be possible for individual Divisions to join. In this case they will be charged on a "cost recovery plus" basis, but remote sites which are close to Universities or IT/CAEs may still find it worthwhile to join. Even if CSIRO is not a part of AARNet, CSIRO-IP will still be connected to it through the 48Kb link to The University of Melbourne. So individual Divisions which join AARNet will have good access to the CSIRO-IP network.

# A Available IP Software

There aren't many computers that don't support TCP/IP. For those not mentioned here talk to the manufacturer.

## A.1 Unix

TCP/IP networking comes bundled with Berkely Unix (BSD 4.2 and 4.3, with BSD 4.4 due out within a year). It does not come with the current variants of System V, but most System V systems come with "BSD enhancements" which include TCP/IP. In future System V release 4 and beyond will include TCP/IP networking along with most other Berkeley features. Until then The Wollongong Group sells TCP/IP for System V unix.

There are a number of additional utilities and protocols which come with BSD-style TCP/IP, but not normally with others. These include the LPD line printer protocol, and the r-series commands: rlogin, rcp, rsh, rexec, rwho. The LPD protocol makes no pretense at portability: it just transfers data files and a BSD-style queue entry file as-is, but still it has been implemented on some other systems. Some of the r-series commands depend on a protection scheme in BSD TCP/IP: programs can't use connection magic numbers below 1024 unless they are run from (or setuid) root. This means that you can rlogin to a machine which trusts your machine without supplying a password (similarly with rcp for copying and rsh for remote execution without having to give a password for the remote machine). The point to be made about all this is: don't tell your machine to trust another unless it is a BSD machine (and be careful of implementations of r-series commands on non-BSD machines: they may be a security problem). [1]

## A.2 VAX/VMS

There are *lots* of implementations of TCP/IP for VMS. Many of them come with mail software of dubious quality. I recommend using the PMDF mail software. It is available for $100 from University Computing Services, The University of Melbourne (attn Chris Chaundy).

Carnegie-Mellon University has modified a version written at Tektronics. This CMU-TEK implementation is available for the cost of distribution. There is already a version licensed to CSIRO, and you can get a copy of that from Bob Smart in DIT (phone 03 347 8644). This software is not as well documented and may not be as robust as commercial products, but if your usage is low or you want to keep your money for other things it is certainly quite a usable product. It comes with source so you can fix things if you can speak Bliss. A particularly useful feature is the support for the LPD protocol which makes it possible to share printers between VMS and unix machines.

There are implementations from Fusion and from Excelan. These have had mixed reviews and may not be easily available in Australia.

The Wollongong Group WIN/VX implementation is basically just a straight port of BSD TCP/IP using the Eunice tools. It is probably a good choice if you have Eunice, or if you don't commands which are unixy rather than vmsish. It claims a server version of NFS, but this is not closely integrated with the file system, so may be inefficient. This is available through Techway in Australia.

The Multinet product is another port of the BSD network software to VMS by David Kashtan and Kenneth Adelman. It used to be from SRI, but is now for a company TVG (not to be confused with TWG). This is in many ways the most sophisticated product. It includes support for some other minor protocols, most notably XNS. They are working on interfacing it to DECwindows (so that you can interface it with other X-windows software). They also have a kernel mode NFS server and are working on an NFS client, though it requires a leap of faith to believe that anyone other than DEC can do the interface to the file system correctly. They support the LPD protocol, r-series commands, and their

[1] The mechanisms for telling a BSD machine to trust another are /etc/hosts.equiv and .rhosts files in user home directories. Look out for + characters in hosts.equiv files shipped with new machines.

neatest trick is that they support DECNET over TCP/IP. They don't have an Australian distributor, and I have been unable to get prices. I believe they are expensive in the US.

Last but not least, DEC now have their own product. This is called UCX which stands very misleadingly for the Ultrix Connection. The initial version only supports FTP, X-windows and a funny sort of NFS server. They are working on telnet. They will probably implement mail as a Message Router interface. Don't consider that, there will be a PMDF interface soon. There is a rumour that UCX will be bundled with VMS 5.2 when that comes out, probably some time in 1990.

## A.3 Macintosh

Macs have their own networking protocol, Appletalk. This runs over the traditional twisted pair cabling and over ethernet, called localtalk and ethertalk respectively. There are at least three products which act as a gateway between a TCP/IP network (running on ethernet) and an appletalk network (localtalk or ethertalk or both):

**Kinetics Fastpath** The original.

**Webster Multigate** Designed at Melbourne University, it allows for four localtalk network plus async/dialup Appletalk. The software is more flexible than the Fastpath.

**Cayman Gatorbox** Less flexible than the others, it has one great selling point: it supports an AUFS to NFS gateway. This means you can access your unix files from the Mac without running any special software on either.

These allow the Macs with only localtalk connections to talk TCP/IP. The Mac software must encapsulate each TCP/IP packet in a special Appletalk packet and send it to the gateway box. The gateway will strip off the Appletalk layer and dump the IP packet on the ethernet. TCP/IP packets from the ethernet bound for a Mac reverse this process. This means that you can run TCP/IP applications from you Mac. The most commonly used is the NCSA Telnet package which allows remote login and FTP file transfer. There is a similar package from Stanford called Mac-IP, but it is not Public Domain and so is harder to get.

Macintoshes with direct ethernet connections can talk TCP/IP directly, but the software to do this is not yet ideal. This should change soon with Apple's release of MacTCP which will also support the traditional localtalk gateway approach.

The appletalk-ethernet gateways also act as localtalk-ethertalk gateways, and also allow TCP/IP hosts to talk Appletalk using a protocol that embeds appletalk packets in TCP/IP. This allows Unix systems to talk to a Laserwriter printer on the appletalk for convenient sharing of that resource. It also allows the unix system to appear as an AUFS Macintosh file server. The CAP (Columbia Appletalk Package) is a public domain system which supports both of these capabilities, and others. It is supplied with the multigate.

## A.4 IBM PCs and Compatibles

There are many implementations of TCP/IP for the PC. I have little experience in this area, but here are some things I have heard of:

KA9Q by Phil Karn is freely available and powerful. It supports the SLIP protocol and can act as a router. It also supports a protocol that runs over amateur packet radio, called AX.25.

IBM has their own implementation.

PC-NFS from Sun includes, as you might imagine, NFS support. So you can access your unix disks as if they were local.

Several implementations, including IBM's, support the tn3270 protocol which allows you to act as a 3270 terminal when connected to an SNA network through a direct connection to an IBM mainframe or through a gateway.

Various PC Lan products include TCP/IP gateways.

[More detail needed in this section]

## A.5 NGENs

Convergent Technology have been promising to release their own version of TCP/IP for the NGEN for a while. A third party, Sirius Systems, has developed a version called INTERNET-CT. This appears to be based on the KA9Q version for the PC.

INTERNET-CT supports ethernet (thin-wire), SLIP using the NGEN's serial lines, and X.25 using the NGEN X.25 product. It can also act as a router, but I don't know how efficiently.

It supports ftp and client telnet. Its telnet can be used with unix systems after suitably updating /etc/termcap. For use with other systems, e.g. VMS, there is a separate product, CT-102, which emulates the VT102 DEC terminal.

It supports smtp mail interfaced to its own very simple mail user agent. One promised advantage of CT's own implementation of TCP/IP for the NGEN is that it will interface to CT-mail.

Another optional product with INTERNET-CT is S-NET. This gives the functionality of CT-NET, but using TCP/IP instead of raw ethernet so that it can work over a wide area network.

Sirius systems claims to be working on an NFS client, and on a 3270 product using the tn3270 protocol. INTERNET-CT and related products are currently distributed in Australia by Sigma Data and by Redlink.

### A.6 IBM Computers

IBM has an implementation for VM, an MVS version is coming. These are rather expensive.

Techway sells the Mitek TCP/IP-SNA gateway. This provides much of the functionality you would want: file transfer and logins in both directions, and mail gatewaying. It can even support SNA LU-6.2 task to task communication with a remote TCP/IP connected machine. Hardware Prices start at $55K with a medium speed SNA link (up to 256Kb) up to $80K for a channel attached version. Software prices typically add another $40K depending on what features are desired.

It is possible to use a IBM RT, or a Sun computer as a gateway, though without the throughput or functionality of the Mitek box.

## B CSIRO Administration Networking

CSIRO Administration needs to support IBM style networking, i.e. SNA, for connection of 3270 terminals and NGENs to the many Administration applications which are implemented in a very IBM-specific way. They have found a way to meet this need with a type of X.25 switch which is optimized for use with SNA-over-X.25.

There are several ways that Administration and Research networking in CSIRO can cooperate. The most obvious is to multiplex lines or take advantage of situations where Telecom charge less for 2nd and subsequent lines between two points.

### B.1 Option 1: Research Use of the Admin Network

A more sophisticated technique is to run TCP/IP over the Admin X.25 network where appropriate. The IP Router box at Clayton could be connected to the admin X.25 switch there (the better IP Routers know how to talk X.25). Then anybody near the site of any of the other X.25 switches could use a connection to that to join the IP network. They could connect another IP router to their nearby switch but that would normally be a waste of firepower. The alternative is to connect a host computer such as a Sun (using Sunlink X.25) or an NGEN (using NGEN X.25 software and the INTERNET-CT implementation of TCP/IP for the NGEN), and let it act as an IP router.

Those responsible for Administration networking have indicated that they are concerned about the possible impact of IP over X.25 on the performance of the SNA use of the X.25 network. They are prepared to accept it on sections of the network where the remote sites on the given link agree, or if some arrangement can be made to give the SNA traffic priority so that the IP only uses the otherwise unused bandwidth.

The suggestions for connecting sites via X.25 could also be used from commercial X.25 services such as Austpac or CSIRONET, but in these cases the usage charges are likely to make this prohibitively expensive.

## B.2 Option 2: Admin Access via the TCP/IP Network

However the IP over X.25 approach has some disadvantages. Not every site will have an Admin X.25 switch: many will just take an SNA link to a centrally located switch. Where Admin network use is at a low level (as it is at DIT in Carlton) it makes more sense to try to do things the other way around. I.e. to provide the ability for sites connected to the CSIRO-IP network to get access to the Admin functionality they need through that network.

There are many ways that 3270 access could be provided, but it is not clear that any of them will be satisfactory. I believe these options need to be investigated.

Access to the GEAC Library system is more clear cut. GEAC have recently released new hardware designed specifically to allow connection from VT100 and VT220 terminals (or emulators thereof) connected directly or through terminal servers. The Library network administration have agreed to buy this new GEAC communication controller and allow it to be connected to the CSIRO-IP network, so every site attached to the CSIRO-IP network will be able to access all the library capabilities without the need for NIMs, GEAC terminals, micronode or any other sort of X.25 hardware.

## B.3 CSIRONET Connectivity

Sites attached to the IP network will be able to access traditional line-mode CSIRONET. A suitably configured terminal server will be placed next to a micronode with some serial lines into the micronode. Users connecting via TELNET to an appropriate address will then be able to obtain the familiar "UNRECOGNISED FIRST RECORD".

## C The Domain Name System

If computers on networks are to have unique names, then we need some way to allocate names to machines. The traditional way is a central registry, but that has become completely impractical for the Internet: it's too big. The solution to this problem is the domain name system.

This is a system of heirarchical names. A central body allocates only top level domains. The current top level domains are all the offical 2 letter country abbreviations (Australia is **au**), plus a group of 3 letter international groupings by type (**edu**, **com**, **gov** and others) which are mostly only used in the United States.

The central body then delegates administration of the top level domains to some responsible person. In fact it delegated most of the 3 letter domains to itself. The **au** domain was delegated to Robert Elz (who works at The University of Melbourne). Robert Elz can create subdomains of **au**. The current ones are **oz.au**, **edu.au**, **gov.au**, **com.au**, **csiro.au** and **telememo.au**. Each of these has its own administrator, though in some cases it is Mr Elz again. For example **edu.au** is administered by Mark Legg of Flinders University. He can then register and approve names such as **latrobe.edu.au**, and La Trobe can do what they like with that. In this way the job of administering an enormous name space is distributed among many people.

Of course the resulting domain names can be long. Most computers provide some way of defining short aliases. So the Radio Physics network would be set up so that most workstations there can use the name **ateles** for the Convex, instead of **ateles.rp.csiro.au**.

The next question is: how does your computer find an IP address, given a host domain name. The traditional mechanism is a file of host-address pairs. On BSD unix this file is /etc/hosts. This has also become hopelessly unwieldy in the Internet. The solution is to distribute the job of keeping this host–address information. The information is kept on a heirarchy of name servers. So when your program calls the host name to address lookup subroutine, it may query a local name server which may pass the query to a root name server which may pass the query to some other subsidiary name server. This process will delay connection, however the name–address translation will be cached locally, so further connections to the same place within a few hours will not result in a new name server request.

## D IP Addresses and Routing

TCP/IP addresses are 32 bit integers. They are usually called host addresses, but this is slightly misleading. The address is actually that of the host-interface pair. So a host which has interfaces to multiple networks will have multiple addresses. Normally a machine with connections to more than one network will be acting as a router between them.

The IP address is traditionally written as 4 decimal numbers separated by dots, representing the 4 8-bit bytes (sometimes called octets). So the number of the machine I am on is 128.250.1.82. The address can be split into two components, a network number and a host number within that network. The split does not occur at a fixed point, instead there are three classes of networks. What class you are in can be determined by the first byte:

1-127: Class A networks. Host number is the last 3 bytes.

128-191: Class B networks. Host number is the last 2 bytes allowing 65534 hosts.

192-254: Class C networks. Host number is last byte allowing 254 hosts.

The number of possible hosts is reduced because host numbers with all 0s or all 1s (in binary) are reserved. As you will see the type of network can actually be determined from the first two bits of the first byte. The network number consists of 1, 2 or 3 bytes respectively. Network numbers is what IP routers know about: they keep tables indicating which networks can be reached through which interfaces. IP routers talk to each other about what networks they know about.

The intention of the IP addressing scheme is that Class C will be used for most networks. Class B is for large ethernets and similar things. Class A for very specialized very large networks. It hasn't quite worked out like that because of the invention of sub-netting.

## D.1 Subnetworks

Routing boxes have to remember where large numbers of networks are. This becomes a performance problem. This is aleviated by subnetworks. The idea is that a large organization (such as a University) is given a Class B network and looks to the outside world (and in particular to the routers) like one large network. Internally it is broken up into subnetworks.

For example the machine I am using is 128.250.1.82. To the outside world it seems to be host 1.82 of network 128.250. But within our Class B network (which is actually that of the University of Melbourne) this is actually host 82 of subnetwork 128.250.1. The way this works is that there is a subnetwork mask, in this case 255.255.255.0. Instead of following the standard class rules to determine what is host number and what is network number, the mask is used. So if my machine sees an IP address 128.250.1.26 then it sees that after applying the mask then it is on the same network, and so sends to it directly. Alternatively if it sees 128.250.2.53 then it will know that this is on another network and look for an appropriate router to send the packet to.

### D.2 Gateways/Routers

...

## E One Person's View of Networking in 1989 and Beyond

The following is a personal view of networking in 1989, how we got there and where we are going. The current situation is a mess and this will allocate the blame in a way which many would undoubtedly dispute. I feel this appendix has to be written or many readers will find it hard to reconcile what is planned with the things they read in Computer magazines and books. Magazines particularly are rarely dispassionate, and frequently they serve up partisan views as if they were fact.

At the end of the seventies the experimental networkers in America moved from host-to-host wide area networking (the original Arpanet running NCP) to LAN-to-LAN internetworking (TCP/IP). At the same time the European telephone monopolies (the PTTs) became seriously interested in providing a packet switch service. However their point of view was significantly out of date. They thought computer networking was about connecting terminals to mainframe computers. They didn't understand how to layer network software: they certainly didn't realise that the network level software should be kept simple with the job of error recovery, defragmentation, packet reordering, session connection management and more being done in the connecting machines. The result of this was the CCITT X.25 standard.

The next development to consider is OSI. The International Standards Organization, ISO, decided to get into the business of standardizing heterogeneous networking. In my opinion standards organizations are most useful when they find de facto standards and tidy them up and publicize them. There was only one possible heterogeneous protocol suite in existence that they could consider using and that was TCP/IP. They decided instead to take on TCP/IP with their own suite of protocols. I believe that the reason for this was political. The ISO Committees were dominated by Europeans, and they felt that adapting an American standard would put the fledgling European Computer industry at a disadvantage. They felt that they could do networking better, and give the European Computer and Communications businesses an advantage. The decision was mostly ignored or criticized in the American networking community (see Padlipsky's book "The Elements of Networking Style"). It certainly looks bad now, and the European computer industry is still fledgling. The difference between the American and European attitudes is that the Americans like to get their hands dirty and experiment and learn from experience. The Europeans frequently imagine that they can bypass that phase and work out how things should be done by pure introspection. Invariably the result has lots of interesting ideas but fails to work as well as imagined in practice: ALGOL 68 and ADA are examples of failure of the European style.

ISO started out by defining carefully what they were going to do: the famous 7 layer OSI Reference Model. This led to the incredible fraud whereby companies would claim that their software complies with the OSI Reference Model. They weren't claiming any OSI interoperability, only that their networking software was organized in layers. A lot of people spent a lot of time talking about OSI's wonderful 7 layer Reference Model — for a long time there wasn't much else to talk about. In fact the OSI Reference Model is flawed: the Presentation Layer would be better if it was not a separate layer, but was an Application Layer service like a number of others that have been defined. The current arrangement makes it hard for an application to do checkpointing conveniently. However by the time this was discovered it was too late to do anything: the 7 layer model was sacred text.

It is important to understand the relationship of OSI to X.25. Many people see X.25 as a necessary part of OSI. This idea was pushed particularly by the PTTs in the early days of OSI. In fact X.25 precedes OSI and does not conform sensibly with the OSI Reference Model in several ways. X.25 is not necessary for OSI which can also run over other network layers, including particularly the ConnectionLess Network Protocol, CLNP. The IP Network boxes being considered for the CSIRO-IP network will also (simultaneously) support CLNP, as will those of AARNet. So the ability to handle OSI networking is guaranteed. CLNP is also called ISO IP, so there will be no need to change the name of the CSIRO-IP network as we move to include OSI capabilities.

OSI was significantly oversold. Key standards have taken much longer than expected to be finalized and implemented. However there are now plenty of OSI products around. FTAM covers File Transfer and Management. Remote logins is covered by VTP, Virtual Terminal Protocol. Electronic mail and related services is handled by X.400. The OSI name is MOTIS, but the CCITT name X.400 seems bound to stick. Unfortunately OSI has run very hard to catch up with the rest of the networking world, only to find that the rest of the world has moved on. Remote terminal operation has been replaced by X windows, a network window system. And people no longer want to transfer files, they want to be able to access a remote file system as if it were local as with TCP/IP NFS or Apple's Appleshare or DEC's DFS and Local Area Vax Cluster software. And OSI doesn't have plans for software to handle network printers as with TCP/IP's LPD or DEC's DQS.

Americans are now much more positive about OSI. DEC plans to base the next phase of DECNET, Phase V, on OSI. In fact Phase V DECNET is now called DECNET/OSI. When Americans talk about OSI, they may consider some support for Connection Oriented Network Service, i.e. X.25, but they mostly imagine a datagram network similar to TCP/IP or DECNET Phase IV. When Europeans talk about OSI they have mostly never heard of ConnectionLess Network Service, and imagine use of X.25 as the Network Layer. Governments in America and Europe have issued documents describing what part of OSI they want supported in computers tendered for future Government contracts. These are called GOSIPs for Government OSI Procurement profiles. At the moment machines conforming to the American GOSIP would not interoperate with machines complying with the various European GOSIPs. There are plans for interconnecting the two with Transport layer gateways, though such a beast would seem to be incompatible with the rules of the OSI Reference Model.

The OSI proponents envisage that it will take over all networking. This seems extremely unlikely. However two events in 1990 will make OSI one of the three major network suites. Those events will be the release of DECNET/OSI and the wide-spread availability of BSD 4.4 Unix with a full OSI capability (and essentially free to Academic Institutions). The other major networking suites, TCP/IP and SNA, will continue to be very important throughout the 1990s. PC Local Area networks will have a large installed base but will access the world through gateways to the major protocols. The factor driving networking advances in the 1990s will be the need to find ways to do something sensible with the enormous bandwidths of fibre-optic communication.

# F Administration and Library Access over TCP/IP at DMT

The Division of Manufacturing Technology have shown that there is a lot of potential for using an IP network to support any sort of network connection that might be required. For the curious here are some details.

In Melbourne DMT have an X.25-capable Micronode. This is connected to a BRIDGE CS-200 Terminal Server with two BSC lines. The BRIDGE Terminal Server is connected to the ethernet and talks the Telnet protocol. The DMT TCP/IP network includes the ethernet and remote ethernets connected to the ethernet by IP routers. Machines on the TCP/IP network can connect to the BSC lines by making a Telnet connection to IP addresses which the Terminal server associates with those lines.

You can nearly connect two BSC lines like this:

[picture of: BSC—Terminal-Server—TCP/IP—Terminal-Server—BSC]

and lock down the terminal link. This system doesn't work because all the BSC polling floods your TCP/IP network.

Murray Jensen's solution is a continuously running program, a daemon, which runs in a host computer on the ethernet with the Terminal server. This talks to a BSC line but slows dow the polling. Things wanting to talk to the BSC line must do so via this daemon, and they do so in a way that doesn't involve any polling exchange. So the Melbourne picture looks like this:

[picture of: Ethernet, 2 Suns, BRIDGE CS-200, Micronode with two lines to BRIDGE box, line going off to Adelaide].

One way to use the BSC lines is with a 3278 emulator that Murray wrote. It can run on any terminal on a BSD unix system. This emulator utility connects to the daemon described above and converses in BSC packets.

Another way to use the BSC lines is used in DMT's Adelaide Laboratory. There they have an NGEN serial port running BSC connected in to a CS-200 terminal server BSC port. A program on a Sun waits for the NGEN to wake up and send some interesting packets: when it does it connects across the Adelaide-Melbourne link and connects via the Sun program to a free BSC line into the micronode.

Murray also has a similar but slightly more complex way of connecting a NIM (and attached GEAC terminals) in Adelaide to an X.25 micronode in Melbourne. This adds more possible points of failure to the GEAC's currently rather delicate network environment, but luckily there are better ways in view for network connection to the GEAC.