TECHNICAL MEMORANDUM

209

Report on the eighth meeting of Member State Computer Representatives, 27-28 September 1994

A. Lea (Compiler)

Operations Department

September 1994

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen

CONTENTS

PART 2: MEMBER STATE CONTRIBUTIONS

PREFACE

The eighth meeting of Member State Computing Representatives took place from 27-28 September 1994 at ECMWF. Fifteen Member State personnel took part, the list of attendees is given in Annex 1.

As in previous meetings, the Head of Computer Division (Geerd-R. Hoffmann) opened by giving an overview of the Computer Division's status and plans. This was followed by each Member State Computing Representative giving a short presentation about the use their particular service makes of ECMWF's computer facilities, The remainder of the meeting was a mixture of discussion sessions on the Centre's computer services and technical presentations about some specific aspects of the Centre's computer facilities. The programme is given in Annex 2.

This report briefly summarises each session, in particular concentrating on the discussions. Part 1 covers ECMWF's contribution, Part 2 the Member State contributions. All the reports in Part 2 have been provided by the Representatives themselves.

PART 1

ECMWF STAFF CONTRIBUTIONS

1. ECMWF'S COMPUTER STATUS AND PLANS - Geerd-R. Hoffmann

1. INTRODUCTION

In figure 1 you see the current ECMWF configuration, including the upgraded IBM Data Handling system and the recently installed CRAY T3D (MPP) system. The compute server is still the CRAY C90, but now augmented by the CRAY T3D. The data server is, as before, the IBM but now it is a 6 processor ES/9000-720 machine.

There has been little change in the staff of the Computer Division since the last meeting. One new staff member (Petra Kogel) has joined the Workstation Support Group, while for one year a trainee analyst (Kemal Dokuyucu) has joined Computer Operations Section from the Turkish Meteorological Service. The two new consultants are Charles Wilkins (IBM support) and Stuart Mitchell (Workstation support).

I am very pleased to report that so far in 1994 we have had only two incidents causing a forecast delay of more than 1 hour. This is much better than we have experienced in previous years. However, the next two years (1995/6) may, unfortunately, be not quite as good, as there will be a lot of upheaval due to replacement of the cooling equipment and installation of the new data handling system in 1995, followed by the CRAY C90 replacement in 1996.

Figure 1 ECMWF computer configuration - 1994

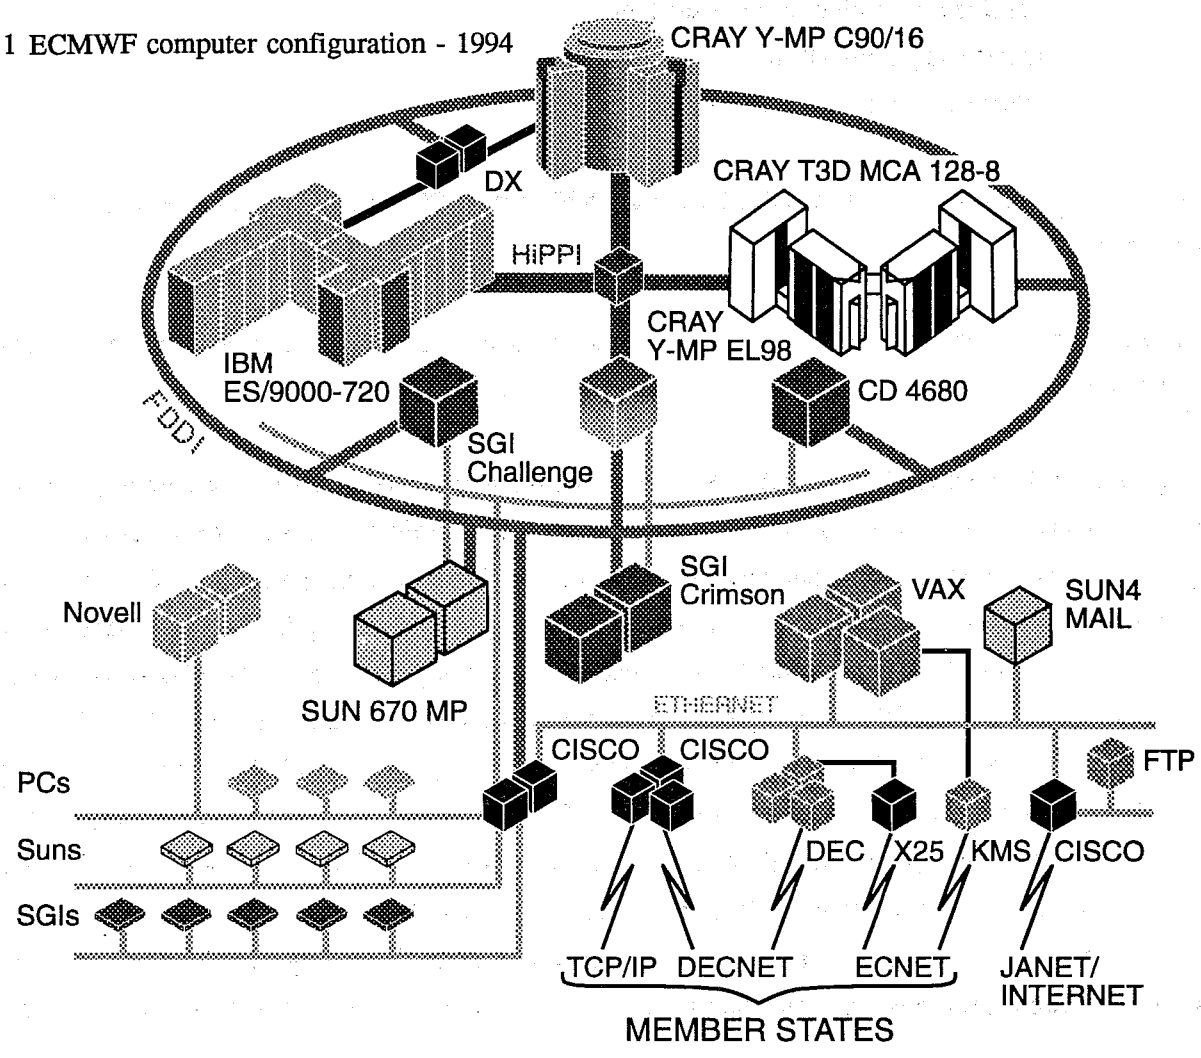

2. CRAY SERVICE

2.1 CRAY C90

The CRAY C90 was recently upgraded to 2 Gbytes of main memory with full banking, this should lead to better job turnaround, especially for fully vectorised jobs. Also, we have replaced some DD62 disks with newer technology DD301 and DA301 sub-systems, the DA301 being RAID type.

The reliability of the C90 has been varied, especially at the beginning of the year and then again when we installed the new memory in July/August. We seem to experience periods when several component failures and other problems occur in quick succession. Then the system stabilises for sometime before another period of instability occurs. On average the overall reliability has not been too good. Hence we have a task force with Cray Research which meets regularly and keeps in contact with the machine developers in the USA. Hopefully, this will result in a more stable system in future.

On a much more positive note, the CPU utilisation is now very good. A year ago we were running at about 60% user time. We then worked hard on this, so that over these past 12 months the user CPU time has risen to over 90%, with the idle and system times now very low. This has been achieved by:

- considerable tuning of key user jobs
- introducing UNICOS 8
- adding the extra 1 Gbyte memory

Overall this has led to the user time rising by around 50% (from 60% utilisation to over 90%), a very good achievement.

The rate of jobs through the system now seems to be acceptable to all, we have not received any complaints in recent months.

To our surprise the introduction of UNICOS 8 led to a reduction in the number of accounting units being delivered, whereas we had expected an increase. We looked very carefully at this and then realised what was happening. To encourage users to multitask their jobs we effectively only charge for one CPU. If a job uses more than one then that is a bonus for the user - his job goes through quicker and is hence charged proportionally less than for a single tasked job. Because most major jobs are multitasked now we found that by reducing the system and idle times significantly more multitasked user jobs were able to go through the machine in a given time. Each job is now being charged on average less than before because the memory integral is less. Also, because the introduction of UNICOS 8 has led to much less system single tasking, again the number of units per job has dropped. To return to the level of units we had expected to deliver in 1995 (as specified in the Technical Advisory Committee paper to be presented to the 20th session of the Technical Advisory Committee later this week) we propose to revalue the unit by 20%, in other words charge each job 20% more than now. Overall, this should mean the normal user is back to where he was before

UNICOS 8 was introduced in terms of accounting. This change in the unit definition will come into effect on the first Monday in 1995 (2 January) when the new allocations are introduced.

2.2 CRAY T3D

The CRAY T3D configuration is shown in figure 2. It consists of a Y-MP2E front end with 0.5 Gbytes main memory and a 128 processor T3D. Each T3D processor is a DEC Alpha chip with 64 Mbytes of memory. The configuration is a 3D torus (hence the name T3D!).

The Y-MP2E provides the disk space (43 Gbytes) plus connections to the outside world.

The machine arrived on 9 August, Provisional Acceptance was passed on 28 August and we are currently towards the end of the Final Acceptance period.[1]

The machine is already being used by Centre staff, Member State projects are now being invited to start via their User Support contacts.

Initial studies have shown that the Centre's IFS model code is rather scalable, but that single node performance on the T3D is rather poor (12 Mflops out of 150 Mflops maximum). Investigations are underway to improve that figure, for example using:

- compiler optimisation
- faster intrinsic functions
- faster FFT
- data read ahead
- code optimisation
- optional 32 bit representation

However, the fact that the Alpha chip has no secondary cache may make it difficult to reach very high single node performance.

3. DATA HANDLING SYSTEM

Figure 3 shows the current configuration following its upgrade earlier this year. This machine now is destined to provide full C90 support until 1996 and then take the MARS load on the C90 successor up to 1998.

---

[1] After the meeting - Final Acceptance was successfully passed on 5 October.

Figure 2 **Cray Y-MP 2E/T3D Configuration - August 1994**

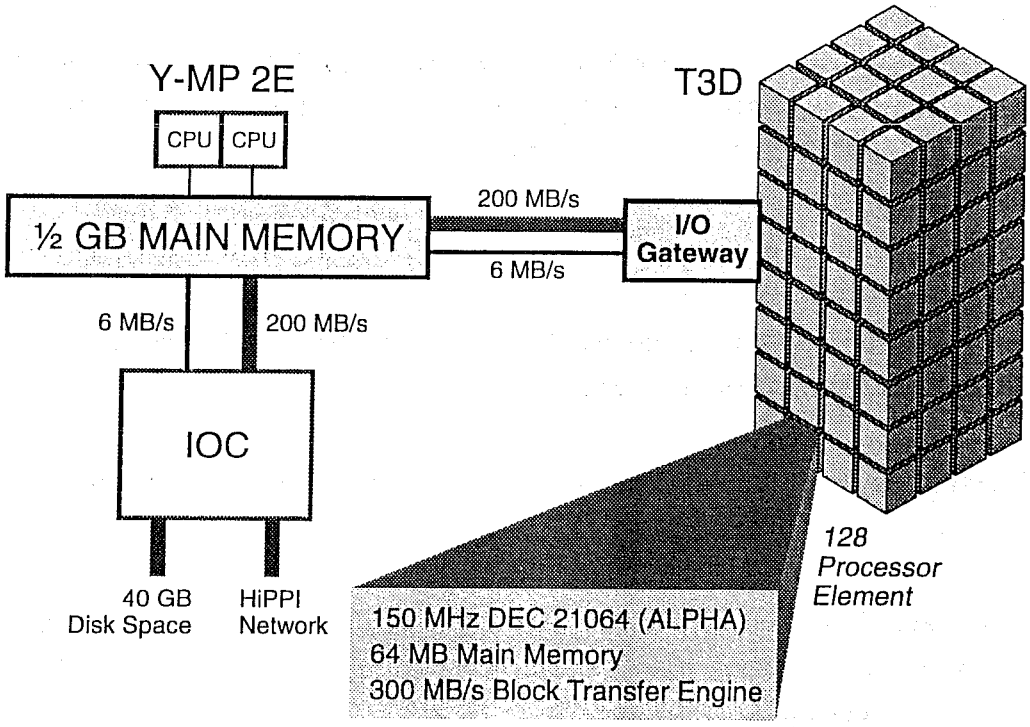

Figure 3 **ECMWF data handling system - IBM ES/9000-720 configuration**

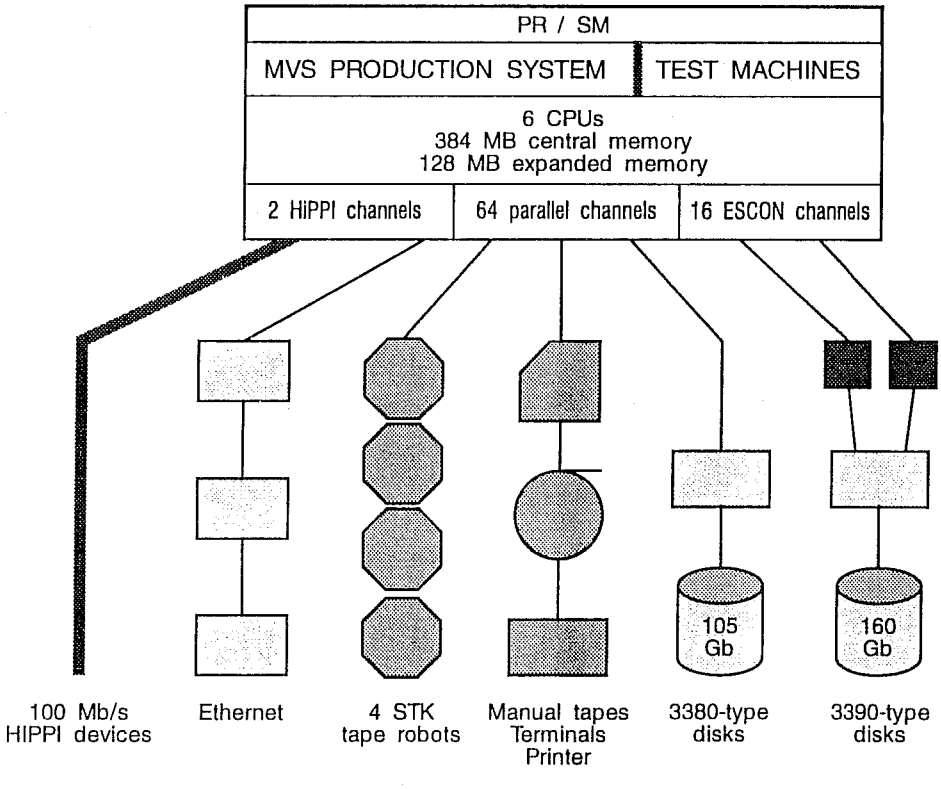

The machine is extremely reliable. The only problem has been when we installed the second HIPPI channel (we believe we are the only site world-wide running two HIPPI channels under MVS/Prism). Overall the IBM is the most reliable system we currently have.

The IBM copes well with the C90 load, there are no longer delays on the C90 due to queues on the IBM.

CFS and MARS continue to run well. CFS now runs with TCP/IP connections to the Cray systems, an in-house development which has led to very good performance. The overall volume of transfers are about 100 Gbytes per day. The total volume of data stored in CFS is now around 21 Tbytes, in about 5 million files. We predict that by 1998 this will rise to about 150 Tbytes of data and 37 million files. Finally, Member State usage is still below the 10% limit allocated to Member States.

4. VAX SYSTEMS

The cluster is now based on two 4100 systems which are much cheaper to maintain that the previous 6000 and 8000 systems. We have one 8250 left to support the only remaining Member State still on ECNET protocols.

Many Member States connect through CISCO routers, running TCP/IP and/or DECNET protocols. Six Member States run DECNET only. Those Member States running TCP/IP can participate in the ETR (Electronic Traffic Routeing) experiment whereby they can transfer data between themselves via the ECMWF links. So far two Member States use this facility, Germany and Switzerland.

Data received from Member States via the NTC links is slowly reducing, while the data sent to Member States on these links is increasing. The reduction in data received is because more and more Member States are switching to TCP/IP links for incoming traffic. However, data dissemination, which is the bulk of the data sent to Member States, still goes via the NTC system and the volume requested is rising continuously.

5. WORKSTATIONS AND THEIR SERVERS

Last year ECMWF selected Silicon Graphics as its main supplier of workstations and servers. Figure 4 shows all the different servers we currently have. Note the server "munin" (alias ecserver) is for Member State users (shared with Operations Department internal users).

Unfortunately, "munin" has suffered some downtime recently. Initially a hardware problem was suspected, but later it turned out to be an operating system problem whereby the kernel was not properly protecting a tasks's memory space. It took a long time to isolate this problem.

A study of "munin's" CPU utilisation shows a typical interactive server pattern, peaks during the working day, with not much activity overnight. On the other hand, some servers used by internal staff (e.g. "hugin") show overnight activity as T21 models are now being run there.

The availability of the CRAY Y-MPEL server has varied. We had a lot of downtime when we installed new disks in May this year, it turned out to be a disaster. Cray reacted quickly and now the availability has been a lot better, although not as good as some other servers. This machine is used for model test runs, cross compilations, testing of sequential mode codes etc. Current user utilisation is around 60%.

The "odin" server has RAID technology disks and hence is used to hold all the home directories. There was some down time at the end of last year but it is now very stable.

Figure 4

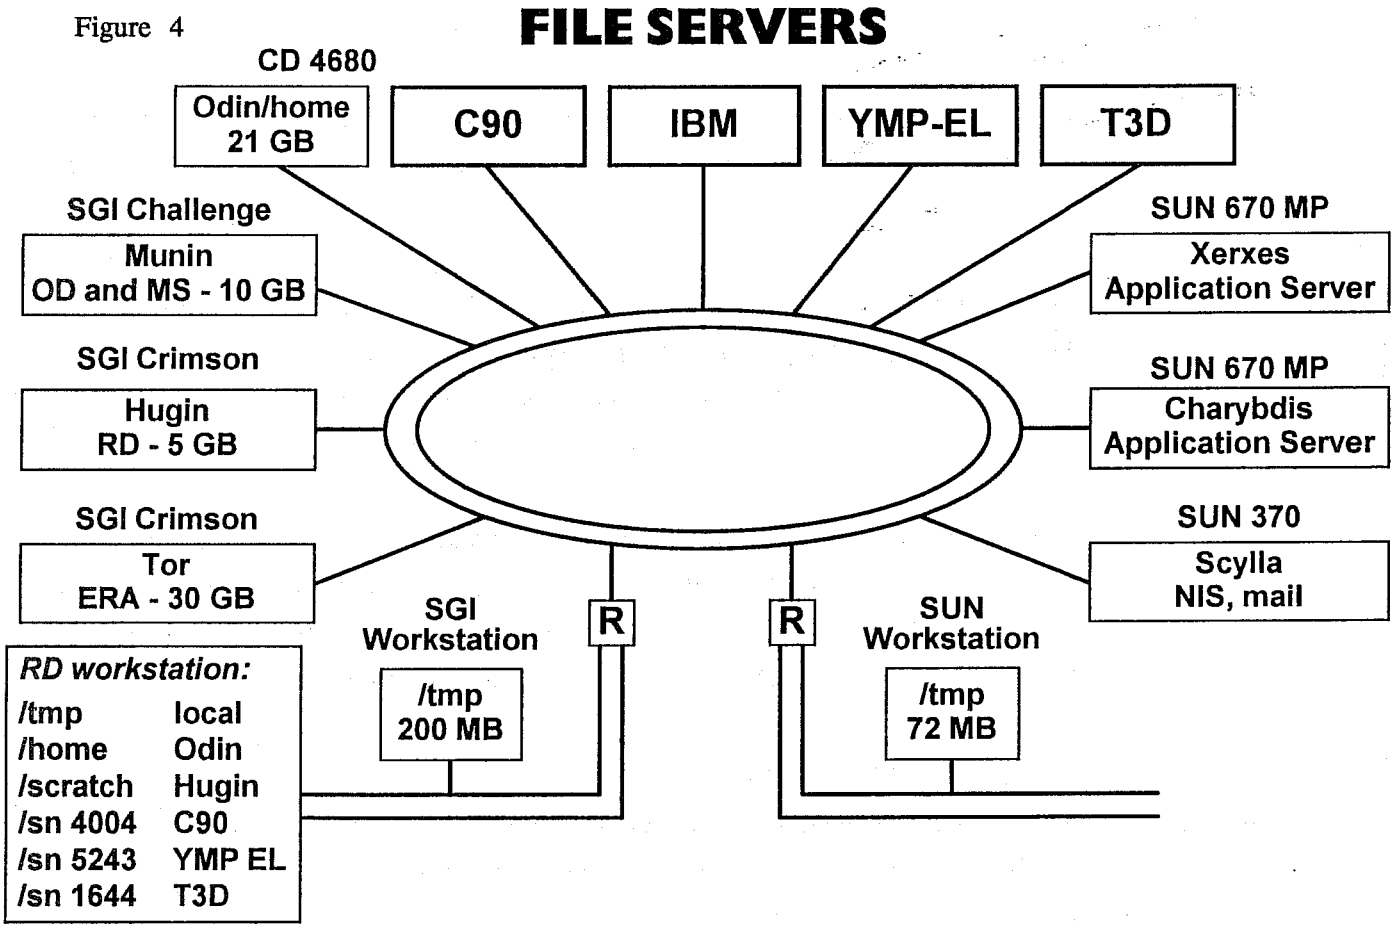

6. LOCAL AREA AND WIDE AREA NETWORKS

The Centre's LAN has two FDDI rings and HIPPI connections including a HIPPI switch, plus many routers and hubs. We are trying hard to make it all reliable, with backups being available to quickly take over if any specific box fails. Currently it is all running reasonably well, with incidents down to one a day or less. That is much better than it was a year or two ago.

The main change in the wide area network has been due to the decision by Council to allow 64 kbps links for all Member States from 1 April 1994, to be paid for by the Centre. All Member States, except Turkey, now have 64 kbps links in operation. Two further links will be installed shortly from our associate Member States, Iceland (64 kbps) and Hungary (9.6 kbps).

France asked the Technical Advisory Committee to recommend to Council that it be permitted to raise its link speed to 128 kbps.

Overall these changes have been very positive and we hope that all Member State users can benefit from the improvements.

7. OTHER SYSTEMS

We have one system (Daphne) which is outside the ECMWF "security fence" to provide external mail, anonymous ftp, etc. It is also used to distribute/collect data from Internet sites.

All printers within ECMWF are now PostScript.

Parallel (MPP) activities include:

GP-MIMD2 — an EU project to show that European MPP systems can work well in an operational environment. A CS-2 machine has been installed in Toulouse for the meteorology community, to show it can work in that environment. Initial problems have meant the project is currently running late.

RAPS (Real Applications on Parallel Systems) is an initiative of some large European computing sites to try to migrate real applications to parallel platforms. The codes are then given to manufacturers to run on their systems. ECMWF has contributed its IFS code, which has now been successfully run on machines from Fujitsu, Convex, Cray, IBM and so. We have got some good feedback, making it a successful project.

PPPE — This is an EU funded project to develop tools for use in a parallel environment. ECMWF has one staff member paid for by this project and it has proved very valuable in feeding information on such tools into our own parallel project work.

8. PLANS

Two ITTs were issued recently to enhance the cooling capacity, the closing date is 3 October. One ITT is to replace the pipework, the second is to replace the chillers. Work will commence later in the year and is to finish around mid-1995. It will lay the foundation if any increase is required in capacity for the C90 replacement.

An ITT will shortly be issued for the Data Handling System replacement. This replacement is required because the IBM cannot support the entire load that will be generated by the C90 replacement, neither can CFS be modified any further to support the data volumes expected. Remember that CFS is a product that is no longer developed commercially. We thus look for a new system based on
Unix and able to support new storage technologies.

In 1995 we will begin by selecting the Data Handling System replacement and getting approval of the various committees and Council, with installation of the first phase in the 4th quarter 1995.

The other major event will be the issuing of the ITT for the C90 replacement in February 1995. This machine will be 5 times the power of the C90, which means in practice a machine with at least 30 to 40 processors, maybe many more. It will require message passing programs to use all the machine efficiently.

The C90 replacement will be installed early in 1996 with a 6 month parallel run with the C90 beginning April 1996. At the same time the new data handling system will be brought into service to support the C90 replacement. The C90 and T3D systems will then be de-installed in October 1996.

Phase 2 of the new data handling system will be installed in 1997 to fully support the C90 replacement. That C90 replacement machine will itself be upgraded in 1998 to double its performance. Finally the IBM ES9000 with a concomitant upgrade of the new Data Handling System will be de-installed in mid-1998.

Towards the end of the four year period (1995-8) we have to consider whether to upgrade our LAN to higher speeds, for example using ATM, fibre links, etc.

Finally, some further items we have to consider in this 4 year period include:

- possibility of 2 Mbit/s links for some Member States;
- different mechanisms to provide our WAN e.g. via managed networks, satellite links, etc.;
- studying pan-European research networks for computer access from Member States, perhaps as an EU 4th Framework project;
- the use of Smart Cards to make remote access more secure;

- remote use of MARS/WS and Metview/WS directly from workstations in the Member States;
- using servers for operational work e.g. pre- and post-processing;
- batch support for server platforms to run such non-vectorising work as MAGICS and MARS;
- introduction of DCE/DFS in place of ecfile on the new data handling system.

QUESTIONS

A. Dickinson — To clarify, are you recommending that all jobs should use multiple CPUs, up to the full 16 available?

G-R. Hoffmann — To multitask your jobs is what we have said all the time. However, to go to 16 CPUs is a bit tricky for efficient scheduling. Perhaps 4 to 6 CPUs per job is a good compromise. Remember, the more CPUs you actually use the bigger the bonus you get from the accounting formula.

S. Pasquini — What is the Centre's future policy concerning DECNET, will it continue to be supported or do you want all Member States to use TCP/IP?

G-R. Hoffmann — There is no plan to get rid of DECNET. However, we will probably freeze the VAX/VMS system at some stage and then not upgrade VMS any more.

R. Rudsar — ETR, what sort of data are you talking about being transferred?

G-R. Hoffmann — Council agreed that for a two year experiment Member States can transfer data amongst themselves, for Research purposes only. It is not to be used for operational purposes, for example it is not a backup for the GTS links between Member States.

B. Barg — The ECSN project has talked about using this facility, but in practice they have not used it.

G-R. Hoffmann — We know. The only users so far are Switzerland and yourself. There is a request to the next Technical Advisory Committee meeting to allow Iceland and Denmark to exchange data.

S. Kruizinga — The transmission of permanent files (such as batch output and SENDTM) has such a high priority on the link that it completely stops any interactive traffic. Is there a solution to this problem?

G-R. Hoffmann — At this time we do not have a solution. It happens at the level of the routers. The only priority differences that can be assigned are where you have different senders/receivers. Because batch output, SENDTM etc. all go through the VAX NTC system, they get the same (high) priority as product data dissemination. Interactive traffic will always be at a lower priority.

W. Zwieflhofer — It is current policy that data dissemination (and hence all other NTC traffic) gets higher priority at the router level than non NTC traffic.

T. Bakker — One solution could be to send data dissemination traffic from one source and batch outputs plus SENDTM from another source.

G-R. Hoffmann — We take note of the problem [NOTE: see the Discussion Session on page 46].

L. Campbell — Are many Member State users running X Windows sessions with munin as a server?

W. Zwieflhofer — Very few. We made some tests on a simulated 64 kbps line which showed that for just character based work (e.g. X terminal) it works reasonably well. However, when used for graphics it can be very slow. Ultimately however, Member State users must judge the response for themselves.

B. Barg — The pan European networks, are these PTT provided or by private suppliers?

G-R. Hoffmann — At the moment we are thinking of research networks, which are funded usually by government subsidies in one way or another. According to our convention we are allowed to use such networks.

B. Barg — In Germany the research networks are overloaded. We tried to send small files to other institutes e.g. Hamburg, over WIN, but they always arrived late, so now we are going back to PTT provided networks.

G-R. Hoffmann — We are not considering replacing our leased lines with the research networks. We are considering a two layer structure, time critical data over the existing 64 kbps links, whereas computer access (which is not so time critical) using a cheaper way that provides more bandwidth but offers no guarantees. Remember, the C90 replacement will offer 5x, then 10x the power of the C90. This means that many Member States will have more capacity available at ECMWF than locally. To use that capacity Member States will need high capacity links.

2. MAGICS AND METVIEW - Jens Daabeck and Elisa Nishimura

1. MAGICS

MAGICS, the Centre's graphics package, has always tried to conform to existing standards, for example, to the WMO GRIB and BUFR standards for data representation, and to ISO GKS and Fortran 77 coding standards. This has made it easier to port it to the several different platforms that are used within the Centre.

Many Member States have a copy of MAGICS, either of release 4.1 or 4.2. Now we are close to releasing version 5 (more specifically 5.0.3). This latest version supports:

- satellite image plotting
- improved colour setting and shading
- shading of land/sea
- stretched/rotated grids

plus a few more minor improvements.

Note that satellite images are repacked and stored in extended GRIB code format. This is an experimental WMO standard which we, and a few other sites, have adopted. We hope that it becomes a full standard shortly. Also, colour control can now be provided by the HSL (Hue, Saturation, Lightness) model, which will be slowly integrated throughout MAGICS to provide a uniform method of colour control.

At the Centre, MAGICS is currently supported on the CRAY C90 and EL systems, plus Silicon Graphics and Sun workstations. These are the systems we foresee we will continue to support. Note that there is no further development of MAGICS on VAX systems, it will be frozen at release 4.2. This is because ECMWF has no internal demand for the latest release 5 to be ported to the VAX systems. Looking to the future, we do not intend to port MAGICS to the C90 replacement. In other words, our policy will be to move graphics applications to the Unix based workstations and not to the supercomputer.

As to the future development of MAGICS, these are the current plans:

- support for the extended analysis feedback (as required by the Research Department)
- incorporate new BUFR decoding software
- use a new spectral to grid interpolation package in the workstation version
- plot data values on the reduced grid
- add contour shading in the spaceview projection
- add enhancements specifically required for Metview, to improve its efficiency through tuning improvements
- update the manuals

2. METVIEW

Metview is an interactive utility to access, manipulate and visualise fields, images and observations. On top of the basic system are special application modules which can provide higher level functions, such as cross-sections.

Metview runs on Unix workstations only and uses a Motif based graphical user interface. The design has used a distributed processing concept i.e. Metview consists of modules which can run on the same CPU or on different CPUs as required.

An internal macro language is available so that whatever can be done interactively can also be done via macros. Thus, if you have a need for a repeated task, or a lengthy, complicated task, you can code it in the macro language and then run it with just one command.

Data access is via MARS/WS. This means that the user does not have to know where the data is, he just asks for what he wants in meteorological terms. Automatically the data will be fetched from wherever it resides, whether it was on disk, in the tape silo, or on archive tape.

The Metview concept is that the user creates icons, an icon holding for example, the definition of a MARS retrieval or of a MAGICS plot specification. Then, using a drag and drop technique, the user can visualise some data (drop a data definition icon on to a window with a map), edit any definition (select an icon and enter an edit phase), and so on.
[At this point in the talk an on-line demonstration was given of the principal features of Metview].
As to the future, ECMWF is working on:

- adding more application modules. Currently there are 4 or 5, there are another 30+ waiting to be ported from Metview/Batch
- providing full support for titles and legends
- adding a colour image editor
- improving support for observations
- providing a batch version
- enhancing the macro facility, so that for example you can "record" a series of interactive actions and then "replay" them.

On the staff side this package is the result of a co-operation project between ECMWF and INPE/CPTEC (Brazil). Also, one staff member from Météo France has been working on the ECMWF side. We believe this product is unique in the meteorological world at the moment because of its icon based (object oriented) structure which gives users a very friendly and flexible interface.

QUESTIONS

L. Campbell — You said MAGICS will not run on the C90 successor. Where will Member State users run it, on their ECMWF server (ecserver)?

J. Daabeck — The only reason why it is not available to Member State users on ecserver at the moment is the current concern about capacity on that machine. Hopefully this can be solved. Anyway, ECMWF realises that it must provide an alternative platform for Member States MAGICS users once the C90 goes.

3. OPERATIONAL SERVICES - RECENT EVENTS - Peter Gray

1. INTRODUCTION

This talk aims to cover the events of the last two years i.e. since your last meeting two years ago. In those two years we have made significant changes to the equipment, in fact every month there is usually some change or other going on in the Computer Hall.

Major equipment that has been removed over the past two years includes the CRAY Y-MP8 and CRAY Y-MP4, the CDC Cyber 962 (when the NOS/VE service terminated in November 1993) and three DEC VAX systems (6210, 6310, 8350 - replaced with two 4100 systems). Major equipment that has been installed includes the CRAY C90, two DEC VAX cluster 4100s, a CRAY Y-MP EL, Silicon Graphics Indigo workstations, Silicon Graphics servers (one Challenge, two Crimsons), and within the last two months the CRAY Y-MP2E with the 128 processor T3D.

Over the same period various upgrades have also been carried out, including the IBM data handling system from an ES9000-580 (3 CPUs) to an ES9000-720 (6 CPUs), the CRAY C90 memory doubling (plus extra disks), upgrading the STK tape drives from 18 to 36 track (including compressing existing 18 track tapes to 36 track tapes), and adding an IOS plus 4 extra processors, disks and a HIPPI channel to the CRAY EL.

On the network side, the Member State network had a very significant upgrade earlier this year when all lines (except Turkey) were changed to 64 kbps. We currently run both DECnet and TCP/IP protocols on those lines. Seven Member States use exclusively DECnet, seven exclusively TCP/IP and three Member States use both. The ECNET protocols, which we started to use in 1978 (before there were any international standards) have now almost gone. Only Austria still uses ECNET protocols and they hope to abandon them shortly. At that stage we can then dispose of the last remaining VAX 8350.

Within the Centre the Local Area Networks (LAN) has changed - we installed a structured cable system (a mixture of fibre optic and twisted pair copper cables) within the Office Block. This was done just before the majority of the SGI workstations were installed, which was fortunate as these workstations would have totally overloaded the old LAN system. Finally, within the Computer Hall a HIPPI network now connects all the mainframes together at 800 Mbits/second.

2. SERVICE IMPROVEMENTS

Looking at the CRAY C90 performance, various user optimisation efforts over the past 1½ years have produced excellent results (see figure 5). By the way, for those who want to optimise their own programs I recommend you read and act upon ECMWF Computer Bulletin B2.2/1 (I/O Optimisation Techniques on the CRAY YMP/C90). You can see the typical improvement achieved by comparing the CPU utilisation charts for 16 February 1993 and 11 May 1994 (figure 6). Then, later, UNICOS8 made another big improvement by reducing the system overhead by at least a half. Finally, the memory upgrade has reduced idle time to around 1-2%. Overall, from the user viewpoint, the Cray service is now much improved

compared to what it was 18 months ago.

The other major improvement, again from the user point of view, is that to the IBM Data Handling System. The three changes mentioned already (doubling the number of CPUs, doubling cartridge storage capacities, installing the HIPPI connection) have also indirectly helped the C90 performance, as C90 programs do not wait as long as they did for data from the IBM.

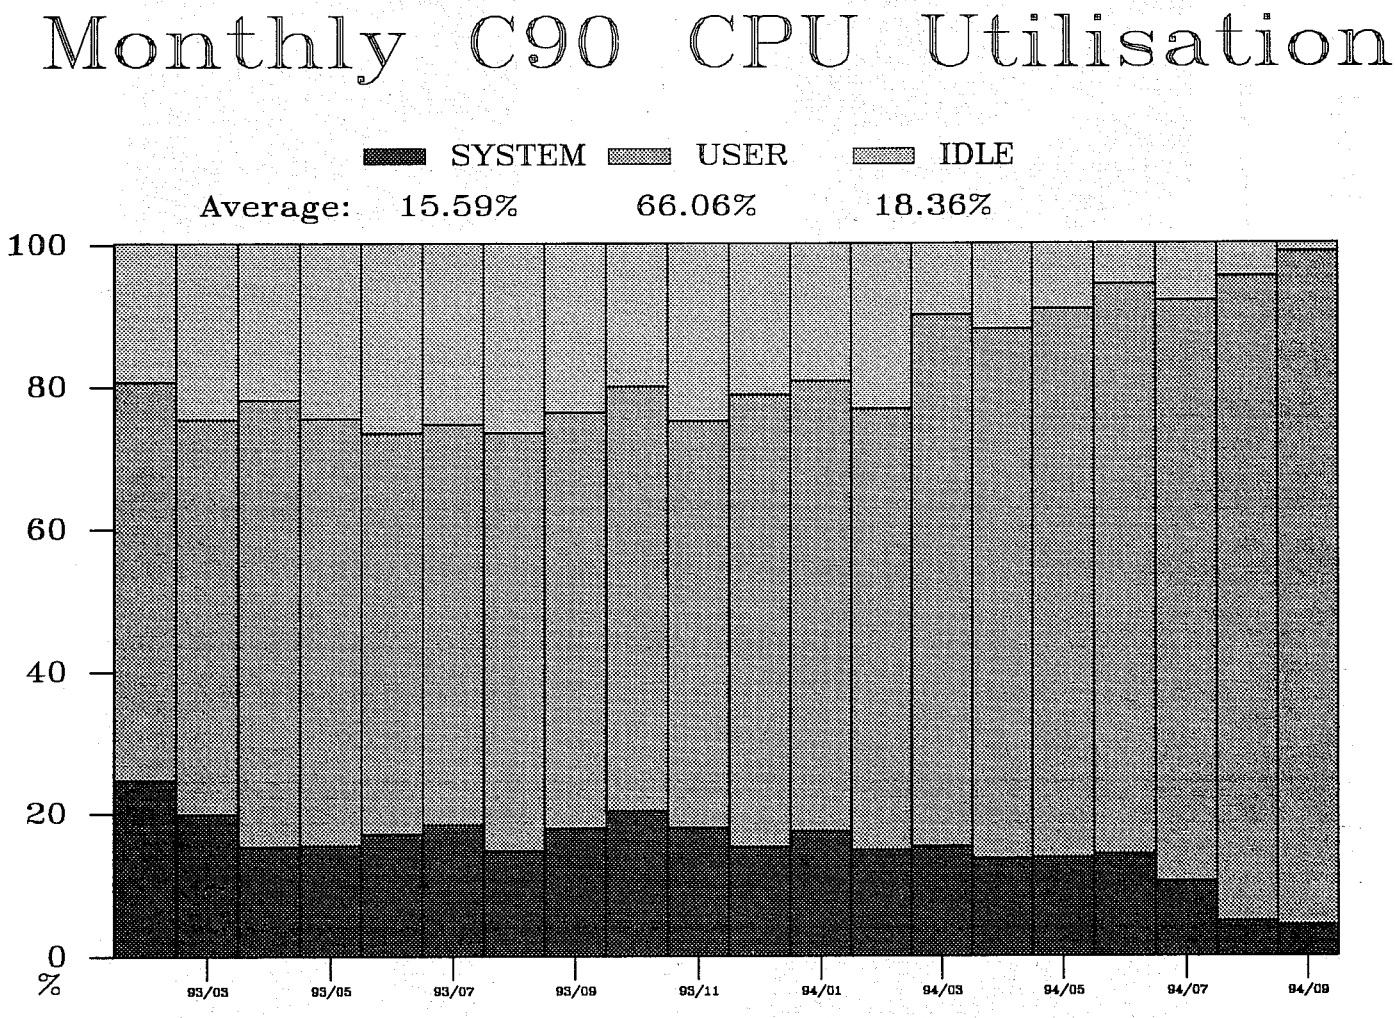

Figure 5

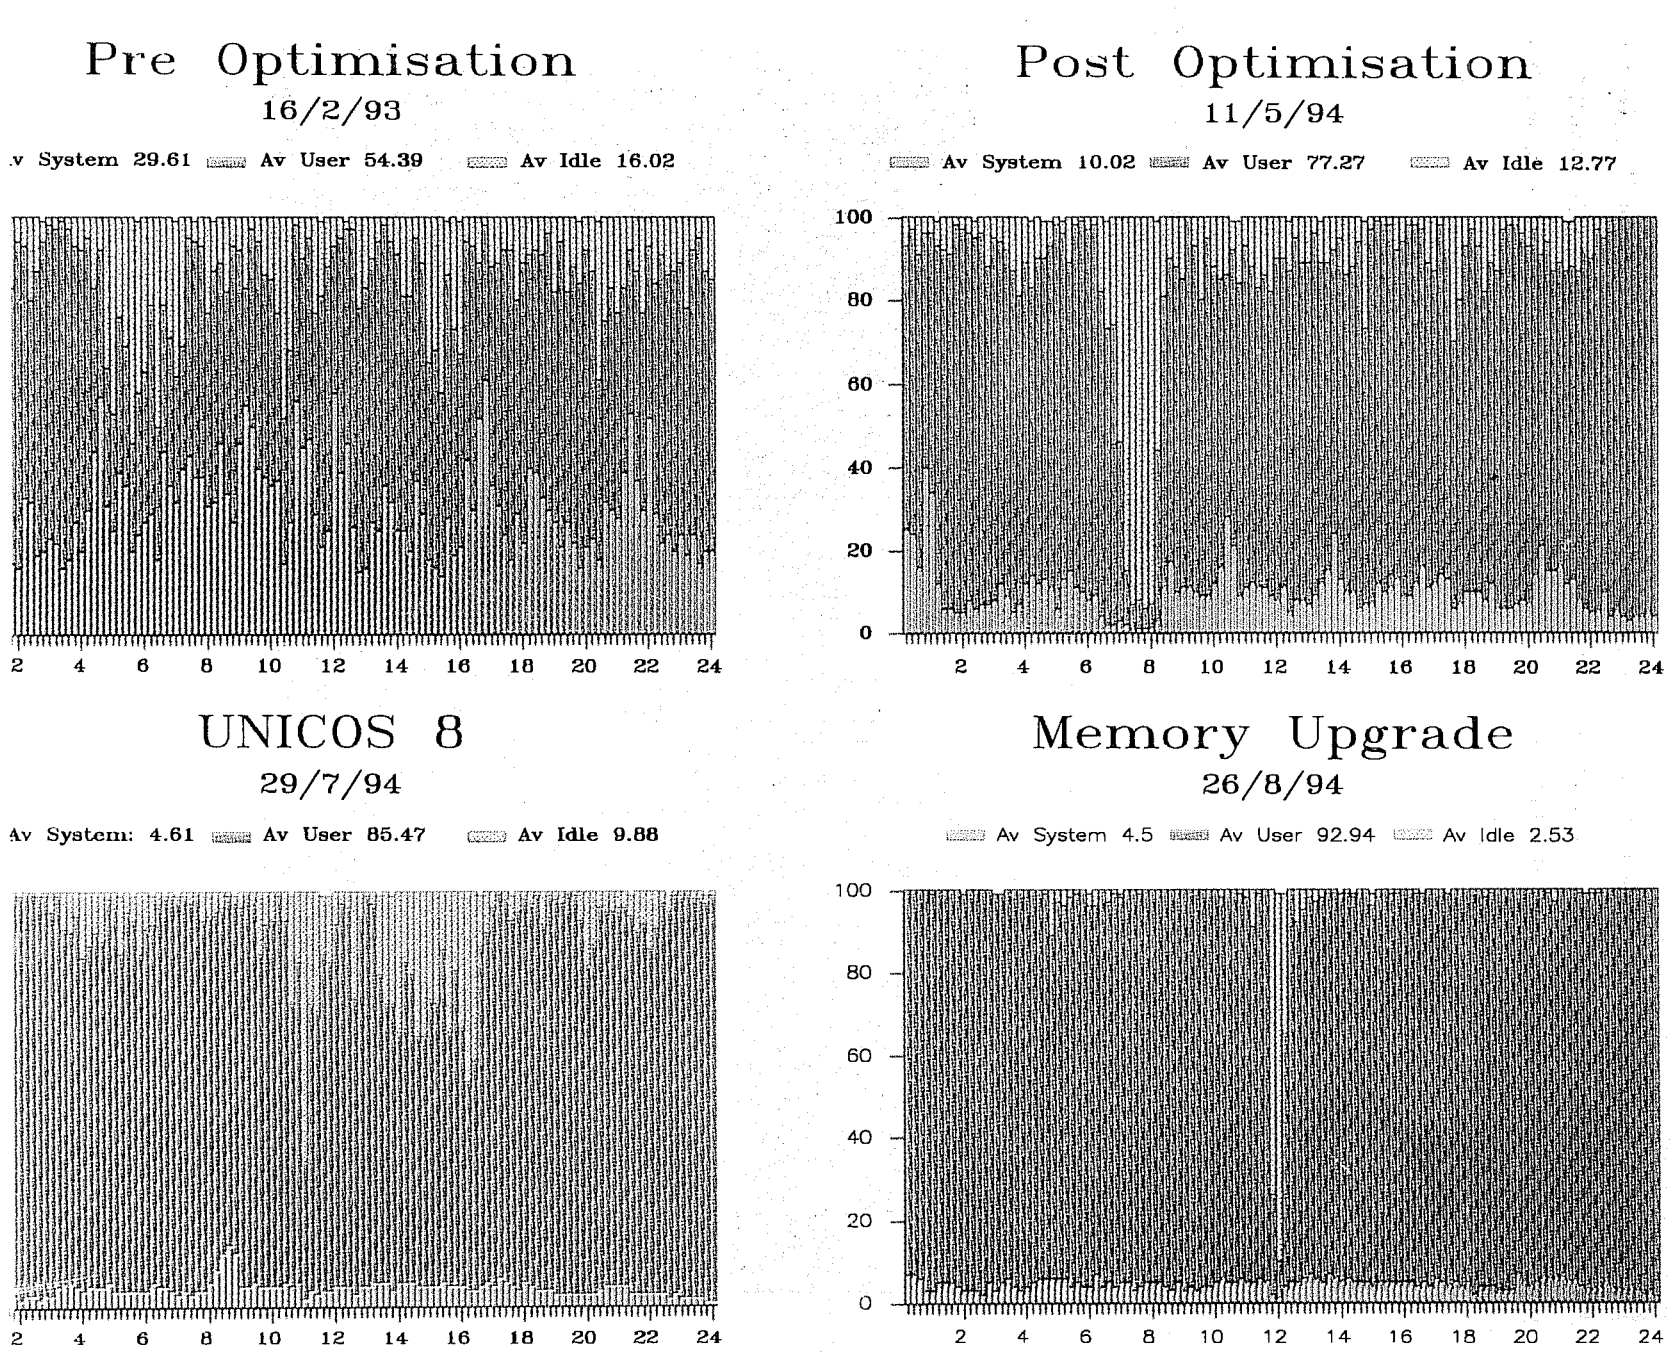
Pre Optimisation
16/2/93
.v System 29.61
Av User 54.39
Av Idle 16.02
2 4 6 8 10 12 14 16 18 20 22 24
Post Optimisation
11/5/94
Av System 10.02
Av User 77.27
Av Idle 12.77
100
80
60
40
20
0
2 4 6 8 10 12 14 16 18 20 22 24
UNICOS 8
29/7/94
Av System: 4.61
Av User 85.47
Av Idle 9.88
2 4 6 8 10 12 14 16 18 20 22 24
Memory Upgrade
26/8/94
Av System 4.5
Av User 92.94
Av Idle 2.53
100
80
60
40
20
0
2 4 6 8 10 12 14 16 18 20 22 24

Figure 6

3. CALL DESK

A change in the organisation came about as a result of setting up a Call Desk to handle and co-ordinate service problems. In the past we have not necessarily handled all problems as well as we could have done. Now the Call Desk is the central focal point for all reported faults connected with the operational service. However, the Call Desk is not a replacement for the User Support function, User Support continue to deal with all programming queries, software assistance, etc. etc. The reported problems come largely from in-house users, although now we are opening it up to Member State users as well. Thus Member State users can contact the Call Desk to report such things as line faults and to enquire about service status, etc. All reported faults are logged, assigned to someone to investigate and answers relayed to the user. A problem continues to be "active" until the user is happy with the solution. The Call Desk is continuously manned, 24 hours a day, by dedicated staff during office hours and by the shift operators at all other times. A database system is used to track all calls and problems, plus produce statistics for management purposes. We handle about 500 problems a month currently.

To contact the Call Desk either telephone: +44 1734 499 303
or send an email message to: call_desk@ecmwf.co.uk or cdk@ecmwf.co.uk

4. IMMEDIATE AND FUTURE PLANS

Recently, financial economies have been achieved on the IBM and STK maintenance contracts, on colour printing (replacing the printer with one that uses much less expensive consumables), on Sun workstation maintenance, and by replacing continuous stationary output with a bound A4 service.

Currently we send the ensemble forecast charts by fax using a manual procedure. Some nights it takes 2 hour, other nights a lot longer. We will shortly install a system that will create fax charts directly from the MAGICS plots and send them automatically to you. This is a Unix based package that will run on one of our Sun systems and will handle 4 fax modems. It will accept either PostScript files, or input in the form of a mail message from an internal user.

We currently have 11 Unix servers that the operators have to monitor. To help them control all these servers we are investigating various monitoring packages that will check each server on a regular basis and report, via one screen, any problems (e.g. no response) or potential problems (e.g. file system over 90% full).

The next big task will be the cooling system changes. These changes are required because:

- we need to ensure there will be sufficient cooling capacity to handle the C90 replacement, whatever it might be
- the existing chillers are reaching the end of their economic life
- we can improve the energy efficiency and so cut the Centre's electricity bill by installing,

for example, better controls plus variable speed pumps, by allowing for higher water temperatures in the winter months and, perhaps, by utilising heat recovery.

The ITT for this exercise has been issued and replies should be received soon. Work will then begin in December 1994 and carry on through to June 1995. As a consequence of this, all users should be aware that there may be substantial downtimes in the spring of 1995, hopefully restricted to weekends.

Overall we will continue to try to improve the service, both from the efficiency point of view and providing better facilities.

4. EXTERNAL NETWORKS - Tony Bakker

1. MEMBER STATE CONNECTIONS

As has been stated before, all Member States (except Turkey) now have 64 kbps links. So far the usage on these links is not that high, except perhaps for France and the UK. Eleven of our Member States use Cisco routers, running mainly TCP/IP protocols, although three use DECnet (Germany, Italy, Finland). Cisco routers handle several high level protocols e.g. TCP/IP, DECnet Phase IV, DECnet Phase V. Below that we currently support three level 2 protocols, X25 (the one we have used for years - very good on poor quality analogue lines), HDLC and PPP (both of which are good for higher quality digital links).

The Domain Name Service (DNS) is an important part of any TCP/IP connection for two reasons:

(i) it allows simple names to be used instead of complex numeric addresses
(ii) it stores information required by the SMTP mail service.

On our TCP/IP Member State links we have configured our DNS with the Member State DNSs so that your users can use our internal host names. These DNS facilities will become more important later as TCP/IP spreads across more and faster networks. Another use of our DNS is to ensure that mail destined for Member States uses our leased line network. Without an internal DNS that mail would automatically go via Internet.

Early this year we upgraded our Cisco routers so that we could standardise on Cisco 4000 routers for all our wide area connections. These upgraded routers now provide more ports per router (6), and allow us to easily and quickly swap one router for another.

With extra memory these routers can support up to 4 Mbits/second on each link. We currently use 4 routers operationally, with a fifth available as backup.

Some Member States use DEC X25 routers, which can only run X25 as their low level protocol. They support line speeds up to 64 kbps only, and have no TCP/IP capability. They do support DECnet Phase V, but in-house experiments have so far not been successful between a DEC X25 router and a Cisco router running DECnet Phase V. Overall I hope that DECnet Phase IV will be all that is required as I believe it will be difficult and complicated to migrate to DECnet Phase V. Finally, note that no access control (security) is possible on these routers, another reason why we believe these routers will eventually need replacing. Cisco routers, on the other hand, provide quite good security features. We will therefore leave the DEC routers running for as long as they are required, but we believe they are a "dead end" for us as far as future developments are concerned.

The Electronic Traffic Routeing (ETR) experiment was approved by Council at its June 1992 meeting, to allow Member States to exchange data amongst themselves across ECMWF links. Monitoring of usage shows a traffic of about 35 Mbytes per day between the two active participants (Germany, Switzerland).

Finally, we will get two new connections shortly from our associate Member States, Iceland and Hungary. The Icelandic 64 kbps line is now under test. When operational, Iceland will join the ETR group so as to be able to exchange data with Denmark. The other new link will be a 9.6 kbps analogue line to Hungary, it is now on order.

2. VAX CLUSTER

As also mentioned before, we have upgraded (and simplified) our VAX cluster. Two VAX 4100s have been installed, replacing some VAX 6000 and 8000 systems. They continue to support the standard NTS services, namely dissemination of ECMWF products, remote job submission, batch output transfer, and file transfer (SENDTM). They also continue to provide the data acquisition and pre-processing services they have always done. However, in principle these two services could transfer to Unix systems if required. It is the Centre's intention to now keep the cluster as is for some 3-4 years at least, although maybe not develop it any further.

3. INTERNET

Our Internet link to ULCC (London University) was upgraded to 256 kbps in January 1994. ULCC is one of the main network nodes for JANET (UK academic network) and hence we connect from there to Internet. The use of this link has grown and grown, especially since the availability Mosaic and the World Wide Web (WWW) service. This link is also part of the GP-MIMD2 project where we need a high speed link to CERFACS (Toulouse). Overall, we may need a 2 Mbit/sec link to ULCC at some stage.

Initially, our Internet link level 2 protocol was X25, but recently we moved to native IP (TCP/IP over HDLC). Again, initially, all Centre machines had Internet access, but with increasing security problems we decided to restrict access to one system only, where we can exercise control more tightly.

We have one external server (Daphne) which is outside our security "firewall", its address being daphne.ecmwf.co.uk. It is used as an external mail server by staff on mission and it also provides an anonymous ftp service to the world.

4. ISDN PILOT PROJECT

On request from Météo France we have run an ISDN experiment, to determine if an ISDN link with our Member States could be used as a backup to the normal leased line, or as additional capacity on an overloaded link. We had to wait some time for the ISDN-2 link from BT, then again had to wait for a suitable interface to Cisco routers hardware. Tests of a digital dial-up service were then set up with France, running the PPP protocol, with CHAP (Challenge Handshake Access Protocol) to provide a level of security. So far we have made the PPP protocol work, but failed to get CHAP functioning correctly. Tests continue still, with the active help of Cisco (UK).

As an aside we have also used the ISDN link to provide an internal test bed to imitate a Member State 64 kbps line over which we could try out various XWindows based interactive services such as Metview.

5. PUBLIC X25 (PSS) LINK

We still have a public X25 link which can be accessed from workstations or the VAX cluster. Its main use at the moment is for Omnet, to access various mail boxes used by the meteorological community. Later it may be possible to access Omnet via Internet though.

5. UNIX SERVICE FOR MEMBER STATES - Richard Fisker

1. MEMBER STATES "ecserver" SYSTEM

The UNIX service for Member States was designed to replace the NOS/VE interactive service that terminated on 30 November 1993. Also at that time we had some other internal requirements to satisfy, hence we purchased a Challenge server from Silicon Graphics to cover all those needs. Its host name is munin with an alias of ecserver. The alias was created and advertised so that later we can move the service to another server without users having to learn a new server name. Thus, Member State users can access ecserver, as well as the CRAY C90 and the T3D (for specific projects only). However, they cannot access any of the other systems at ECMWF.

The specific hardware used currently for ecserver is a Silicon Graphics Challenge (with dual MIPS R4400 150 MHz processors), 384 Mbytes of memory, 14 Gbytes of disk space and one connection to the FDDI ring.

This Challenge server, in addition to providing a Member State service, is also used for:

- SGI compute server for ECMWF Operations Department users
- Internet gateway for all users (alias igate)
- MARS cache server for workstation users
- ClearCase server for ECMWF Operations Department users, ClearCase being a source code control system (it replaces the NOS/VE source Code Utility SCU)
- file server for IRIX 5.2 workstations. There are few such workstations currently, but they will grow in number shortly as we migrate all the SGI workstations to IRIX 5.2
- Empress database host for Computer Operations statistics (REPGEN)

As stated before, ecserver has 14 Gbytes of disk space, split as follows:

• system (including backup and test systems)	3.8 Gbytes
• /tmp (control by a Select/Delete mechanism)	0.5 Gbytes
• /usr/local	1.2 Gbytes
• MARS cache	4.8 Gbytes
• scratch space for large temporary files for ECMWF Operations Department users (again controlled by a Select/Delete mechanism)	1.8 Gbytes
• ClearCase database	1.9 Gbytes

We keep all home files for users on a Control Data 4680 system (host name odin), which has two sets of RAID disks. Member State users have part of a 3Gbyte partition for their home directories. In addition, these file systems are fully backed up (to CFS). Hence with both the RAID disk technology and the CFS backups, we hope we can provide a secure location for all your home directories. To ensure a few users do

not take all the available space a quota system is in operation (typically 10 Mbytes per user).

The CRAY C90 /ec and /tmp file systems are mounted on ecserver, where they are known as:

/unicos/ec	or	/sn4004/ec
/unicos/tmp	or	/sn4004/tmp

Thus users can quickly access/modify C90 files without having to log on to the C90 or run a batch job.

2. SERVICES PROVIDED

Any user with a basic VT terminal, or VT emulator, can log on to ecserver and work in character mode. The standard Unix facilities are available, including:

- shells - Bourne, C, Korn
- editors - vi, (plus emacs for compatibility purposes only)
- file viewers - more, less
- RQS for job submission to the CRAY C90

Most ECMWF users are now used to using XWindows based interfaces. However, not all popular packages are yet available on Silicon Graphics systems (e.g. WordPerfect, Island Graphics), hence we provide, for internal users, the capability to run such packages on a Sun server. As mentioned in a previous talk, we have tried an experiment whereby we tested the use of X based services across a "pretend" Member State 64 kbps link. Although packages which rely upon heavy graphics usage (e.g. Metview, Mosaic) ran poorly, others such as our X based editor (called ie) ran quite happily.

3. SERVICE AVAILABILITY

Recently, we have had a period of poor availability on ecserver, which mostly occurred after we upgraded its operating system to IRIX 5.2. Several crashes, originally diagnosed as hardware errors, turned out to be due to a monitoring package we were evaluating. Another set of problems, reported as "stack underflow/overflow" occurred mostly when we ran Metview. This turned out to be an SGI IRIX problem which occurred when the system attempted to write a core file to an NFS mounted file system. We are awaiting a fix from SGI for this, in the meantime we have set "coredumpsize" to zero.

Occasionally we get automounter problems which are still under investigation. We may get round them by using a public domain version of automounter known as "AMD". Incidentally, if anyone has experience with AMD we would be interested to hear from them.

At the moment it is not possible to submit jobs to the CRAY T3D from ecserver as RQS does not support the extra parameters required by the T3D. We are thus planning to evaluate CraySoft's Network Queuing Environment (NQE) software as a replacement for RQS. NQE could also be used to provide a batch service on ecserver, something we want to do in the long run. This would then allow us to offload some of the smaller jobs from the CRAY C90.

6 CRAY T3D SERVICE - Neil Storer

1. CONFIGURATION

The "T3D" in fact consists of a CRAY YMP-2E that front ends the actual T3D. Connected to the Y-MP 2E is an IOS, this IOS containing 4 IOPs. These 4 IOPs provide the following connections:

IOP-0	operators workstation and FDDI link
IOP-1 & 2	16 DD62 disks, providing ~40 Gbytes of disk space
IOP-3	HIPPI connection to the other major system via an NSC switch

The YMP-E itself has two CPUs and 64 Mwords (512 Mbytes) of main memory. Of this 64 Mwords about 40 Mwords is available to user jobs, the rest being used for the operating system and some ldcache space. Note there is no SSD on this machine.

The T3D consists of 64 nodes, with 4 spares. Each node (see figure 7) has two processor elements (PEs) connected to a Block Transfer engine (BLT) and to a network interface. Each node is connected to its 6 neighbours. At the edges nodes are connected to their equivalent neighbours on the opposite edge, creating in effect a 3 dimensional torus. Hence the origin of the name T3D (3D Torus). In addition there is an extra node which is slightly different from the others, it has a 200 Mbyte/sec connection to the Y-MP 2E and hence is known as the I/O gateway (IOG) node.

Effectively, one has two machines. Some people have the misconception that everything runs on the T3D itself. This is not true. Tasks such as I/O to disk, compilation, loading etc. take place on the Y-MP 2E. Note that the T3D is entirely dependent on the Y-MP 2E, if the 2E crashes then everything on the T3D is lost. Conversely, if the T3D crashes the Y-MP 2E can continue to operate quite happily.

The basic T3D resource that the user requests is a number of nodes. Users must request these in powers of 2 i.e. 1,2,4,8 etc. If you literally request just one PE you will be given one node (two PEs) and then the unused PE is not available to anyone else. In geometric terms you are given nodes that form simple "boxes" (in 3D). This can mean that if your job wants 16 nodes, then although 16 nodes may be available, if they are fragmented into several smaller "boxes" your job will not be scheduled. To be scheduled there must be available the requested number of nodes in one rectangular box.

In more detail, each node consists of two PEs (see figure 7), each PE having one DEC Alpha 21064 chip (rated at a theoretical 150 Mflops) plus 64 Mbytes of memory. Memory access from a PE to its own local memory is reasonably fast. However, access from a PE to memory in another PE must go through the network (at 300 Mbytes/sec), this slows down memory access considerably. However, as each node connects to 6 neighbours and all these connections can be doing memory accesses at once, a total throughput of 1.8 Gbytes/second is possible, which is quite high for MPP systems.

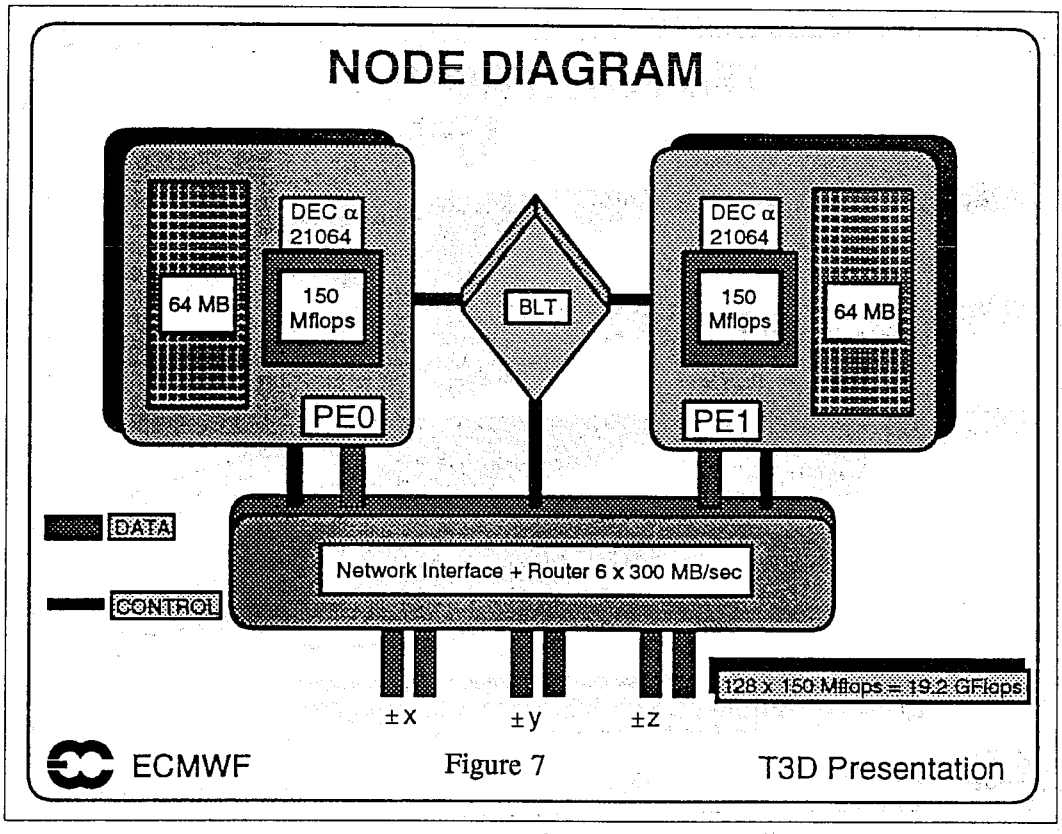

Figure 7

In theory each PE can run at 150 Mflops, there are 128 PEs in our T3D, hence there is a theoretical maximum performance of 128 x 150 Mflops, or 19.2 Gflops. That is an impressive figure, more powerful than our C90! In practice, based on experience with our IFS code, users only see around 10 to 12 Mflops per PE, making the full machine equivalent to a 1 to 1.5 Gflop system.

The main difference between the T3D and CRI's other machines is that on the T3D each PE has its own local memory, which can be accessed (albeit slowly) by other PEs. On CRI's other systems (e.g. the C90) the memory is truly global. Any of the C90's 16 CPUs can access any part of memory equally quickly. This then is the main difference between MPP (Massively Parallel Processor) systems such as the T3D and PVP (Parallel Vector Processor) systems such as the C90.

2. T3D JOBS

Your job runs initially on the YMP 2E and asks for T3D resources (a number of nodes). It is sharing the system with other jobs. Look now at figure 8 where we have three jobs running in parallel.

The first job running on the Y-MP 2E creates a T3D binary (a.out) and requests 64 PEs, he gets 32 nodes in a rectangular "box" as shown in figure 8. The second job asks for 32 PEs (16 nodes), the third job also asks for 32 PEs and gets the remaining block. If the third job had asked for 64 PEs it would have had to wait, as there are not 64 free PEs available.

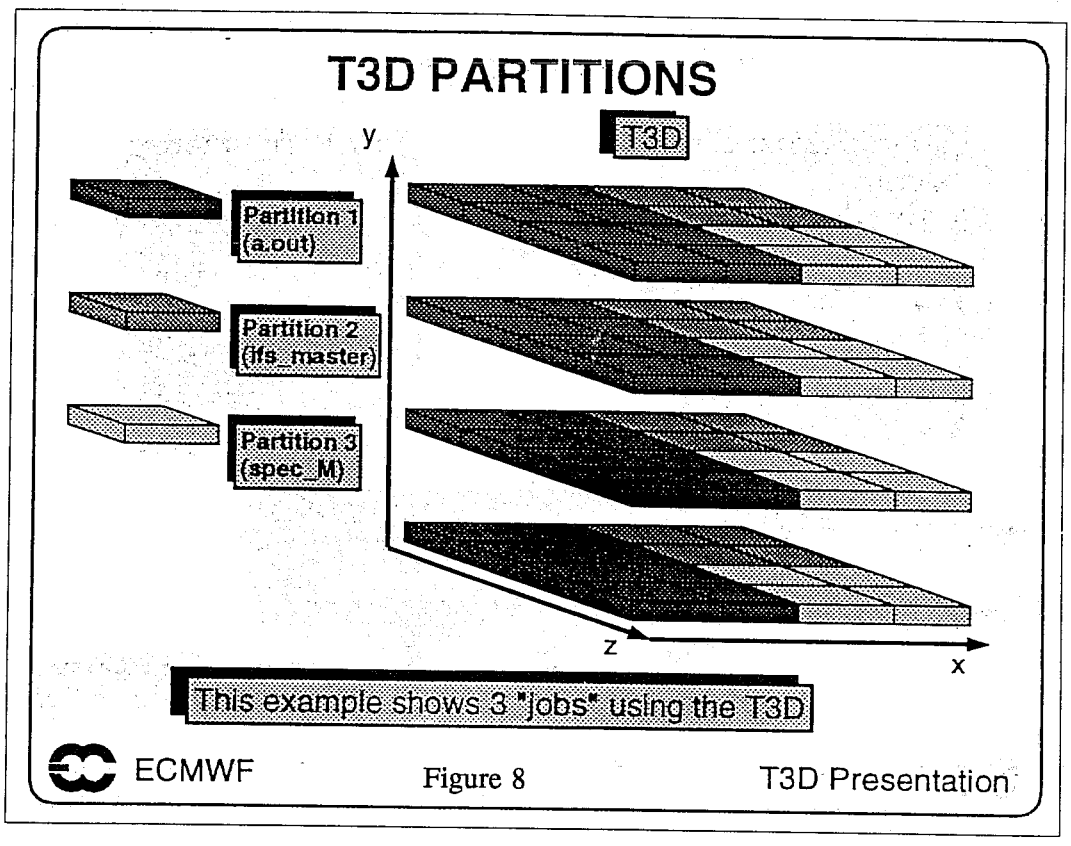

Let's look at the first job in more detail, as shown in figure 9. Note there are two sets of NQS (QSUB) parameters, the first set asks for Y-MP 2E resources, the second set is new and asks for T3D resources:

mpp_p	maximum number of PEs required (must be multiple of 2)
mpp_t	total time, in seconds, for the entire job
mpp_t	total time for any one process

These QSUB parameters are interpreted on the machine the job is coming from. At the moment NQS/RQS on the C90, EL, VAXes and 2E itself can handle these new parameters, but RQS on the workstations cannot. We are evaluating a new product, NQE, on the workstations that will handle these new parameters.

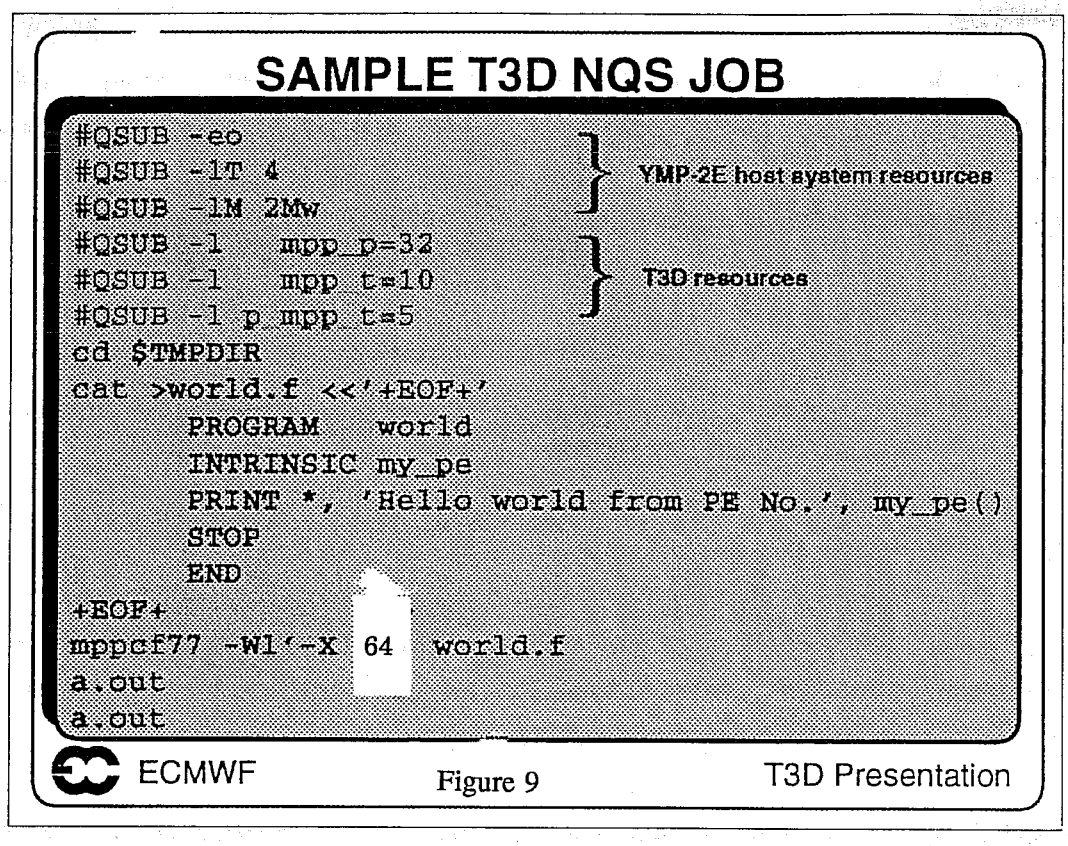

All commands from cd down to mppcf77 run on the 2E. Note the compiler name is mppcf77, this produces a binary for the T3D. The normal compiler, cf77, would produce a binary a.out that will only run on the 2E itself. Note the -X parameter on the mppcf77 command that specifies the number of PEs required (64 in this example). Note that the time limit on the T3D is elapsed time, not CPU time.

In more detail, when your job on the 2E requests T3D resources, a process called "mppexec" starts on the 2E. In turn, mppexec copies your binary a.out, via the Unicos 8 operating system, to the requested number of PEs. The process mppexec is in fact a multitasked program, there being 1 task for every 4 PEs. Hence, if you ask for all 128 PEs the mppexec process becomes quite highly multi-tasked.

Each PE has a very small operating system running on it, called Unicos Max (Micro Kernal Architecture). This operating system is very basic, in fact whenever it needs a normal operating system service (e.g. give me the time of day) it has to go to Unicos 8. The way it is done is rather tortuous. Look at figure 10. If your binary a.out running on a PE wants to know the time of day, it sends a system call to Unicos Max. Unicos Max, in turn, then connects, through Unicos 8, to mppexec running on the 2E and passes it the "time of day" request. Now mppexec asks Unicos 8 the time of day. When it received the reply it is passed back to the relevant PE, to its Unicos Max and finally a.out receives the time of day. This procedure is followed whenever a.out requires a system function, in particular this is how I/O is done. Hence a word of warning, try to avoid as much as possible any system calls from within your executing T3D program. Have the collection of all such information, including reading any input files, done on the 2E before extending your binary on the T3D. As an aside, this is why mppexec on the 2E is multitasked. If it were not, and many PEs issued system requests simultaneously, then the bottleneck of all such requests funnelling through one task on the 2E would be intolerable.

Various job queues have been created on the 2E, some with T3D resources, others without. A "qstat -m" will show you the current queue structure. Note this may alter over time as we gain more experience with the running of the machine. At this time, about the only control the operators have over a T3D job is through the queue definitions and their limits.

Job scheduling considerations are fairly rudimentary:

- partitions are available, each providing 2^N PEs in "rectangular boxes", where N=0,1,2, etc.
- PEs cannot be shared, once a set of PEs has been allocated to your job, you hold them for the entire duration of the job, whether you actually use them or not
- partitions cannot currently be suspended or check pointed, the only way a job stops is when it terminates itself, or it hits a time limit, or it is killed. In a future release of the operating system there will be an operator facility to roll out a T3D application
- jobs using the T3D may need to be killed and then rerun to allow operational work to execute. Thus to avoid your job being rerun from the beginning if it is killed, it is a good idea to build periodic restart points into it
- T3D time limits are elapsed wall clock times, not active CPU time. Thus, as said before, avoid as much as possible using any facility from the T3D that takes a variable amount of real time to complete e.g. any I/O, MARS, or ECFILE requests etc.

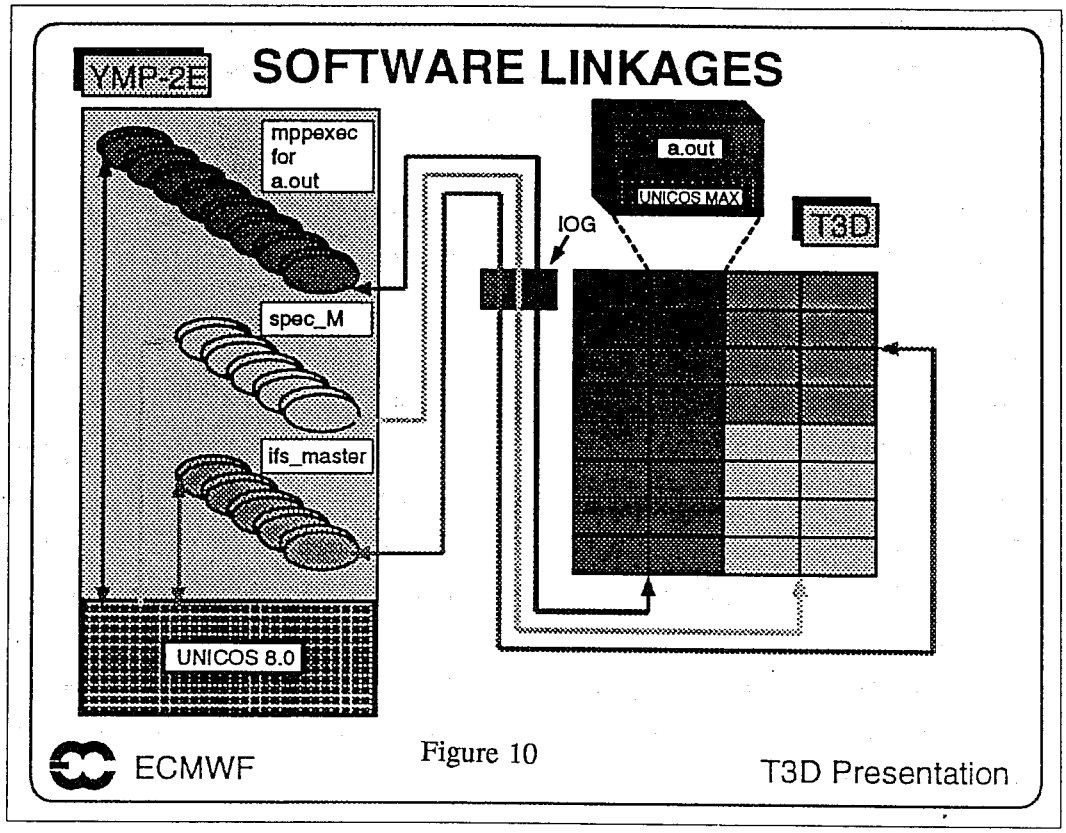

Figure 10

3. T3D PROGRAMMING

Like most MPP vendors, Cray provides two basic ways to program the T3D. One is Cray proprietary and is called CRAFT (CRay Adaptive ForTran), the other is non-proprietary and is called message passing.

Programs written in CRAFT will only run on Cray hardware. The actual style uses commands which look like Fortran comments, in a similar fashion to the way in which autotasking is done. **Note that it has been decided within ECMWF that we will not use CRAFT.** This is because the C90 replacement may well not be Cray compatible and then any programs written in CRAFT will cease to run and have to be rewritten. On the other hand, message passing is non-proprietary and hence programs written using it can be run on many different vendors equipment. ECMWF uses the message passing system known as PARMACS. The IFS model, using PARMACS, has already been run on probably ten different architectures.

Message passing consists either of macros or subroutine calls and is similar to doing I/O. Other message passing systems exist, in particular PVM (in fact PARMACS uses PVM at its lower levels). An emerging standard MPI (Message Passing Interface) is also slowly coming into existence.

Under PARMACS, a program is split into two portions, a host part which acts as "controller" and a node part where the work is done. Often the host part will run on the 2E, while the node part is copied to however many nodes (PEs) you ask for. The host is usually very simple, it starts by splitting the input data amongst all the PEs, waits while the PEs execute, then it collects together the final results and terminates. The nodes perform the actual calculation, including using PARMACS calls to move data from one PE to another as it is required. The "PARMACS portion" of a program mostly consists of macros which SEND and RECEIVE data to/from PEs. One PE SENDs data from its memory to a second PE. The second PE issues a call to RECEIVE the data and temporarily stores it in its own memory.

QUESTIONS

J. Greenaway — Could you comment on the different arithmetics used on the T3D and the Y-MP 2E?

N. Storer — The YMP-2E uses a Cray proprietary arithmetic (floating point) format, whereas the T3D uses IEEE floating point format. Thus, if you produce a file on the T3D that is to be processed on the Y-MP 2E (or C90 or EL come to that) you need to convert the floating point number format. The easiest way is to use a parameter on the assign statement to cause the T3D program to write CRAY-format floating point numbers. Alternatively, there are conversion packages available which you can call.

S. Pasquini — So the recommended way is to collect all one's input data first on the 2E, launch the T3D program, then, when it terminates, use the 2E to write the data?

N. Storer

Yes. Beware that as each PE creates output data it should be written as a separate output file, one per PE. The 2E is then used to finally concentrate the output files into one file at the end. Otherwise, if every PE writes to the same file, then, unless the program used AQIO to write a "direct access" file, the data will arrive in that file in random order. No two runs would produce the same output file! That could be very confusing.

Note that at the moment, Member States who wish to use the T3D must submit a written request to ECMWF. A User Support person will then be assigned to assist in getting the work running and seeing that it continues to run smoothly. This will last to the end of this year. Then, as from the beginning of 1995, allocations will be available on the T3D for each Member State and Member State users can use it in a similar fashion to projects that run on the C90 currently.

J. Greenaway

When your a.out finishes on the T3D, can you then release your PEs for other jobs to use?

N. Storer

This in fact happens, but NQS scheduling may refrain from starting any other "T3D jobs". A utility is available called "qrelmpp".

D. Lucas

What do the new QSUB parameters look like in the VAX format?

N. Storer

Not sure, I will have to check. Anyway, remember that VAX users can now use the format shown in figure 9.

7. DISSEMINATION AND MARS UPDATE - John Hennessy

1. DISSEMINATION

Originally dissemination was in GRID and ECMWF bit codes, but now increasingly GRIB code is used. In fact, it is policy that all new enhancements, new products, etc. will only be in GRIB code in future. On the GTS, products are sent both in GRID and GRIB codes.

There has been an explosive growth in demand from Member States for products over the past 12 months. Today we send out 60,000 products each night on your 64 kbps links, this amounts to about 334 Mbytes of data. You may be interested to know that when we started in 1979, on the first night we sent out 27 products on two 50 baud lines! These figures include all product groups: atmospheric forecasts (midnight and 12Z), the global and mediterranean wave models, and the ensemble forecasts. Note that as of July this year, the global wave model went to 1.5 x 1.5° resolution.

With recent enhancements, dissemination in GRIB now includes the following:

- anyone can now define any area they require
- the data may be on pressure or model levels, including surface data
- any parameter listed in the manual is supported
- spectral or grid point data at any resolution can be requested
- "hollow box" boundaries are available for limited area model input use
- users can define their own bit map to select particular points
- the area that may be selected now includes a sub-area for quasi-regular Gaussian grids
- a maximum message length of 15000 octets means that larger products have to be split using BLOK code. However, at CBS this year BLOK "died". In future to concatenate parts of a GRIB message, WMO will recommend use of the TXX option. ECMWF will continue to support BLOK for those Member States that still use it.
- we see the future as going in the direction of ftp, TCP/IP etc., and this will permit the dissemination of pure binary files (i.e. data with no structure or embedded BLOK code). This facility is currently under test and will be discussed at the Technical Advisory Committee this week.

- a choice now exists for packing accuracy.

These new options are selected as follows:

FILE	= BINARY	(BIN)	for binary files
	= VMS BLOCKED	(VMS)	for current files (default)
ACCURACY	= n		n is the number of bits required
	= NORMAL		accuracy as now, default
REPRES	= QG		for quasi-regular Gaussian grid
GRID	= 160		

On 23 August 1994 we changed the GRIB code definition we use here at ECMWF. We added in some extra bytes in section 1, reserved for local use within the GRIB code This section is already used in the Ensemble products. It will also be used in the Wave products.

We also changed to complex packing of spherical harmonics, because we found with the simple packing that at T213 resolution some parameters were no longer at the required accuracy.

All in all these changes should have been transparent to users. The archived data and the data produced every night by the model already contains all these changes. So far dissemination has not been affected. However, we would like to introduce into the dissemination the new Section 1 as soon as possible as there are several advantages, in particular both ourselves and Member States have had problems distinguishing between various products in the past (e.g. test suite runs from operational runs).

The changes will be in octets 41 to 52 of Section 1 (the area reserved for local use). In more detail:

Octet 41 a local definition number

1. product labelled with MARS parameters (class, stream etc.). This is already used in the daily analysis and forecast, ensemble forecasts, and will be used in the wave products

2. clusters of ensemble forecasts (mean, standard deviation, etc.)

5. forecast probabilities (temperature, rainfall)

Other numbers are also used internally.

Octet 42 the class of the data which, as far as Member States are concerned, will always be 1. Research uses 2, while 3 means re-analysis (this is a MARS parameter also).

Octet 43 type e.g.

1. First guess
2. Analysis
3. Initial Analysis
9. Forecast
10. Control Forecast
11. Perturbed Forecast
14. Cluster Means
15. Cluster Standard Deviations
16. Forecast Probabilities

You will see that 10 onwards are for ensemble forecasts

Octets 44-45 Stream. There are quite a number of these, those relevant to you being:

1025 daily operational (atmosphere model)
1035 ensemble forecasts
1045 wave products

You can use these numbers to distinguish easily between products

Octets 46-49 Version number - 2 for the short cut off 00 UTC analysis, 1 for every other product, 11 and above for test suites and parallel runs, etc.

Octet 50 the ensemble forecast number (0 if it is not an ensemble forecast)

Octet 51 the number of forecasts in the ensemble (0 if it is not an ensemble forecast)

Octet 52 Always 0

2. MARS

I have already mentioned above some of the changes that have recently taken place e.g. GRIB code changes, addition of local headers, complex packing of spherical harmonics. As a transition aid, for those who cannot immediately handle these changes, we have temporarily introduced a new MARS option -G. This will convert the complex packed spherical harmonics to the simple format used before. It will also remove the local information from Section 1. However, we would like people to move to the new options as soon as possible. Bear in mind the "temporary" option -Z, introduced in 1991, to bridge the change from GRIB edition 0 to edition 1, is still in use today! A recent check through the logs showed at least half a dozen

users still using -Z. We hope that a similar experience does not happen with the -G parameter.

Some of the data from the Re-analysis project has already been stored in MARS. It took quite a lot of work, especially to pack a month's worth of analyses into one CFS file. This was required in order to cut down the growth in the number of CFS files.

The MARSINT interpolation package is being rewritten. It is already being used in the dissemination and will soon be available to in-house users.

We plan to produce pure binary target files. At the moment any MARS file retrieved on the Cray is in Cray blocked format. We want to give the option to have the file in the same (binary) format as binary dissemination files. We have routines available that will read messages from binary GRIB or BUFR files. These routines are well documented and are available to you now.

Finally, something which will involve the MARS team in a lot of work is the replacement of the CFS and IBM Data Handling System over the next four years.

Data which has recently been added to MARS includes:

- monthly mean data:
  - Analysis data from July 1985
  - Forecast data from August 1985
  - There is no initialised surface data
  - The normal conversion retrieval facilities are available e.g. sub-areas

To retrieve monthly mean data use:

STREAM	= MONTHLY [MO]	
DATE	= yymm00	the 00 as the day indicates monthly mean

This information will be included in the next edition of the MARS guide.

3. WORKSTATION MARS

This is an experimental version of MARS which runs on a workstation rather than a Cray system. It is, in fact, a new MARS program. Bear in mind the Cray version is over 10 years old and a lot has happened in software design techniques in those 10 years. We therefore propose to take advantage of these advances, and also to take advantage of Unix system facilities. The existing MARS is written in Fortran, the new one is written in C. The new one will also be ported to the Cray, together with the new MARSINT package.

This new version will access a variety of databases, starting with data on the user's local disk, then the Empress relational database which contains the MARS cache (as a caching system), the Cray Fields Database

(FDB), and, as a last resort, the IBM data handling system archive. When we get a new data handling system it will become just another source of data to this version of MARS.

Finally, I would like to deal with some of the points raised in the Computing Representatives own presentations.

S. Pasquini — Could the 00 UTC products be sent out earlier than currently?

J. Hennessy — When this project was set up by Council it was to provide boundary conditions for Limited Area Models (LAM). It was stated in the Council paper that the products would be useable up to 08.00. Hence, it is not late if the products are delivered before 08.00 UTC. Looking at the actual times the products have been available in the last 6 months, on only 3-4 occasions have the last products not been available before 05.15 UTC. Bearing in mind there is a 3 hour cut off, this is a good record for a 00 UTC based forecast. Because of all the work involved in actually producing this forecast it is therefore difficult to get it out any earlier.

N. Olsen — Why is it not possible to repeat a single product?

J. Hennessy — One used to be able to, but unfortunately it is no longer possible. In the original "Product Catalogue" every product had a 6 digit number. Each Member State pre-defined its area, grid-mesh, etc. from a prepared list and there was no flexibility. With a catalogue number it was then trivial to repeat a product. However, with the introduction of the new flexible dissemination system, Member States can now prepare any request they like, a different area and grid for every product if they wish. With this flexibility there is no numbering scheme possible. The only way, therefore, to repeat a product is for the Member State to submit a full repeat "request" command specifying area, grid, etc. It is a little bit pointless for ECMWF to develop such a "request" command as it is, in fact, just a MARS request followed by an ftp, SENDTM, or whatever. Therefore, we suggest that you use the MARS retrieval mechanism for single product repeat requests.

B. Barg — Requests for test data sets are at present difficult to handle.

J. Hennessy — There are a number of options:

(i) set up a special "Group Code" - this is quick to do;

(ii) use the "delay transmission" feature, so that such test products would be held back until after the normal dissemination has completed. Then use the

repeat facility to get that file of test product;

(iii) it was suggested to use MARS.

I am not too happy about that last idea, at the moment, because there are a number of differences between MARS files and dissemination products. Dissemination allows you to use bit maps, frames for boundary values, sub areas for a quasi-regular grid. None of these are in MARS. We do plan to introduce these features into MARS so that MARS products and dissemination products will be identical. Hence this problem will disappear over the next year I hope. However, in the short term I suggest you ask us to set up a group code (option (i) above).

B. Barg

We propose that you add to MARS a Member States local GRIB use definition:

Octet 41 local definition number: 192 - 255

Octet 42 National information

This would allow Member States to define local information which does not conflict with ECMWF usage.

J. Hennessy

Local use of GRIB code, Section 1 octet 5 identifies the Centre - 98 identifies ECMWF, 78 identifies DWD, for example. ECMWF has defined what it will use this local part for, DWD is free to define whatever they want to use it for. It has been suggested that we reserve some global numbers for Member States. I would like to think about this to see what can be done. A possibility may be to use octet 26 to define a sub centre. This may give us a way to solve this problem.

QUESTIONS

S. Pasquini

ECMWF's software, provided to Member States, will reconstruct large files from the BLOK code. Will this software be updated to include the new WMO standard method of transmitting large files?

John Hennessy

With binary files you will not need any deblocking software at all. BLOK was always an experiment. WMO want eventually to go to binary file transmission without any blocking required. This is the way we see the future too.

S. Kruizinga

Could you identify, to the relevant Computing Representative, who is still using -Z? Maybe users are still using it without realising it.

J. Hennessy

We will do.

B. Barg — When do you plan to make pure binary target files available from MARS?

J. Hennessy — Late this year, early next year.

S. Pasquini — Do you foresee that workstation MARS will be easily transportable to another platform e.g. IBM or DEC?

J. Hennessy — Bearing in mind the way the new MARS is constructed, we cannot see any reasons why it cannot be easily interfaced to many different databases. You have to write a "database module" which translates the request into whatever is required to access the data on your particular machine.

S. Pasquini — If you access Empress using SQL, is it possible to use the same SQL commands to access another database which recognises SQL?

J. Hennessy — We do not use SQL. We use the C level interface into Empress which is much more efficient.

S. Kruizinga — Some important changes to dissemination are coming, for example, the change from BLOK to binary. If one wants to test such changes it can be difficult. It would be useful, therefore, if test dissemination files (in the new format) could be made available - say in ecfile - for anyone to pick up so he can test his software. Is this possible?

J. Hennessy — Yes, use the Cray utility "uscpblock". Retrieve the file from MARS in Cray blocked format, run it through uscpblock and you have a pure binary file immediately. Alternatively we could set up a test file if required.

8. USER SUPPORT SERVICES - Andrew Lea

1. REGISTRATION AND ACCOUNTING

Since the last meeting we have simplified the annual registration procedure. No longer do you have to re-register each user at the beginning of the year. Existing users are now just carried over from year to year. From our point of view it all seemed to go well, we had no adverse comments from yourselves.

Geerd Hoffmann, in his "Status and Plans" review talk, has given the reasons why we propose to make a change to the C90 account unit definition next year. It amounts to changing the overall multiplying factor from 0.00284 to 0.00341, an increase of 20%. This proposal will go before the Technical Advisory Committee meeting later this week.

Beginning next year each Member State will have an allocation on the CRAY T3D which you can allocate to your projects as you wish in the normal way. The number of units available for all users will be 10,000 per day. However, at this time the account unit for the CRAY T3D has not yet been fully defined. We are currently running with a test formula based on the following two parts:

Y-MP 2E	-	definition the same as for the C90
T3D	-	the number of PEs requested multiplied by the elapsed (wall clock) time

These two parts will then be combined in some ratio yet to be decided, and finally multiplied by a normalisation factor to convert to 10,000 units per day.

Every 4 weeks you receive from us, through the post, a list of your current C90 projects and their cumulative usage through the year. However, most Member States can now access files here at the Centre through ftp, SENDTM, etc.; plus more and more of you have interactive access. Thus we would like to propose that in future we send these reports in some electronic form, perhaps via email, or in a similar fashion to news sheets.

2. DOCUMENTATION

ECMWF's policy is to provide more and more of its documentation as on-line files. The current system, echelp, is based on the Internet product gopher/xgopher. It access a series of directories and files which contain much information about ECMWF services e.g. computer bulletins, news sheets, write-ups, etc. We will continue with this policy and, as before, will be reducing the number of Computer Bulletins still further as they are replaced by other forms of on-line information.

Within the Centre we are beginning to experiment with Mosaic, which again some of you may be familiar with via Internet. As well as displaying documents, graphics, etc. it allows the development of hyperlinks between documents, giving the user the ability to get more information on a specific topic via those links. Mosaic is X windows based and hence, for Member States to use it, it comes back to the question of how

well an Xbased product will work across a 64 kbps link.

Manufacturers are providing more and more of their manuals on-line:

- Silicon Graphics have a system called Insight, which is a browser for searching through on-line manuals;

- Cray have experimented with on-line documentation. Currently they provide Doc-view, but at the next release of Unicos (level 8.0.3) Doc-view is replaced by Cray-doc.

Again, as an experiment Cray has provided some of their T3D manuals as PostScript files. We have copies which can be made available to any registered user of the ECMWF system.

Finally, it should be remembered that as far as copyright is concerned, on-line manuals carry the same copyright protection as do the paper manuals. Thus, again, these on-line manuals should not be passed to people who are not registered users of the ECMWF computer systems.

3. TRAINING

ECMWF will present another series of Member State computer training courses in Spring 1995, probably in March.

It has always been the policy of ECMWF to provide training only in those aspects of the computer service that are not covered by commercially available courses. For example, we teach local facilities such as MARS and ecfile, but we have never taught Fortran as this topic is covered by many commercial training companies. For the last three years we have run a Unix course, initially for in-house staff but later for Member States. We now find that more and more Unix courses are available commercially, hence we propose to drop the one week Unix course from our next set of courses.

Finally, note that ECMWF is adding another one week course to the set, namely on MPP (T3D) programming.

4. OTHER TOPICS

The only other item to bring to your attention is the article in the summer edition of the ECMWF Newsletter (No. 66 pp42-44) concerning the mark 16 release of NAG. Again, there is quite a collection of new routines in this release, plus a few existing routines being deleted.

QUESTIONS

A. Dickinson — Are you going to base all future on-line documentation on Mosaic?

A. Lea — It is a possibility. Currently, the on-line files are plain text files, viewed with gopher or xgopher. The question is one of man-power to go through our existing files and create the hyper-link structure.

On-line documentation is clearly popular. An analysis of access to our present on-line documentation showed that two thirds of all documents had been accessed by just over half our internal users (based on records collected over a 2 month period). Thus, as said before, we will continue to move in this direction and, if we can find the manpower to prepare a hypertext version, then that will be of benefit to all.

9. SECURITY ISSUES - Walter Zwieflhofer

1. RECENT SECURITY ADVISORIES

ECMWF, and probably most Member States, monitor various sources of security advisories, including CERT, CIAC, NASIRC, NASA and USENET. Of these, probably the most useful one is USENET which we find, on a daily basis, gives lots of useful hints.

To see the degree of the problem, CERT alone has issued 16 alerts in the past 12 months. Of these, 8 applied to one or more ECMWF Unix system. I will discuss three of these as examples of the serious security threats we are faced with.

1.1 Sendmail

This was a particularly nasty problem as it allowed a hacker to breach a firewall system. A mail would be sent to a mail daemon, which would fool the daemon into executing a script contained in that mail message. The script would be executed under any user id except root.

1.2 Network monitoring attacks

A feature called "etherfind" on SunOS systems allows one to monitor network packets. It is used a lot within ECMWF. A hacker would use one of the other known bugs to become root on a system, then would modify certain programs such as ps (to prevent the hacker's program being displayed by ps), login (to ensure he could always get back into the system), and netstat (displays active network connections). This way, the hacker would monitor the system for passwords completely invisibly. The etherfind program would be used to capture the first 128 bytes of any network connection, including of course the user's password. CIAC estimated that by March this year this technique had compromised over 100,000 user accounts worldwide.

1.3 "8LGM" advisories

A group, including previously convicted hackers, sent out mail messages containing so-called "advisories" in the same style as bodies like CERT. The first set of mail they issued comprised 4 advisories, 3 of which allowed one to become root, and in those advisories gave full details (in an easy to use cut and paste style) of what to do to break into a site.

Several of these examples would have worked on ECMWF's Member State server (ecserver) and so they show how important it is to provide more secure access to ECMWF's system, hence the smartcard proposal (see below).

2. EXTERNAL ftp SERVER

Several Member States already have a similar system to that being brought into service at ECMWF. Just a word of warning therefore.

At ECMWF we have deliberately not allowed anonymous ftp write access. This is because those sites that do allow such access very soon end up as a store for illegal software, where third parties use it to exchange software that should be licensed but is not. Thus, if your site is preparing a similar system, be aware of this point.

3. IMPORTING SOFTWARE AND DATA

More and more of our users are discovering how useful Internet can be. There are many very good programs which are either free or available for a small fee. However, one day a program will be imported that contains a virus, trojan horse, or other hidden malicious software. Two points therefore to be aware of:

- we attempt always to only import source code which is then compiled here. This habit in the Unix world of always exporting/importing source code has meant it has been much more difficult for virus writers etc. to gain a foothold, as against the PC world where binaries are the usual form of importing software.

- data can be dangerous! A PostScript file (containing a picture or document) is itself a program which gets executed in a PostScript printer for example. If you use a PostScript viewer then you are executing the PostScript "program" on your own system (or perhaps on a server). The PostScript language includes statements that open and write files, hence the "program" can delete files and/or introduce viruses. Thus, we have changed the default within Ghostscript (which is used to view PostScript files) to incorporate the mandatory parameter -dSAFER. This forbids the opening of files from within PostScript other than in read only mode. This we have done site-wide and is maybe something you might like to consider also.

4. SMART CARD PROJECT

We have recently run an experiment whereby we tried to evaluate whether a smart card authentication system could be used in ECMWF's batch job environment. This experiment was required because all commercially available systems of this nature work in an interactive environment, but no-one sells a system explicitly for the batch environment.

We decided to use a smart card system called SecurID. A small, 25 card, system from Security Dynamics was installed, together with a public domain version of NQS (source code from CERN). We then wrote local code to combine these two packages. Problems were experienced initially, but with excellent support from Security Dynamics we managed to overcome them. Finally, the experimental system was used by Denmark and The Netherlands and found to work satisfactorily.

The first test was of batch job submission. Here one has NQS running both in the Member State and on an ECMWF system (a gateway machine). The Member State user would submit his job via a command called, for example, ecqsub that operates (from a user point of view), similar to the normal qsub utility. Then ecqsub would prompt the user for the password from the smart card. As the job enters the ECMWF gateway

that smart card password is checked and, if valid, the job is passed on to the target machine for execution. If the password is not valid the user receives a mail message telling him the job was not submitted because of password failure.

We are now considering providing a second version of ecqsub that does not use NQS in the Member State. Instead, it would establish, from the Member State, a direct connection to the ECMWF NQS gateway. At that time the user is prompted for his smart card password. If it is wrong the user is informed immediately. Because there is no NQS at the Member State end, there is no spooling in this option. Hence, if the link is down the user cannot submit work and have it automatically sent once the link returns.

Where NQS is available at both ends of the link, the output will normally be automatically returned to the place from where the job was submitted. This, however, raises security problems. Hence we have modified ECMWF's NQS so that a job's output is returned to the ECMWF gateway machine rather than to the Member State. When the output is complete the gateway will pick it up and will copy it to the Member State submitting machine, via another utility (called eccopy for example).

The most important service to protect is interactive access, as in terms of security, it is the most dangerous. Member State users will telnet to one ECMWF server (e.g. ecserver) and enter the smart card password when prompted. Once accepted by ecserver the user can then telnet to any other permitted system (e.g. C90). The same mechanism will apply to ftp access also.

Unattended remote batch job submission has proved the most difficult to deal with satisfactorily. This is where batch jobs are submitted to ECMWF by a program running on a Member State system, rather than a user. There is a strong requirement for this from some Member States. One answer is a ticket system whereby the user initially gains a ticket or token (using his smart card) which is then valid for, say, 2 days. Any job quoting that ticket will then be accepted by ECMWF without further authentication. This, however, considerably weakens the security offered by smart cards. Hence it should be restricted to a very few user identifiers from specific Member State hosts only. The Member State Computing Representative would control this facility. In particular, they would have a utility which would turn on/off this unattended remote batch facility for specific user/host combinations. It would thus be the responsibility of the Computing Representative to ensure this facility was not abused.

Based on the above experiment a proposal is to be discussed at this week's Technical Advisory Committee meeting. If approved, we propose to buy the system towards the end of this year and provide an interactive service within about 2 months. Then ftp and batch would follow some 3 months later, the whole system being fully functioning by mid-1995. We will not touch the VAX/VMS service in general, but may include VMS interactive access into the smart card system later if required.

QUESTIONS AND COMMENTS

S. Kruizinga — We have installed this experimental system for batch submission. It works well and works fast. It was not realised in the beginning that the card generated valid numbers for 5 minutes rather than just one minute. We had some problems whereby the mail message was not received properly if the passcode was wrong. However, I feel that users will avoid using it wherever possible, and will switch to interactive use of ecserver, submitting their batch jobs from there. Maybe users will also get round the unattended batch problem by running on ecserver and submitting their jobs from a program running on that machine. That means their output will stay on ecserver. Can the user's ecserver program then invoke eccopy to return the output to his own system?

W. Zwieflhofer — Yes, we see no problem with this idea e.g. using crontab on ecserver to submit, say, C90 jobs at specific times. Also, eccopy would be available to users directly to use (replacing SENDTM). As an aside, this may solve the SENDTM priority problem currently experienced by The Netherlands, as eccopy files would not have the high router priority associated with data dissemination files.

N. Olsen — We are very pleased with the system in Denmark. In fact, we would like it to be implemented within the Danish Meteorological service as well, to control access to our own systems.

A. Dickinson — The UKMO already has a smart card system. If ECMWF chooses the same system will one card work for access to both ECMWF and UKMO machines?

W. Zwieflhofer — This was discussed at length with Denmark. In theory it is possible. The SecurID card can be used to select one of three different seed values, generating three different number streams. Thus it could be used for accessing up to three different sites, which clearly the users would prefer (rather than having three different cards). However, the user does have to remember which seed stream to use for which system. We will investigate this idea further to see how feasible it is in practice.

D. Birman — Unattended remote batch - french users have a requirement to access MARS at ECMWF from within jobs running on machines at Météo France. The ticket based system could be a good short term solution for them, but a better long term solution would be automatic access to MARS, outside the smart card system. This could probably reduce the need for ticket based usage considerably.

W. Zwieflhofer — I agree. Where access is restricted to one application only (e.g. MARS retrieval) then it could well be outside the smart card system. We will look into it.

G-R. Hoffmann — There is a discussion within ECMWF about remote MARS access from Member States, where a MARS client would run on your machine locally in Toulouse and contact the ECMWF MARS server to get the data. Some questions still exist about security and the licensing of data. Then it is a manpower issue to find the effort to implement it. Further into the future one could envisage DCE/DFS direct access to our file servers and hence to MARS data. It is something we are planning to look at in the longer term.

10. GENERAL DISCUSSION - G-R. Hoffmann

G-R. Hoffmann — The Netherlands raised the problem of SENDTM files having the same high router priority as data dissemination (DD) files and so killing interactive response while they are being transmitted. Now the SENDTM utility does not know which type of link a Member State has, whether it is a DECnet or TCP/IP link. SENDTM just sends the file to the VAX where it is treated as a normal DD file. It is, therefore, not obvious how we can change the priority. We will continue to look into it.

S. Kruizinga — In the Netherlands we would be perfectly happy to use "eccopy".

G-R. Hoffmann — Of course users in each Member State know what type of link they have, hence if they use ftp (or ftua) that would also solve the problem. In general, it is difficult to modify SENDTM.

G-R. Hoffmann — The next item is the suggestion from Sweden that the Member State Computing Representatives meeting be held annually. Those Technical Advisory Committee representatives here at this meeting are now aware of this proposal and it can be brought to the attention of the Technical Advisory Committee. However, what is the view of all those Member State Computing Representatives here today?
[At this point a show of hands indicated that three Member State Computing Representatives out of fifteen were in favour of annual meetings].

S. Pasquini — The Member State Computing Representatives meeting conveys detailed technical information to the Member State users, whereas the Technical Advisory Committee deals more with policy. However, I would be in favour of an annual meeting if it stimulated Member State users to use ECMWF's services more. The policy in the Italian Meteorological service is to run ECMWF's products and services as much as possible, hence these meeting are important in learning as much as possible about ECMWF's services. However, at these meetings I do not learn a great deal about how to use these services, I just get the information. More information is learnt through ECMWF's training courses, workshops and seminars. However, ECMWF staff are busy with many things and their resources are limited. I believe that we should try as much as possible to communicate ECMWF's services to Member States and, therefore, the Member States Computing Representative's meeting is one possible mechanism, but not the only possibility.

L. Campbell — The risk with annual meetings is that there may be in some years very little change to the service. There have in the past been such years. Thus, if one ties oneself too rigidly to annual meetings there may in some years be very little to discuss.

R. Rudsar

Maybe the length of the meeting should reflect the amount of change. There have been quite a few changes this time and therefore this meeting has been rather intense.

S. Pasquini

At the Technical Advisory Committee there is often little chance to ask technical questions about an issue under discussion, especially if the question is really due to a lack of technical knowledge. However, the Member State Computing Representative meeting could be a forum, before any Technical Advisory Committee meeting, to discuss technical issues in more depth. Two examples come to mind:

(i) the use of SQL and commercial databases, in place of a locally developed data base. SQL is a standard, but the MARS interface to Empress apparently does not use it, for performance reasons.

(ii) the use of CRAFT versus PARMACS on the T3D.

G-R. Hoffmann

Regarding SQL and Empress, it is a Centre's management decision that if a commercial database was used then the interface should either follow the standard or be very quickly convertible to that standard (for example if the commercial database supplier went bankrupt). We believe that is the case for the MARS Empress interface.

On the issue of CRAFT on the T3D, Council made it clear that in any future acquisition of an MPP system, it must not be a requirement that CRAFT be available. Hence Parmacs is the chosen interface. CRAFT is there, you can use it, but you must be aware that it could well not be available on the next system. Nor will it be a requirement that it be available on the next system. Therefore you realise the risk that you take if you use it.

Going back to the question of frequency of these meetings, I gather from the views expressed that there is only a minority in favour of annual meetings, but the majority wish to continue as now, namely to hold meetings on an event driven basis. On such a basis the next meeting would probably be in early summer 1996 when the new data handling system was operational and the C90 replacement had been installed.

The third item to be raised concerns the length of these meetings. I appreciate the comment that the 10 minutes assigned for the Computing Representatives own contributions is a bit tight. Hence a full two day meeting might be better. The question is really do you prefer two complete days (9-5 each day) or have it spread over 3 days (starting after lunch on day 1 and ending before lunch on day 3)? The

latter may make travelling easier but it means you are away from your own work one extra day.

[At this point a show of hands indicated:

- two full days 2
- two days spread over three 11]

Thus we will go for the second option, two days spread over three.

B. Barg — I am confused, am I a Computing Representative or a Computer Representative?

A. Lea — I apologise! I am the one responsible for all this confusion, as I can never remember whether it is Computer or Computing Representatives. Hence I went to look in the original paper that set up the scheme of Member State representatives and there it states that you are Computing Representatives. I shall have to try to remember this in future!

G-R. Hoffmann — As there are no further questions I would like now to close the meeting. Before I do, I wish to thank the staff of ECMWF who have been involved in this meetings and, in particular, to User Support for organising this meeting so well. It has been a pleasure to meet you all again and we appreciate all your comments and criticisms which we will take careful note of. It just now remains for me to close the meeting and wish you all a safe journey home.

EIGHTH MEMBER STATE
COMPUTER REPRESENTATIVES' MEETING

E C M W F
Shinfield Park, Reading, U.K.

27-28 September 1994

==========

PARTICIPANTS

AUSTRIA	Dr. G. Wihl
BELGIUM	Mrs. L. Frappez
DENMARK	Mr. Niels Olsen
FINLAND	Mr. T. Hopeakoski
FRANCE	Mr. D. Birman
GERMANY	Dr. B. Barg
IRELAND	Mr. L. Campbell
ITALY	Dr. S. Pasquini
NETHERLANDS	Mr. S. Kruizinga
NORWAY	Mrs. R. Rudsar
PORTUGAL	Mr. C. Fernandes
SPAIN	Mr. E. Monreal
SWEDEN	Mr. S. Orrhagen
SWITZERLAND	Mr. B. Bachofner
UNITED KINGDOM	Dr. A. Dickinson

P R O G R A M M E

EIGHTH MEMBER STATE COMPUTER REPRESENTATIVES' MEETING

E C M W F
Shinfield Park, Reading, U.K.

27-28 September 1994

Classroom

==========

Tuesday, 27 September

09.00	Welcome	G.-R. Hoffmann (Chairman)
	ECMWF's computer status and plans	
10.15	*C O F F E E*	
10.45	Member States presentations	

Each representative will be asked to speak for a maximum of 10 minutes, outlining their Member State's involvement (actual or planned) in the computer service at ECMWF. This should include:

* diagram of own computer equipment, and of connection to ECMWF
* projects run at ECMWF
* experience using ECMWF computers, including suggestions and queries regarding the present service
* plans (involving ECMWF usage over next couple of years)

12.45 *L U N C H* **)

**) During this lunchbreak machine room tours will be arranged if required. Interested participants should contact their User Support Contact Point.

Tuesday, 27 September (cont.)

14.00	MAGICS update and METVIEW/ws	Jens Daabeck
14.30	Operational services (incl. Call Desk)	Peter Gray
15.00	*C O F F E E*	
15.30	External networks	Tony Bakker
16.00	UNIX service for Member States	Richard Fisker
16.30	ATM (general presentation)	Luong Le (Alcatel)
18.00	***COCKTAIL PARTY***	

* * * * * * * * * * * *

Wednesday, 28 September

09.00	Cray T3D service	Neil Storer
10.00	*C O F F E E*	
10.30	MARS and dissemination update	John Hennessy
11.00	User Support services	Andrew Lea
11.30	Security issues	Walter Zwieflhofer
12.00	Discussion session	
12.30	***F I N I S H***	

PART 2

MEMBER STATE COMPUTING REPRESENTATIVES' CONTRIBUTIONS

ZAMG ZENTRALANSTALT FÜR METEOROLOGIE UND GEODYNAMIK

1.a.) C O M P U T E R E Q U I P M E N T Fig.1

Central Systems at Dept.ADP :

Database System (Dual System)
- Sun SPARCserver 690MP (4 CPU's,128 MB Memory)
  - Disk Unit (40GB, IPI dual port, switchable)
  - CD-ROM (644 GB)
  - Tape (8mm, Exabyte, 2.3 GB)
- Sun SPARCserver 690MP (4 CPU's, 64 MB Memory
  - Disk Unit (40GB, IPI dual port, switchable)
  - CD-ROM (644 GB)
  - Tape (Gogatape, DAT 1.0 GB)

Interactive Server :
- Sun SPARCcenter 2000 (4 CPU's 512 MB Memory)
  - Disk Unit (Differential SCSI, 10 GB)
  - CD-ROM (644 MB)
  - Tape (8mm, 5 GB)

Batch Server :
- Sun SPARCcenter 2000 (4 CPU's 512 MB Memory)
  - Disk Unit (Differential SCSI, 10 GB)
  - CD-ROM (644 MB)
  - Tape (8mm, 5 GB)

ECMWF-Server :
- Sun SPARCcenter 10 Model 30 (64 MB Memory)
  - Disk Unit (SCSI, 2.8 GB)
  - Floppy Disk (3.5")

GTS-Server :
- Sun SPARCcenter 10 Model 40 (64 MB Memory)
  - Disk Unit (SCSI, 2.8 GB)
  - Floppy Disk (3.5")

ECMWF/GTS-Server :
- Sun SPARCcenter 10 Model 40 (64 MB Memory)
  - Disk Unit (SCSI, 2.8 GB)
  - Floppy Disk (3.5")

Mass Storage Management :
- Control Data CD 4460 (128 MB Memory)
  - Disk Unit (Differential SCSI, 10 GB)
  - Tape (150 MB, 0.25")
  - Tape (0.5", 9track,, 6250/3000/1600/800 bpi)
  - Roboter (2 Drives,48 Slots, VHS, 14.6 GB/Cartridge)
  - Optical Disk Robotrer (4 Drives, 144 Slots, rewritable magnetooptical Disk, 650 MB/Cartdirge)

Print-Server :
- Sun SPARCstation IPX (24 MB Memory)
  - Disk Unit (SCSI 424 MB)
  - Floppy Disk (3.5")

Domain-,Network-,Mail-Server :
- Sun SPARC IPC (32 MB Memory)
  - Disk Unit (SCSI, 1.8 GB)
  - Floppy Disk (3.5")
  - CD-OM (644 MB)
  - Tape (8mm, 150 MB)

Equipment for software developement :

Sun SPARCstation 2 GX+ (32 MB Memory)
- Disk Unit (SCSI, 861 MB)
- Floppy Disk (3.5")

Sun SPARC IPC/GX (36 MB Memory)
- Disk Unit (SCSI, 861 MB)
- CD-ROM (644 MB)
- Tape (8 mm, 150 MB)

Sun SPARC IPC (24 MB Memory)
- Disk Unit (SCSI, 514 MB)
- Floppy Disk (3.5")
- DC-ROM (644 MB)
- Tape (8 mm, 150 MB)

ending at end of 1994 :

Control Data Cyber 180-860A (1 CPU, 20 PPU, 32 MB Memory)
- Disk Unit (7.5 GB)
- Tapes (4 Drives, 0.5", 9-track, 1600/6250 bpi)

Control Data Cyner 180-830A (16 MB Memory)
- Disk Unit (1.4 GB)
- Tape (0.5", 9-track, 1600/6250 bpi)

At other Departements :

Dept. SYNOPTIC :

6 Sun SPARCstation (several kinds)
- Disk Units (SCSI, several capacity)
- Floppy Disks (3.5")

ending at end of 1994 :

Control Data Cyber 18-10M (64 MB Memory)
- Disk Unit (8.8 MB)
- Tape (0.5", 9-track, 800 bpi)

DEC Micro VAX II-VAXSTATION (9 MB Memory)
- Disk Unit (159 MB)
- Tape (0.5", 95 MB streamer cartridge)

Dept.CLIMATOLOGY :

2 Sun SPARCstation 2GX (32 MB Memory)
- Disk Unit (SCSI, 1.5 GB and 861 MB resp.)

Dept.ENVIRONMENT :

2 Sun SPARCstation 2 and 10 resp, GX+ (32 and 64 MB Memory resp.)
- Disk Unit (SCSI, 1.5 GB and 848 MB resp.)
- Floppy Disk (3.5")

Dept.GEPHYSICS :

Sun SPARCstation 10 GX (128 MB Memory)
- Disk Unit (SCSI, 2.9 GB)
- Floppy Disk (3.5")

Dept.TECHNICS :

Sun SPARCstation 10 GX+ (96 MB Memory)
- Disk Unit (SCSI, 2.5 GB)

Regional Plants at Innsbruck, Klagenfurt and Salzburg (each equipped with) :

Sun SPARCstation 2 GX (32MB Memory)
Disk Unit (SCSI, 1.5 GB)
Floopy Disk (3.5")
Tape (Gigatape, DAT, 1.0 GB)
Sun SPARCstation 2 GX (32 MB Memory)
Disk Unit (SCSI, 861 MB)
Floppy Disk (3.5")
CD-ROM (644 MB)
Tape (8 mm, 150 MB)

Institute Meteorology and Geophysics University of Vienna :

7 Sun Sparcstations,
1 Textronics XD 8810

Further equipment:

PC's	286,386,486,Atari,McIntosh
Printer	Centronics, Genicom, Epson, NEC, Itoh
	A3: WDV Thermotransfer, IBM Colorjet
	A3: QMS Laser, Inkjet
	A4: HP Laserjet II, III,IV
	A4: Kyocera 1500
Plotter	A0: Calcomp 68436, Benson 1645R
Terminals	a-n: DEC, Qume, CD monochrom
	X-Terminals (monochrom, color)

1.b) S O F T W A R E

Sun Systems :

Operating System	SunOS 4.1.x and SOLARIS 2.3 (UNIX)
Compiler :	Fortran 77
	C, ANSI C
Graphics	Sun GKS
	Xelion GKS
	MAGICS
	PV-Wave
Libraries	IMSL
	NAg Mark 15/16
Database	SYBASE

Control Data Cyber 860/830 :

Operating System	NOS/VE 1.5.3
	VX/VE 2.5.3 (UNIX)
Compiler	Fortran 77
	Pascal Level 1 ISO 7185
	C
	Cybil (Cyber Implementation Language)
Graphics	GK-2000 GKS Level 2B
	Erlgraph
	PV-Wave
Libraries	IMSL
	NAg Mark 15/16

Control Data Cyber 18 :

Operating System	MSOS

DEC Micro VAX :

Operating System	VMS 5.0
Compiler	Fortran 77
Graphics	Digital GKS Level 2C
	GTS-GRAL GKS Level 2B
	MAGICS

Personal Computer :

Operating System	MS DOS, MS Windows, Interactive Unix, (TOS, McOS)
Compiler	Fortran , Turbo Pascal , C , C++ , Basic
Applications	MS Word, Word for Windows
	MS Works, Works for Windows
	Excel , Lotus 123
	MS Powerpoint
	Corel Draw
	Ventura Publisher
Network	PCNFS

1.c) N E T W O R K **Fig.2**

WAN : Vienna ; Innsbruck , Klagenfurt , Salzburg operated by ZAMG
this network should be operated by Austro Control Ltd in 1995

LAN ZAMG Vienna : TCP/IP (Telnet,FTP,SMTP,NFS) and Control Data CDCNET
6 logical segments of ETHERNET
1 FDDI segment

3 CISCO Routers (7000,7010,4000)
2 Network Peripherals FDDI Workgroup Concentrators
8 Control Data Communication Processors
Repeater , Optical Links

switched lines :

kbps	end of line
128	Univ.Vienna,Inst.Meteorl.Silbergasse
64	ECMWF (splittet do Cyber (FTP/EEP) and Sun-Network (TCP/IP))
64	University of Vienna : ACONET , Internet (Domain : zamg.ac.at)
64	Austro Control Ltd Central
64	Austro Control Ltd Airport
19.2	Geological Survey
4.8	Slowenia
4.8	Slowenia
2.4	Goverment (environmental purposes)
1.2	County of Lower Austria : Highway

dial lines :

9.6	Computerdial/Weatherphone
9.6	Danube Power Plant Corp

DATEX-P (packet switched network):

9.6	16 virtuel connections

LAN ZAMG Innsbruck (Aerport) : Ethernet
direct plugged to local Router of Austro Control Ltd

LAN ZAMG Klagenfurt (Airport) : Ethernet
direct plugged to local Router of Austro Control Ltd

LAN ZAMG Salzburg : Ethernet
direct connection to local Router of Austro Control Ltd (Aerport)
1 CISCO Router 2500
switched line :

kbps	end of line
64	Austro Control Ltd Salzburg (remote)

2.) Daily computed results based on ECMWF products
(normally based on 12z data, updated with 00z data)

Model output diagnosis (MOD) , organised in 2 blocks.
These fields can be shown in graphic-mode on X-Terminals or Plotter either as single pictures, or combined with each other, or combined with other fields as synoptic charts or satellite pictures, or in cross sections showing the situation between choosen points within Europe. **Fig.3,4**

1) PIB (primary information block) :
with direct model output (DMO) and derived quantities
(dispatched within ZAMG and to Austro Control Ltd, etc)
- absolute topography 1000 hPa
- absolute topography 500 hPa
- advection of absolute vorticity 500 hPa
- equivalent potential temperature 500/850 hPa
- thermal frontal parameter **Fig.5**
- thermal advection of potential temperature 500/1000 hPa
- humidity index 500/850 hPa
- prectipitation

2) SIB (secondary information block)
with direct model output (GMO) and further derived quantities :
(dispatched within ZAMG only)
- temperature 850 hPa
- vertical velocity 850, 700, 500 hPa
- ECMWF cloud cover
- relative vorticity 500 hPa
- shearing term of vorticity 500, 300 hPa
- twisting term of vorticity 500 hPa
- isotachs 300 hPa
- divergency 1000, 850, 700, 500, 300 hPa
- stratching and shearing deformation 500, 300 hPa
- potential temperature 500/850 hPa
- thermal vorticity including twisting term 500/850 hPa
- relative humidity 850, 700 hPa
- advection of relative hunidity 850, 700 hPa
- time till condensation 850, 700 hPa
- Showalter index 850, 700 hPa
- frontogenetic parameter

Further special purposes (from short-term up to medium-range forecast) :

Trajectories at 1000, 850, 700, 500, 300 hPa Fig.6
from fixed and special choosen locations in and around Austria for transport and environmental (air pollution purposes)
QFA and weather forecast at distinct points Fig.7,8
for energy companies, avelange forecasts, traffic, newspapers, broadcast and TV, telephone-information service, agrarmeteorology etc.
Model output data verification on 31 synoptic stations (T2m, amount of precipitation, cloudiness, humidity index, 1000 hPa wind)

3) **Special Projects running at ECMWF :**

- "Diagnostic tests of the ECMWF physics package "

- "Diabating heating of the global atmosphere"

(both by M.Hantel,University of Vienna, Institute of Meteorology and Geophysics, Dept. of theoretical Meteorology)

4) **Other scientific work using ECMWF data :**

- at ZAMG :
  Dept. SYNOPTIC, Working Group of Satellite Meteorology Fig.9-11
  combining satllite and aerosonde data with ECMWF fields
  Dept. ENVIRONMENT : trajectory-modeling
- at University of Vienna : Institute of Meteorology and Geophysics
  Dept. of theoretical Meteorology (M.Hantel)
  Applied Analytical Meteorology (P.Kahlig) :
  Special: T.Haiden : "Modelling orographic upslope rain" Fig.12,13
  presented at ITAM94,Lindau
  Dept. of general Meteorology (R.Steinacker)
  Dept. of environmental Meteorology (H.Kromp-Kolb)
- at University of Innsbruck : Institute of Meteorology and Geophysics
  Dept. of theoretical Meteorology (H.Pichler)

5) **Plans for further use of ECMWF data :**

As at the end of 1994 the old Cybers will be out of use work is done to move over for all applications to the new SYBASE-Database and Sun-environment based on the Operating System SunOS and SOLARIS resp.
One main step in this direction was to cancel the dissemination of ECGRID-Data and use only GRIB Data. To test this new dataset we splitted the line to receive both datasets, the old one in direct connection to the Cybers 860/830 (FTP/EEP), and the new one to the Sun-world (TCP/IP).
When we finished our tests and wanted to cancel the old dataset colleagues unfortunately found out that they have had five more important applications to migrate. We hope that this work is done at the mid of October.
Then we want to enlarge our normal dissemination set to fullfill the requirements by the Austrian Environmental Authorities which are very interested in special modelling of trajectories to have a meteorological instrument in the **Federal Warning an Alarm Plan .** Fig 14
The Ensemble Forecast System will be used in digital form and not by FAX.

Fig. 1

Hardwarekonfiguration ZAMG

Stand: Juni 1994

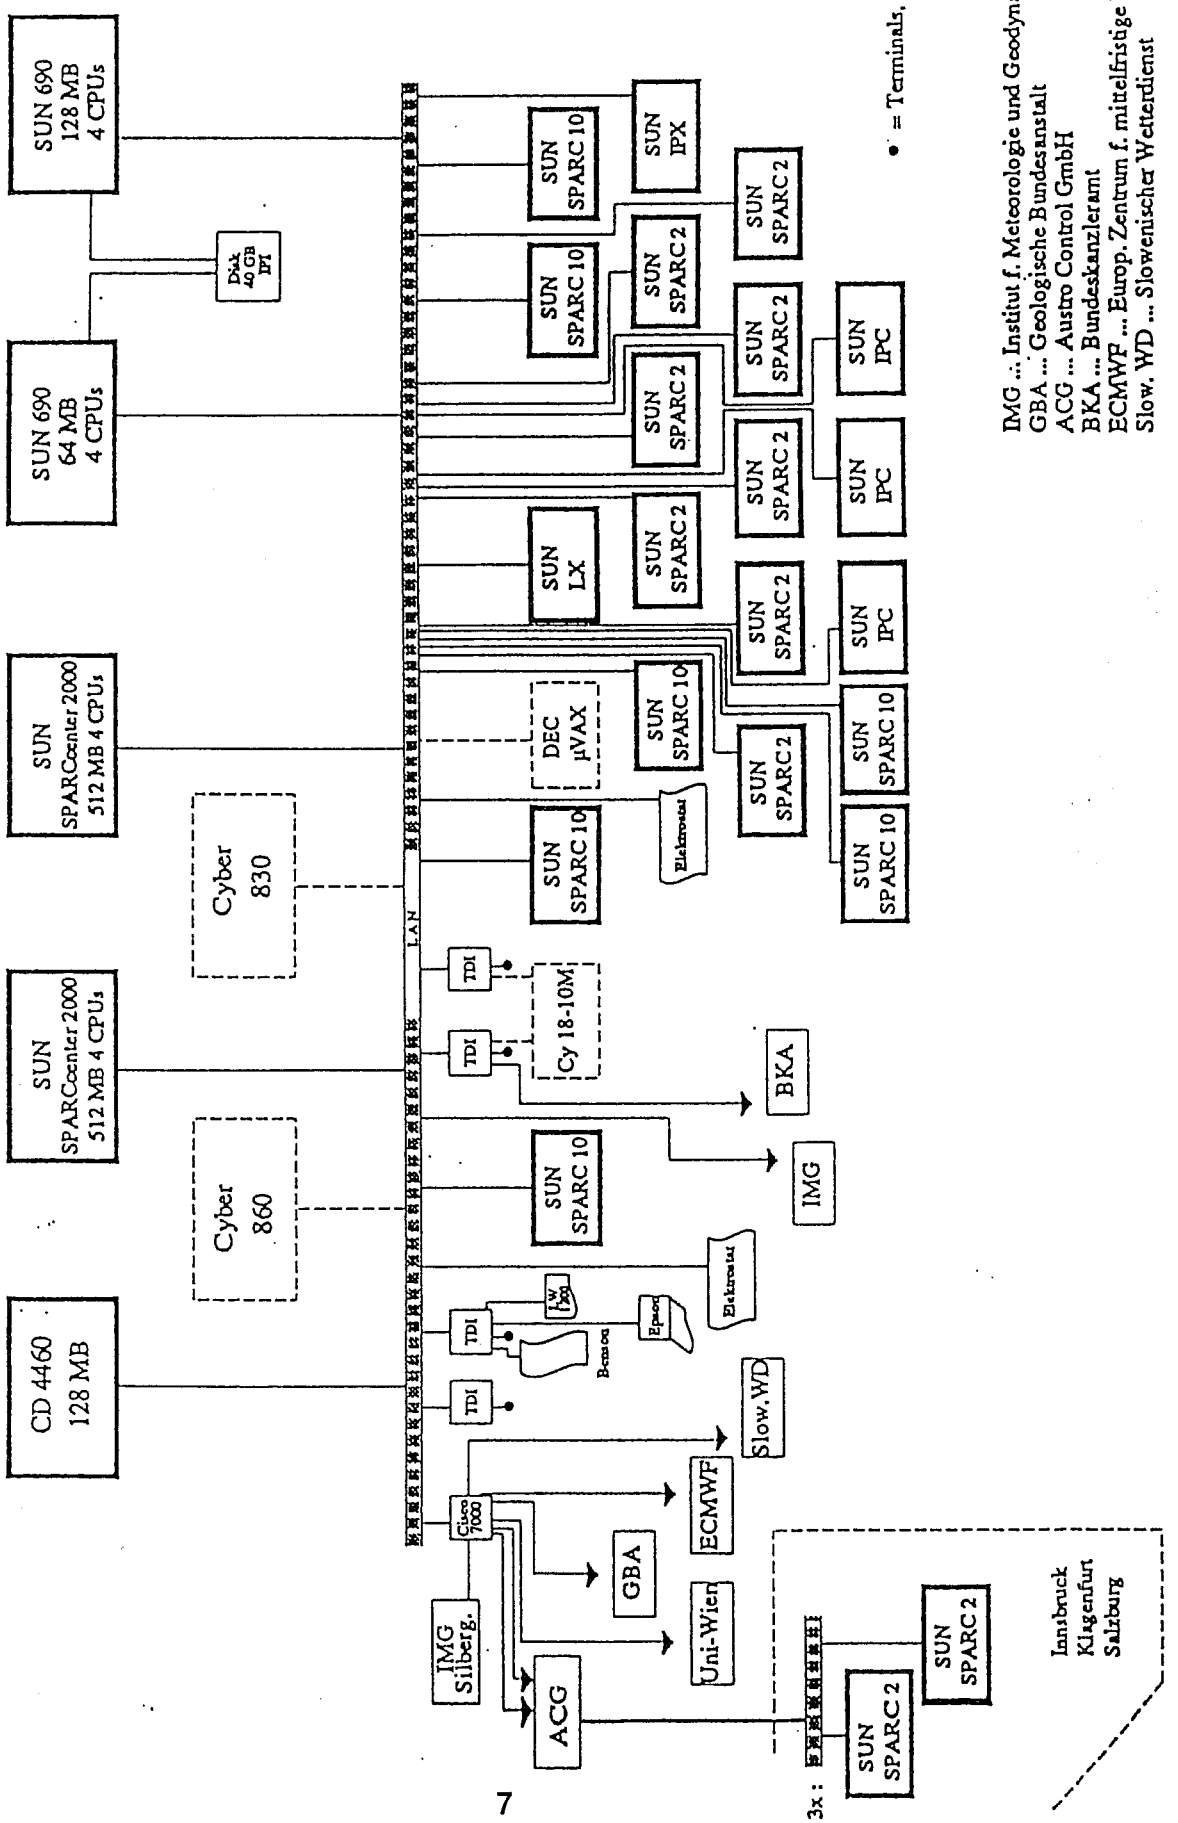

Fig. 2

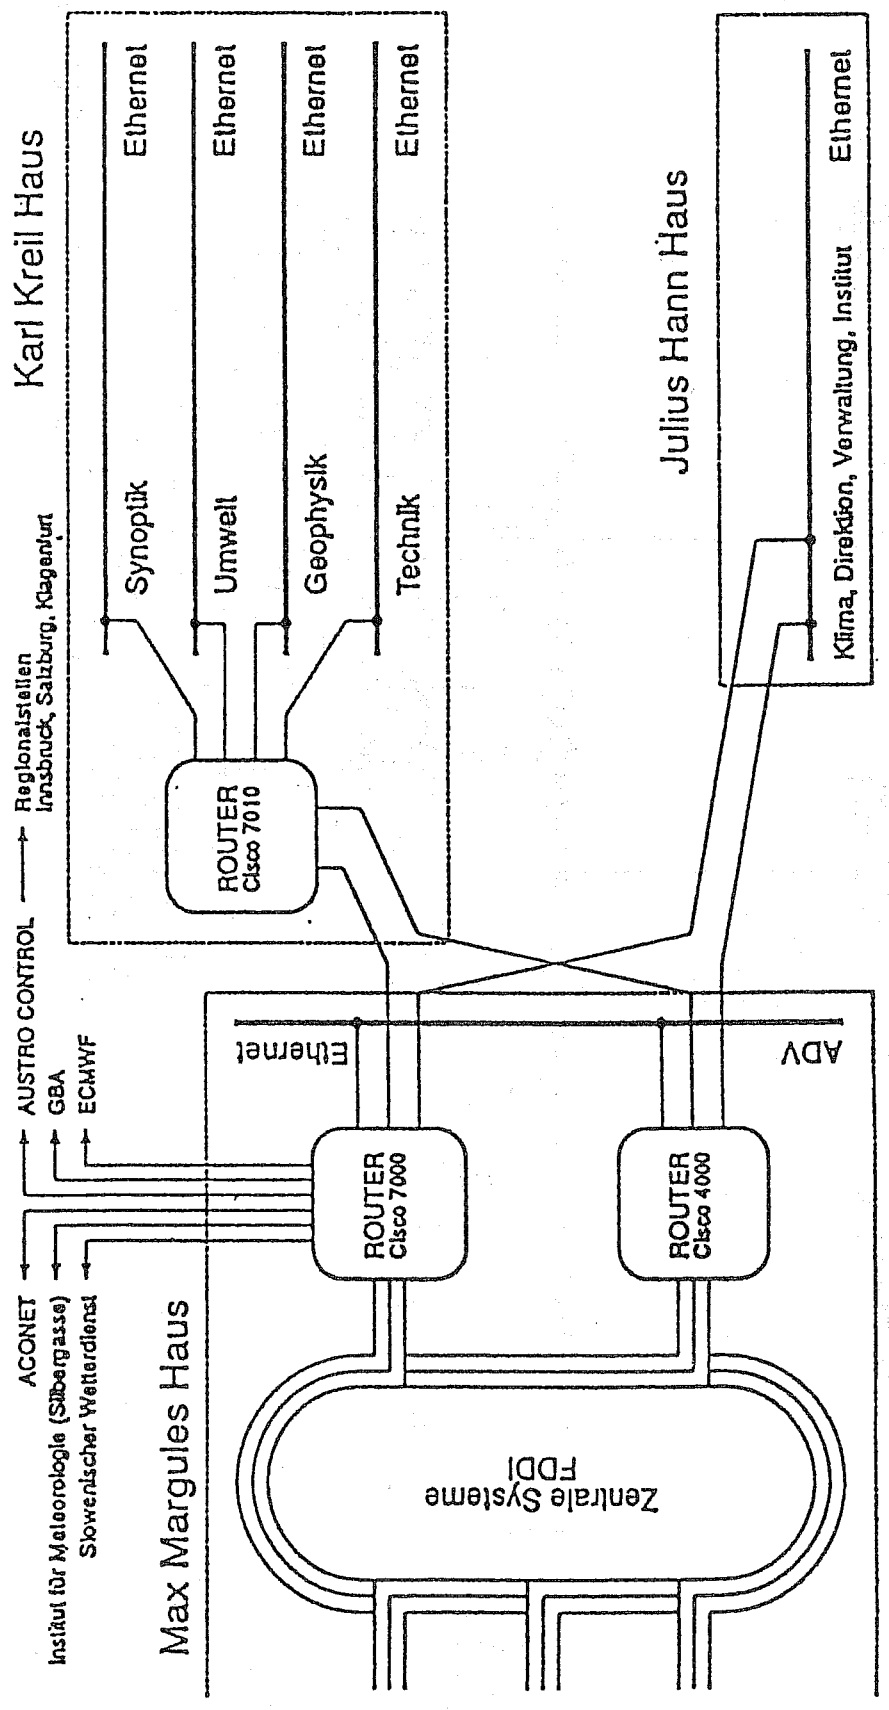

ZAMG/6-1994

Fig. 3

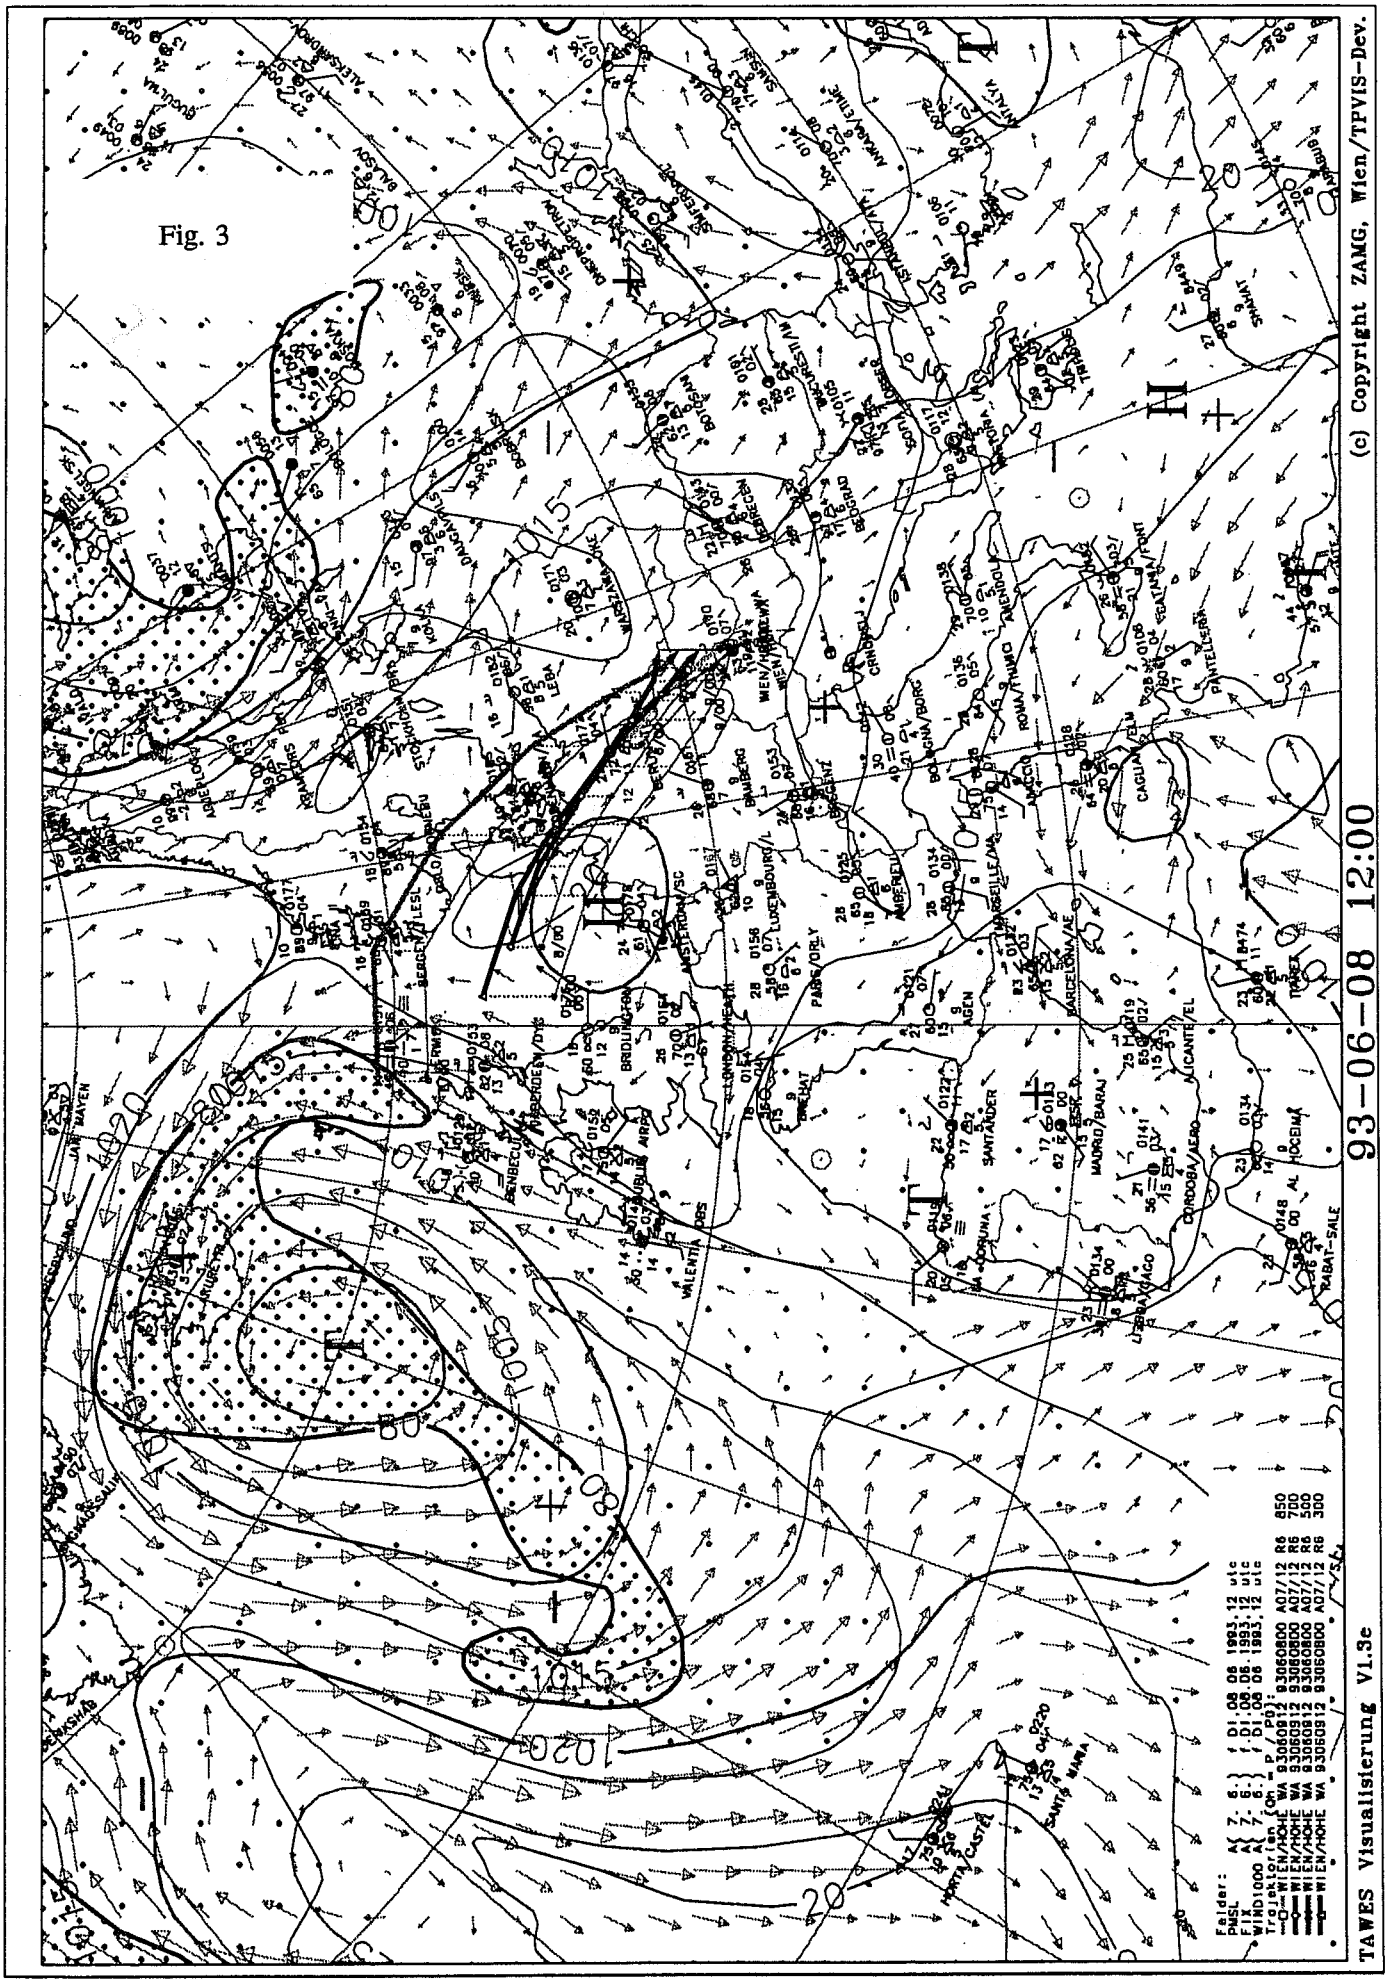
93-06-08 12:00
(c) Copyright ZAMG, Wien/TPVIS-Dev.
TAWES Visualisierung V1.3e
Felder:
MSL
FI1000
WIND1000
Trajektorien (Om_P/P0):
WIEN/HOHE WA 930609 12
WIEN/HOHE WA 930608 00
WIEN/HOHE WA 930608 12
WIEN/HOHE WA 930609 12

Fig. 4

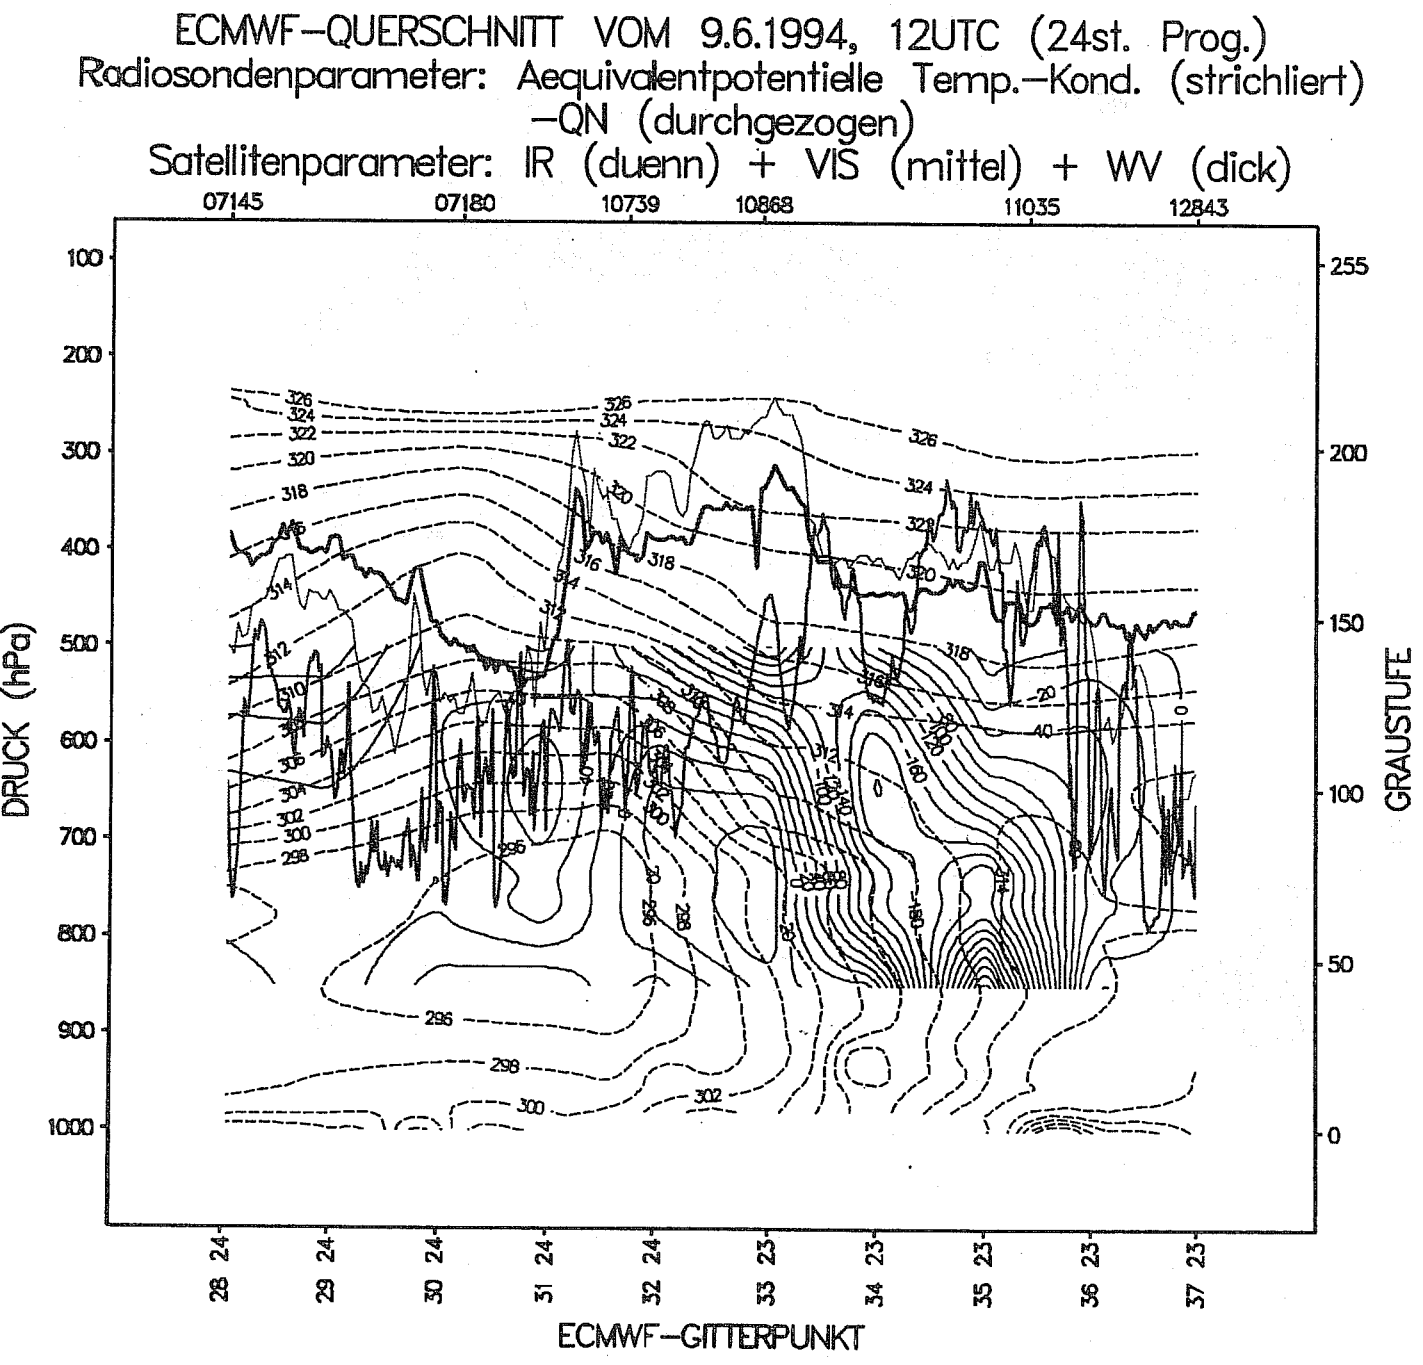
ECMWF–QUERSCHNITT VOM 9.6.1994, 12UTC (24st. Prog.)
Radiosondenparameter: Aequivalentpotentielle Temp.–Kond. (strichliert)
–QN (durchgezogen)
Satellitenparameter: IR (duenn) + VIS (mittel) + WV (dick)
07145
07180
10739
10868
11035
12843
DRUCK (hPa)
GRAUSTUFE
ECMWF–GITTERPUNKT

Fig. 5

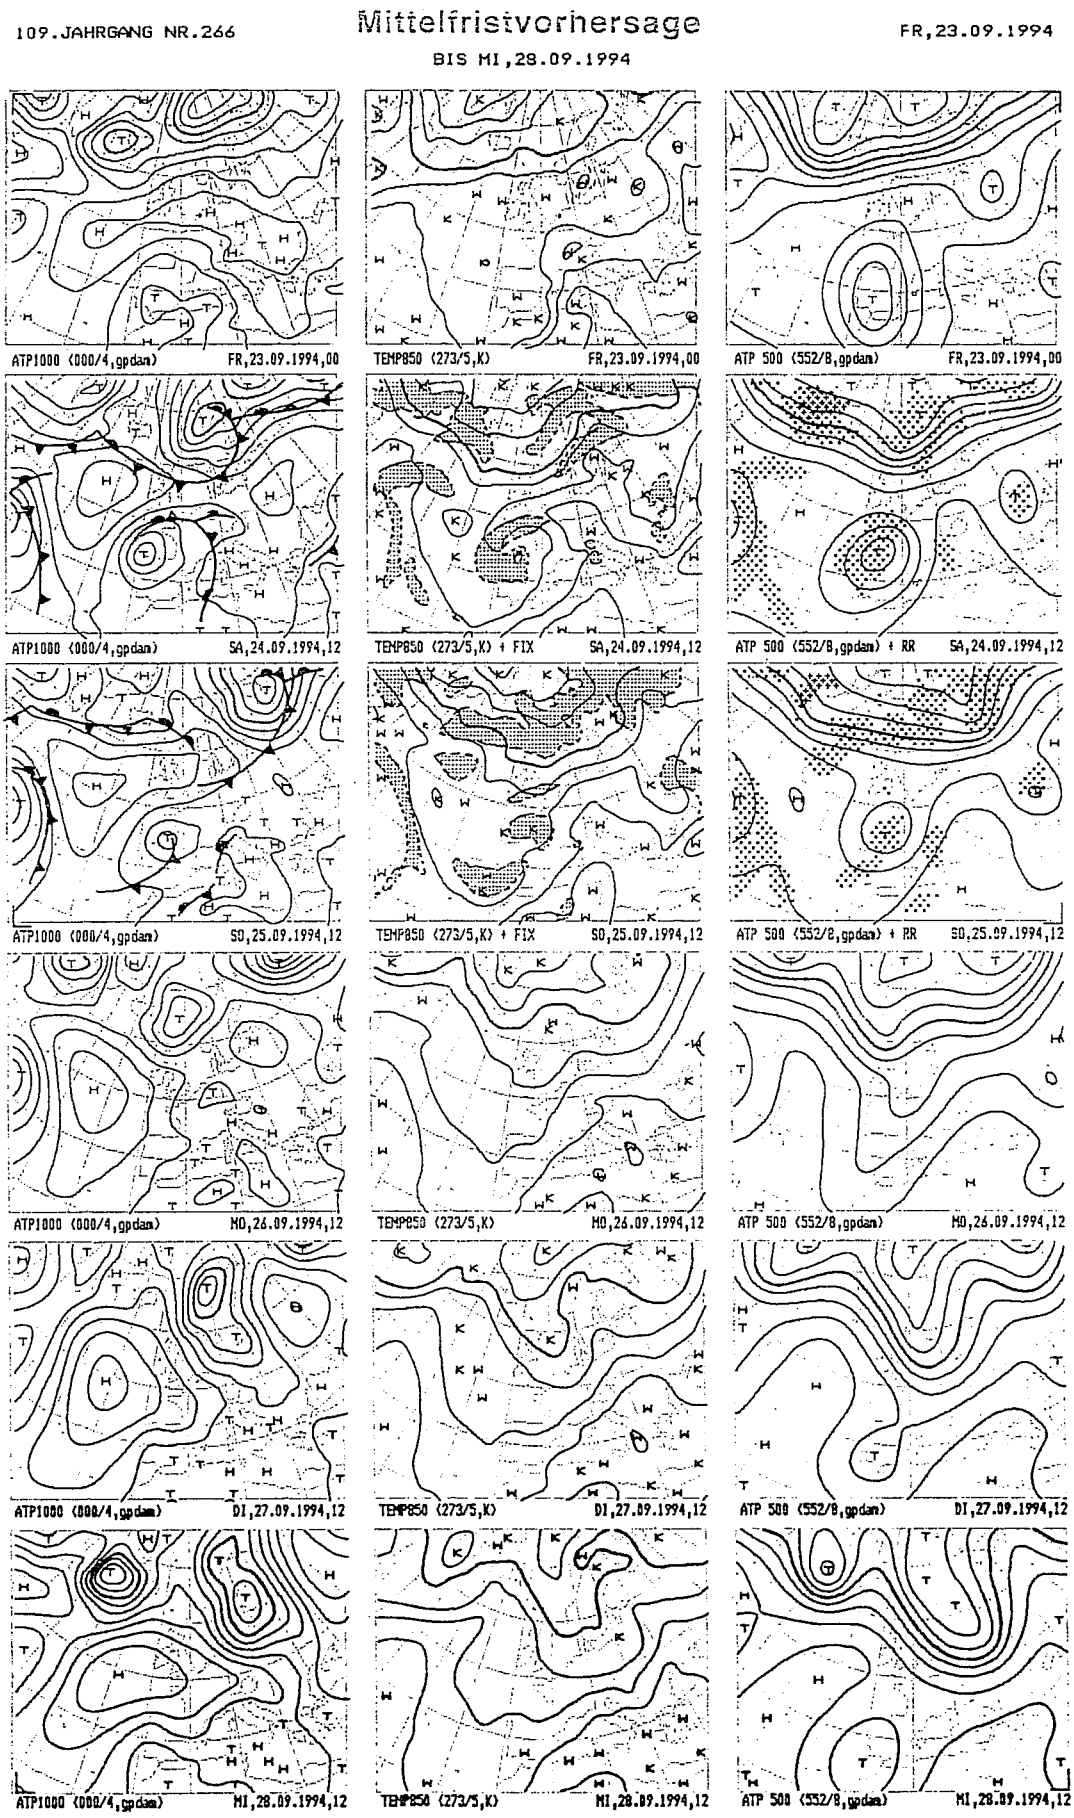
109.JAHRGANG NR.266
Mittelfristvorhersage
BIS MI,28.09.1994
FR,23.09.1994
ATP1000 (000/4,gpdam) FR,23.09.1994,00
TEMP850 (273/5,K) FR,23.09.1994,00
ATP 500 (552/8,gpdam) FR,23.09.1994,00
ATP1000 (000/4,gpdam) SA,24.09.1994,12
TEMP850 (273/5,K) + FIX SA,24.09.1994,12
ATP 500 (552/8,gpdam) + RR SA,24.09.1994,12
ATP1000 (000/4,gpdam) SO,25.09.1994,12
TEMP850 (273/5,K) + FIX SO,25.09.1994,12
ATP 500 (552/8,gpdam) + RR SO,25.09.1994,12
ATP1000 (000/4,gpdam) MO,26.09.1994,12
TEMP850 (273/5,K) MO,26.09.1994,12
ATP 500 (552/8,gpdam) MO,26.09.1994,12
ATP1000 (000/4,gpdam) DI,27.09.1994,12
TEMP850 (273/5,K) DI,27.09.1994,12
ATP 500 (552/8,gpdam) DI,27.09.1994,12
ATP1000 (000/4,gpdam) MI,28.09.1994,12
TEMP850 (273/5,K) MI,28.09.1994,12
ATP 500 (552/8,gpdam) MI,28.09.1994,12

Fig. 6

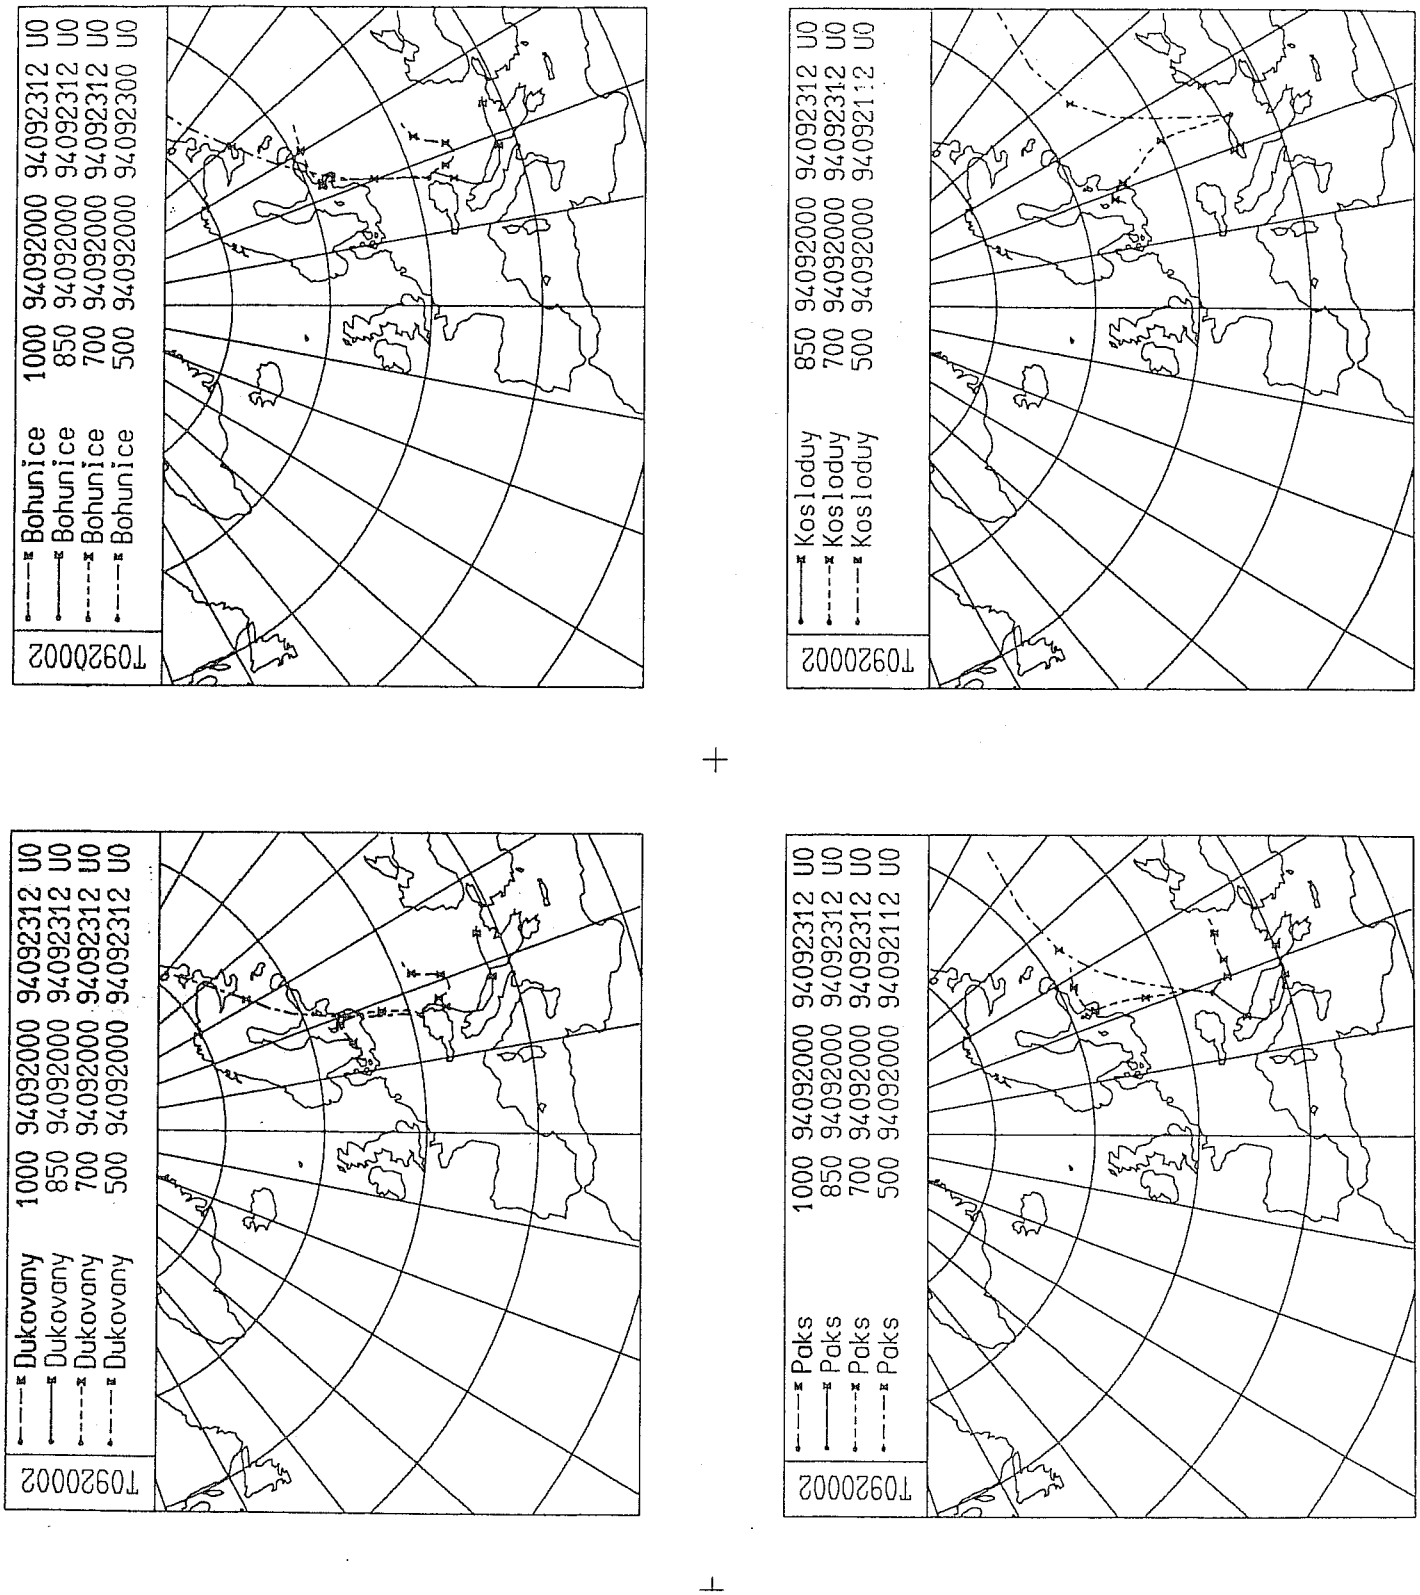
Bohunice 1000 94092000 94092312 U0
Bohunice 850 94092000 94092312 U0
Bohunice 700 94092000 94092312 U0
Bohunice 500 94092000 94092300 U0
T0920002
Kosloduy 850 94092000 94092312 U0
Kosloduy 700 94092000 94092312 U0
Kosloduy 500 94092000 94092112 U0
T0920002
Dukovany 1000 94092000 94092312 U0
Dukovany 850 94092000 94092312 U0
Dukovany 700 94092000 94092312 U0
Dukovany 500 94092000 94092312 U0
T0920002
Paks 1000 94092000 94092312 U0
Paks 850 94092000 94092312 U0
Paks 700 94092000 94092312 U0
Paks 500 94092000 94092112 U0
T0920002
1994-09-20
12:03:51

Fig. 7

Der STANDARD

Wetter für morgen (Österreich und Umgebung)

vom. Mo. 19. 9. 1994

für die Zeitung am. Di. 20. 9. 1994

Muster:

"Wetter" °C

°C "häufigste" Mittagstemperatur

ausgegeben von der ZAMG, 1190 Wien
(Meteorologe Tel. 36 44 53-2310 DW)

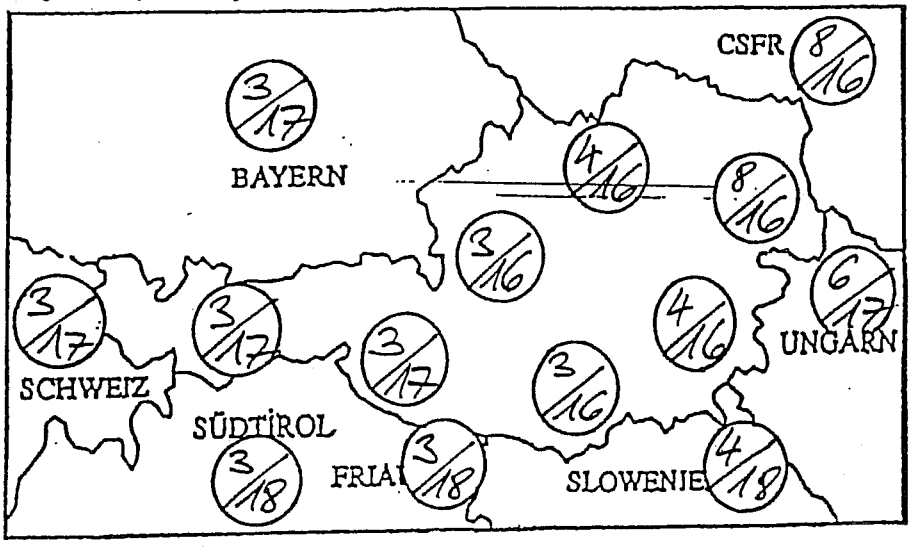

1=wolkenlos 4=st.bewölkt 8=Regenschauer

2=heiter 6=Regen 9=Gewitter

3=wolkig/mäßig bew. 7=Schneefall 11=Nebel

Wettertrend für WIEN in den nächsten Tagen

Vorhersage für	Dienstag 20.	Mittwoch 21.	Donnerst. 22.	Freitag 23.	Samstag 24.	Sonntag 25.	Montag 26.
Maxima	16	16	22	23	24	22	20
Minima	10	11	10	12	14	14	14
Gesamt-wettertrend	4	6	2	2	2	8	8

Fig. 8

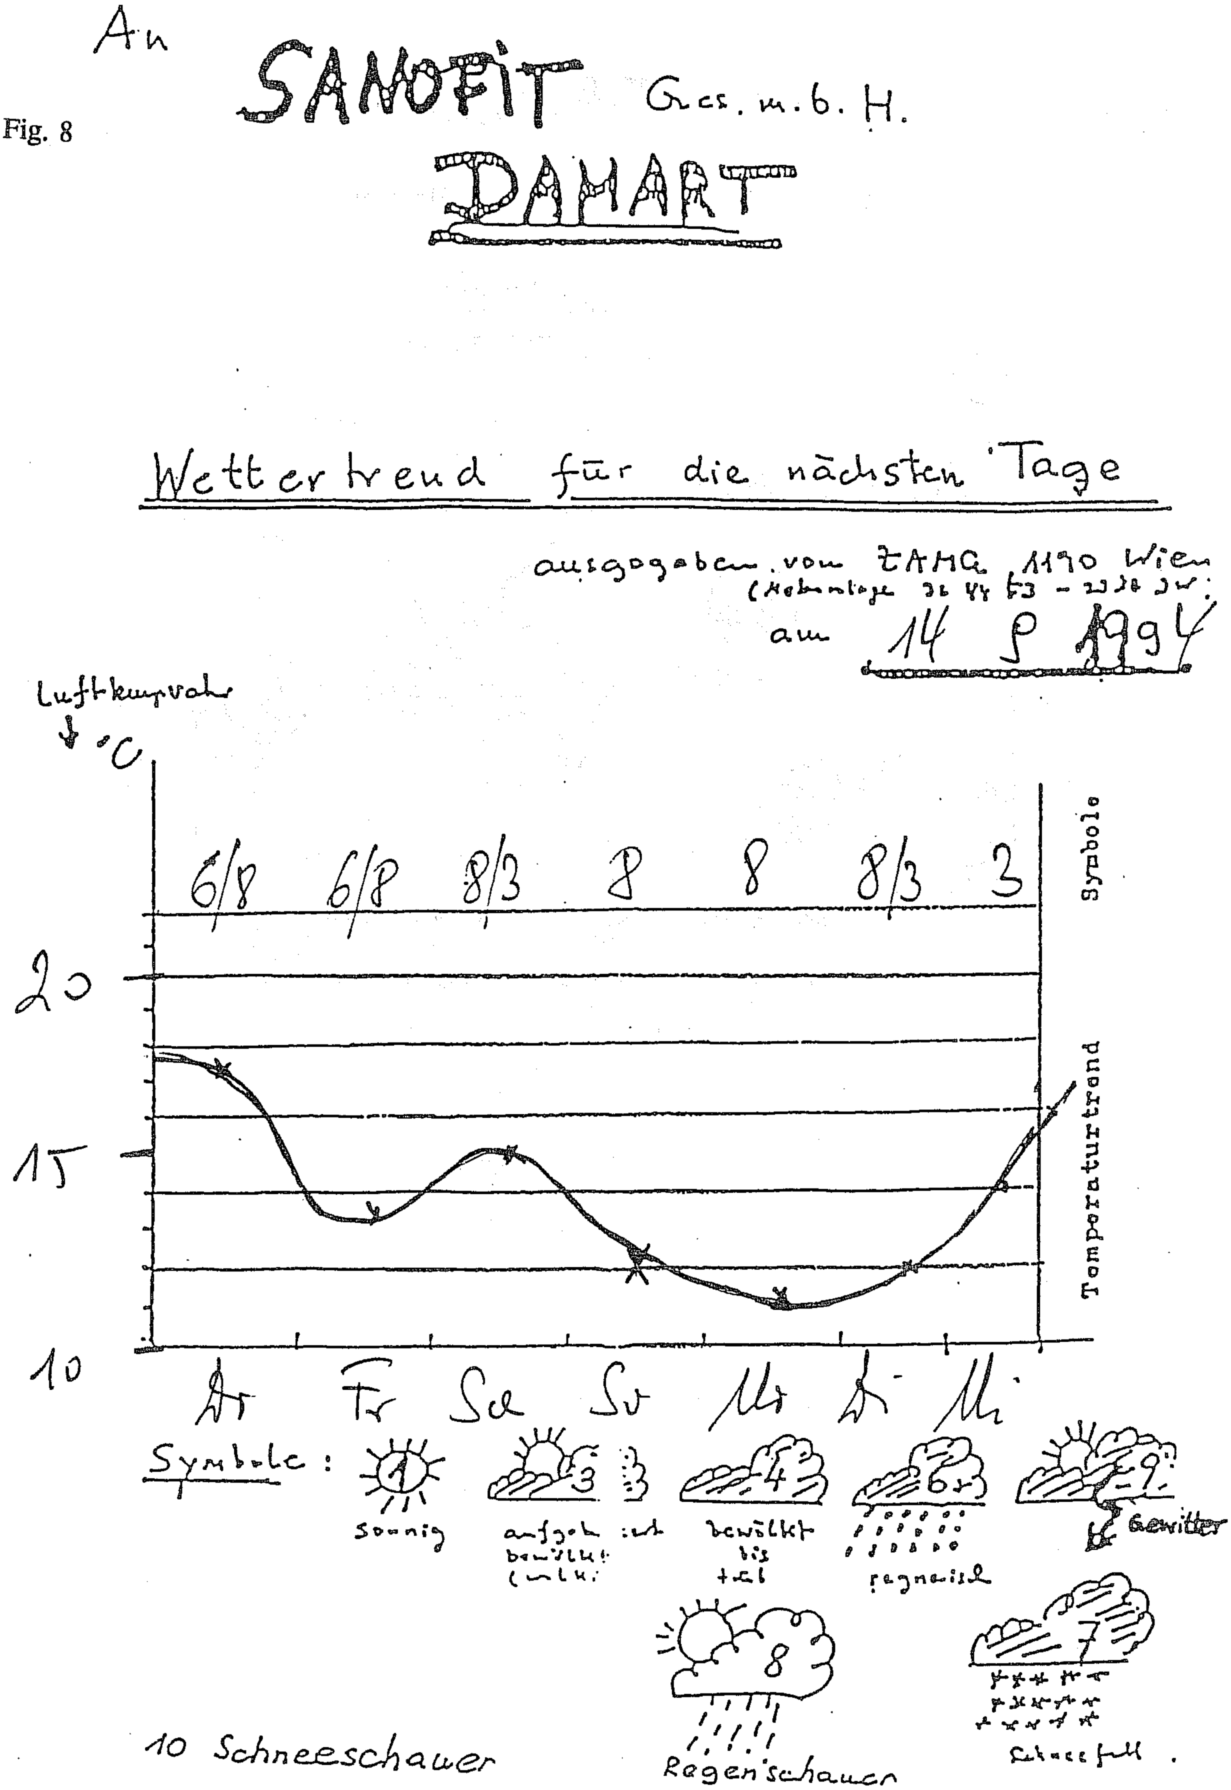
An
SANOFIT Ges. m. b. H.
DAMART
Wettertrend für die nächsten Tage
ausgegeben von ZAMG, 1190 Wien
am 14 8 1994
°C
20
15
10
6/8 6/8 8/3 8 8 8/3 3
Symbole
Temperaturtrend
Do Fr Sa So Mo Di Mi
Symbole:
1
sonnig
3
4
bewölkt
6
regnerisch
9
Gewitter
8
Regenschauer
7
Schneefall
10 Schneeschauer

Fig. 9

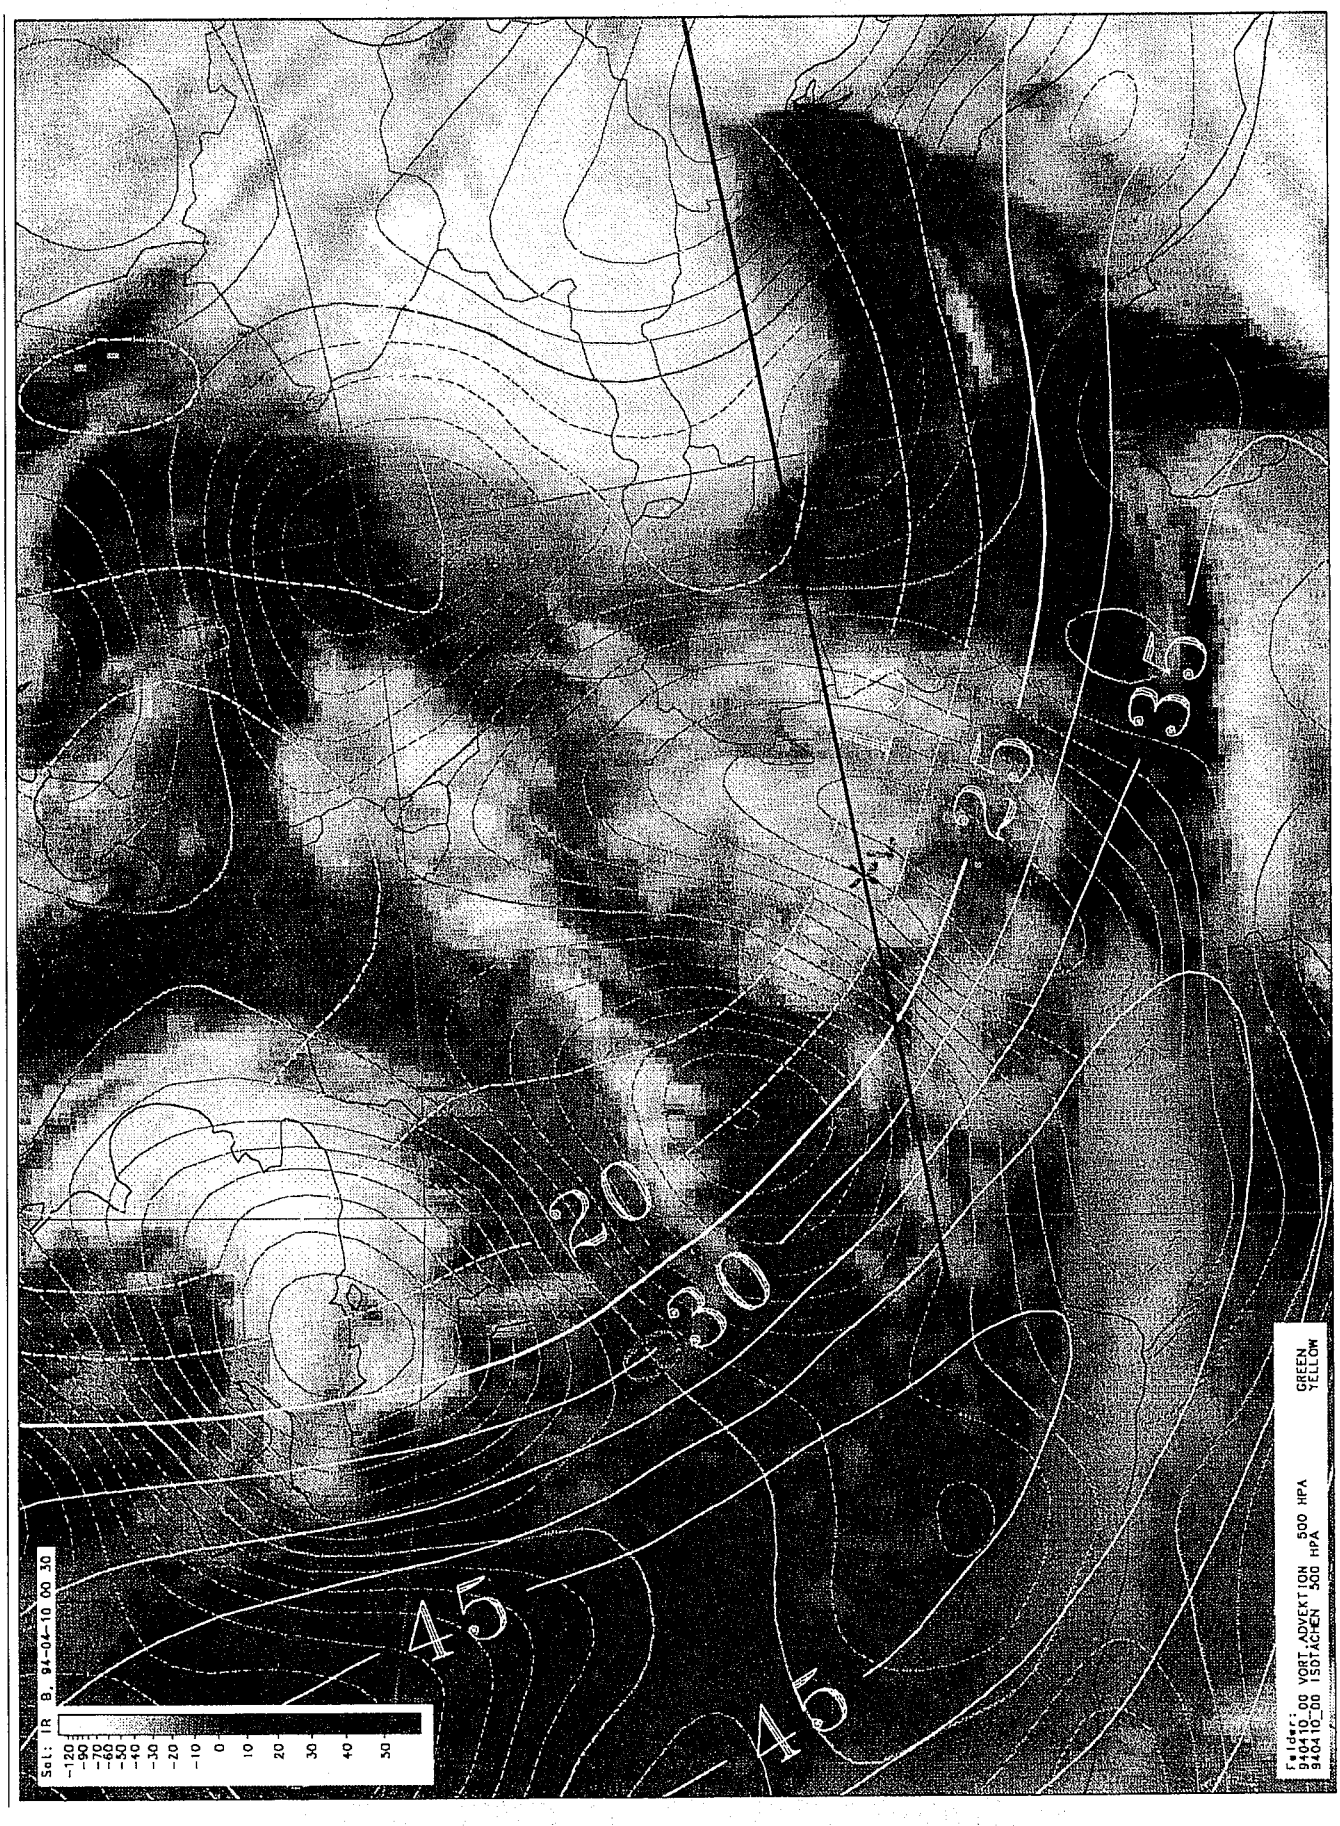
Sat: IR B, 94-04-10 00 30
-120
-90
-70
-60
-50
-40
-30
-20
-10
0
10
20
30
40
50
45
20
30
25
35
45
Folder:
9404I0_00 VORT_ADVEKTION 500 HPA
9404I0_00 ISOTACHEN 500 HPA
GREEN
YELLOW
TPVIS V1.5a
Plot-Datum:
94-08-17 12:16
(c) Copyright ZAMG, Wien/Satellit

Fig. 10

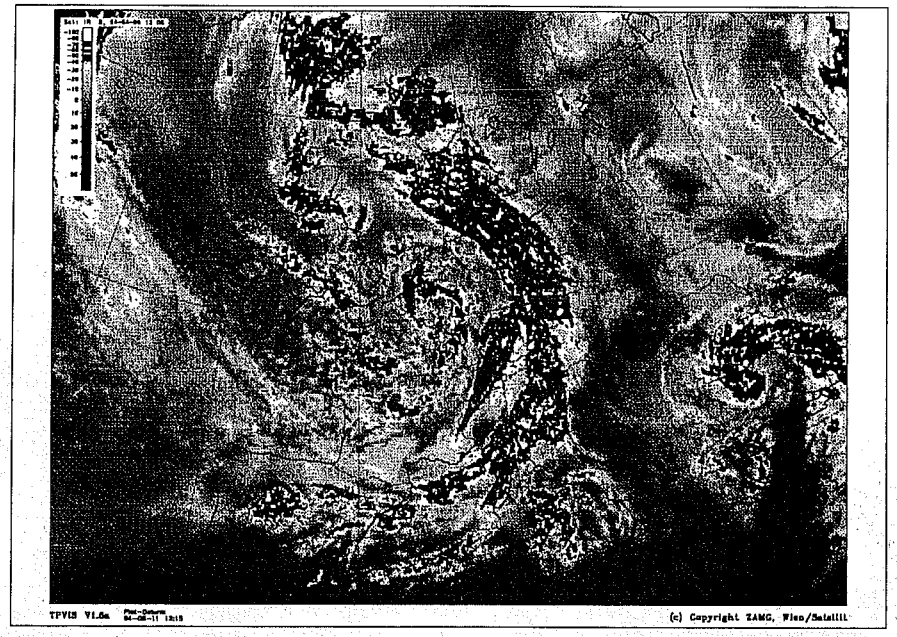

IR

Water Vapour + Rad. Stations

Line of Cross Section

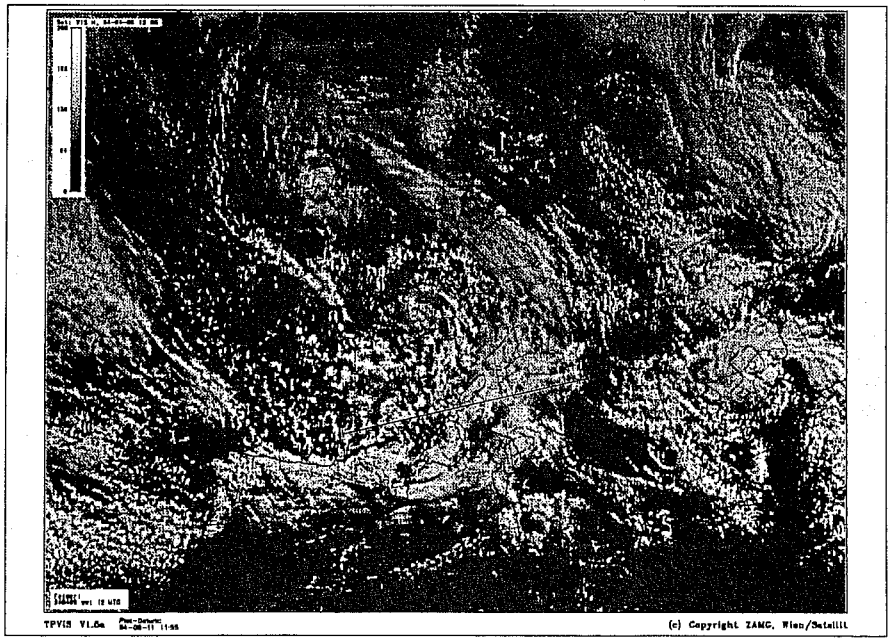

Visible rouge + WW

Line of Cross Section

Fig. 11 Radiosonde cross section, 10.4.1994, 00UTC
Equivalent-potential temperature (black), relative humidity (green)
Orange line: IR signals

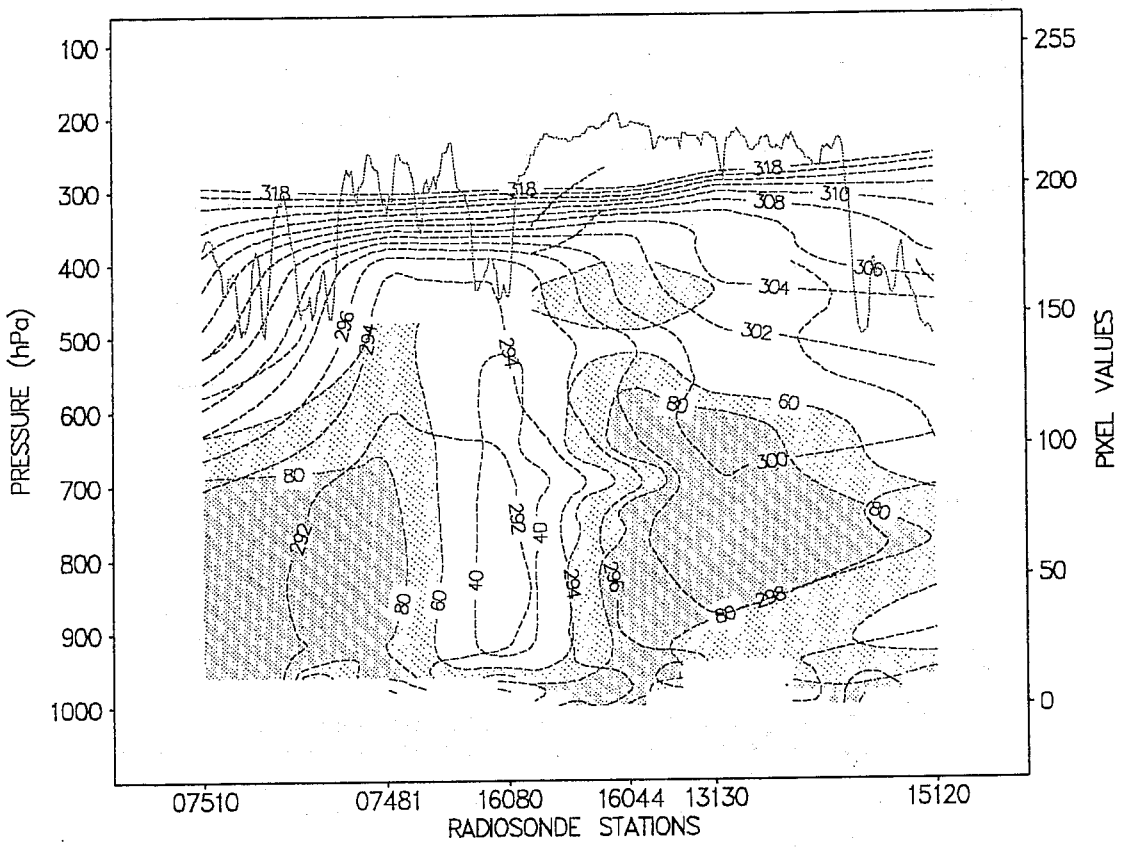

ECMWF cross section, 10.4.1994, 00UTC
Temperature advection (blue)
Equivalent-potential temperature from radiosonde (black)

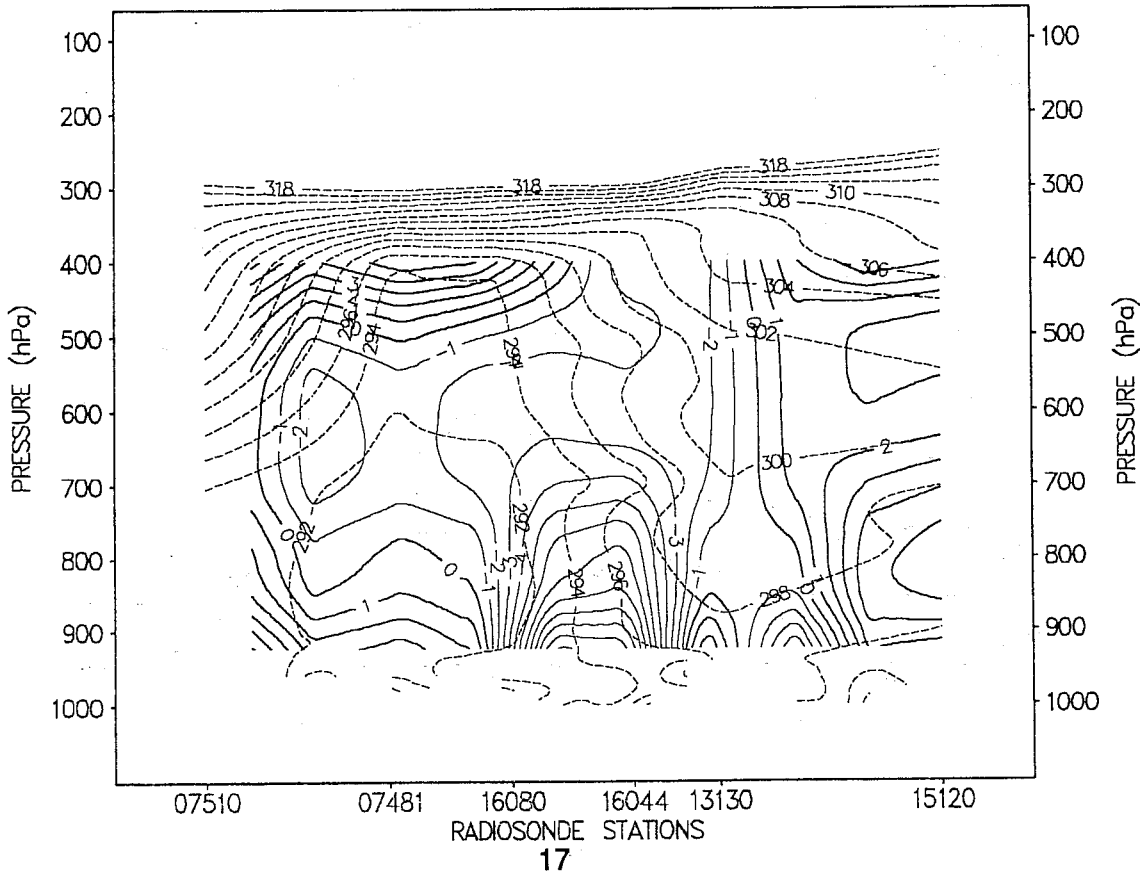

Fig. 12

Haiden 1994

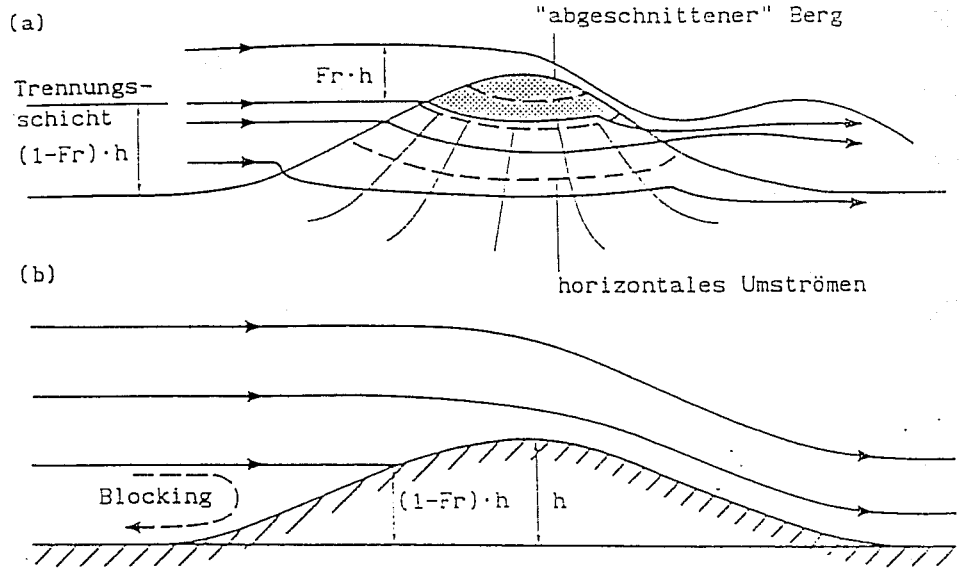

Abb.2.19: Strömungsverhalten bei sehr stabilen Verhältnissen (schematisch); (a): dreidimensionale Strömung über einen einzelstehenden Berg, (b): zweidimensionale Strömung über einen langgestreckten Bergrücken (nach Carruthers und Hunt, 1990).

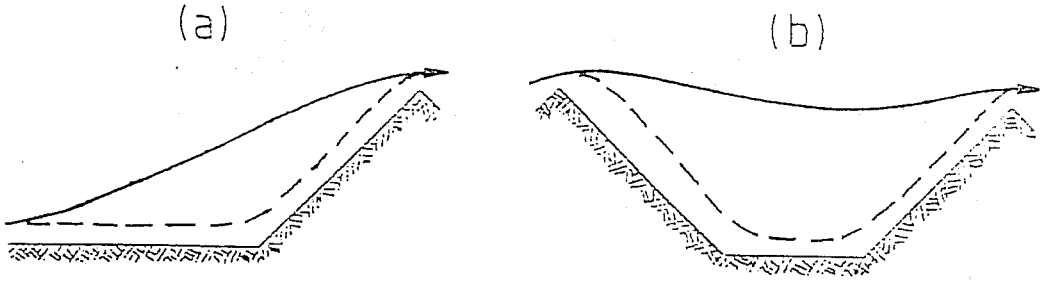

Abb.2.20: Strömungsverhältnisse mit (———) und ohne (- - -) Berücksichtigung von Blocking-Effekten (schematisch). (a):Gebirgsrand, (b): Tal innerhalb des Gebirges.

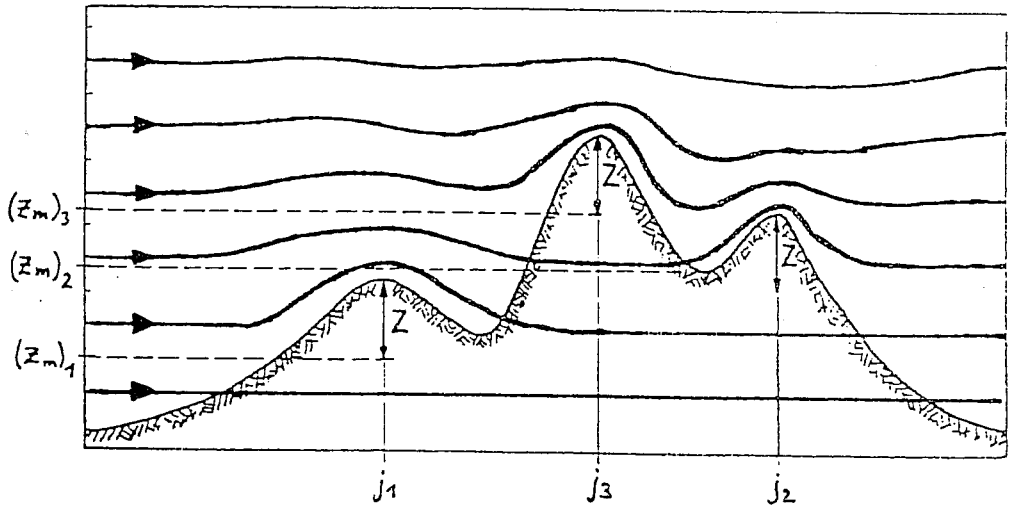

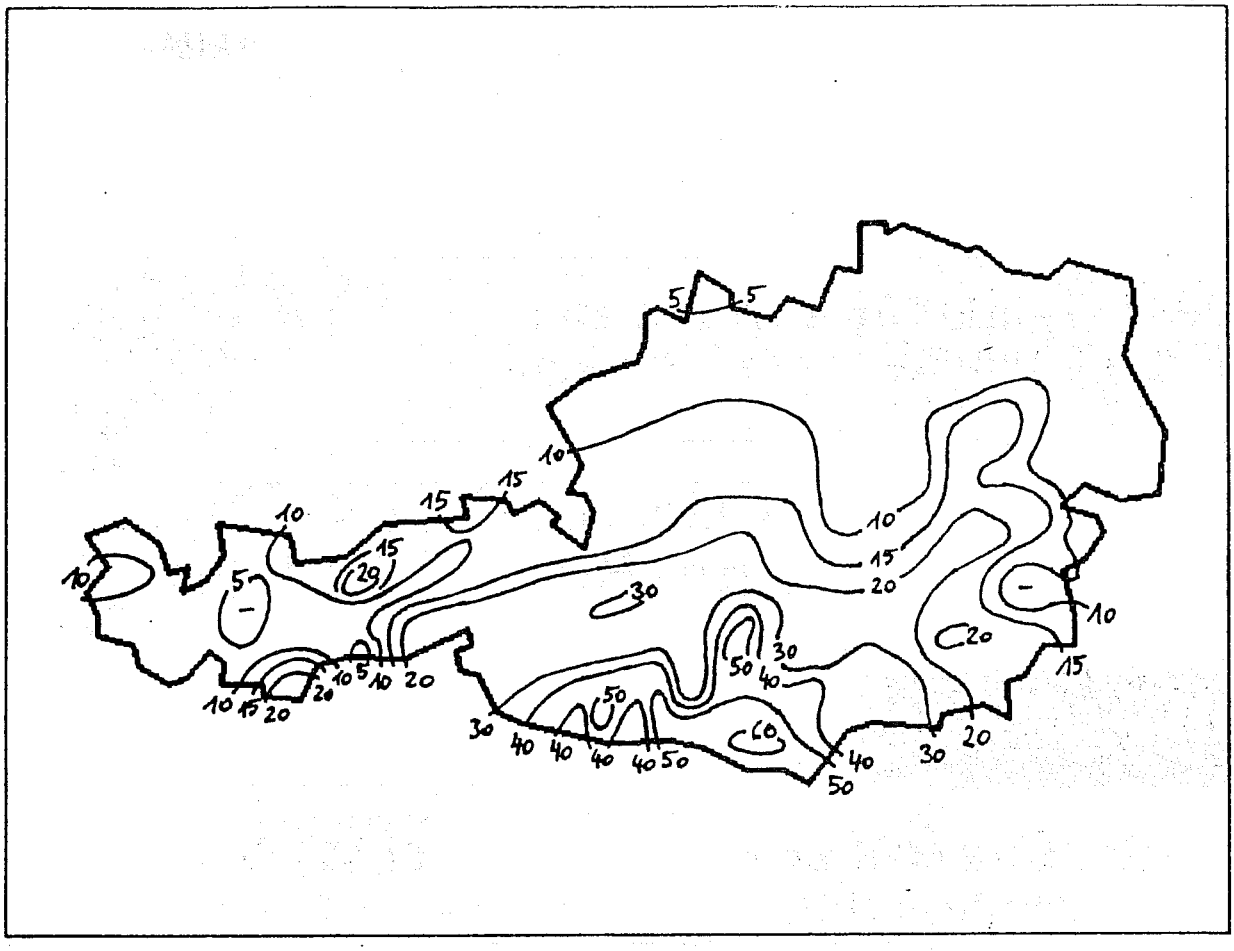

22.10.1993 00 UTC (±6h)

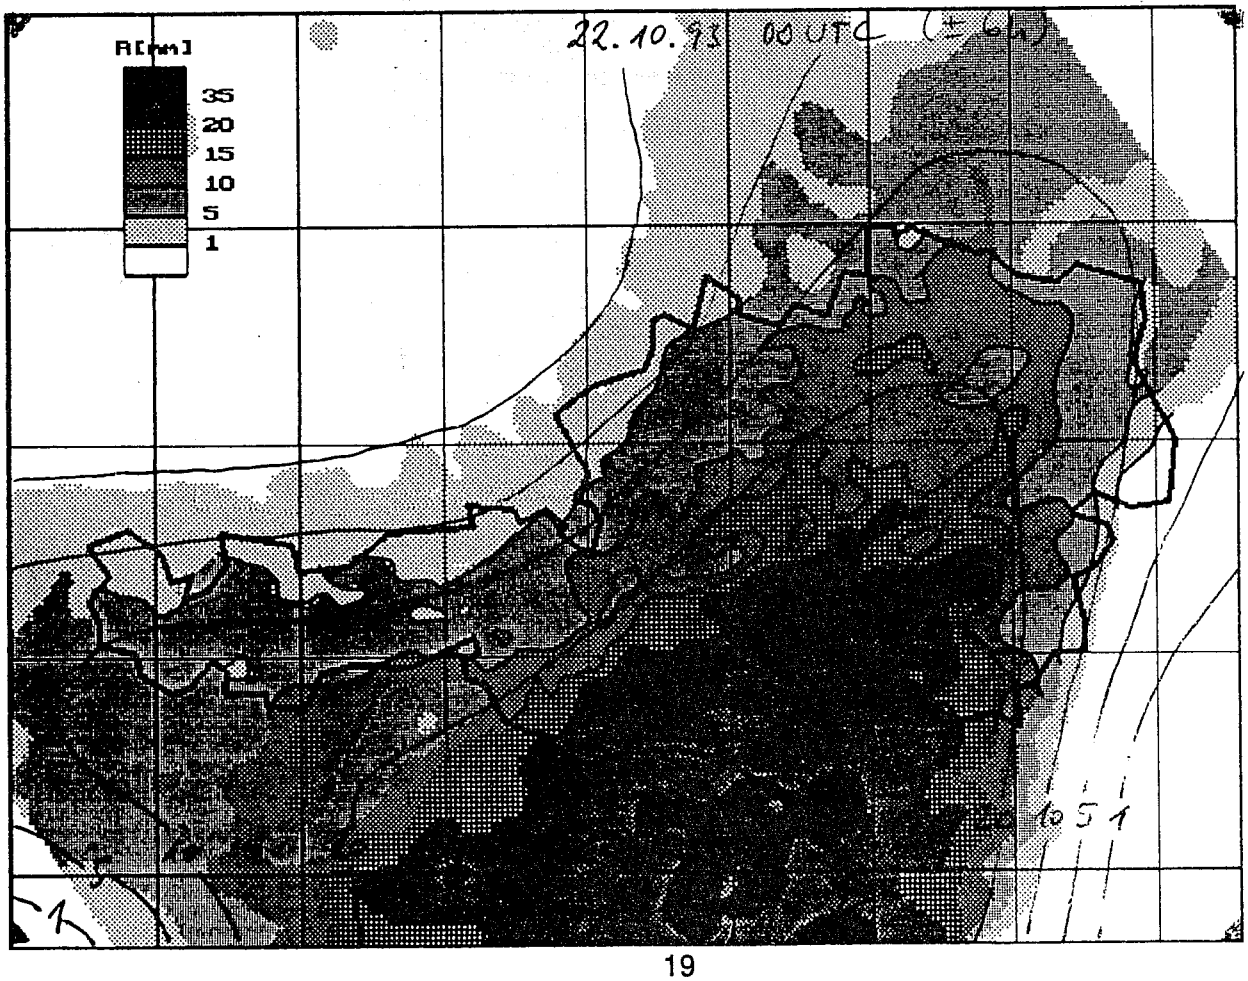

Fig. 14

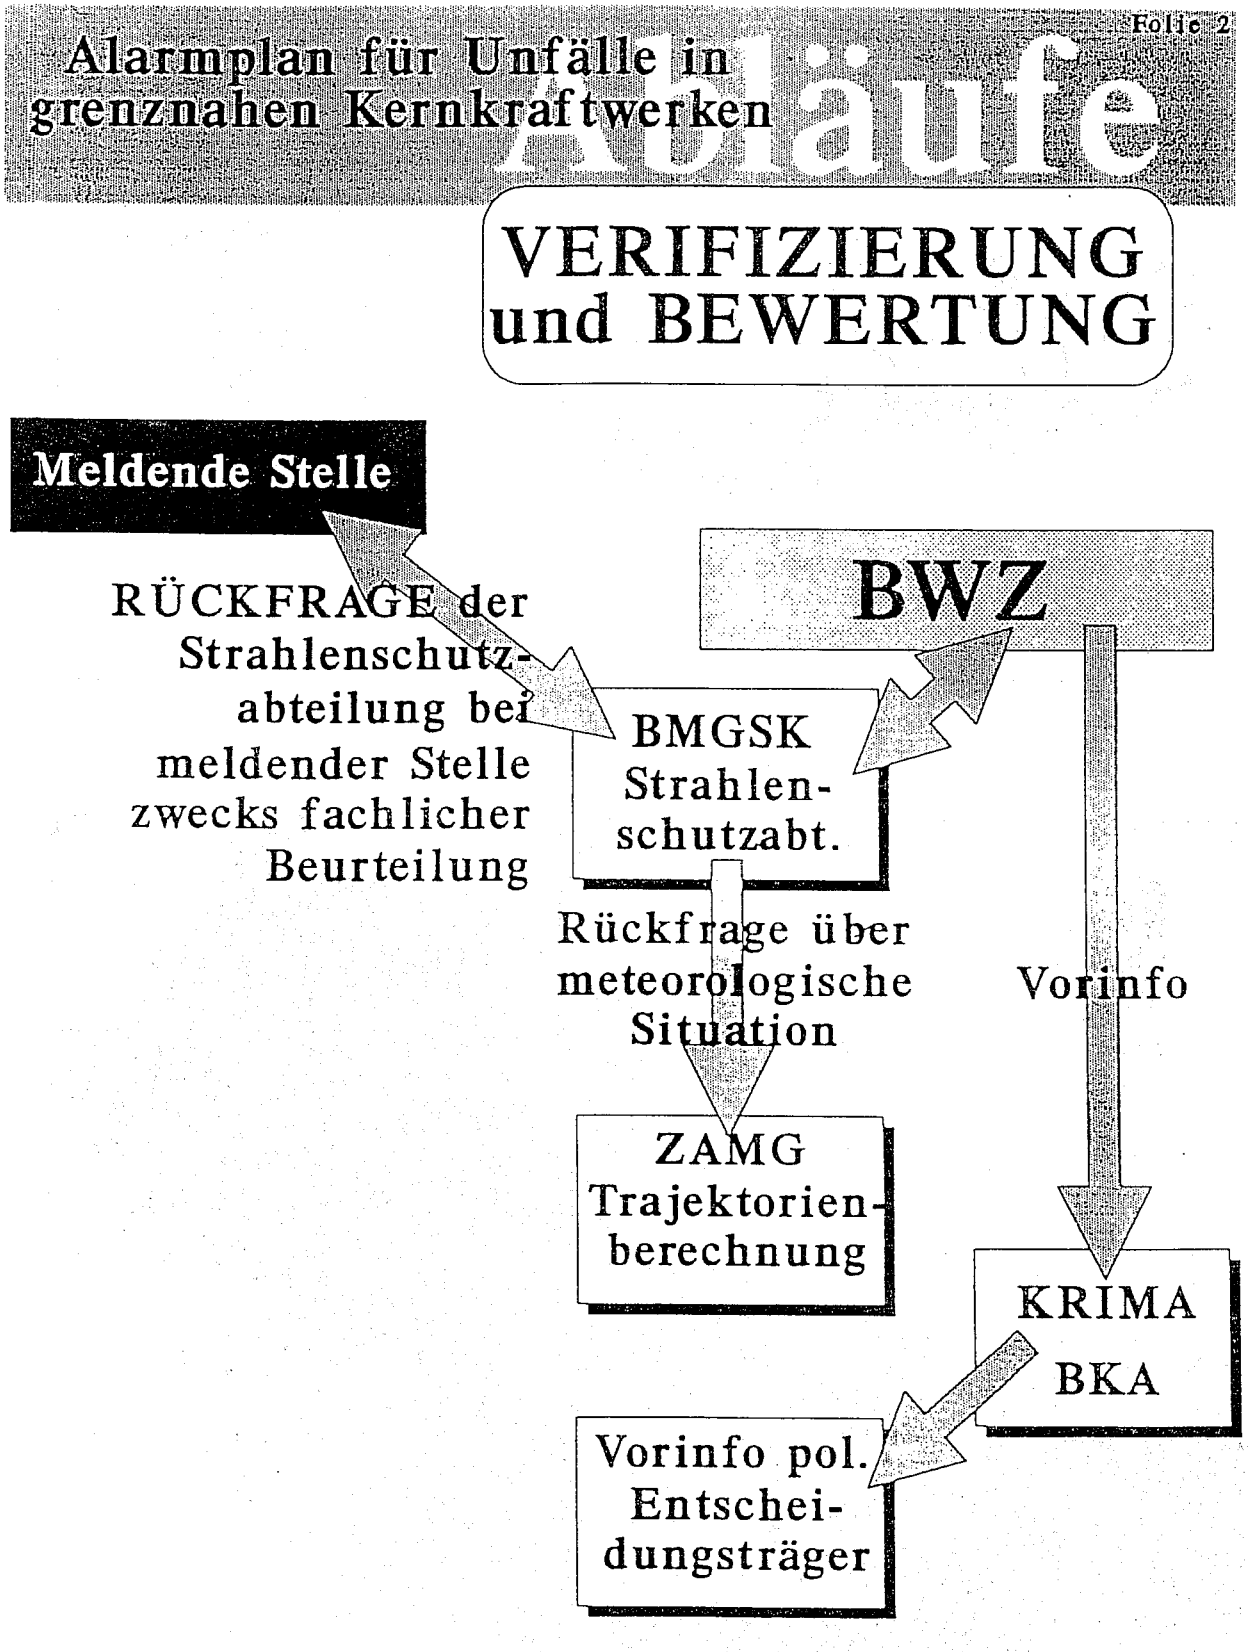
Folie 2
Alarmplan für Unfälle in grenznahen Kernkraftwerken
Abläufe
VERIFIZIERUNG und BEWERTUNG
Meldende Stelle
RÜCKFRAGE der Strahlenschutz-abteilung bei meldender Stelle zwecks fachlicher Beurteilung
BWZ
BMGSK Strahlen-schutzabt.
Rückfrage über meteorologische Situation
Vorinfo
ZAMG Trajektorien-berechnung
KRIMA
BKA
Vorinfo pol. Entschei-dungsträger

Royal Meteorological Institute Belgium 27-09-1994.

Hardware and software installation

The computer installation and the budget are shared by 3 institutes:

- Royal Meteorological Institute (RMI),
- Royal Observatory of Belgium (ROB),
- Institute for Space Aeronomy (ISA).

The infrastructure is operational since mid 1990. The principle was to allocate one departmental unit to each institute and to let one or several number crunching units accessible by the users of the three institutes.

The equipment has been somewhat modified since 1992.

At RMI, we had 2 Apollo DN10000 (one monoprocessor and one biprocessor) as departmental and number crunching units. They were replaced this year by one HP 9000/800 G70: a multiuser system with 2 PA-RISC cpu's, 96 MHz clock, 248 kB I/D cache, CM 96 MB, MM 5 GB, 1 DAT, 1 CD-ROM.

Three number crunching units are accessible in the other institutes for the RMI users:

- 1 Convex C1 at ISA,
- 2 Apollo DN10000 (each with 2 PRISM processors) at ISA and ROB.

A number of workstations are located in different sections of the RMI. Several workstations with different functions were added to give the present configuration:

- 1 communication server: HP 750-66 (1 PA-RISC cpu, 66 MHz, CM 64 MB, MM 2.6 GB, 1 DAT, 1 CD-ROM),

- 1 routine forecast server: HP 755-99 (1 PA-RISC cpu, 99 MHz, CM 128 MB, MM 4 GB, 1 DAT, 1 CD-ROM),

- 1 graphic server: HP 730-66 (1 PA-RISC cpu, 66 MHz, CM 64 MB, MM 1.8 GB),

- 1 server provisionally dedicated to archive conversion: HP 735-125 (1 PA-RISC cpu, 125 MHz, CM 176 MB, MM 5 GB [1 internal disc, 4 external discs], 1 DAT).

Several X-terminals were recently introduced in the network: 3 HP Entria, 4 HP Envizex.

The printing and drawing units consist of:

. line printer HP 2567B (1200 lpm)
. HP laserjet series II
. HP laserjet 4

. Tektronix phaser III
. Versatec 8900

Nearly all these units are connected to an Ethernet LAN.

SOFTWARE

- For the Apollo stations: OS is Domain (environment: AEGIS, BSD 4.3 ans SYS V). 90% is SYS V. Korn Shell.

- For the HP stations: OS is HP-UX 9. Korn Shell.

The backups are performed by HP Omniback through the LAN for the UNIX stations. Omniback uses DDS or TCP/IP protocol.

On the LAN, the basic protocol is Ethernet 802.3.

- between the Apollo stations: all the commands are network wide under DDS protocol. Domain contains the functionalities of TCP/IP and NFS.

- for HP stations, UNIX and Xenix PC's: TCP/IP + NFS;

- for DOS PC's: Domain Personal Computer Interconnect (DPCI) will be soon replaced by PC/TCP.

External communications

- with GTS: data arrive now under X25 protocol through a 14.4 kbps line and a MEGAPACK X25 switch. The data arrive on the communication server. A backup is configured on the routine forecast server.

- with ECMWF: data arrive now on the communication server through a CISCO router via a 64 kbps liaison. Protocol is X25 with TCP/IP above.
The speed was increased from 9600 bps to 64 kbps in last June. Since we have this high speed facility, we have strongly encouraged our users of ECMWF computers to use the interactive and FTP facilities instead of remote batch. Since 10th September 1994 is the remote job facility from Belgium to ECMWF no more available.

- with belnet: belgian research network between universities and research centers. The control is located into the VUB (Vrije Universiteit van Brussel). It runs under HDLC at 64 kbps speed and is connected to Internet. It is a gateway to Internet.

- with the belgian public network (DCS): through an HP station of ISA.

Libraries

- mathematical: Naglib 13.0 g2.0 on Apollo DN10000 and workstations;

- graphical: Uniras 6.1 f (GKS) on Apollo workstations;

- relational DB: Ingres 6.2 on Apollo DN10000 and client on Apollo workstations via NCS.

Plans

The contract of the Apollo number crunching units, the departmental units and some workstations comes to end in mid 95. An invitation to tender has been prepared for two distinct contracts: one for a common number crunching unit for the three institutes and another one for the departmental units.

For the end users, the following option has now been choosen: acquire X-terminals connected to more powerfull servers rather than increase the number of small stations.

Projects run at ECMWF

At the present time, two projects are using the computing facility of ECMWF. The other retrieve MARS data.

- Development of an atmospheric transport model. This model is intended to become the central part of a future Chemical Transport Model (CTM) of the atmosphere. An offline version of this CTM is coupled with the ECMWF forecast T63 model: dynamical fields calculated by NWP model are read as input by the offline version of the CTM. This model was first used to investigate the formation of polar stratospheric clouds.

- Predictability experiments on a simplified thermal convection model. The base equation is the extended Kuramoto-Sivashinsky equation integrated with rigid boundary conditions. The resolution leads to chaotic solution. The aim was to study the error growth and the spatial properties of this error.

LAN and WAN of KMI-IRM (for internal use only) 21/09/94

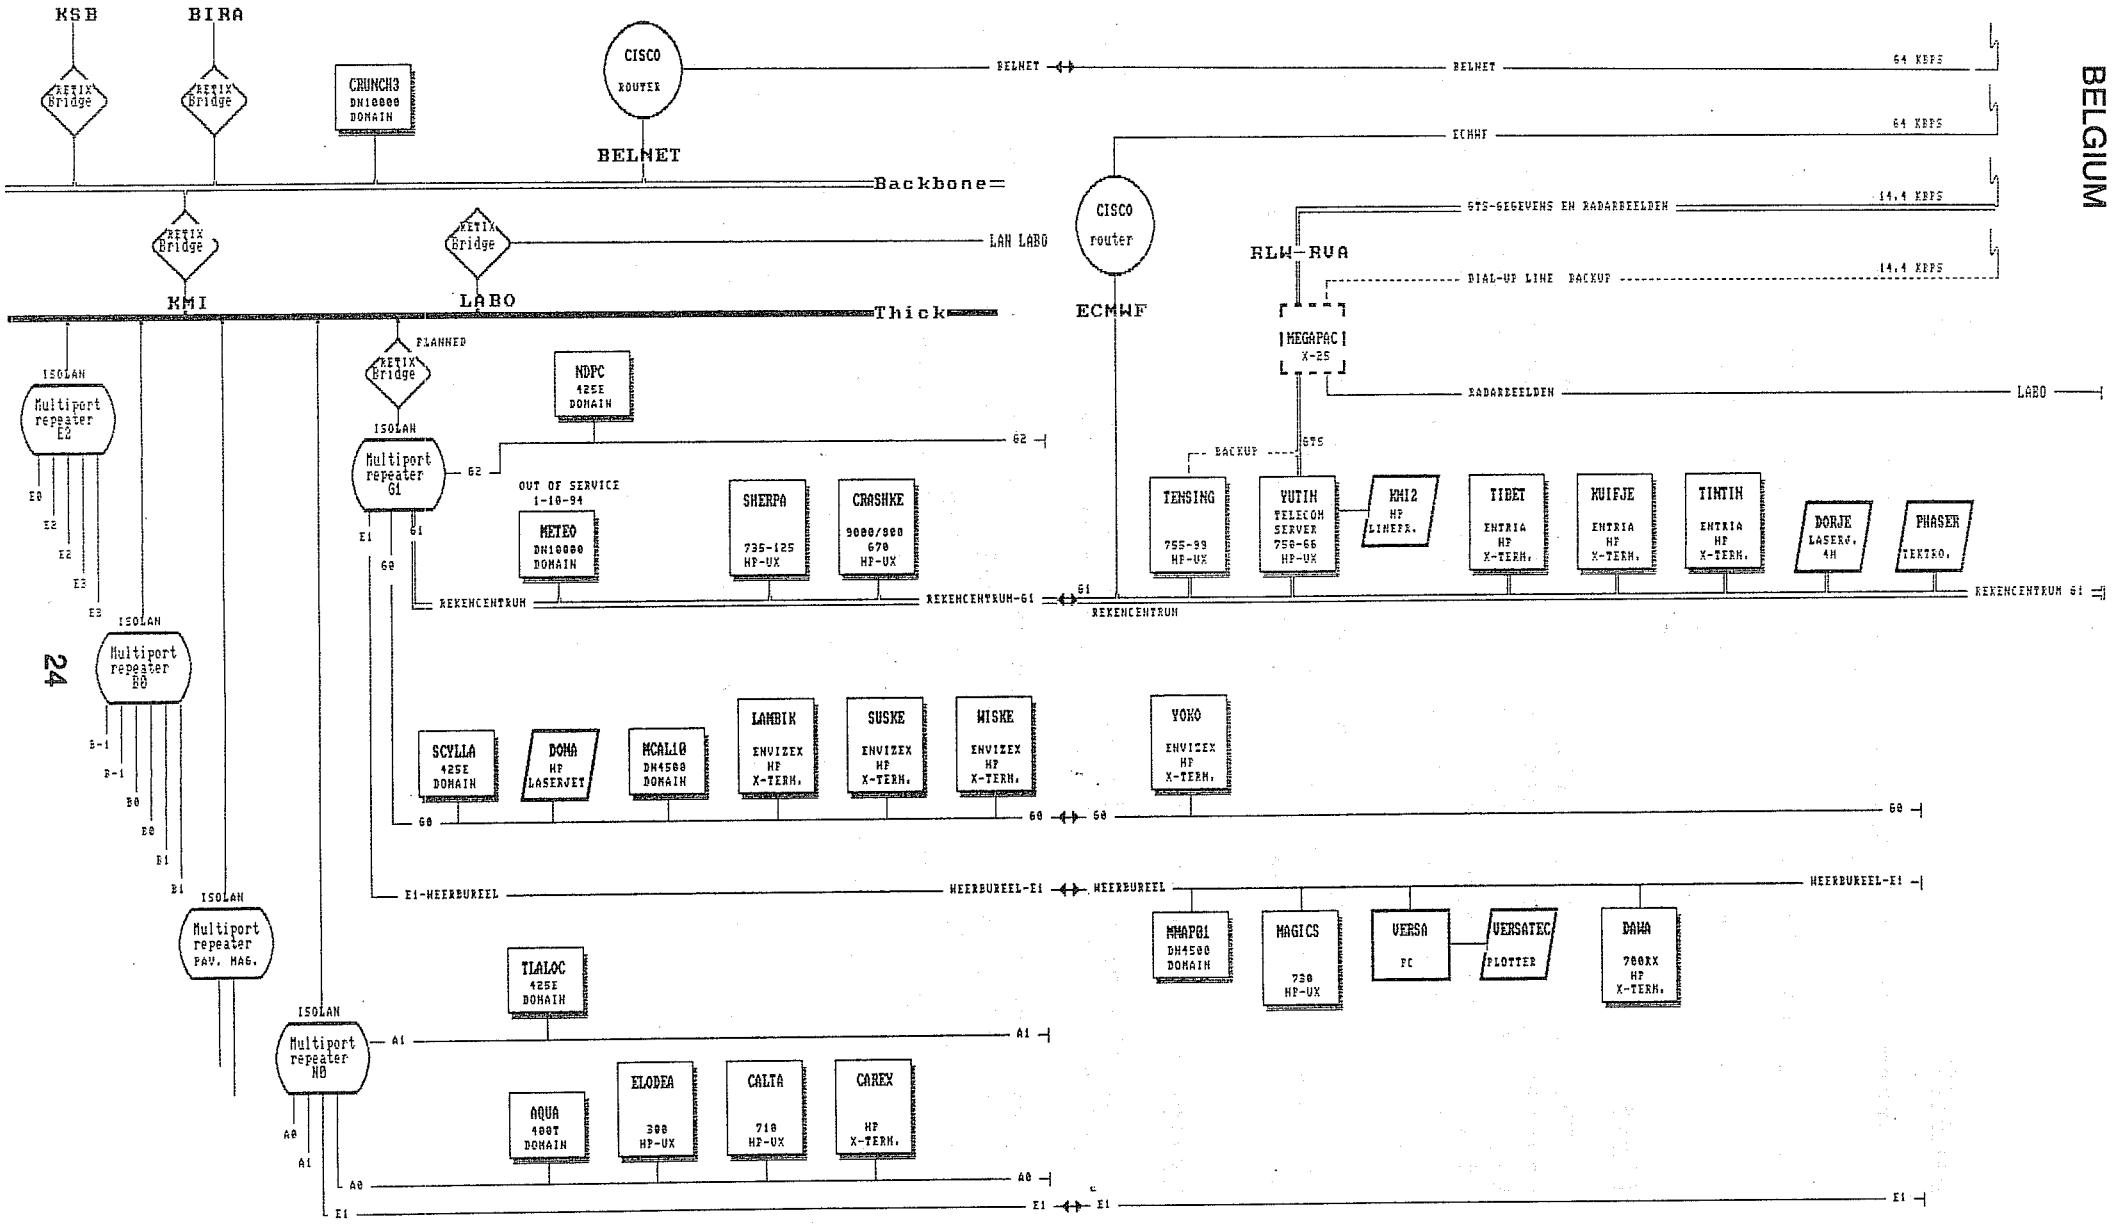

BELGIUM

24

The computer system at DMI

The computer system at the Danish Meteorological Institute has changed a lot in recent years. From a system based on a mainframe computer a distributed computer system has been developed.

The mainframe computers in the DMI system are two VAX-6410 computers coupled into a cluster and two Convex vector-computers, a C3220 and a C3880. In addition to these computers there are approximately 200 workstations i. e. Sun's and SGI's, and approximately 200 PC's.

As you can see in the diagram, the network is based on a segmented 10 Mbit ethernet coupled by multiport bridges connected via a 100 Mbit backbone net.

The communication system connected to the GTS-lines and the domestic lines are handled by one of the Vaxes, while the decoding system is running on the other Vax.

The decoding system, based on the ECMWF Data_Acquisition and Pre_Processing system, has been heavily modified and new programs have been added e. g. message-switch, bulletin-generation and a program that searches for bulletins in data not recognized by the Data_Acquisition system.

Bulletins handled in the decoding system are distributed by the message-switch to the Danish airports, to Greenland and to the GTS via the ESWI line.

All data received by the decoding system are routed to workstations placed in the weather department. The workstations placed in the weather department and in the airports are used by the forecaster.

Bulletins including observations are decoded and reformatted into BUFR, while GRID-bulletins are converted into GRIB. The BUFR-data are collected on the Vax-system four times a day and sent to the C3880, and the GRIB-data are sent to the GRIB-database that is placed on both Convex machines and on two SGI-machines.

Plotting of received observations are mainly done by CalComp penplotters. In Copenhagen we use two plotting systems: the Vax-system for presentation of synoptic observations, and the Sun-system for presentation of upper-air-, synoptic- and metar-observations. In the airports the Sun-system is used.

The disseminated analysis and prognosis from ECMWF are primarily received at the C3880. When data have been extracted into the different fields, the data are put into the GRIB-database and distributed to the other Convex and to the two SGI-machines and stored in the GRIB-database.

Some of the GRIB-data are used for plotting of prognostic maps or for other purposes in the daily forecasting suite. The plots are generated in Post-Script and sent to QMS laser printers using the formats A4 and A3. Other GRIB-data are used as boundary values for the HIRLAM analysis and prognosis system.

The prognosis is made for an area around Denmark and an area around Greenland. The prognosis data are stored in the GRIB-database on all machines with a GRIB-database. After a few days the prognosis data are stored on the metrum-storage, a storage system that is driven by the software UNITRE.

Data used by external customers are put on our commservers. The external users defined in the usercatalogue on the commservers gets the data by calling the

system via the public telecommunications network. It is not possible for external users to pass through the commserver into the internal net.

We have implemented some security measures. The users have to change passwords with regular intervals and admittance to the system from external users can only be accepted through the Convex 3880 or through the firewall-machine. In the future we are planning to prevent admittance to our system via the Convex machine.

Submission of jobs are still done from the Vax-system. In this connection we are looking forward to a system that does not involve the Vax machinery.

The remote jobs running at ECMWF are to a large extent data retrieval from the MARS archives.

The research department has made some experiments with the HIRLAM model, and the ECMWF computers have been used in calculating trajectories for the stratosphere.

In the future the main use of the ECMWF computers will be retrieval of data from the MARS archive, but there is a possibility that a research project about ozone will be initiated.

The need for retransmission of disseminated products has increased because of changes in our setup and because of growing needs for retrieval of BLOCK-data. In this connection we would like to ask if it would be possible to develop a system of requesting retransmission of one special field or area.

Danmarks Meteorologiske Institut

Netværks topologi

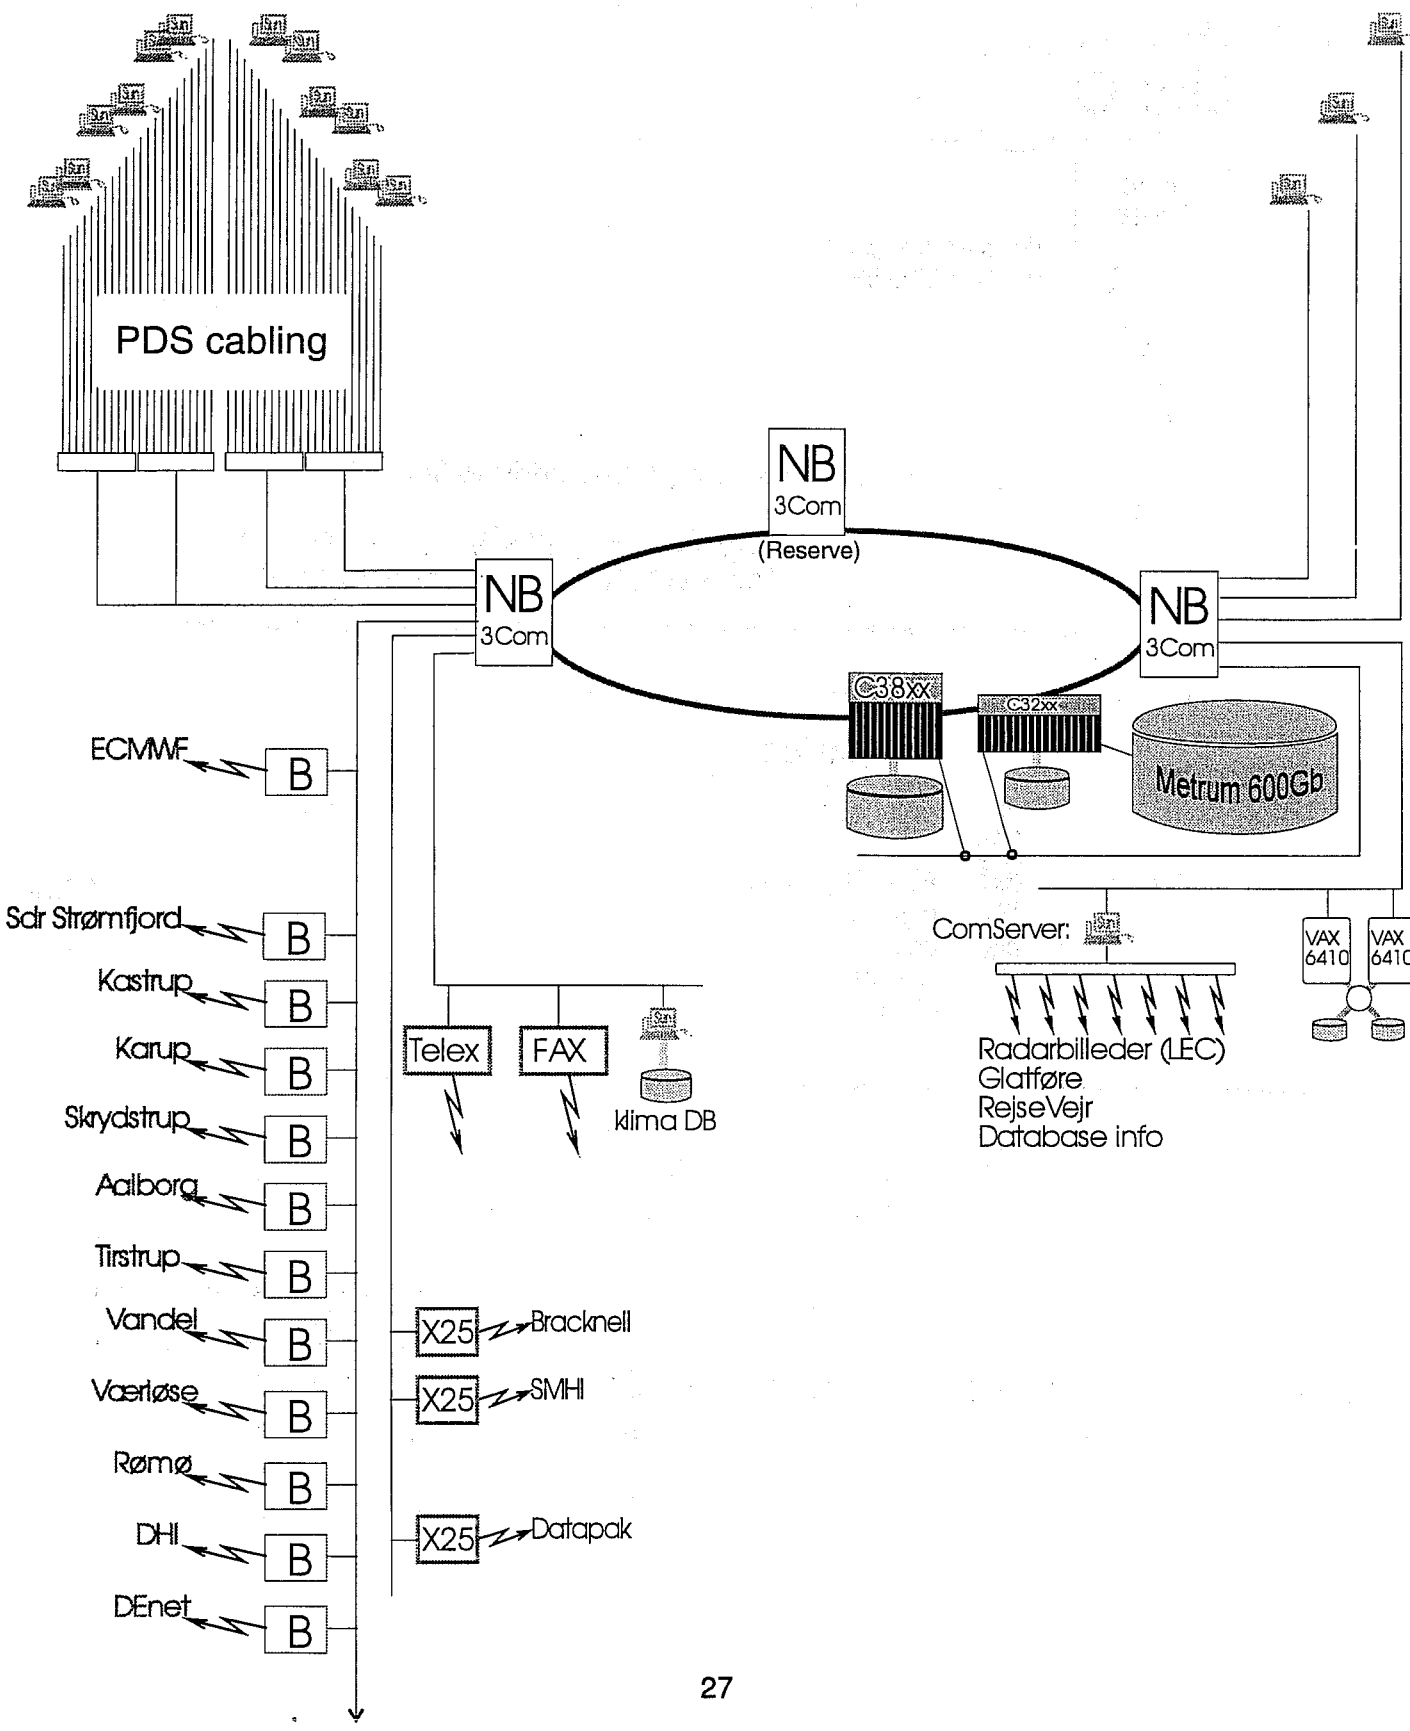

Danmarks Meteorologiske Institut

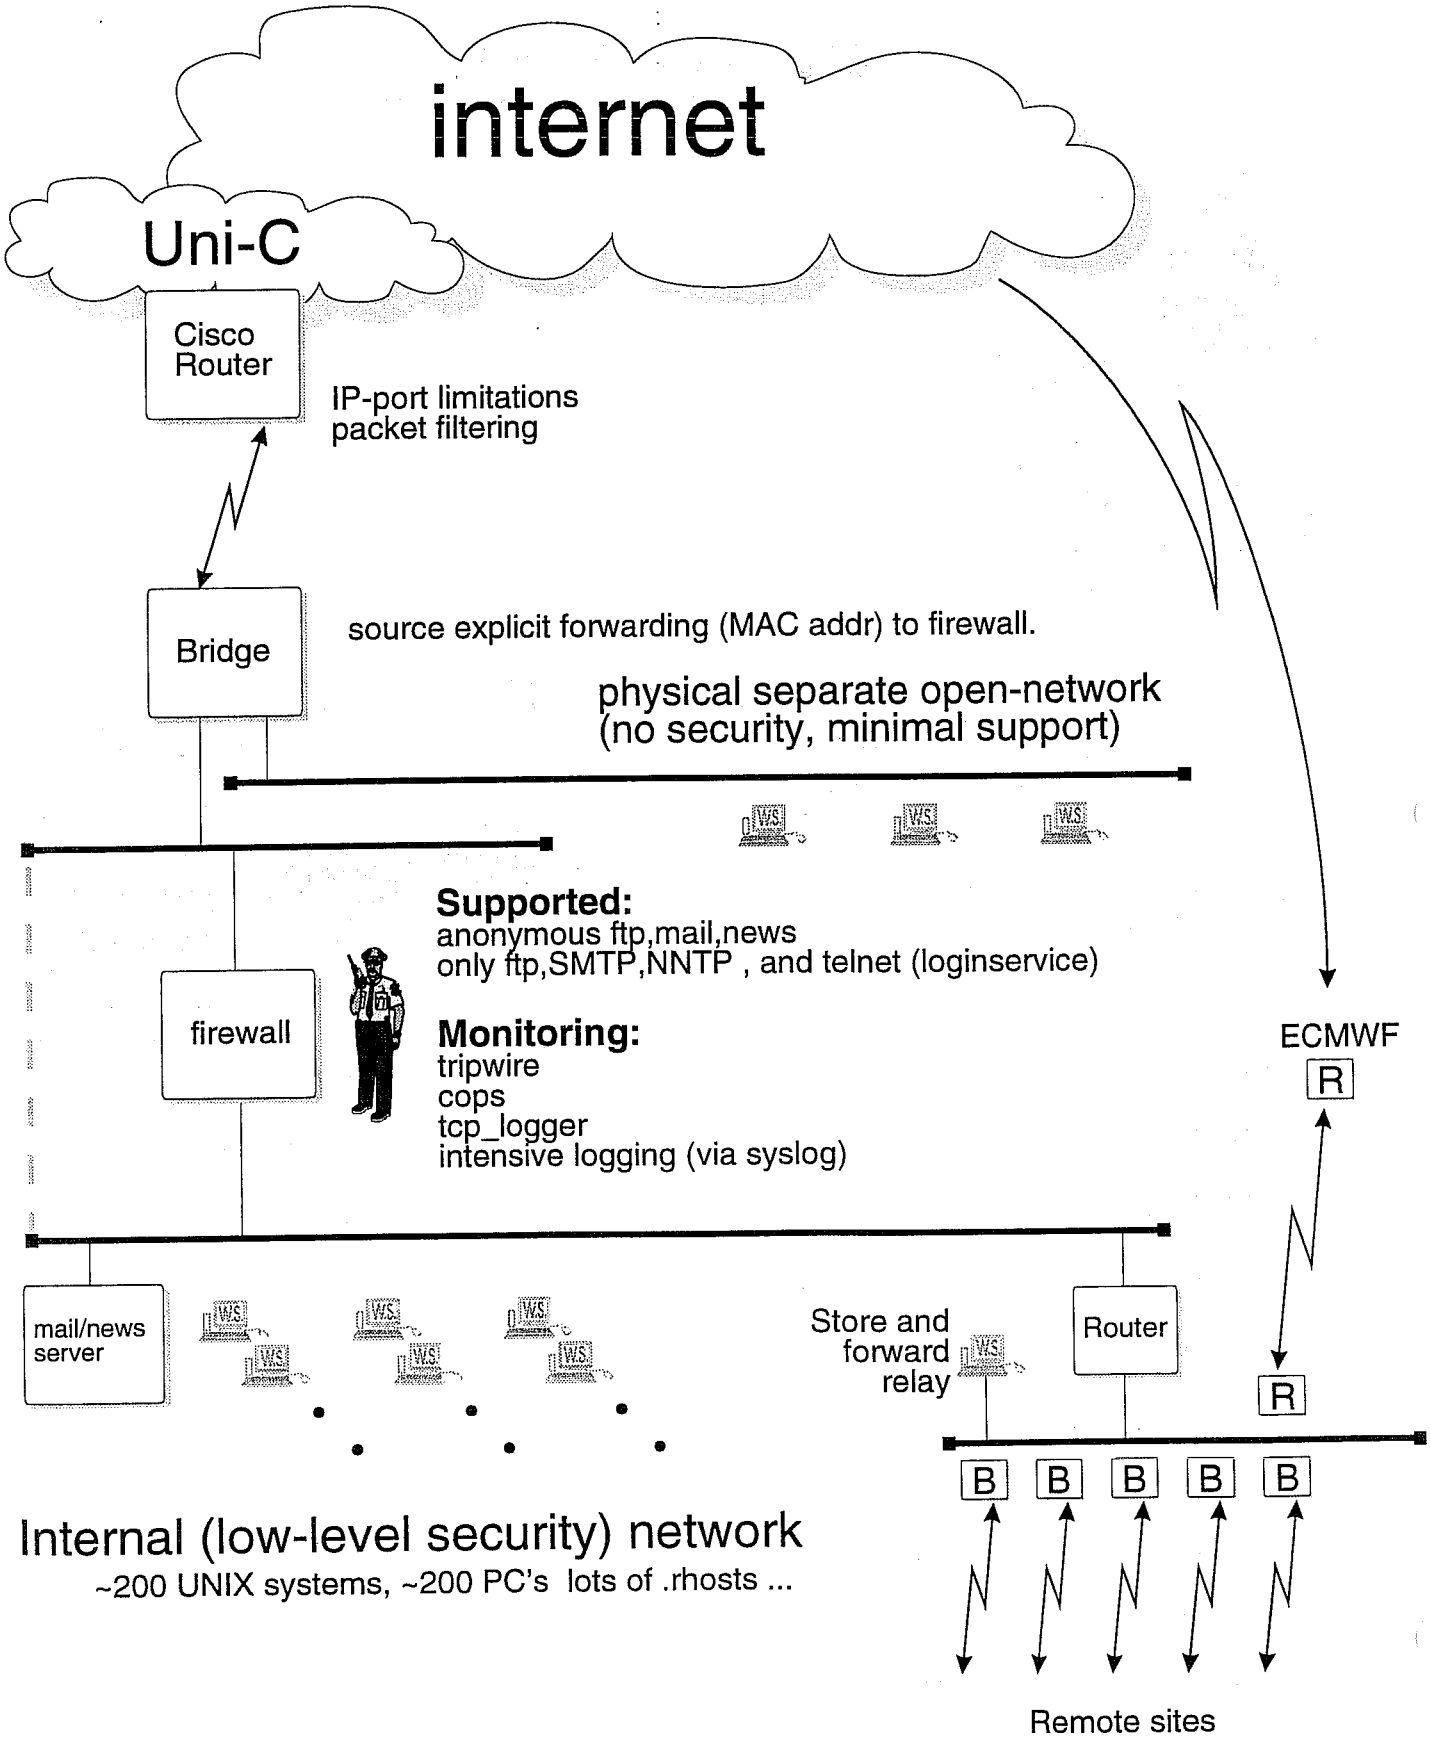

COMMENTS AND QUESTIONS

J. Hennessy — The idea of requesting retransmission of one special field or area makes sense. We will look into it.

B. Barg — DWD is interested in the same idea.

COMPUTER NETWORK AT FMI

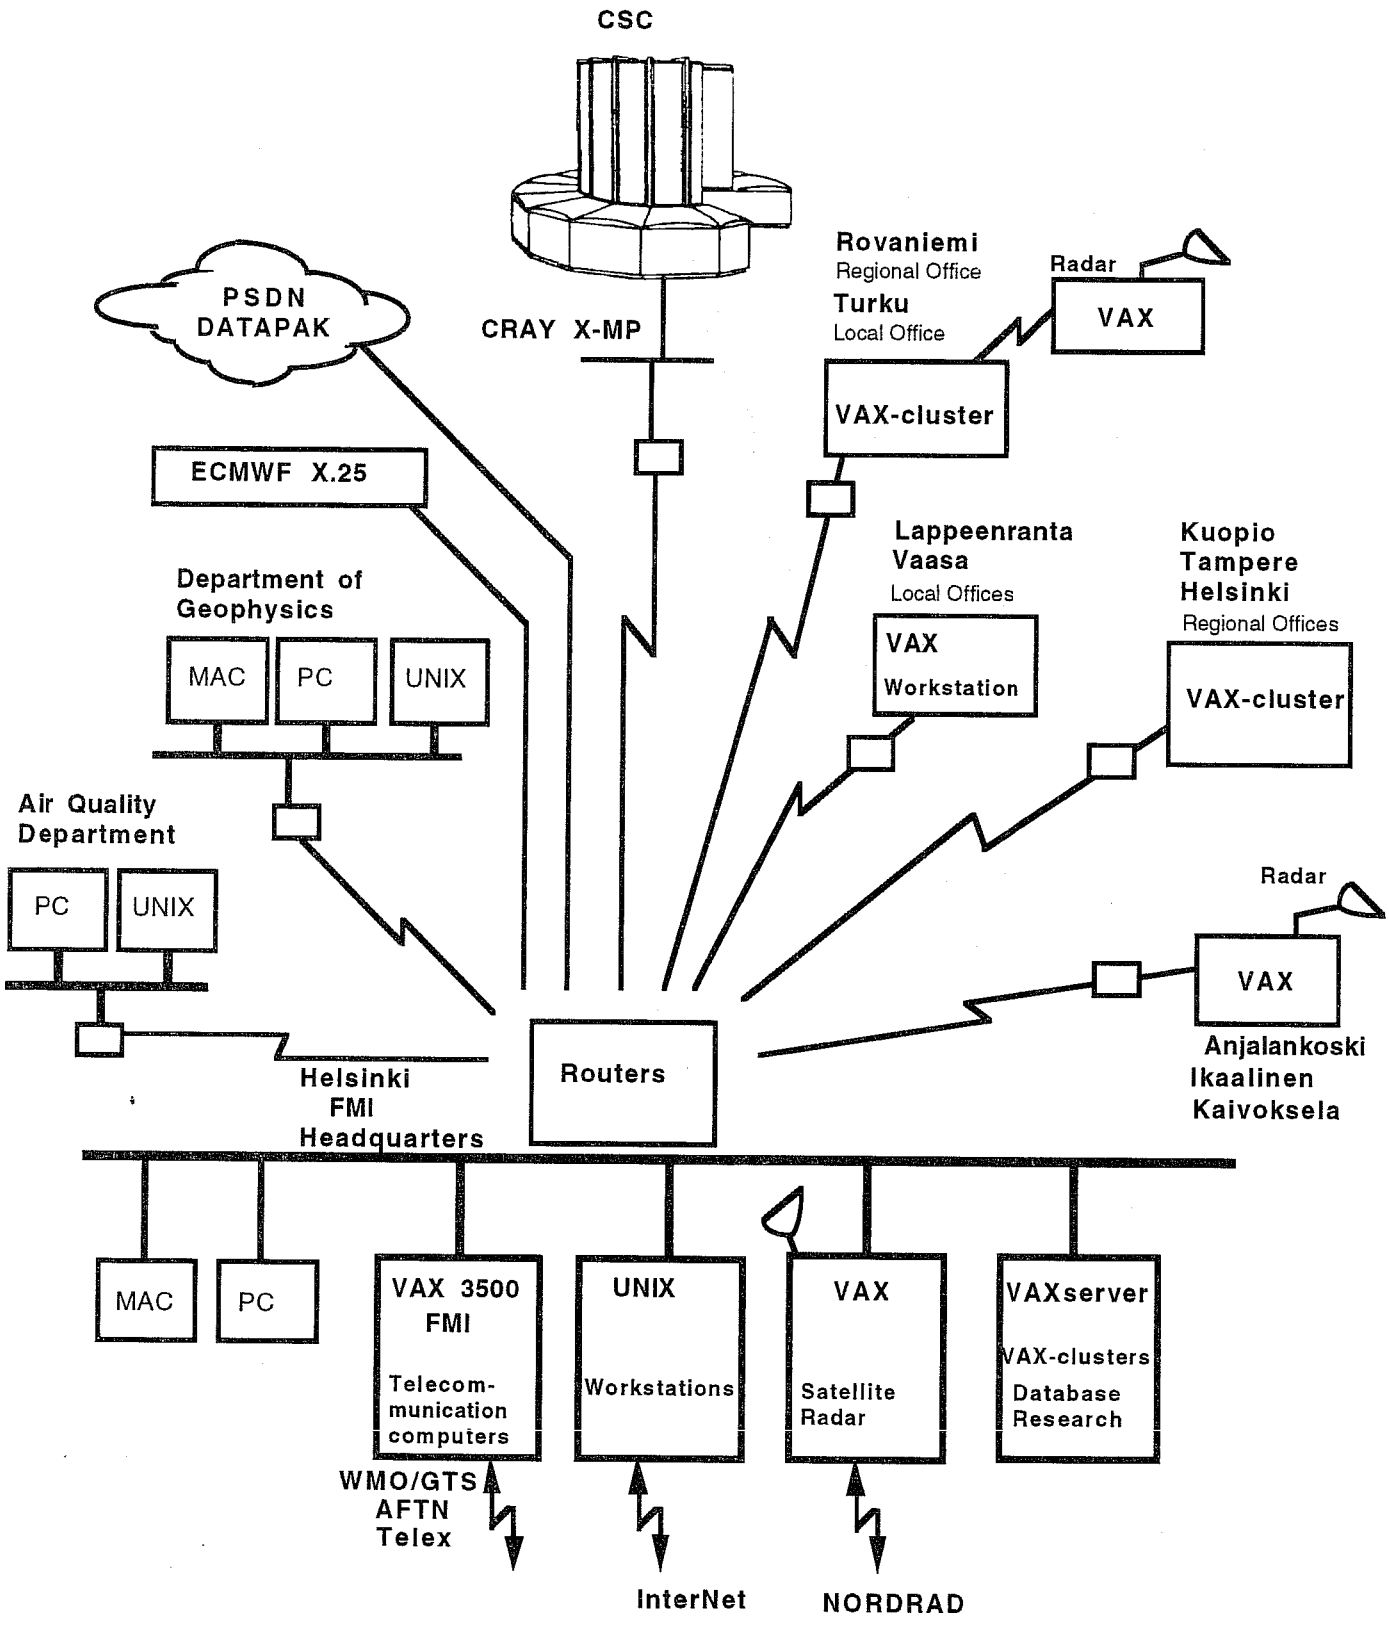

DATA PROCESSING DIVISION SERVICES

DATA COMMUNICATION SERVICES

- FMINET
- Connections to other networks
- Telex, Telefax
- Collection and dissemination of observation data

COMPUTER SERVICES

- ECMWF
- Supercomputer Cray
- VAX- and UNIX- computers
- Micros

8.9.94/KSo

DATA PROCESSING DIVISION SERVICES

TECHNICAL CONSULTATIONS

- Graphical systems
- Database systems
- Meteorological systems
  - Meteorological workstation
  - Videotex-service
  - Radars
  - Satellites
  - Soundings

8.9.94/KSo

COMPUTER NETWORK AT FMI

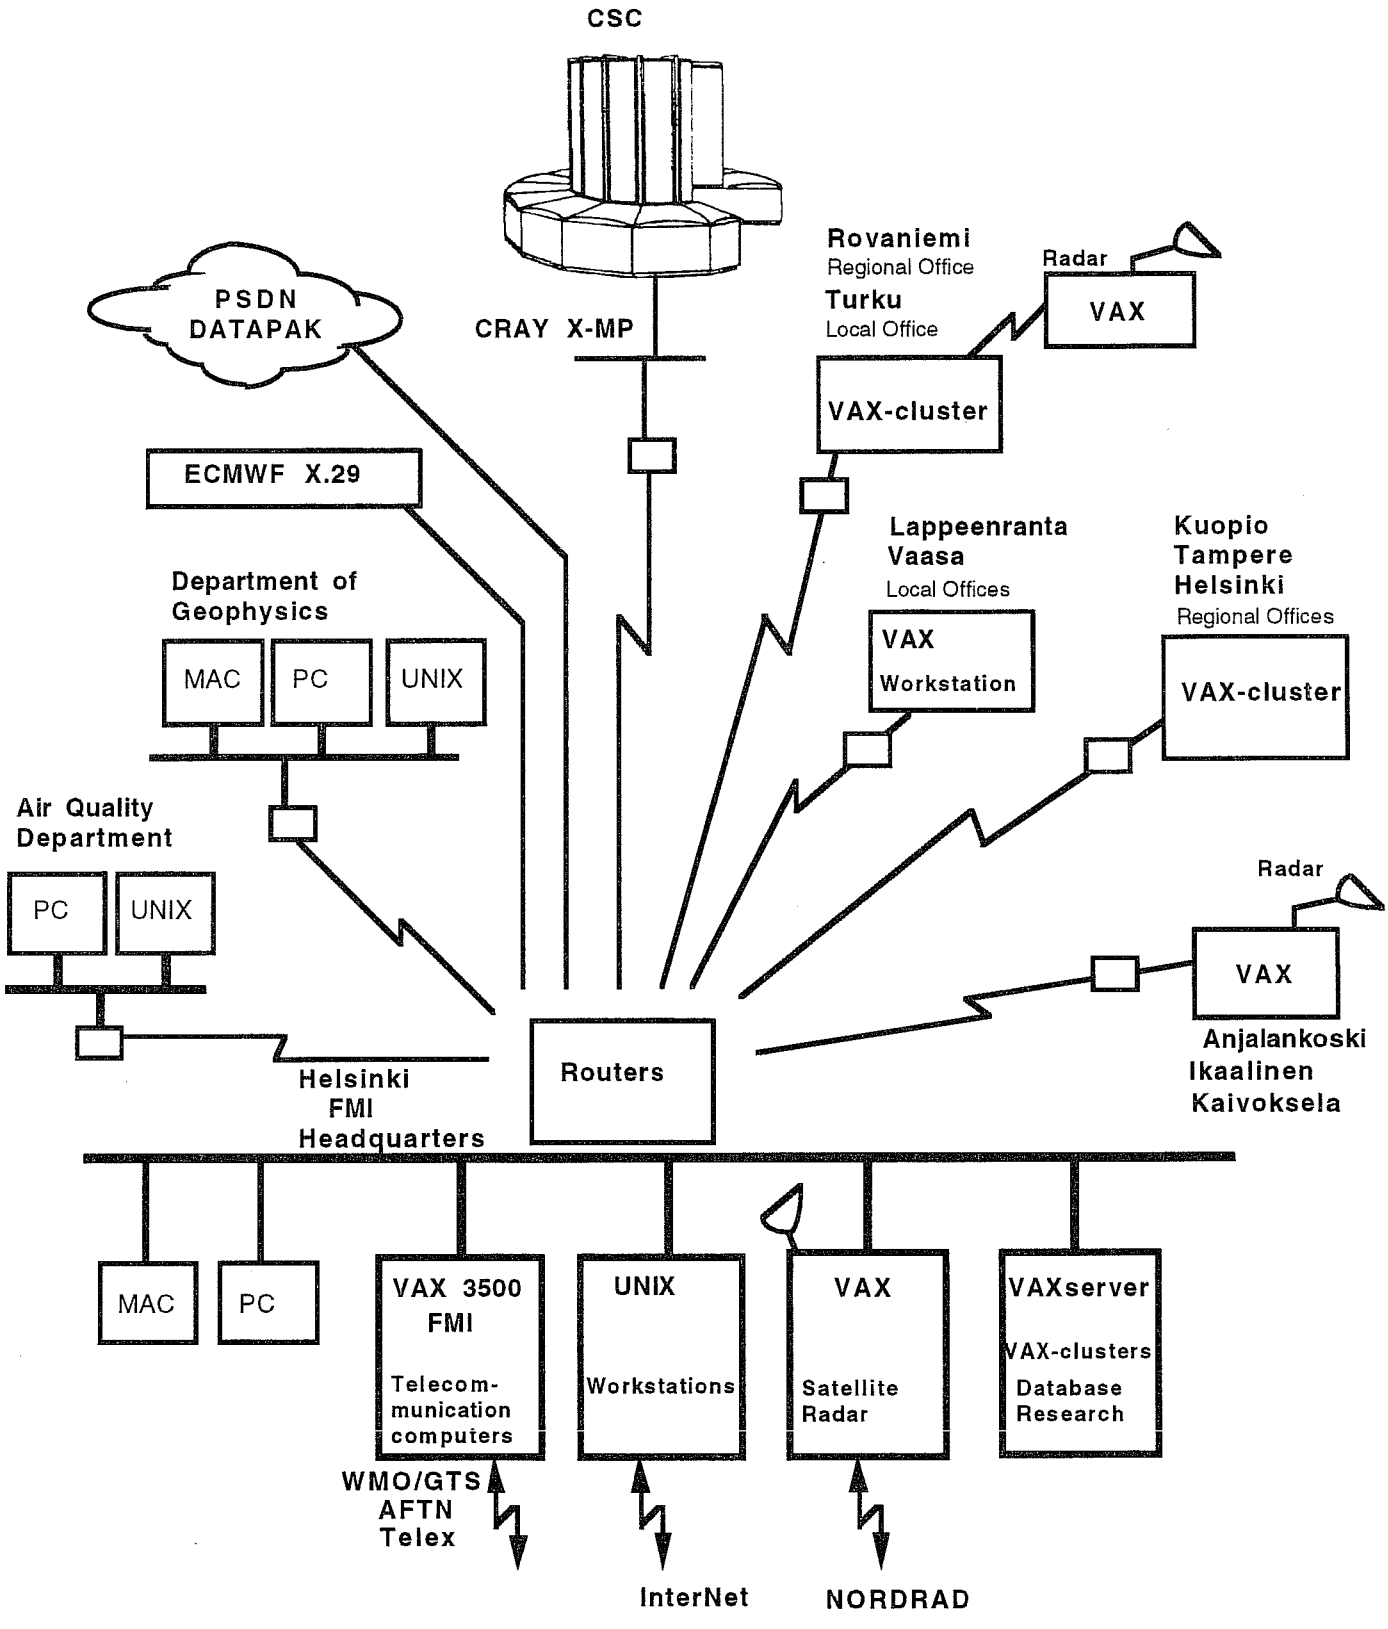

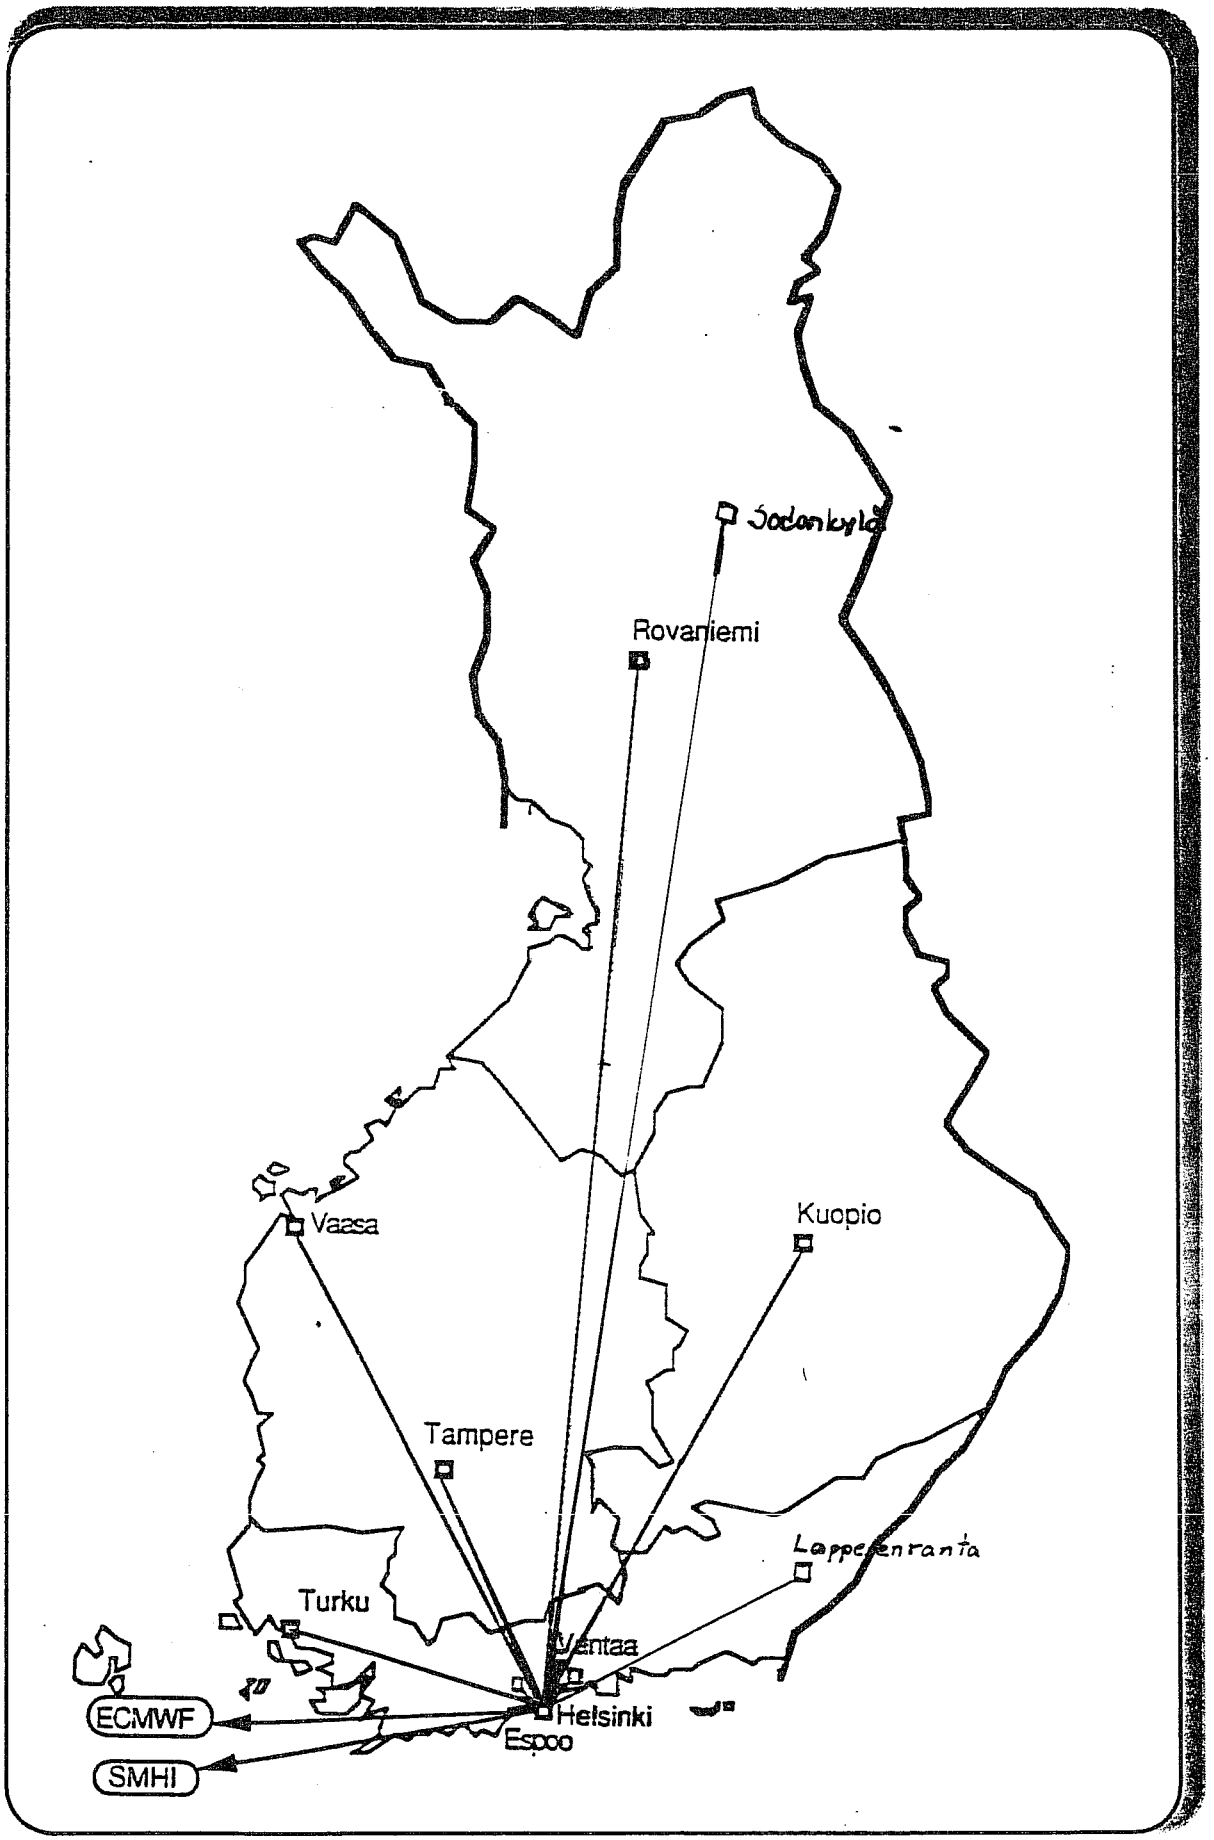
Sodankylä
Rovaniemi
Vaasa
Kuopio
Tampere
Lappeenranta
Turku
Vantaa
ECMWF
Helsinki
Espoo
SMHI

DATA AVAILABLE ON METEOROLOGICAL WORKSTATION

- METEOSAT IR images
- SYNOP observations
- TEMP observations
- radar images
- numerical forecasts
- trajectories
- lightning observations

23.9.1994/KSo

Finnish Meteorological Institute

PROJECTS RUN AT ECMWF

1. Retrieval of data from MARS
2. Development of small scale forecasting models
3. Application of the urban air quality model to European sites
4. Ocean wave model
5. Trajectory models
6. ECMWF preprocessing system
7. Ozone research
8. The Nordic HIRLAM model

COMMENTS AND QUESTIONS

S. Pasquini — I see you use VAX equipment for your GTS telecommunication, probably using the VAX/VMS operating system. Do you have plans to move to OSF1?

T. Hopeakoski — We have plans, but have not decided when to do it yet.

Member State
Computing Representatives Meeting
27 - 28 September 1994

Computers and networks at Météo-France

Number crunchers
The main computer system consists of a 4 processors Cray Y-MP C98, with 2GBytes of main memory, 4GBytes of SSD and 50GBytes of disk space.
It is used for running the operationnal model and for users developments. (twice a day)

The second computer system is a Cray-2 (4 CPUs, 2GBytes of memory and 15GBytes of disk space).
It is dedicated to climate research (with no operationnal use).

Other computers
Two Control Data systems (Cyber 960 and Cyber 860) will be replaced by a set of Hewlett-Packard servers (G50, T500) running HP-UX during the first quarter of '95.
These computers are used :
- to run the operationnal production suite (collecting observations and store them in a data base, pre-processing of data, starting the model on Cray and post-processing of results) ;
- to provide user access to data and ;
- as a development tool for users.

Two Control Data systems are dedicated to a specific use :
- a CD4680 running ORACLE for a climate data base ;
- a CD4360 acting as a network and mail manager.

File server
The file server is made of a Convex C3820 (2 CPUs, 1GByte of memory) with 160GBytes of disks, of which 140GBytes are acting as a disk-cache, and 2 StorageTek silos (ACS 4400) with 4 drives on each. The total capacity of the silos is about 5-6TBytes (with about 600MBytes per cartridge).
The software used is Unitree (version 1.7.6).
The use of the system today is over 2.8TBytes for 830 000 files.

Network
The main network equipment consists of an FDDI ring connecting together ethernet lans and computers. A HiPPI switch (NSC PS32) is used for the access of supercomputers to the file server.

Links to the outside :
- the Transmet system, which is used for collecting (gathering) observed meteorological data (GTS, nationwide...) ;
- dedicated lines (ECMWF, CERFACS, regional services...) ;
- TRANSPAC (french X25 network) ;
- RENATER (french part of the Internet).

Connections from the outside are filtered by a gateway (TIS Toolkit). We have studies for installing a strong authentication system.

Projects run at ECMWF

For France, there are 22 registrated project, of which about 12 are really active.

FRAUIOPG : observation program (FASTEX)
ECMWF resources usage : MARS retrievals, few CPU

FRENUMET : tropical meteorology
ECMWF resources usage : MARS retrievals

FRSCFKAL : developments related to EPS
ECMWF resources usage : MARS retrievals, few CPU

FRSCPMAR : marine prediction, wave prediction
ECMWF resources usage : MARS retrievals, few CPU

FRSCTRAN : various use of observed data
ECMWF resources usage : MARS retrievals

FRTOASPI : nowcasting
ECMWF resources usage : MARS retrievals

FRTOCOPT : TOGA-COARE experiment
ECMWF resources usage : MARS retrievals (automated qsubs from France)

FRTODCLI : validation of models outputs ; climate simulations
ECMWF resources usage : MARS retrievals, CPU intensive

FRTOMAD : data assimilation
ECMWF resources usage : MARS retrievals

FRTOMESO : observation program (FASTEX), climatological studies, mesoscale studies,
land/atmosphere parametrisations
ECMWF resources usage : MARS retrievals (automated qsubs from France)

FRTOSNOW : snow cover simulation on the Alps
ECMWF resources usage : MARS retrievals, few CPU (automated qsubs from ecserver), ecfile

ECSN participation.

Usage of ECMWF resources(today and for the next future)

MARS is the most used resource, as all projects more or less use it. And this will continue during the next two years.
From the point of view of the most intensive users of MARS, this system seems to use a significant amount of CPU (this is one of the reasons for *unattended* remote batch from France). Nevertheless, the general impression about MARS is very positive.

The **CRAY CPU** is used for MARS retrievals post-treatment and, very intensively, for large climate simulations (10 years, T106). This should also continue during next two years.

We hope we'll be able to make a large use of the recently installed **CRAY T3D** during the next two years, for the testing of a parallelized ARPEGE (french flavour of IFS). This should probably generate data to be kept in ECFILE and trafic over the network.

Suggestions and queries

For users, the only problem is not due to ECMWF. It is interactive response time and time of files transfers caused by a heavy trafic over the 64Kb/s telecommunication line between ECMWF and Météo-France. We have set up priorities, we have to begin to use FTA for files transfers, and we are requesting the authorisation of increasing the speed of the line up to 128Kb/s.

Questions of users are about the availability of MAGICS 5.0 and a C++ compiler on UNICOS.

Other questions are Computing Representative questions, related to users administration.
- It could be useful to be able to revalidate passwords on ecserver.
- We need a tool to control ecfile use by our users, as acct_admin give us the control of Cray SBU use.

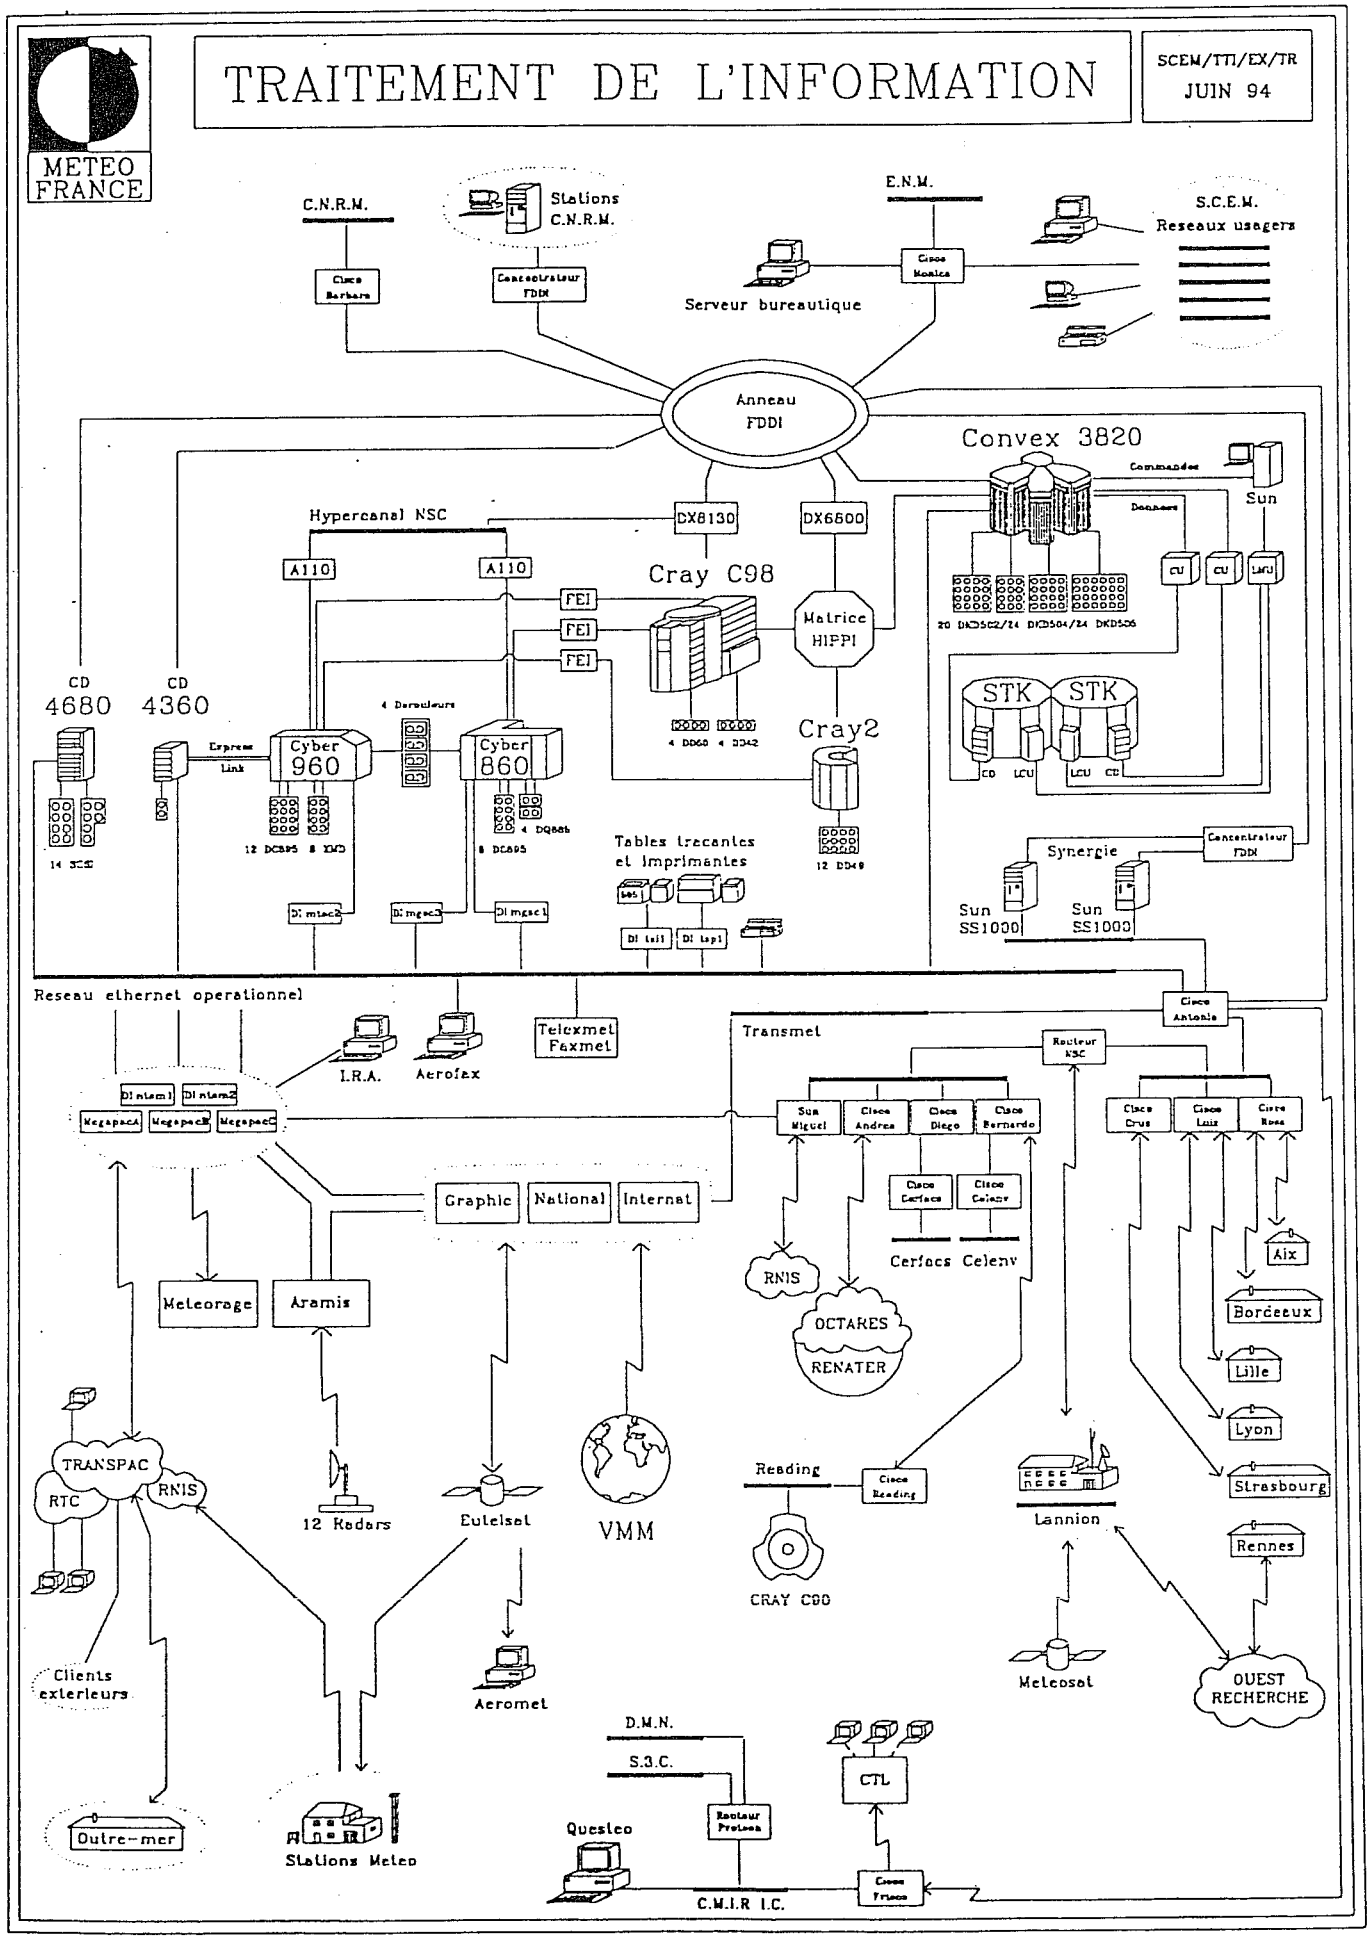
METEO FRANCE
TRAITEMENT DE L'INFORMATION
SCEM/TTI/EX/TR
JUIN 94
C.N.R.M.
Stations C.N.R.M.
E.N.M.
S.C.E.M. Reseaux usagers
Serveur bureautique
Anneau FDDI
Convex 3820
Sun
Hypercanal NSC
DX8130
DX6600
A110
A110
Cray C98
FEI
FEI
FEI
Matrice HIPPI
CD 4680
CD 4360
Express Link
Cyber 960
Cyber 860
Cray2
STK
STK
Tables traçantes et imprimantes
Synergie
Sun SS1000
Sun SS1000
Reseau ethernet operationnel
I.R.A.
Aerofax
Telexmet Faxmet
Transmet
Graphic
National
Internat
Meteorage
Aramis
RNIS
OCTARES
RENATER
Cerfacs Celenv
Aix
Bordeaux
Lille
Lyon
Strasbourg
Rennes
Lannion
Reading
CRAY C90
TRANSPAC
RTC
RNIS
12 Radars
Eutelsat
VMM
Meteosat
OUEST RECHERCHE
Clients exterieurs
Aeromet
Outre-mer
Stations Meteo
D.M.N.
S.3.C.
Quesico
CTL
C.M.I.R I.C.

COMMENTS AND QUESTIONS

G-R. Hoffmann — We will look into both the queries you have raised at the end of your presentation regarding the revalidation of user passwords on ecserver, and a tool to list usage of ecfile by your users.

GERMANY

Deutscher Wetterdienst
Dr. B. Barg

Member States Computing Representatives Meeting,
27-28th September 1994

1. Computer equipment and conncections to ECMWF

There have only been minor changes to the computer environment at the meteorological computer center since the last Computing Representatives' Meeting in September 1992.

The main computer system is shown in Figure 1 and consists of

- one CRAY Y-MP 4/432
- two CYBER (1x CY 180/860 A, 1x Cy 2000)
- two CDC 4680

The CRAY is used for operational work and model development, the Cyber 2000 for pre- and postprocessing, the Cyber 860 for program development, and the two CDC 4680 are used for the migration of all programs - running on the Cyber-equipment under the system NOS/VE into the UNIX-world - and their operational usage. The goal is to finish the usage of the Cyber-equipment in 1997.

The telecommunications computer system is shown in Figure 2. It consists mainly of a VAX-cluster (VAX 8550) which handles the communications.

The link to the ECMWF is shown in some more detail in Figure 3. The line to the ECMWF is rated at 64 kbit/s and runs DECNET or TCP/IP under X.25.

In Figure 4 and 5 is shown the operational data-transfer and the massage routing between ECMWF and DKRZ, ECMWF and ESOC and DWD and SMA.

In the operational work is used DECNET under X.25, in all other cases TCP/IP under X.25.

2. Projects run at ECMWF

There are at present 5 programs which include approximately 30 users.

demppt22	- a small project mainly for referencing and testing purposes These include • running the DWD-Benchmark-Suite • developing and testing benchmark programs for parallel systems • reference site for the CRAY Y-MP4 at the DWD (compiler, libraries and tools)
dewflspm	- spectral (global) modelling
dewf2clm	- regional climate modelling
dewf3eum	- development of a regional forecast model
dewk7wzn	- within this project the global precipitation center in Offenbach gets results of the forecasts of precipitation done by the ECMWF

3. Experience using ECMWF computers

In general there are no difficulties in using the computer equipment at the ECMWF.

Concerning the dissemination system our experiences indicate two minor problems.

1. To request a test data set is at present difficult to handle.

   The first step would be the inclusion of the request in a routine data set. Receiving the routine data set, one has to sort out the test data before the routine data are ready for use. This would mean that a special program is needed for the handling of data (and unfortunately we don't have a program to manage such cases in a simple and general way yet).

   To solve the problem we see two possibilities, Figure 6. Our first proposal is the creation of a special "Group Code" (or something

equivalent) for a **single** test dissemination outside of the routine data dissemination. In this case obviously both test data and routine data would have the same format.

The second proposal would be to try to find a solution within MARS. At present MARS has less facilities than the data dissemination system. We would like to propose that an advanced MARS should allow e. g. output of fields covering a limited area on the globe and/or different types of GRID. This also could solve our problem.

2. In some parts of the documentation of data dissemination there are entries without any explanation.

   Before going into details it is to mention that in general the documentation of data dissemination is good and sufficient.

   An example for an entry without explanation one can find on p. 29 in the Met. Bull. M 3.1 (1), "The dissemination of ECMWF products to Member States", OPTIONS = DELAY, there is a reference to a repeat product facility which is never described in this Met. Bull.

   In other parts of the documentation, explanations are not up to date (e. g. again on p. 29 in the Met. Bull M 3.1 (1), "The dissemination of ECMWF products to Member States", FORMAT = BUFR is missing, but is already active for at least one year).

Finally I would draw your attention to a rather different problem.

The Centre is using (and e. g. for ensemble forecasts already disseminating) GRIB, Section 1, Octet 41 and following, "Reserved for originating centre use" in a form as defined in the MARS User Guide, appendix F, p. 115 and 116.

At the DWD we also include local information in GRIB. However, switching over to GRIB edition No. 1 our local information should go to the same position in the GRIB code as the ECMWF local extensions. Therefore, we would like to make sure that we will be able to distinguish between ECMWF and DWD parts for local use. We think this might be a more general problem, and would like to propose the Centre should restrict itself by adding the following to the MARS User Guide, Appendix F:

Proposal, add to MARS User Guide, Appendix F (Figure 7)

ECMWF Member States local GRIB use definition

Octet Number	Contents
41	Local definition number: 192 to 255 NOTE: These local definition numbers will never be used to define ECMWF local extensions to Section 1, GRIB code. They are free to be defined and used by Member States.
42 - nn	National information.

This would allow Member States to define local information which do not conflict with ECMWF usage.

4. Plans

The plans consist in a more extensive usage of ECMWF computer equipment for developing and testing programs by carrying out (for example):

- tests to parallelizing 1D-, 2D-, 3D-models and special numerical algorithms (local advection method, Helmholtz-solver)
- tests in the framework of the development of the GM_E and LM by using C90 and T3D.

This will require:

- Installation of DWD-libraries on ECMWF computers where testing will take place
- libraries should be available as permanent files.

Heutige Konfiguration

Fig. 1

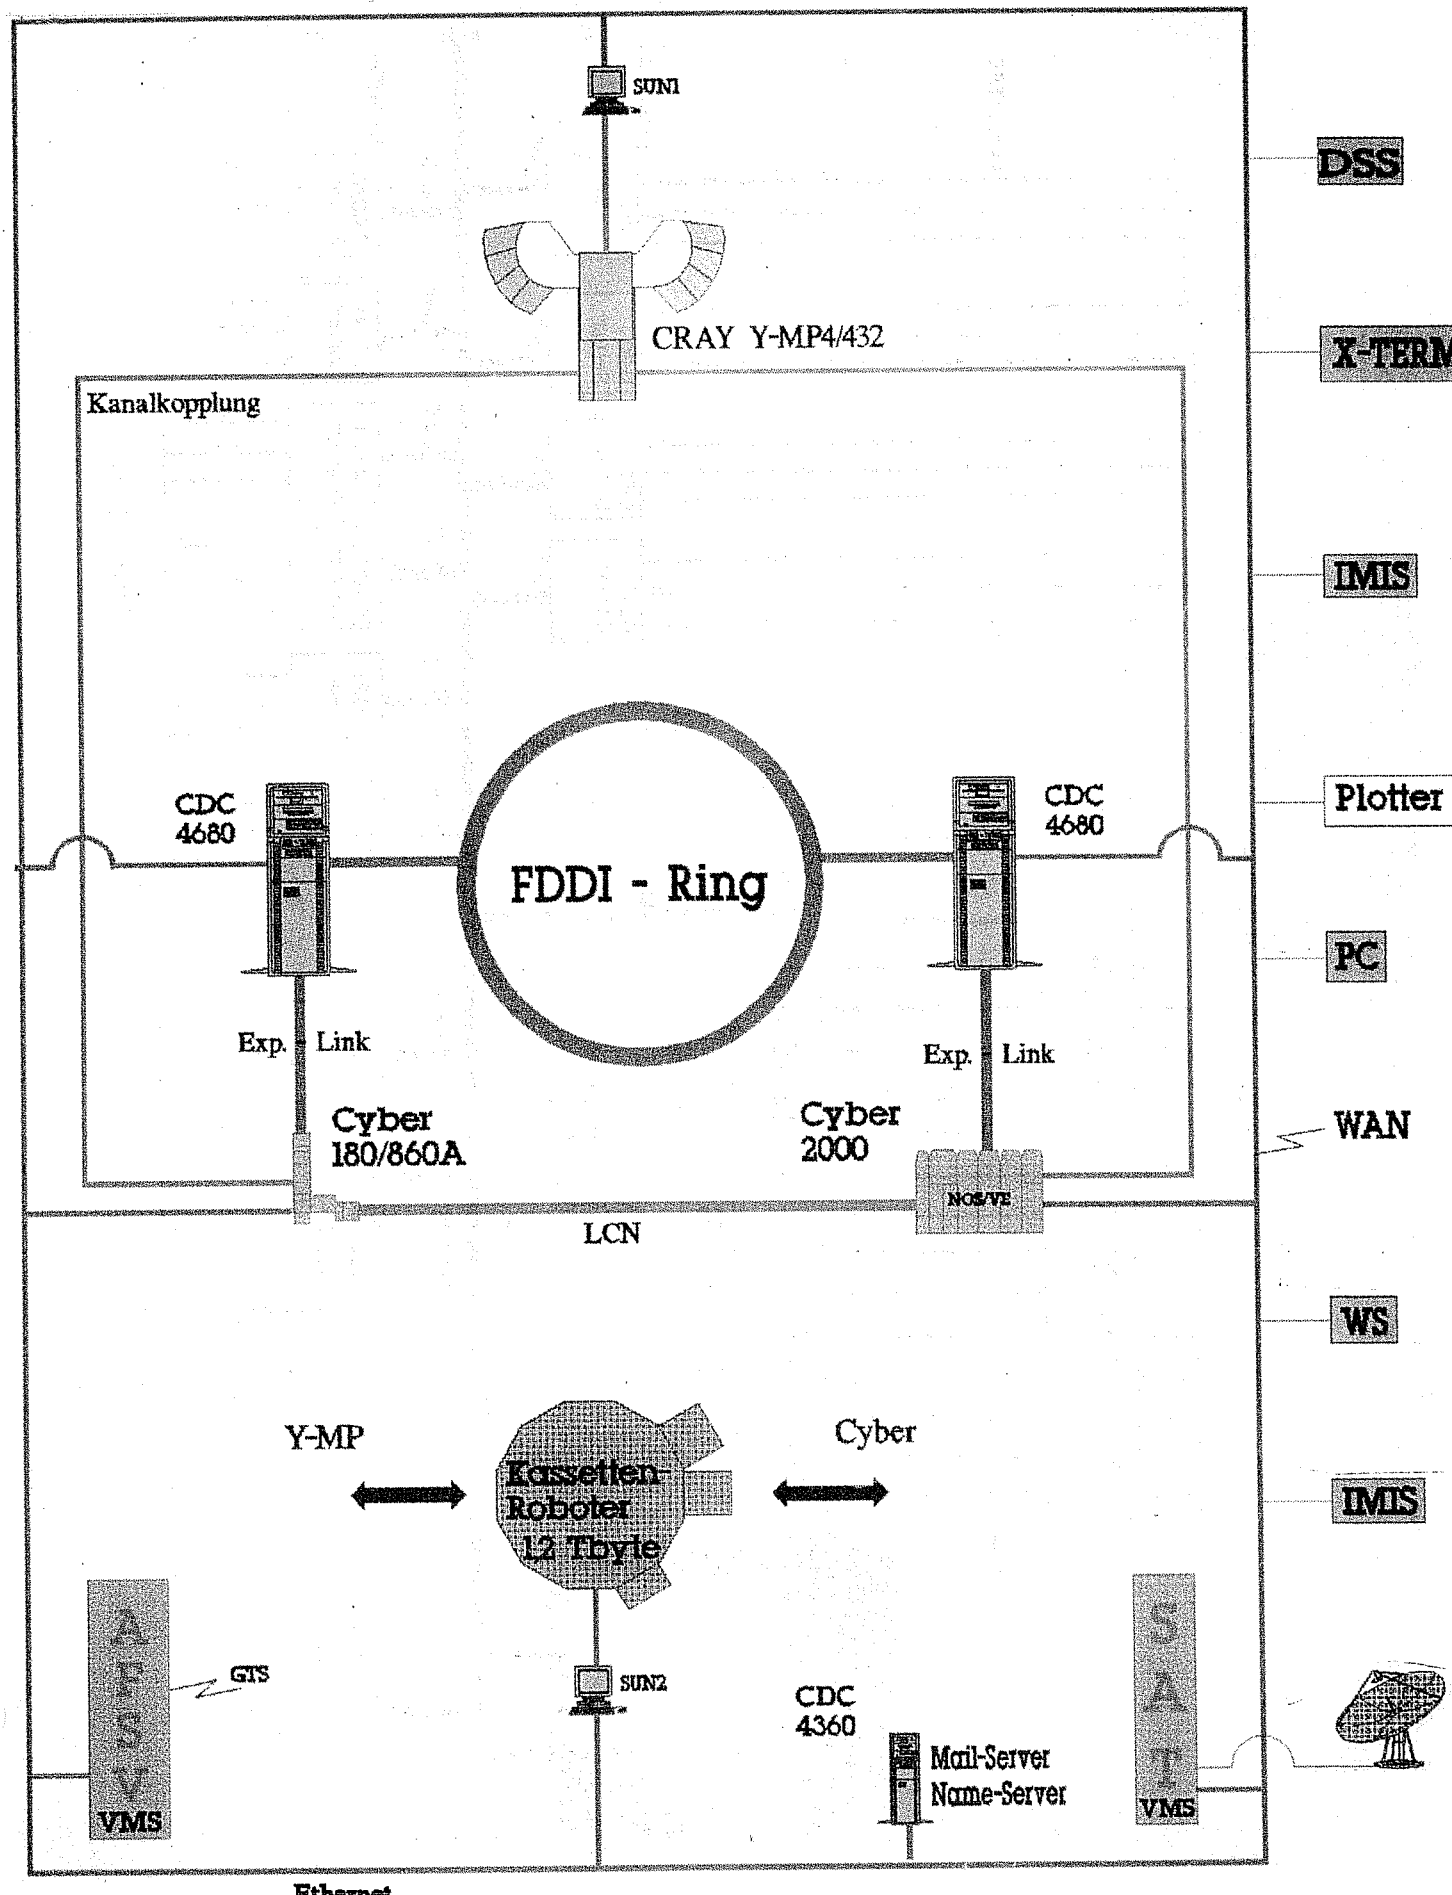

Fig. 2 AFSV des Deutschen Wetterdienstes

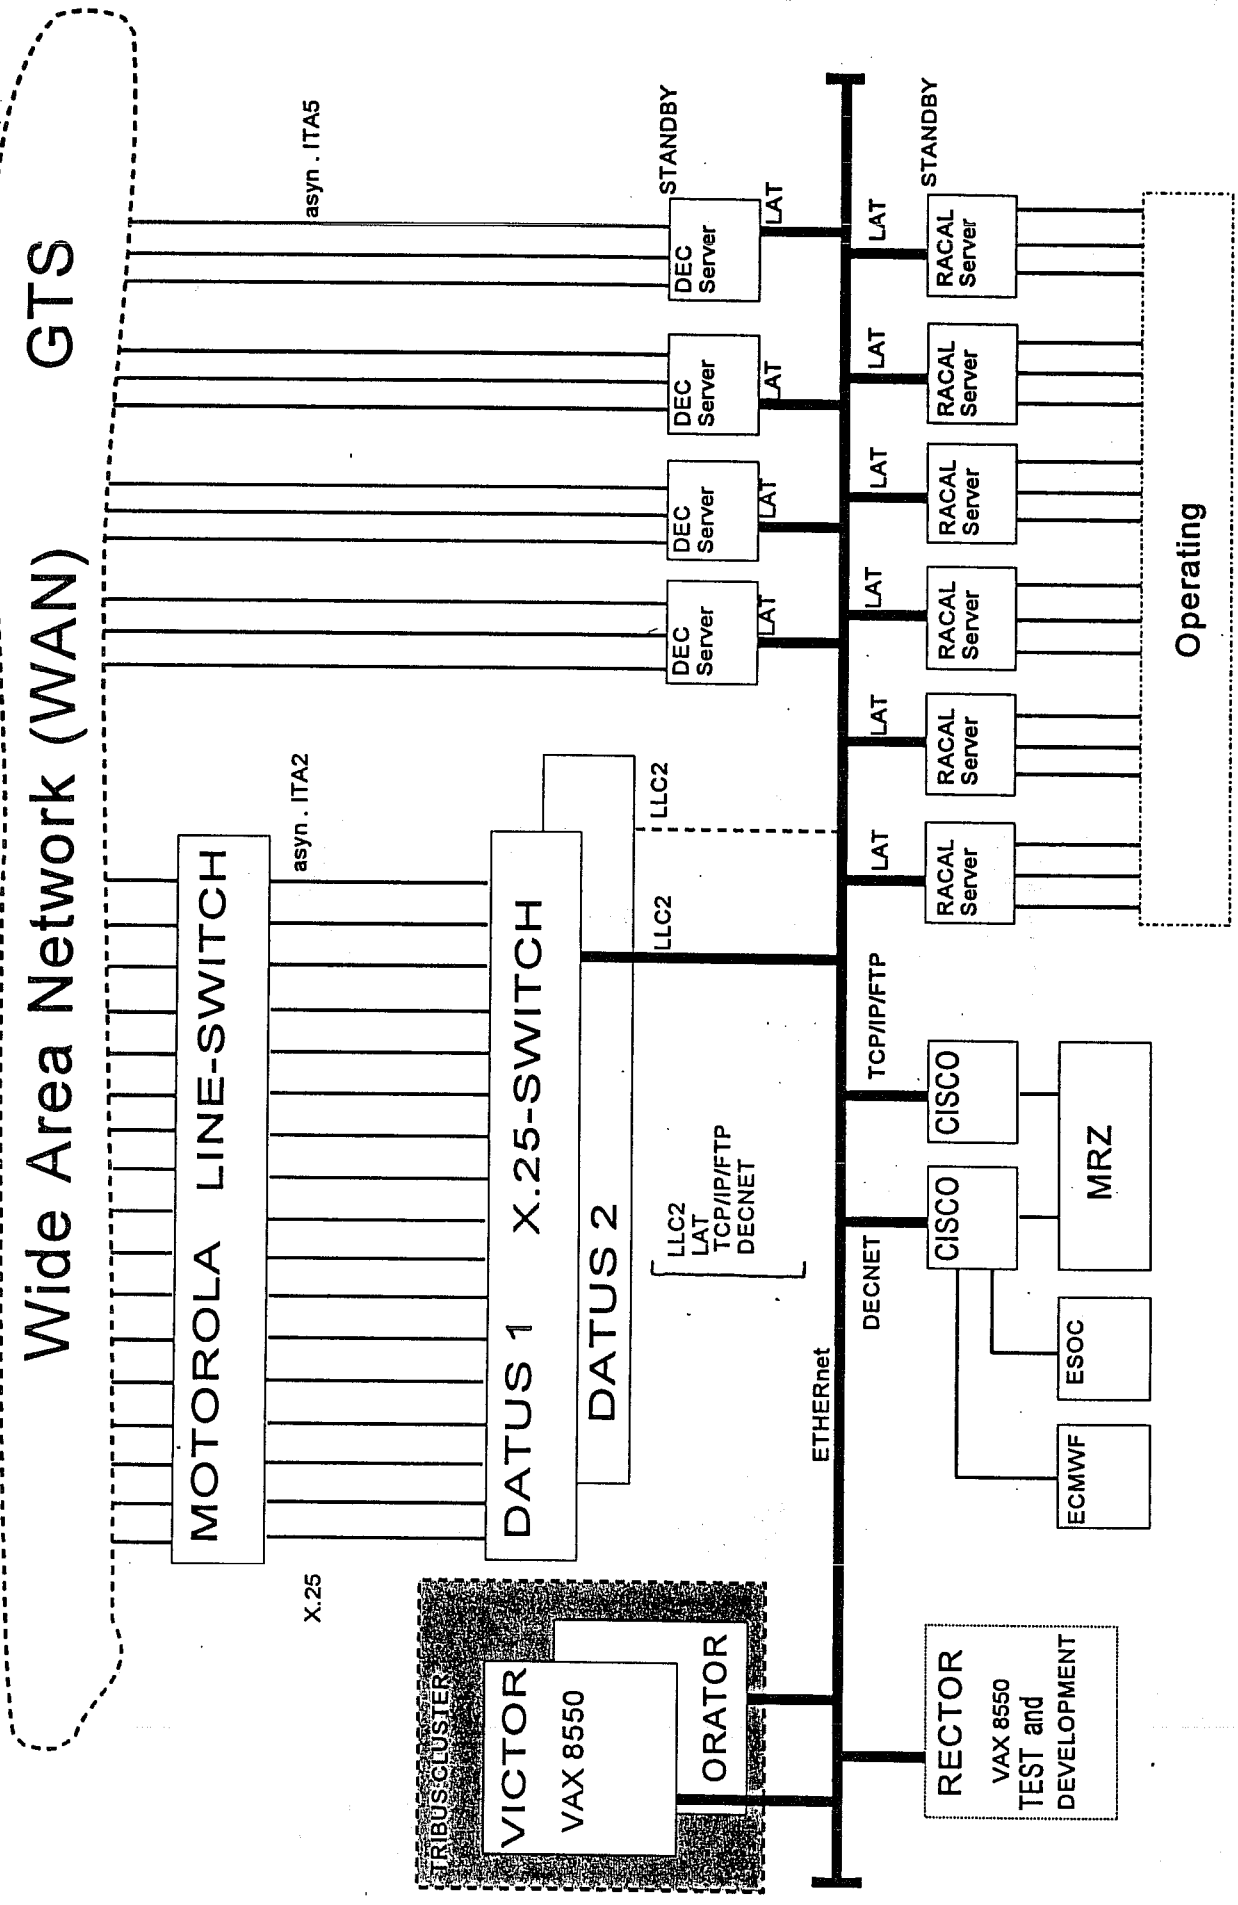

Fig. 3

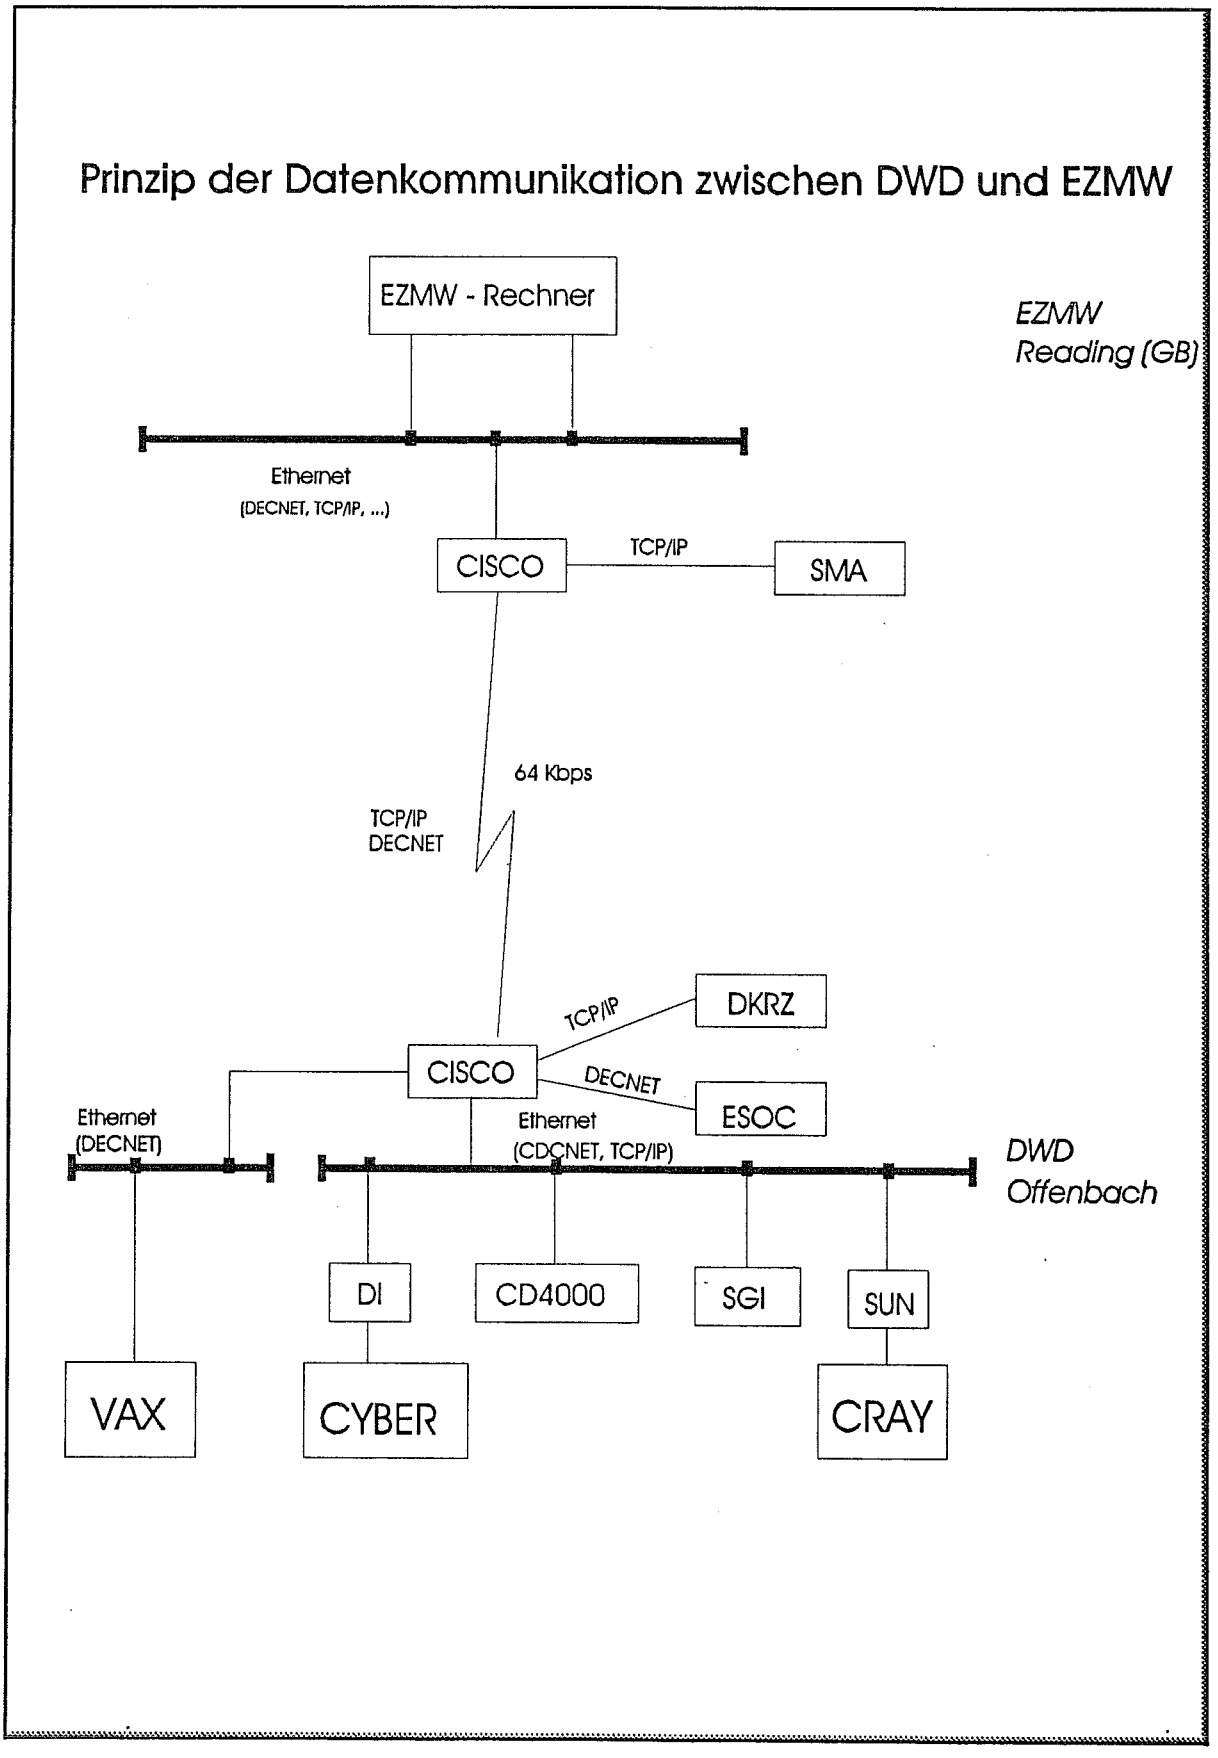
Prinzip der Datenkommunikation zwischen DWD und EZMW
EZMW - Rechner
EZMW
Reading (GB)
Ethernet
(DECNET, TCP/IP, ...)
CISCO
TCP/IP
SMA
64 Kbps
TCP/IP
DECNET
TCP/IP
DKRZ
CISCO
DECNET
ESOC
Ethernet
(DECNET)
Ethernet
(CDCNET, TCP/IP)
DWD
Offenbach
DI
CD4000
SGI
SUN
VAX
CYBER
CRAY

Fig. 4

Data Communication Between DWD and EZMW

I. Operational data transmission

computer sytems:	EZMW - Vax-Cluster - AFSV (Vax-Cluster)
protocols:	DECnet (via X.25)
data quantity:	75 Mbyte/day

II. Message Routing

A. EZMW - DWD - DKRZ

protocols:	TCP/IP (via X.25)
data quantity:	2 MByte/day
remarks:	connection between DWD and DKRZ uses WiN

B. EZMW - DWD - ESOC

protocols:	DECnet
data quantity:	?
remarks:	connection between DWD and ESOC uses 9,6 kbyte/s DDC

C. DWD - EZMW - SMA

protokolle:	TCP/IP (via X.25)
data quantity:	35 Mbyte/day
transmission times:	3.00 / 15.00 (UTC)

Fig. 5

III. R&D - user

protocols:	TCP/IP (via X.25)
data quantity:	approx. 2,5 Mbyte/day (varying)
time of use:	office service period

IV. Additional used services:

A. exchange of electronic mail
B. direct connection of both organizations "domain name services"

Fig. 6

DWD

First proposal:

Creation of a special "Group Code" for a single test dissemination outside of the routine data dissemination

Test data and routine data having the same format

Second proposal, problem probably to be solved within MARS

At present MARS has less facilities than the data dissemination system

An advanced MARS allowing e.g.

- limited area output,
- different grids

could also solve the problem

\wp_text\boe\ezmw\940909.crm, 13.09.1994

Fig. 7

DWD

ECMWF is using GRIB, Section 1, starting with Octet 41 for local use.

This information is already included in ensemble forecast data.

Proposal, add to MARS User Guide, Appendix F:

ECMWF Member States local GRIB use definition

Octet Number	Contents
41	Local definition number: 192 to 255
	NOTE: These local definition numbers will never be used to define ECMWF local extensions to Section 1, GRIB code. They are free to be defined and used by Member States.
42 - nn	National information.

This would allow Member States to define local information which do not conflict with ECMWF usage.

\wp_text\boe\ezmw\940909.crm, 13.09.1994

COMMENTS AND QUESTIONS

J. Hennessy — In general we will discuss your queries in the "MARS and Dissemination" session later in this meeting. However, I believe the problems regarding documentation have already been solved.

Irish Meteorological Service

1. Computer Configuration

The computer configuration at the Irish Meteorological Service is shown in Fig.1.

- The Communications VAX Cluster, comprising two VAX 4000/200's and a microVAX 3100, handles communications and data pre-processing.

- The HIRLAM Numerical Weather Prediction model is currently run in test mode on an SGI Challenge L server, while graphical services are provided by two MIPS servers. A new graphical display system is under development and will run on a series of SGI Indy workstations.

- Radar data is processed on a VAXserver 3300 and displayed on VAXstation 3100 workstations.

- Office services are provided by a Novell Netware PC network.

- The Climatological Database System, based on Ingres RDBMS, comprises a DECsystem 5000/200 server and several Decstation 5000/120 client workstations.

- The Ethernet network supports both DECnet and TCP/IP protocols. The connection to ECMWF is a 64000 baud link via an X25 router.

2. Use of ECMWF Computers

The ECMWF computer system is used for data retrieval from the MARS archive, mainly in support of a number of operational and research projects. A trajectory model is run on the Cray. A significant fraction of our computer allocation is now committed to the ocean wave project, the 00 UTC forecast project and the European Climate Support Network project.

3. Future Plans

The current routine operational dissemination of ECMWF products consists of a mix of data in ECBIT and GRIB codes. We plan to receive all routine dissemination products in GRIB code within the coming year.

We have no specific plans for new projects utilizing the computer resources at the Centre.

Irish Meteorological Service

Computer and LAN Configuration

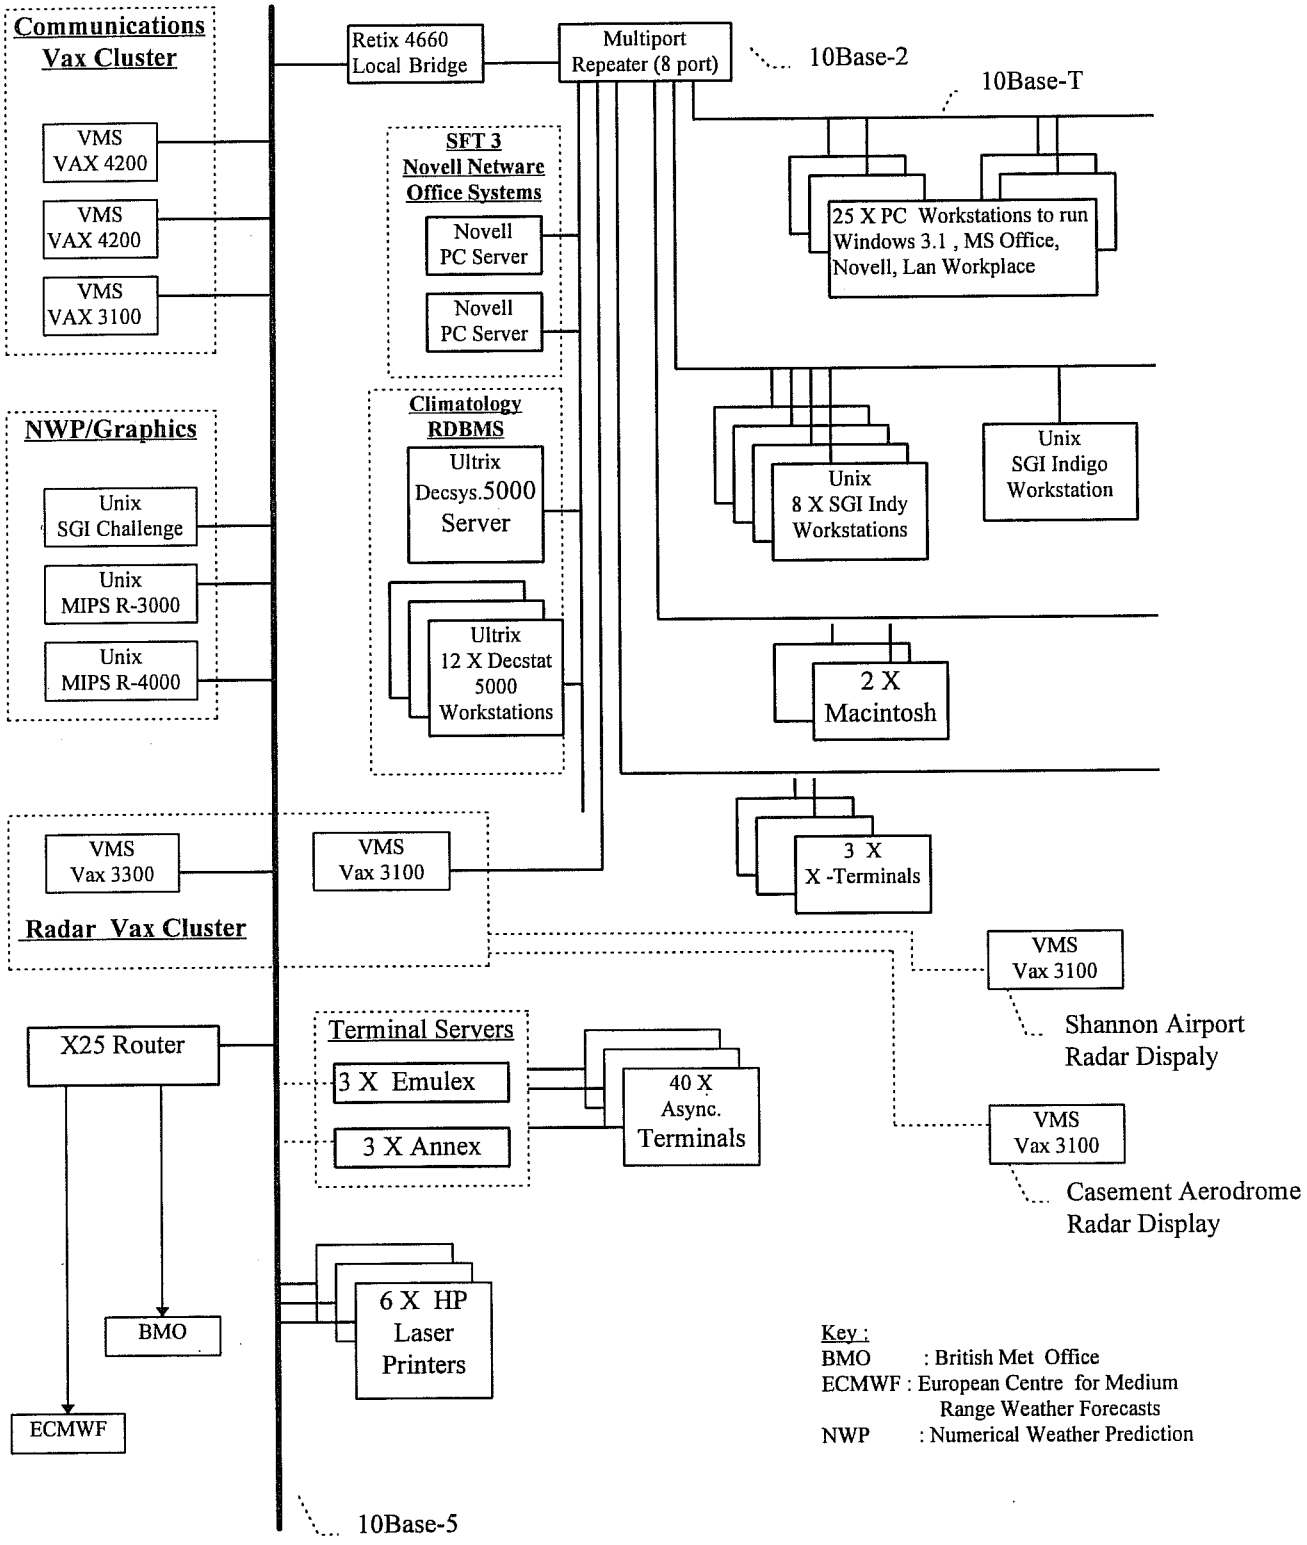

Figure 1.

COMMENTS AND QUESTIONS

B. Barg — Are your climatological people satisfied or not with the Ingres database system?

L. Campbell — Yes, they are satisfied. Just one point I would like to make, it took longer to learn to use it than expected.

ITALIAN METEOROLOGICAL SERVICE

Computing representative: S. Pasquini.

1. Computer equipment and ECMWF connection.

The main components of the central computer system are shown in figures 1 and 2. From a functional point of view it can be divided in the following way:

a. Front-end area.

It handles the telecommunications concerning both the WMO/GTS circuits and the ICAO (MOTNE, AFTN) networks. It is based on two IBM S/1 minicomputers, one operational the other in standby. The system is able of an automatic switchover in case of failure. The Message Switching software has been developed by IBM. Because of the IBM S/1 computing power limitations some links are now handled by two 8250 Vax systems (Bracknell ecc...). A completely new Message Switching system based on updated DEC equipment running a Unix-like operating system and application software provided by Global Weather Dynamics will be installed early next year in the new building of the Operational Branch in Pratica di Mare.

b. Host area.

The IBM 3090 12 E (O.S.: MVS/XA) performs the operational work concerning the analysis and forecast model, post processing, graphical applications and meteorological database. The IBM 4381 P13 provides interactive services and it is used for research&development and as back-up.
There is a procurement action (open tender) aimed to replace the mainframe systems with a distributed networked workstations system.

c. Back-end area.

It deals with the access of remote users to the meteorological database (inquiry) and to the interactive graphical services.

d. ECMWF connection.

It is based on various Vax systems and, at the moment, it handles the 64 kbps link based on DECNET/NTS software and Cisco routers. Moreover the vax systems handle two PSS connections with the Italian Public Switched Data Network called ITAPAC (one for ERSA and the other for internal use in the Met Service) and the links with the computer networks of national agencies (CNR, Worldlab, Enea, MAF) to allow the access to the ECMWF computer systems and meteorological archives by remote users.
In the future the Vaxs will be used to support the message switching software developed by DEC/GWDI in order to replace the IBM S/1 computers. The 64 kbps link is also used as back-up of the ERS-1 data transmission to ECMWF.

e. MDD area.

It deals with the dissemination of meteorological data via the Meteosat satellite mainly to African and Middle East countries.

f. Satellite area.

It deals with the reception of data and images from the meteorological satellites Meteosat and TIROS. Moreover the computer system, based on a DEC AXP 4600 system, provides image animation and TOVS retrieval (fig.3). The satellite area is connected via Ethernet with the meteorological radar images processing system that will be operational next year.

2. Projects, experience and plans.

The main usage of ECMWF services has been the retrieval of MARS data associated with the decoding software and MAGICS. ECMWF data in GRIB form are disseminated in real time to users for their operational work (civil emergencies, agricolture, pollution ecc...). Italy has joined the 00 UTC and ocean waves optional projects: relevant data are also disseminated to users in real time. Some user has requested if it is possible to have the 00 UTC data earlier. The number of Italian remote users (via the IMS) will increase as soon as new computer links are established.

The current projects are:

DIAGNOSTIC STUDIES
WAVE MODELLING
MESOSCALE MODELLING
POSTPROCESSING

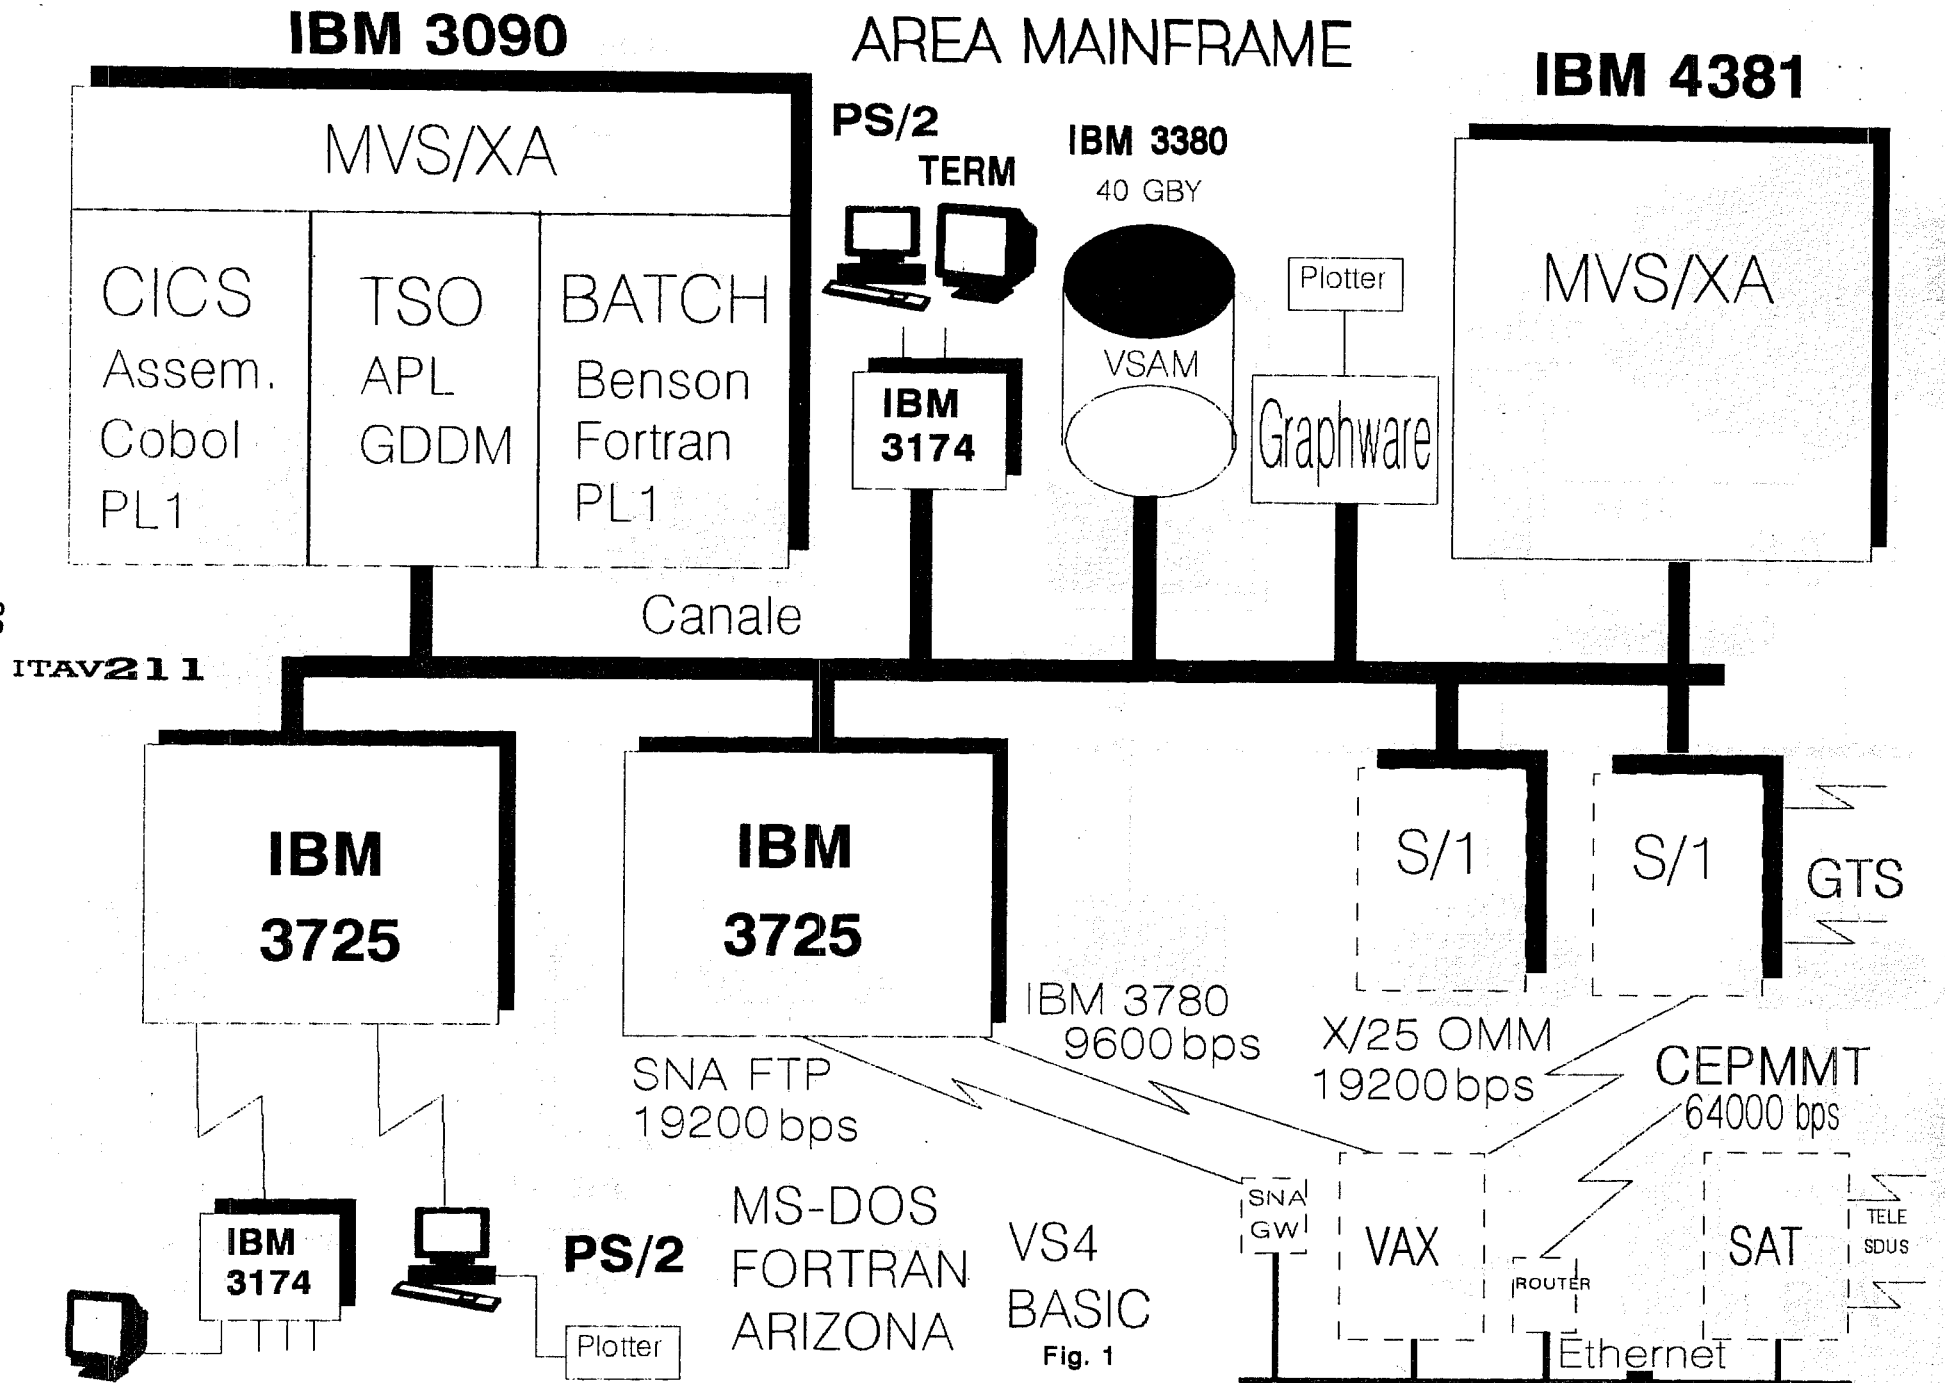
IBM 3090
AREA MAINFRAME
IBM 4381
MVS/XA
CICS
Assem.
Cobol
PL1
TSO
APL
GDDM
BATCH
Benson
Fortran
PL1
PS/2
TERM
IBM 3380
40 GBY
VSAM
Plotter
Graphware
IBM 3174
MVS/XA
Canale
ITAV211
IBM 3725
IBM 3725
S/1
S/1
GTS
IBM 3780
9600 bps
X/25 OMM
19200 bps
CEPMMT
64000 bps
SNA FTP
19200 bps
IBM 3174
PS/2
MS-DOS
FORTRAN
ARIZONA
VS4
BASIC
Plotter
SNA GW
VAX
ROUTER
SAT
TELE SDUS
Ethernet

Fig. 1

MEMBER STATE COMPUTER REPRESENTATIVES MEETING

27-28 September 1994

The Netherlands

1. Local computer equipment.

In figure 1 a diagram of the computer equipment at KNMI is given. The right part of this figure shows the local network at De Bilt. At the left the connections with the regional stations are depicted.

In the last year the local facilities at De Bilt were expanded with a Storage Tek tape robot. This robot is controled by FileServ running on the Convex C240. Forthcoming expansions are the installation of a FDDI backbone network and the installation of a Communication Platform (CoPla) which will handle all incoming and outgoing traffic e.g. VideoTex services, dial-in connections, anonymous ftp and WWW services like Gopher and Mosaic.

In the diagram personal computing facilities are indicated. Currently, in the Research Department alone, about 80 PC's and 30 Workstations are connected to the network.

2. Use of computing facilities at ECMWF.

Currently there are 50 Dutch users registrated on the C90 at ECMWF of which about 30 are active. The tasks performed range from simple MARS-access to the running of complex tracer programmes supporting air pollution research.

For the access to the computing facilities two methods are available: the first is based on telnet for interactive access and ftp for file transfer, the second method is based on Remote Batch facilities. During the last year about 6000 jobs were submitted using the last facility. With regard to telnet and ftp there is no information available.

Of the Dutch account of 130 kUnits about 50 kUnits are allocated to Optional Projects like the Wave model, the 00 UTC run and the reanalysis project. By the end of September already 55 kUnits of the remaining 80 kUnits were used.

In general the Dutch users are satisfied with the facilities offered. However, a clear problem for the interactive users either at the ECSERVER or the C90 is the fact that file-transfers initiated by SENDTM on the C90 can block the interactive connection for substantial periods, leading sometimes to spontaneous logouts etc.

THE NETHERLANDS

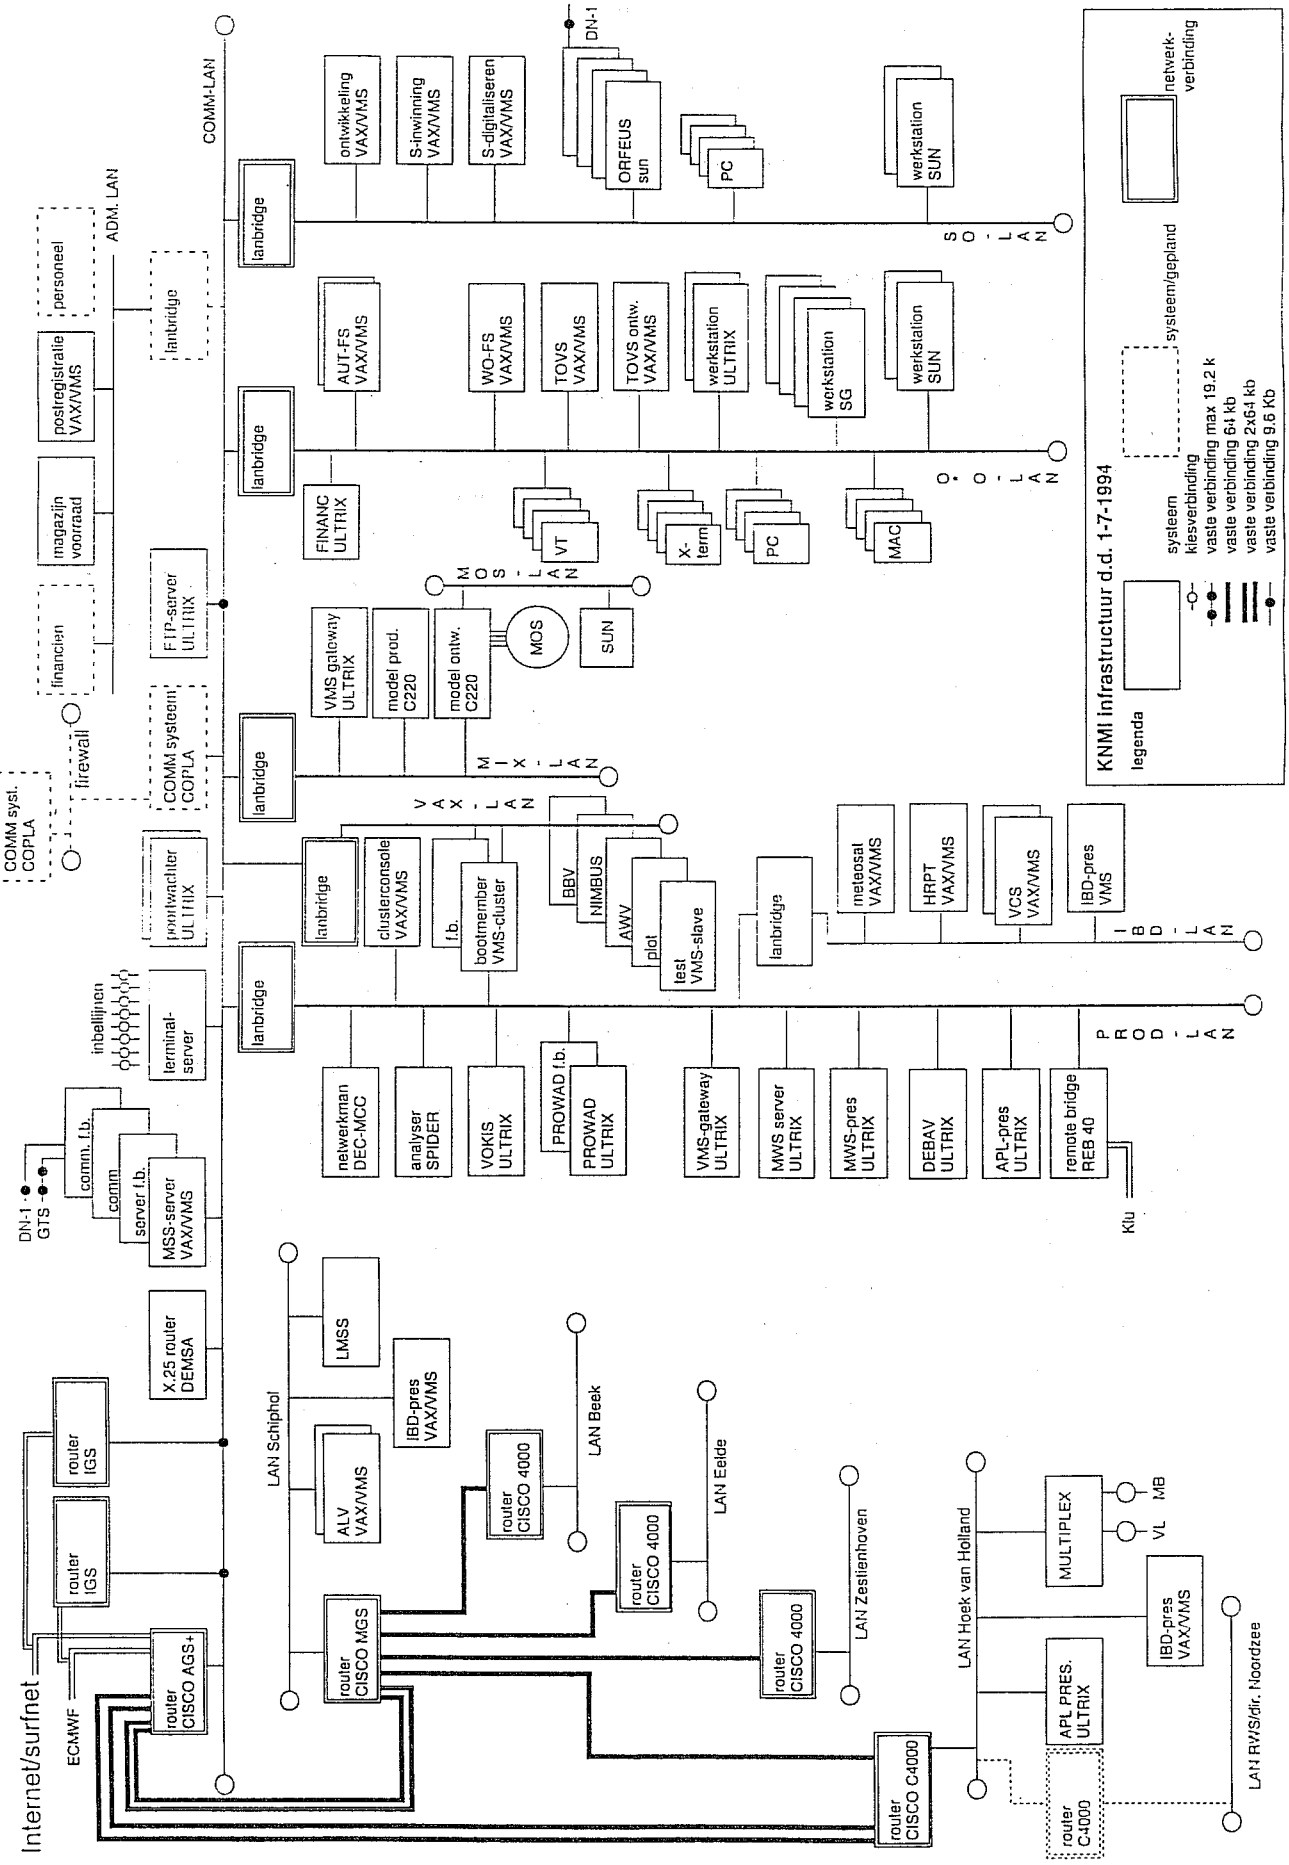

COMMENTS AND QUESTIONS

G-R. Hoffmann — The question regarding interactive response is being covered elsewhere.

Rebecca Rudsar
DNMI, Olso, Norway September 1994

THE COMPUTING ENVIRONMENT AT DNMI

Computer Resources

Computers :

Tandem CLX820
- dual CPU, 2 * 8 Mbyte
- disk system, a mirrored pair of 648 Mbyte each

Silicon Graphics 4D/420
- 2 CPU's
- memory 128 Mbyte
- disk storage capacity 20 Gbyte

Silicon Graphics 4D/340
- 4 CPU's
- memory 128 Mbyte
- disk storage capacity 14 Gbyte

Silicon Graphics POWER Challenge XL
- 4 (R4400) CPU's. These will be replaced by 2 TFP (R8000) CPU's in October 1994.
- memory 512 Mbyte
- disk storage capacity 6 Gbyte

VAX4000-200
- memory 32 Mbyte
- disk storage capacity 1.35 Gbyte

VAX3300
- memory 20 Mbyte
- disk storage capacity 570 Mbyte

VAX4000-200
- memory 32 MBYTE
- disk storage capacity 2.85 Gbyte

Norsk Data (NORD) * 7
- total online disk storage capacity approx. 4.5 Gbyte (interchangeable disks)

Workstations * 60
- Silicon Graphics Indigos and Indys, Sun SPARC stations, IBM RS6000, VAX3100, IVAS(I2S)

Terminals / PCs
- approx. 250

Networks :

Cosmos
- connecting NORD computers

Ethernet
- connecting all computers and workstations and several PCs.

Novell
- connecting some PCs. This is being phased out.

Graphical Devices:

- Hewlett Packard pen plotters

- A3/A4 laser printers

The Tandem ext10 has a dual CPU and mirrored disk system and manages the communication for data acquisition and routing of observations.
Two IBM RS6000 computers, with the AIX operating system, are used to extend the communication capabilities of the CLX.
Observation data is transferred to a NORD computer and to a Unix computer for decoding.
The decoding programs for SYNOP, TEMP and PILOT obtained from ECMWF are run on a Unix platform. These are running as test versions as the code necessary for decoding National Groups has not yet been written.

VAX4000-200 is used for communication with ECMWF. The telecommunication link was upgraded to 64 Kbps in May 1994.
This computer is used for running the Nordic Radar system, NORDRAD, which is a communication system for the distribution and collection of radar data between Nordic countries. It is also used for processing Meteosat DCP data.

VAX3300 is connected to the radar unit and is used for communication to the previously mentioned VAX4000-200.

VAX4000-200 is used for the MISAT satellite system which processes digital image data from Meteosat and NOAA.

The NORD computers are the oldest part of the system and are gradually being replaced.

Oracle database software is installed on the S.G. 4D/420. The design and implementation of the historical database is progressing.

The S.G. 4D/340 is connected via the University network to a CRAY YMP/464 situated at Trondheim. The CRAY YMP has 4 processers, 64 MWord memory and 64 Gbyte of mass storage. The telecommunication line (Supernet) from the University of Oslo to CRAY has a theoretical speed of 34 Mbps. DNMI has a 10 Mbps connection to Supernet. The TCP/IP protocol is used for data transfer.

The Norwegian Limited Area Model for Numerical Weather Prediction has been run on the CRAY since May 1990. The Maritime Prediction models for ocean waves and storm surge are also run on the Cray-YMP as part of the operational suite.

All pre- and post- processing is at present performed on the S.G. 4D/340. The Supervisor Monitor S cheduler (SMS) and X Command and Display program (XCDP), developed at ECMWF, are used to control and monitor the operational suite.

Routers have been installed at the forecasting centres at Bergen and Tromsø and at the airports at Oslo-Fornebu and Bergen-Flesland. The tele communication links are 64 Kbps. ISDN links are used as backup. Two Indy workstations have been installed at each of the centres at Bergen and Tromsø. For backup purposes the workstations have identical file systems for operational products. They function as file servers as well as being used for some graphics for the meteorologists. They also function as file servers for the PCs.
Fornebu and Flesland each have an RS6000 as fileserver. An RS6000 will be installed at Bodø in the near future.
Distribution of products uses the TCP/IP protocol.

The latest aquisition, the S.G. POWER Challenge computer, has until now been used for testing . The pre- and post- processing will be moved to this computer in October and the S.G. 4D/340 will then function as an NFS file server. If the CRAY is unavailable the POWER Challenge computer is used to run a simpler version of the Limited Area Model.

ECMWF Products

Disseminated data from the operational forecast model and the global wave model are received from ECMWF. This data amounts to approx.

20 Mbyte per day. The data is transferred from the VAX computer to a Unix file server.

Dissemination data received from ECMWF is converted from GRIB format and placed in our present fields database.
The data is then accessible by the graphics package which has been developed at DNMI.

The data is also used
1) for general forecasting by the forecasting department.

2) as boundary values for the Norwegian limited area models.

3) as backup for the Norwegian limited area models.

4) as input to the maritime and air pollution models.

5) as input to a ship routing program for the Pacific.

6) the Norwegian Institute for Air Research still receives ECMWF data on a regular basis. The data is utilized in the European Arctic Stratospheric Ozone Experiment.

Data retrieval from MARS is used for research projects.

Planned Projects at ECMWF

Norway is a member of the HIRLAM project. Programs and data are transferred to DNMI and experiments are run on the Cray at Trondheim The results of these experiments are exchanged via the official HIRLAM

libraries at ECMWF.
A half time research position at DNMI will be used for system mantainance of the HIRLAM reference system at ECMWF. The intention is to build up a database which contains forecasts and observations for a time period of four weeks. This database will be used by the HIRLAM staff for verification of test experiments against the present version of the HIRLAM system.
DNMI is involved, together with Meteo-France, in the development of a non-hydrostatic fine resolution model. Testing of the model will proceed during the next six months.

The project "MPP in Numerical Weather Prediction" was initiated last year. So far a parallel version of the dynamical / physical part of the semi-implicit HIRLAM2 system has been implemented. It is possible to switch between versions for PVM, Intel NX or CRAY T3D SHMEM.
Further work includes :
- investigating the input/output capabilities of the T3D.
- optimizing the iterative Helmholtz solver for the T3D.
Testing on the CRAY T3D at ECMWF will begin as soon as it becomes available.

DATANET - DNMI

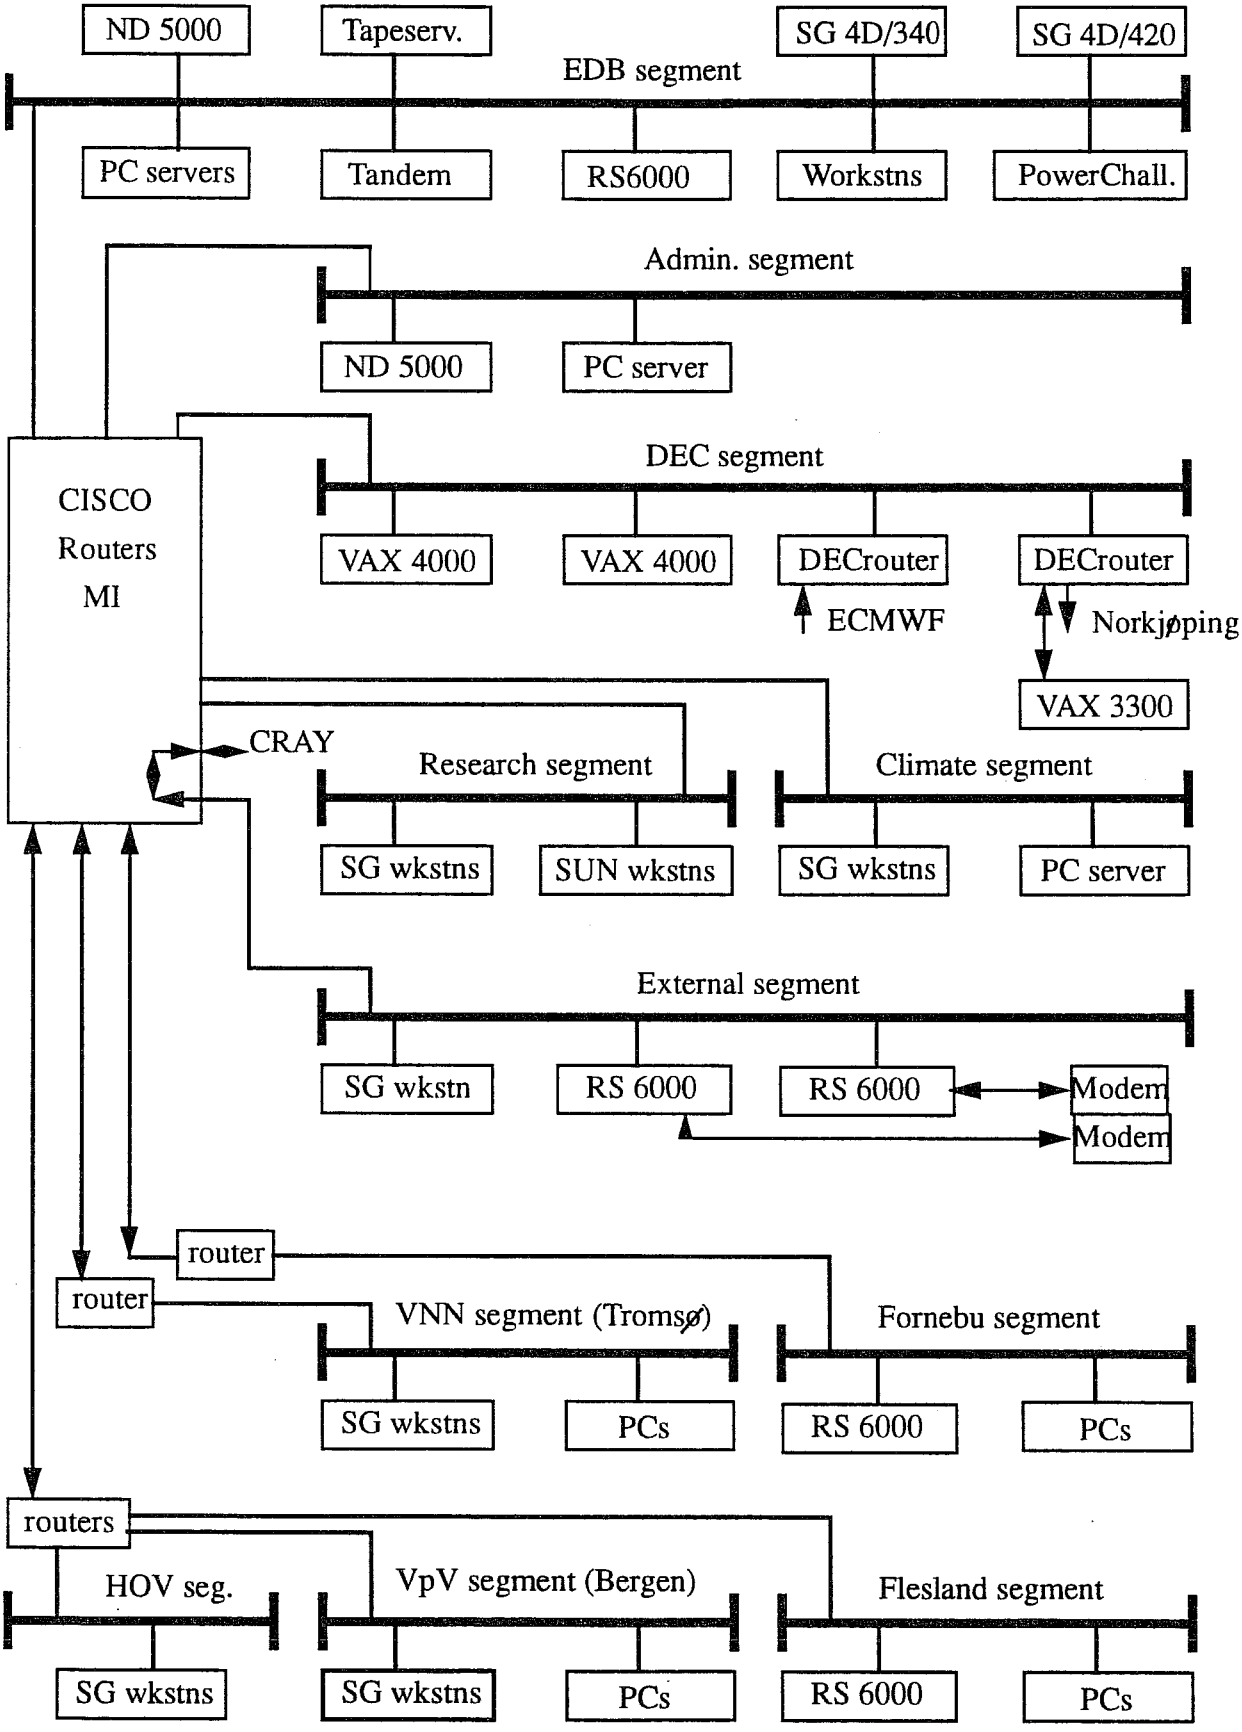

Sept.1994

COMPUTER NETWORK AT PORTUGUESE METEO INSTITUTE (IM)

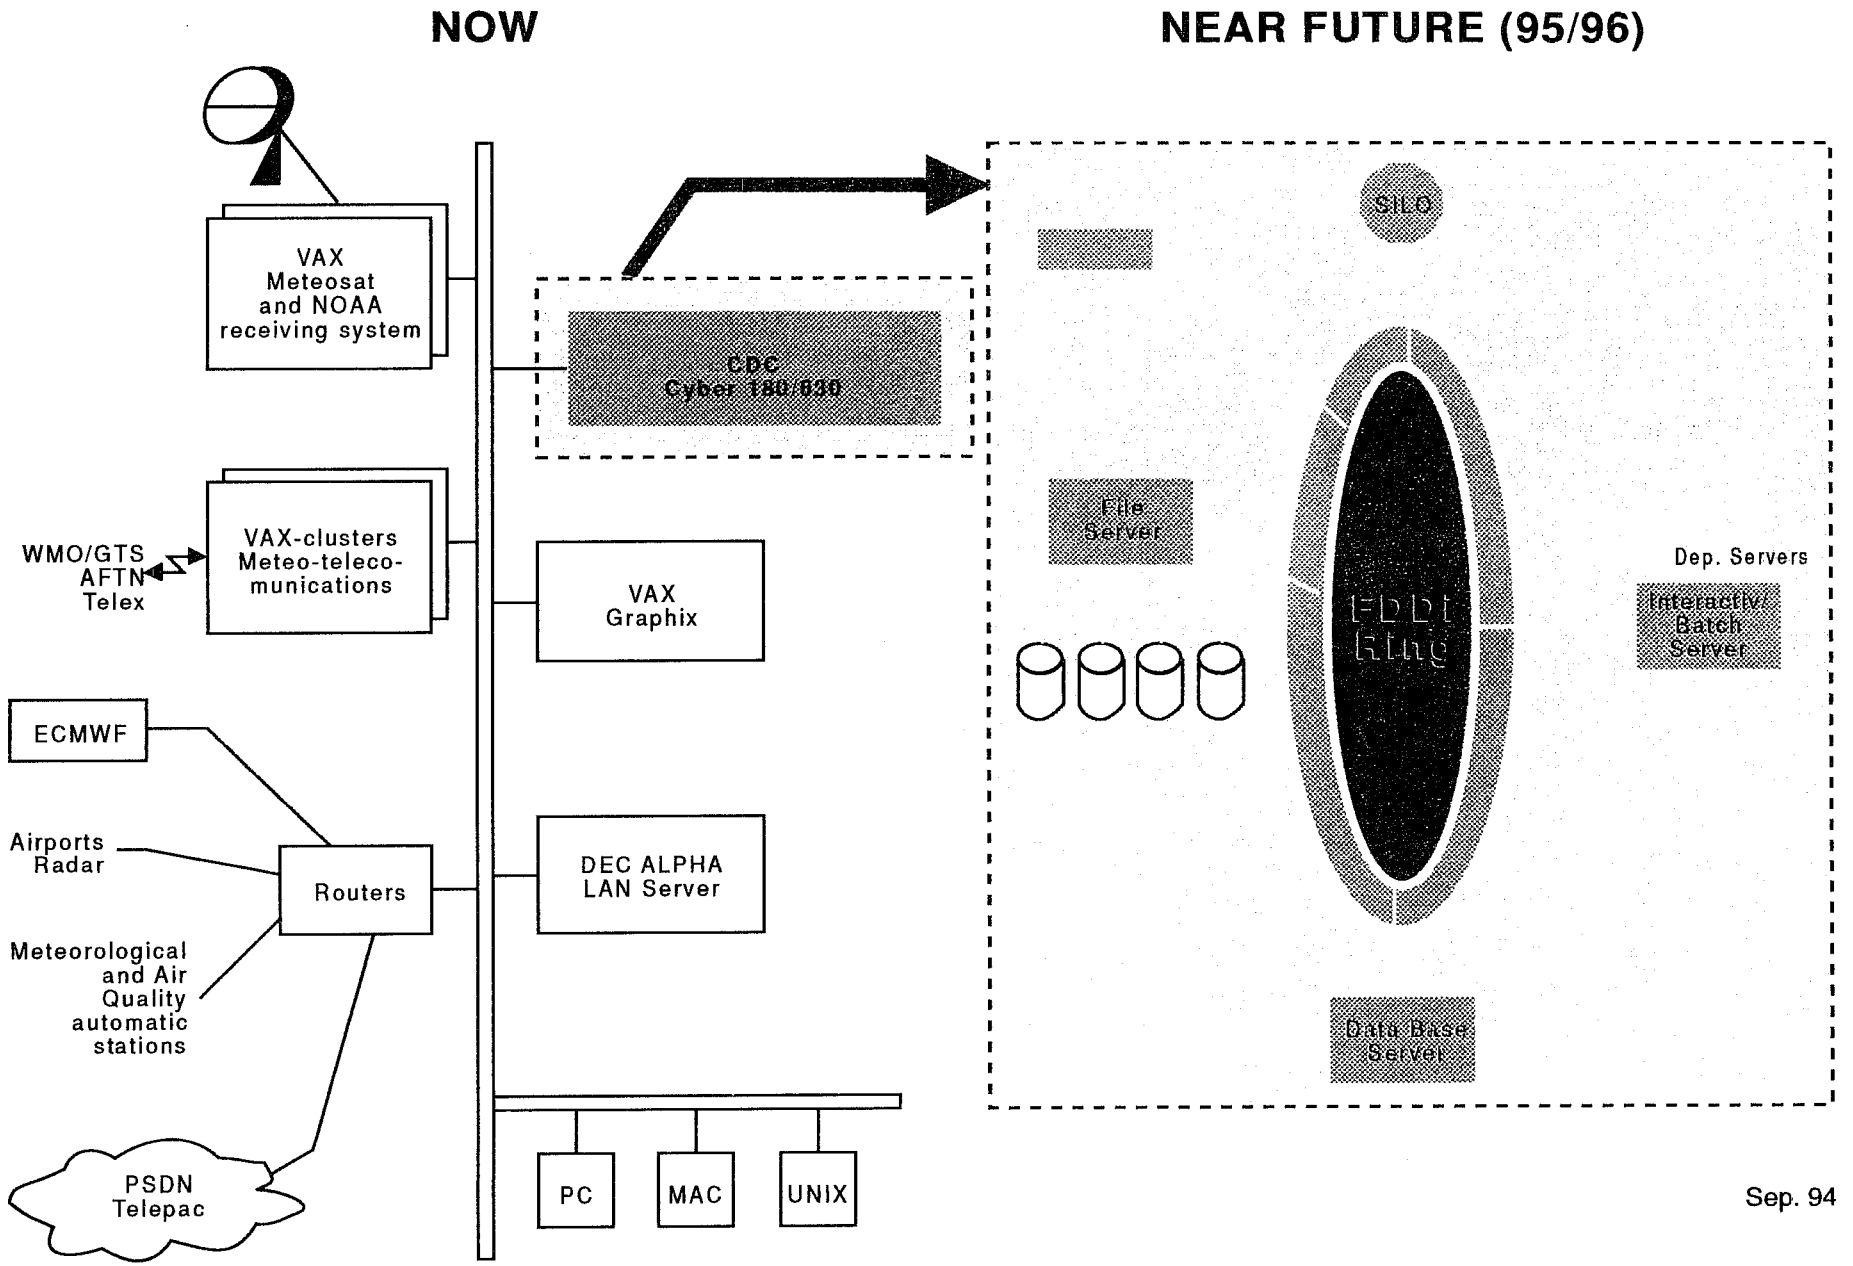

Sep. 94

Main Meteorological Applications at IM

- ***Telecomunications***

  Messages Switching on a VAX cluster, supplied by EDS (UK), handles WMO/GTS traffic, AFTN, telex and PSDN.

- ***Satellite subsystem***

  METEOSAT and NOAA receiving streams on VAX systems and HP ws, supplied by VCS (Germany).

- ***Graphics***

  ECMWF MAGICS V4.2 on VAX/VMS with PS and HPGL printers.

Main Meteorological Applications at IM

- Pre and postprocessing software

  ECMWF decoding software with ISAM file structures on VAX/VMS, interfaced with our Bulletin Data Base;
  ECMWF BUFR and GRIB decoding and encoding package.

- Climatological Archives

  Our national meteorological information is, presently, organized using Cyber File manager structures; due to the relatively small number of stations (30 synoptic, 100 climatological and 750 udometric stations) it is possible to have "on-line" all our information since 1941.

Main Meteorological Applications at IM

- Numerical Weather Prediction

  Presently we do not run any numeric atmosferic models, relying on ECMWF and WMO Regional Centres (Bracknell and Toulose) disseminations. Daily we run, for inshell waters, a local wave model.

IM - ECMWF link

- We have recently moved to a 64 K link based on DECnet with DECWanrouter; no problems were experienced when we switched to this type of link.
- Our main usage of ECMWF CRAY is to access MARS archives, and the users are now quite familiar with the necessary retrieval procedures.
- We foresee, during 1995, to share this link with TCP/IP on UNIX systems.
- On training, we are looking forward for the organization, by ECMWF, of courses in the usage of MPP systems.

INSTITUTO NACIONAL DE METEOROLOGÍA, Spain

Computing Representative: E. Monreal

Member States Computing Representatives' Meeting, 27-28th September, 1994

1. Computer equipment and connection to ECMWF

The main computer equipment and networks installed at INM are depicted in figure 1. There have been several major changes to the computing environment at our Institute over the past two years. The main computer system now consists of a 4 processor CRAY Y-MP C94A, 1 Gbyte memory, 1 Gbyte SSD and 42 Gbytes disk storage, connected to the FUJITSU M-382 through a channel to channel link.

The INM Limited Area Model runs now on the CRAY C94A although the pre- and post-processing are still being done on the FUJITSU M-382. A new Numerical Weather Prediction system, the HIRLAM II model, has been succesfully implemented on the CRAY C94A. Verification of this new model is now being carried out and it is hoped to be in production before the end of 1994. The CRAY C94A is also used for climate studies in cooperation with Climate Research Centers.

The FUJITSU M-382 serves as a destination for all data types concentrated by the different systems and supports the major Data Bases. The major aplication still relying on the FUJITSU is the MCIDAS system, which deals with satellite images (SAIDAS) and serves as operational workstation for forecasters.

A large increase together with a wider use of TCP/IP networks and UNIX based systems, has led to the most outstanding change to the computing environment at INM. Since 1993 almost all the systems are connected to a Local Area Network that consists of several Ethernet networks linked by routers and a small FDDI ring which connects the routers and the CRAY C94A. At the Regional Centers there are also Ethernet networks linked through 9600 bps dedicated lines. Part of the MCIDAS terminals have been yet replaced by PS2 workstations connected to Ethernet networks using TCP/IP protocol.

A new message switching hardware based on a dual DECsystem 5900 runnig ULTRIX operating system (UNIX from DEC) will replace next October the old SINAT system which deals with GTS communication and supports AFTN, MOTNE, etc. It is connected to the Local Area Network using TCP/IP and has enhanced telecommunication features such as facsimil transmission.

Several UNIX based workstations (SUN IPC, IPX and SparcClassic) and Servers (SUN 1000 and SUN 10) were installed last year. They have been our first experience in a distributed computing environment. MAGICS is the main software running on these stations.

INM connection to ECMWF has been involved in another major change. Since May 1994 there is a new 64 kbps TCP/IP line between ECMWF and INM that became operational at the end of July when the old 9600 bps line was over. This line links our Local Area Network to ECMWF Local Area Networ through cisco routers and facilitates interactive (telnet) access as well as file transfer (ftp) to the SGI server "munin", the CRAY Y-MP C90 and the VAX Cluster. Access to this line has been restricted at INM to the UNIX server SUN 1000 due to security reasons.

Figure 2 shows a logical diagram of INM connection to ECMWF. ECMWF dissemination data is received in GRIB code on the SUN 1000 server via ftp from the VAX Cluster. This data is then transferred to the CRAY C94A to be used as boundary conditions for the HIRLAM model and to the FUJITSU M-382 where is decoded, archived and then plotted and ingested into MCIDAS. Our users are able to access ECMWF computing resources only from an interactive session on the SUN 1000 server. Remote job submision to ECMWF CRAY C90 is done via ftp to VAX Cluster.

2. Projetcs run at ECMWF and experience using ECMWF computers

At present, our Institute is not running any project at ECMWF. Most of the ECMWF computers use is for MARS data retrieval and decoding. MARS queries are mainly model outputs not being received via automatic dissemination. A small use of the ECMWF computing resources is due to the use of the HIRLAM Reference System.

Because of the new 64 kbps line better performance it is likely to be an increase of the MARS data being requested as well as an increase of the ECMWF computing resources. At present, this increase cannot be evaluated due to the short time this line has been operational.

We are very satisfied with the new ECMWF computing services and resources we have now access to, in particular the local online documentation has been found very useful.

3. Future plans

Our plans for the next two years envisage the replacement of the FUJITSU M-382, the adquisition of a file server and a mass storage system. Migration of MCIDAS system to UNIX workstations and servers in a distributed environment will be carried out during the next couple of years.

There is not any special project involving the use of ECMWF computers for this period so it is hoped to use ECMWF computing resources in the same manner as it has been described above.

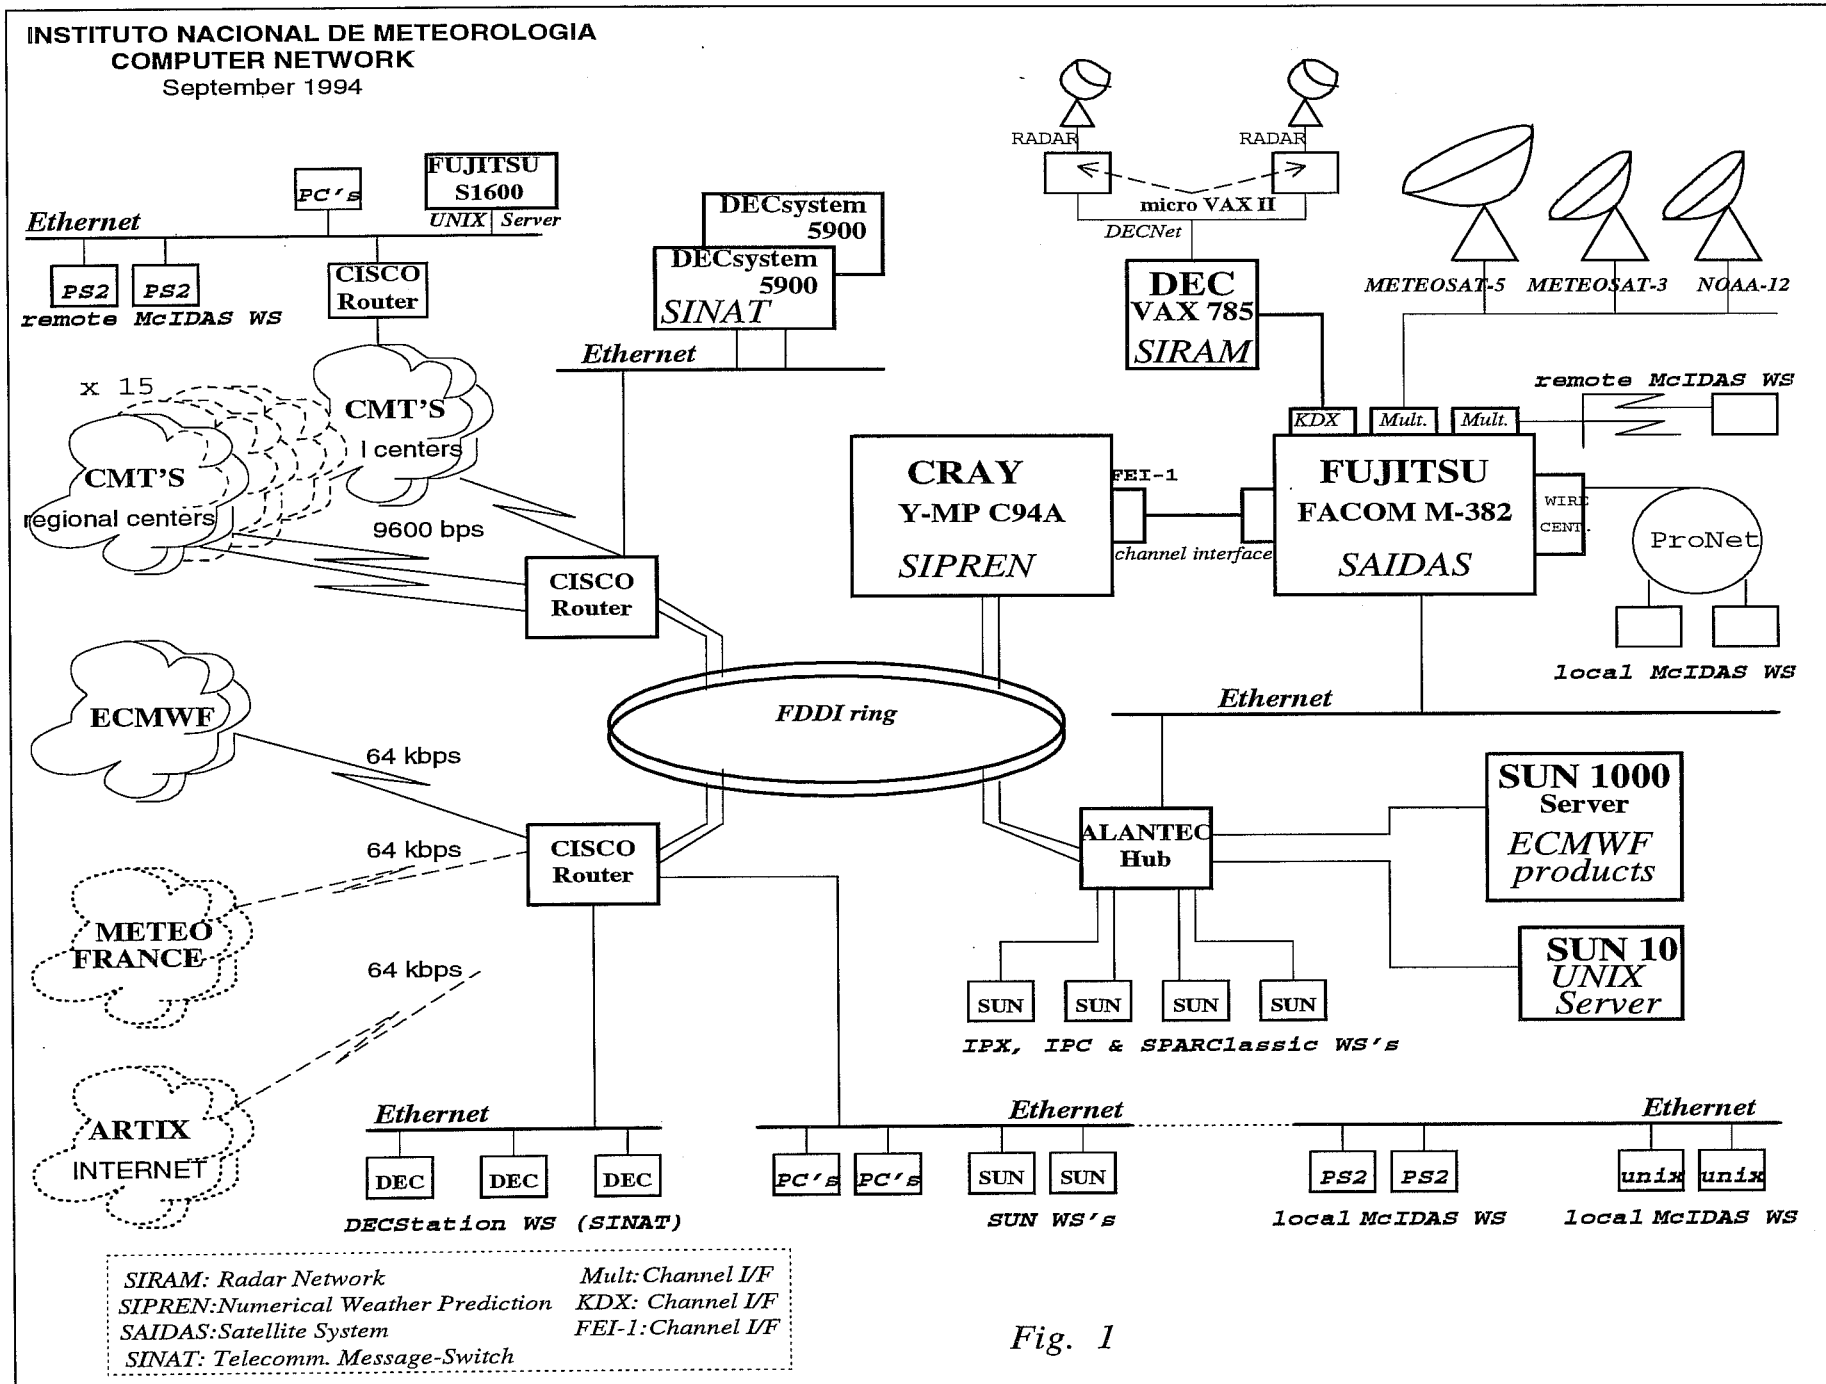
INSTITUTO NACIONAL DE METEOROLOGIA
COMPUTER NETWORK
September 1994
Ethernet
PC's
FUJITSU S1600
UNIX Server
PS2
PS2
remote McIDAS WS
CISCO Router
x 15
CMT'S
centers
CMT'S
regional centers
9600 bps
CISCO Router
DECsystem 5900
DECsystem 5900
SINAT
Ethernet
RADAR
RADAR
micro VAX II
DECNet
DEC VAX 785
SIRAM
METEOSAT-5
METEOSAT-3
NOAA-12
remote McIDAS WS
KDX
Mult.
Mult.
CRAY
Y-MP C94A
SIPREN
FEI-1
channel interface
FUJITSU
FACOM M-382
SAIDAS
WIRE CENT.
ProNet
local McIDAS WS
ECMWF
64 kbps
FDDI ring
Ethernet
SUN 1000
Server
ECMWF
products
64 kbps
CISCO Router
ALANTEC
Hub
METEO
FRANCE
64 kbps
SUN 10
UNIX
Server
SUN
SUN
SUN
SUN
IPX, IPC & SPARClassic WS's
ARTIX
INTERNET
Ethernet
Ethernet
Ethernet
DEC
DEC
DEC
DECStation WS (SINAT)
PC's
PC's
SUN
SUN
SUN WS's
PS2
PS2
local McIDAS WS
unix
unix
local McIDAS WS
SIRAM: Radar Network
SIPREN:Numerical Weather Prediction
SAIDAS:Satellite System
SINAT: Telecomm. Message-Switch
Mult:Channel I/F
KDX: Channel I/F
FEI-1:Channel I/F

Fig. 1

INSTITUTO NACIONAL DE METEOROLOGIA
INM ECMWF LINK
September 1994

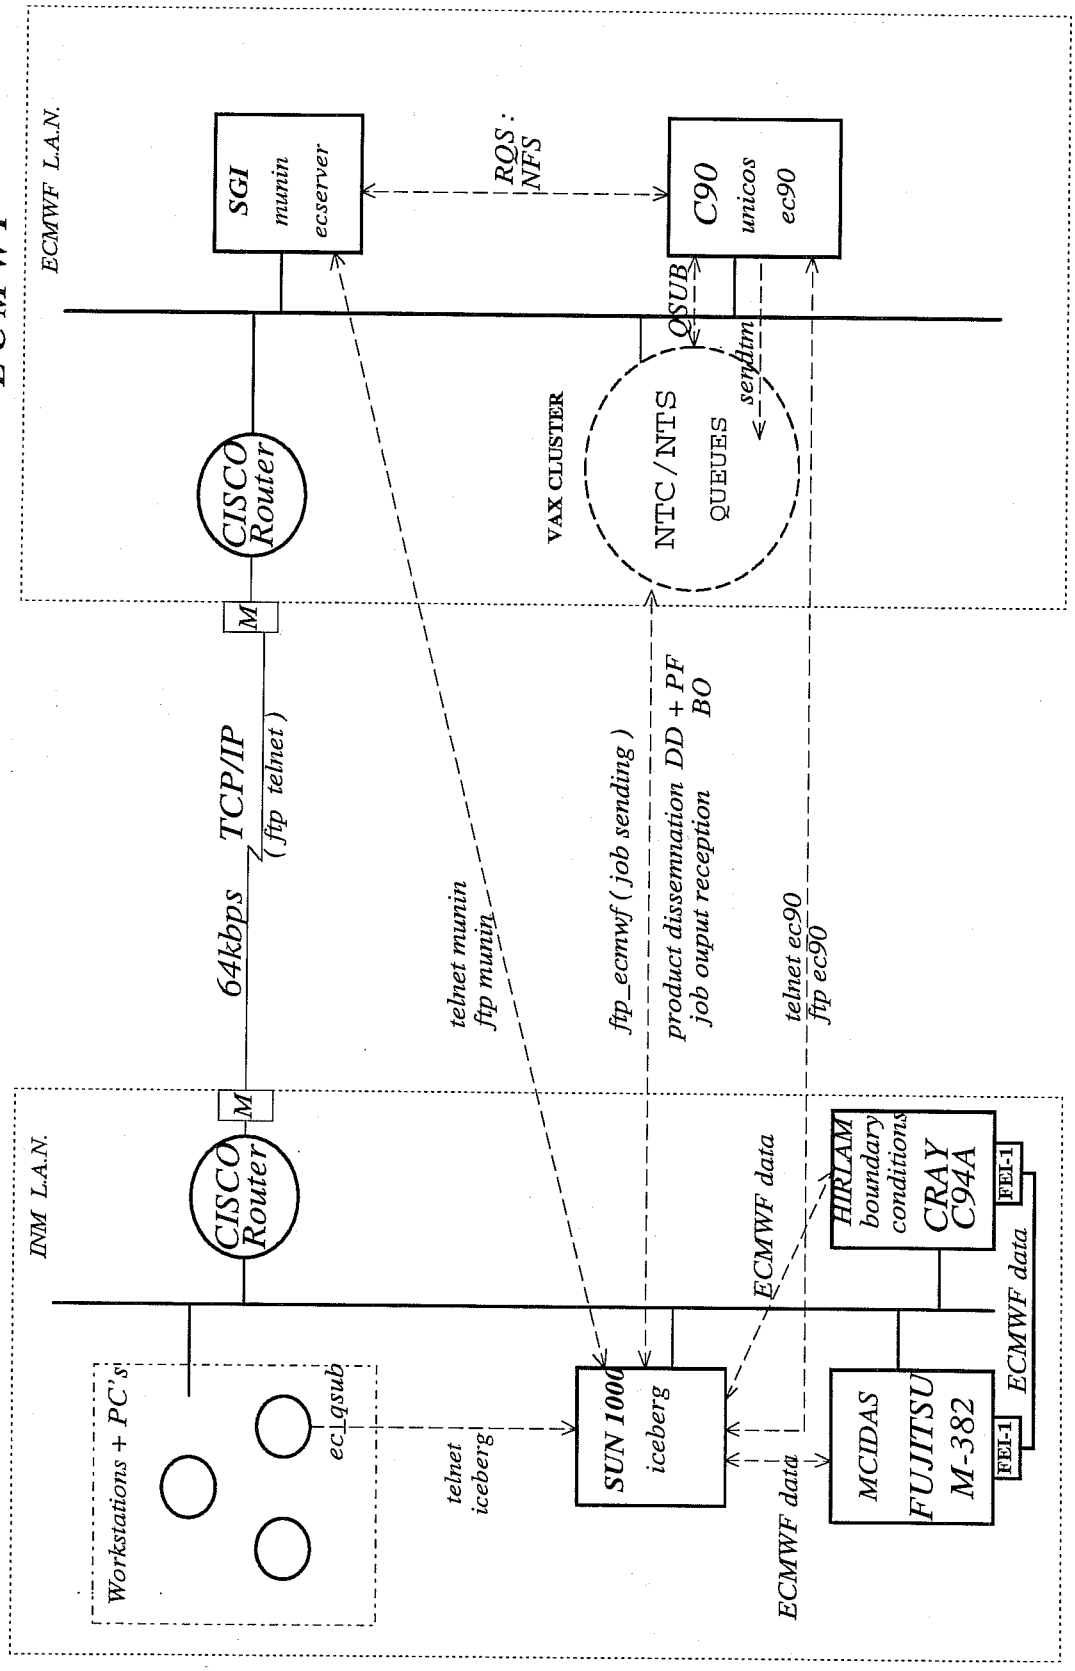

Fig. 2

COMMENTS AND QUESTIONS

S. Pasquini — You are installing new telecommunications equipment? Will it involve new software?

E. Monreal — Yes, we are installing new hardware. It will come into operation next month. The software will not change, we are only changing the hardware.

Sweden

A BRIEF OVERVIEW OF SWEDISH ACTIVITIES IN RELATION TO ECMWF.

i) The computer equipment at SMHI:

The main computer system is based on Digital Equipment VAXs.
At location in Norrköping there are three VAX clusters. One mixed interconnect Vax cluster (MIVC), using FDDI and ethernet, including one combined CI- and DSSI-cluster, comprising of one VAX 6000-610, one VAX 6000-510 and one VAX 6000-410 with a total disk capacity of 40 Gbyte and one Vax 4000-200, one µVax 3100 and two Vaxstation 3100 with local disks. This cluster provides the means to access the main disk-system and a variety of different supporting equipment. Within this cluster there are also equipment for backup and archiving, two cassettetape station TF857 with 7 tapes each, total storage capacity 36.4 Gbytes and a jukebox system with magnetoptical disks, storage capacity 33 Gbytes. The second cluster is a real-time DSSI-cluster comprising of one Vax 4000-300 and two Vax 4000-200. This cluster is the basis for a system that produces customer oriented products for distribution. The third cluster is a cluster with a servernode (Vax 4000-200) supporting 12 Vaxstation 4000.
As decnet node for communication with ECMWF there is a X25-Router (DECnis600) serving. Presently the line-speed is 64 kbits.
There is also a Vax 4000-300 to serve the satellite receiving and processing system. This Vax-system also supports a system for production of various image products. It is also equipped with two very powerful systems for image processing. The image processing systems are to be replaced in the very near future by DEC Alpha AXP-based platforms on which a system for processing of satellite data will be implemented.
A µVax 3100 provides the means for distributing data to customers via modems located on a Decserver 200 on ethernet. It supports around 600 clients with 10,000 products per day.
On site there is a system of a Unix-server (Decstation 5000 Model 240) supporting presently 25 Decstation 5000 of various models all running the DEC version of Unix namely Ultrix.
Since last time an effort has been made to introduce DEC Alpha AXP-based systems into the local environment. There are presently one DEC 3000/400 (dedicated to the Geographical Information System, Arc/Info) and five DEC 3000/300X workstations all running Unix OSF/1. There is also one workstation DEC 3000/300LX running OpenVMS.
SMHI has the responsability for running the Swedish system connecting to the GTS. For this purpose there is powerful Tandem CLX740 system installed called Metcom.
On site there is also a double HP1000 system mainly for data collection from GTS via a direct link to the Metcom system and plotting of observations in near real time. This system also provides the means for collecting data from automatic observation stations.
There is also a remote system, a Convex 3840, a four CPU vectormachine, used to run the local implementation of the Hirlam model. This system is connected to the main SMHI system by five 64 kbits lines.
At the regional centers there are µVax II systems connected to the central system by means of decnet using bridges interconnecting the local area networks using 64 kbits lines. On all systems there are Teragon image-processing systems connected to the Vaxs. These systems are in the process of being phased out. They are presently being replaced with systems based on µVax 3100-80 and for presentation purposes with powerful PCs with 21'-displays. At the regional centers there are also HP 1000 systems connected to GTS for reception and plotting of observations.

SWEDEN

Since the time of the last meeting a system for distribution of meteorological data to the Baltic countries has been developed and put into production. The communication is to a Vax in Riga from which the products are distributed onwards to identical systems in Tallinn and Vilnius. The Vaxes are connected to PC-based presentation systems.
Another system that has been developed since last time is a system for automatic distribution of and handling of on call requests for faxproducts and products to PCs. This system is supported by a Vaxstation 4000.
One thing that could be specially noted is that all production of plotted charts has been moved from pen-plotters and raster-plotters to A3 laserwriters.
There are presently five Vax systems supporting weather-radars, one in Norrköping, one in Stockholm, one in Göteborg, one on the island of Gotland and one in Leksand. There is also connections between Norrköping and the military weather radars in four different locations (Karlskrona, Luleå, Hudiksvall and Örnsköldsvik) and between Malmö and the Kastrup weather radar in the vicinity of Copenhagen. A new system that connects the radars in Sweden, Norway and Finland has been installed. This system connects concentrator nodes in the three countries, which in turn will be connected to the computers of each radar system.
All SMHI Vax systems are supported by local area networks (FDDI and/or ethernet).
For further details see the attached diagram outlining the SMHI computer network.

ii) Projects run at ECMWF:

- High resolution limited area model development
- Research on regional transport
- Aerodynamics and airpollution modelling
- Extraction of data for operational usage
- Hydrodynamic models
- Trajectory studies
- Atmospheric chemistry

iii) Experience using ECMWF computers:

Usage of ECMWF computer resources for HIRLAM development.

The joint Nordic-Dutch-Irish HIRLAM system is now operational in Finland, Denmark and Sweden on local supercomputers. In parallel with the routine operations, further development of the system is taking place on the ECMWF Cray system. The basic idea of using the ECMWF computer system for the HIRLAM work is to keep a "master" version of the system available for all group members. This master system is used for research and development, and as master copy for local installations, benchmarks etc.

Usage of ECMWF computer resources for trajectory calculation.

In cooperation with the University of Virginia, USA, trajectory calculations are carried out for a research project on long range transport of soil particles.

• I suggest that the Member State Computing Representative meeting will become a yearly event because I find it difficult to act as representative when I feel that I lack the close contact with and the good knowledge of operational work at ECMWF that I feel needed. This is probably partly because I am not a regular user of the ECMWF computer system.

iv) Plans:

We will start to use the interactive facility in the very near future.
Apart from that I am not well informed on the plans of ECMWF usage.

SMHI NETWORK

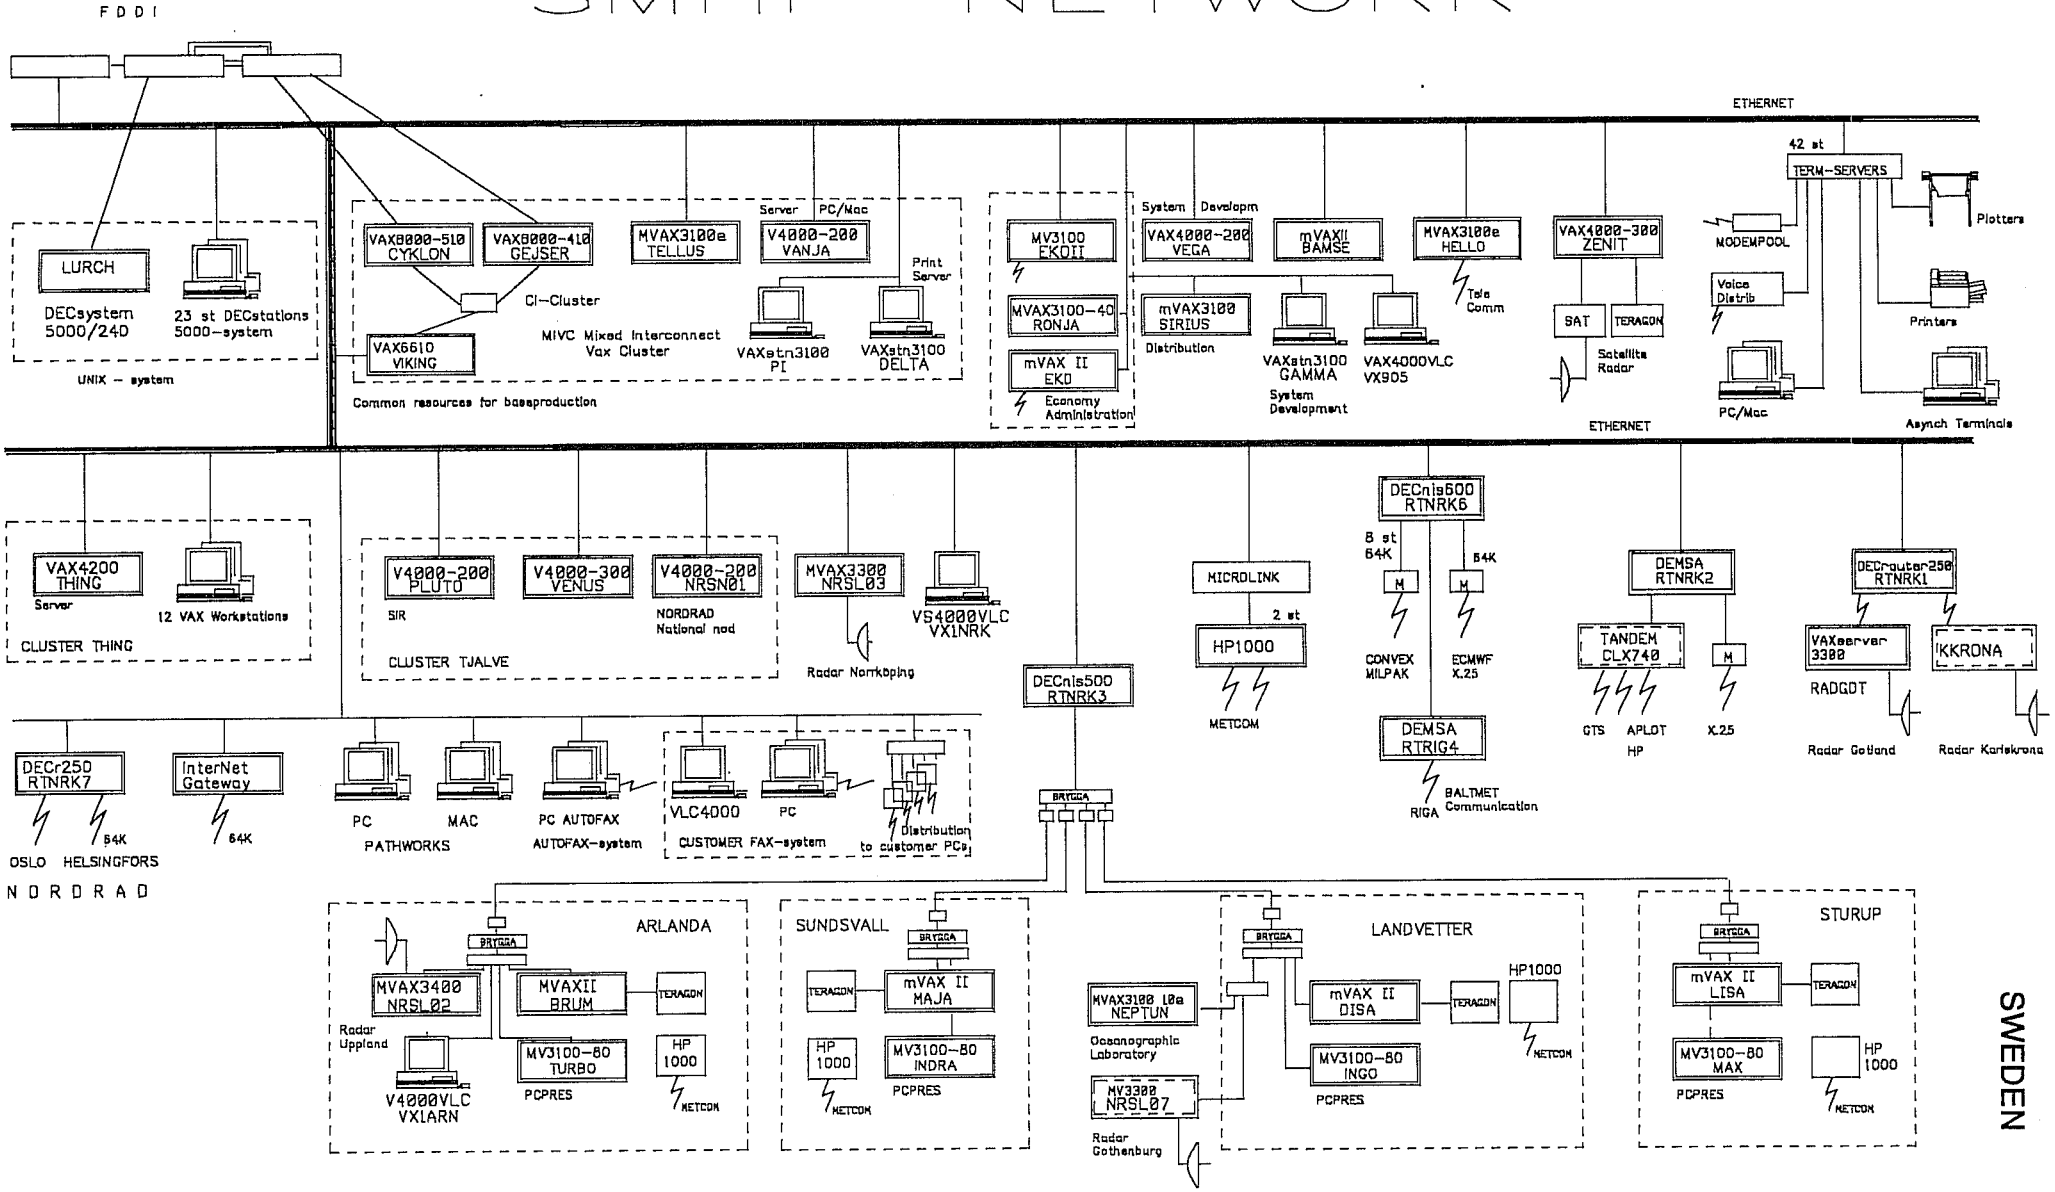

SMHI-ldc 1994-09-22
ROLF STRANDBERG

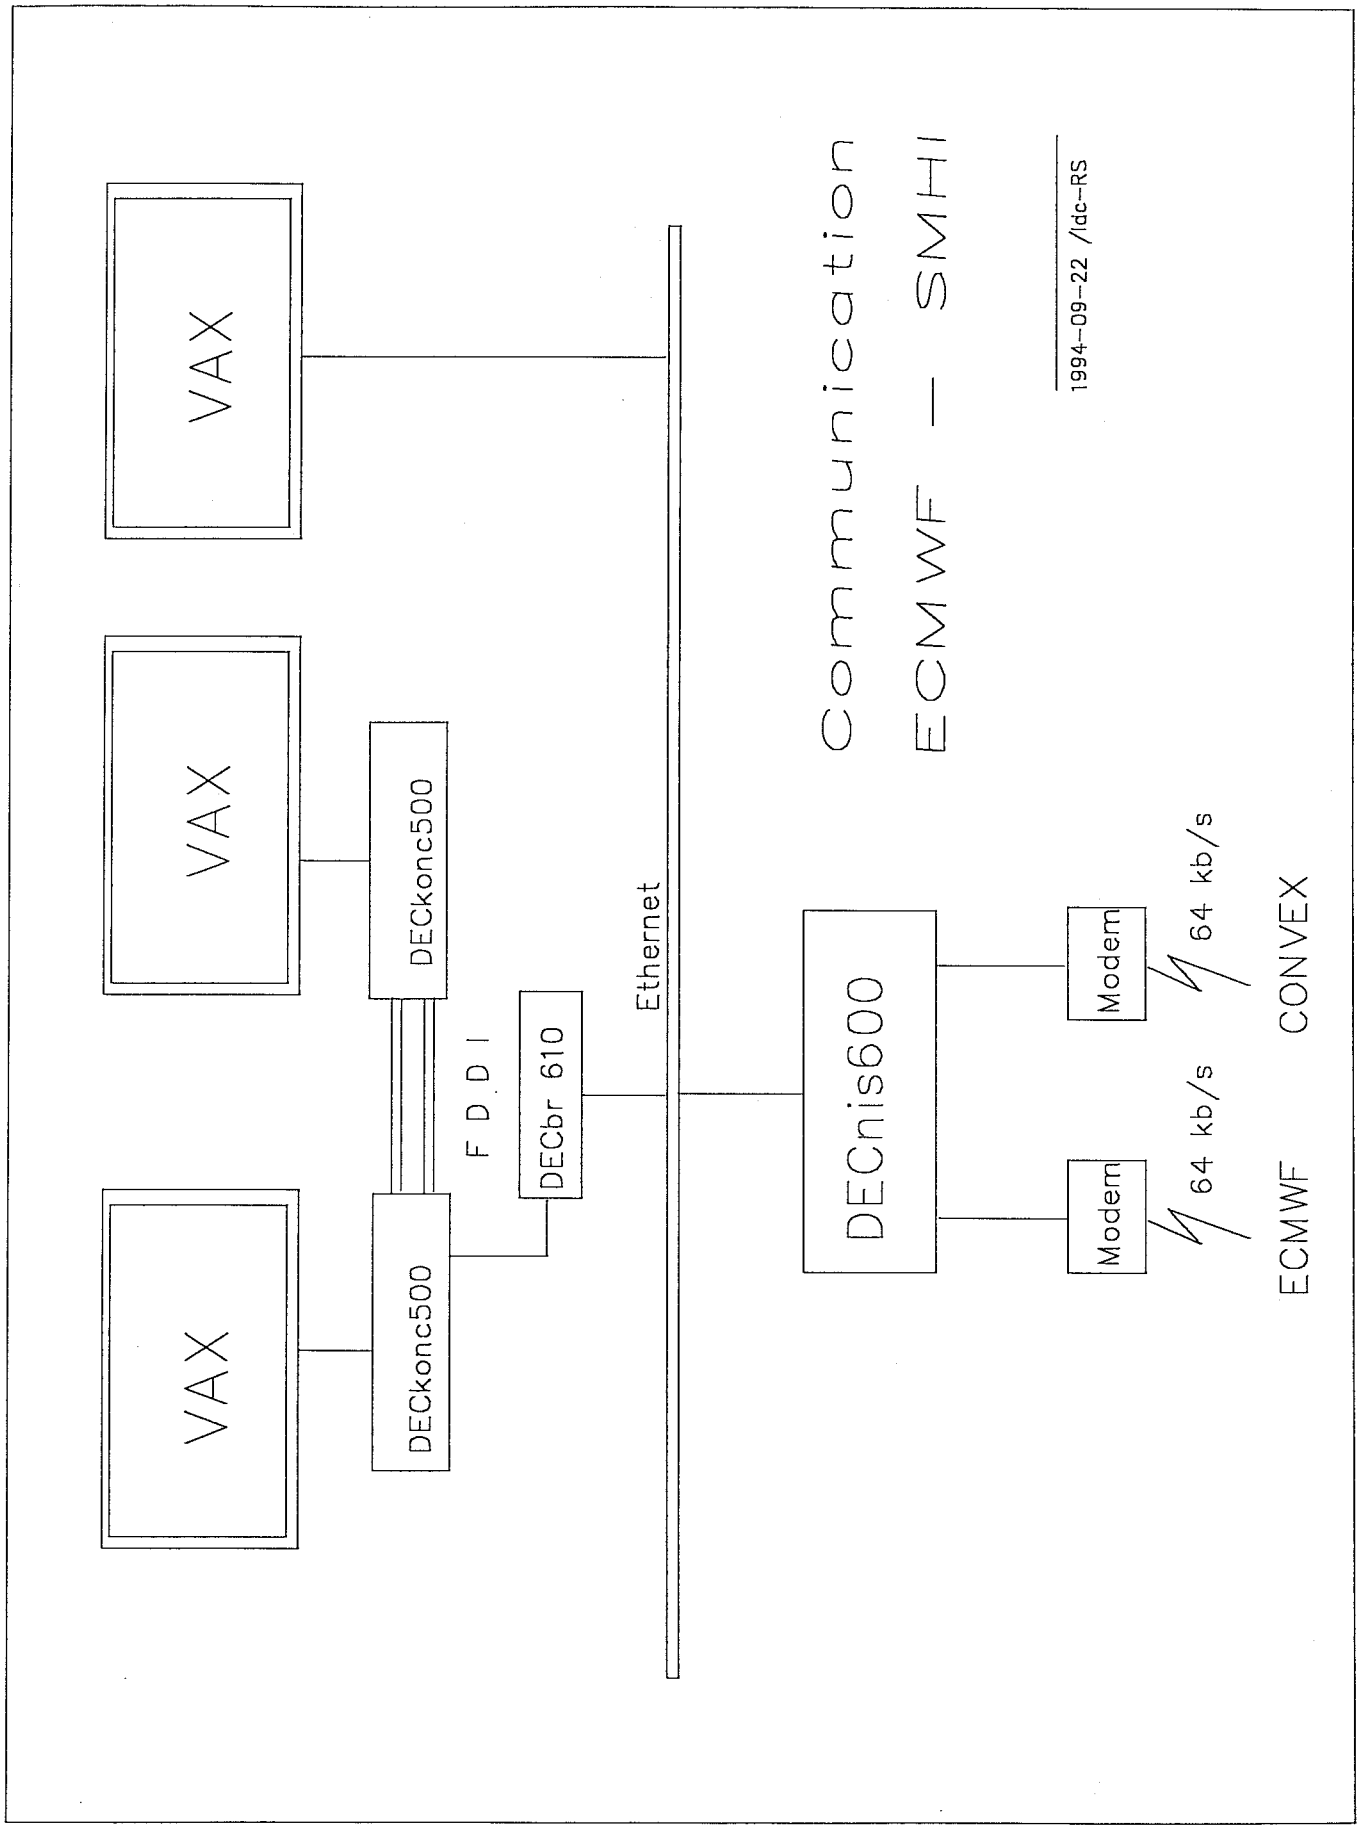
VAX
VAX
VAX
DECkonc500
DECkonc500
F D D I
DECbr 610
Ethernet
DECnis600
Modem
Modem
64 kb/s
64 kb/s
ECMWF
CONVEX
Communication
ECMWF — SMHI
1994-09-22 /ldc-RS

Swiss Meteorological Institute
B. Bachofner
Member States Computing Representatives' Meeting, 27-28 September 1994

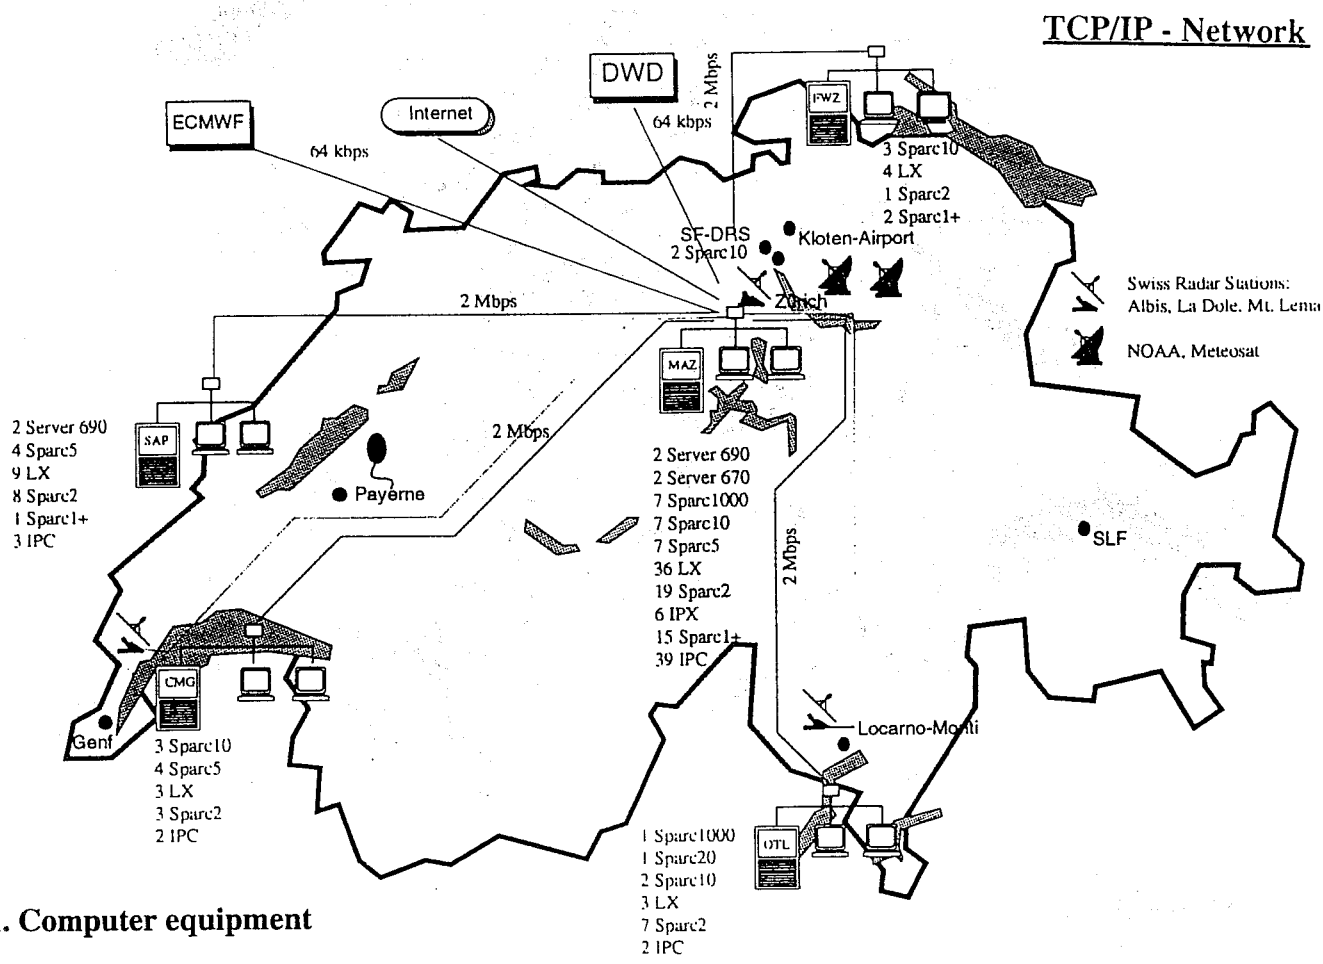

1. Computer equipment

The eight years old Siemens Meteor system has to be replaced until the end of 1996. The function will be taken over gradually by individual UNIX-systems. Each staff-member has his own SUN-workstation, which serves as a multifunctional work environment. Office automation with word-processor and spread-sheet can be applied as well as technical-scientific and graphical applications. All workstations are connected over a 10BaseT-Network to the NIS-Server which is connected to the server-network. Via 2 Mbps-links and CISCO-Routers we have been installed our annex-institutes Locarno, Payerne, Geneva and Kloten.

Two SUN Sparc1000 function as servers for central data, data base and back-up. With a SUN SparcCenter2000 we will install the archive-database.

All communication functions for MOTNE, WMO, AFTN, X25 and our automatic weather station network, are connected by the Global Telecommunication System of Netsys. The frontend is served by two SUN Sparc1000. These systems generates weatherproducts and send them via standard communicationlinks to our clients.

For the wind measuring network two VAX 3400 are in operation, which are periodically connect the 40 measurementstations over the public telefonnetwork.

Two Sparc1000 are the composite-servers for the Swiss Radar Network with 3 radarstations. With a NOAA- and a Meteosat-antenna we receive satellitepictures which will be transferred by WAN and will be visible on all SUN-workstations.

All servers are connected by 10BaseT-Network.

Connections to ECMWF, SF-DRS (Swiss Television), NAZ (Nationale Alarmzentrale) and ETH (Federal Institute of Technology) are made over CISCO-Router (external).

The connection to ETH-Cray allows us to run the numerical models twice a day.

2. Automatic Weatherstation Network

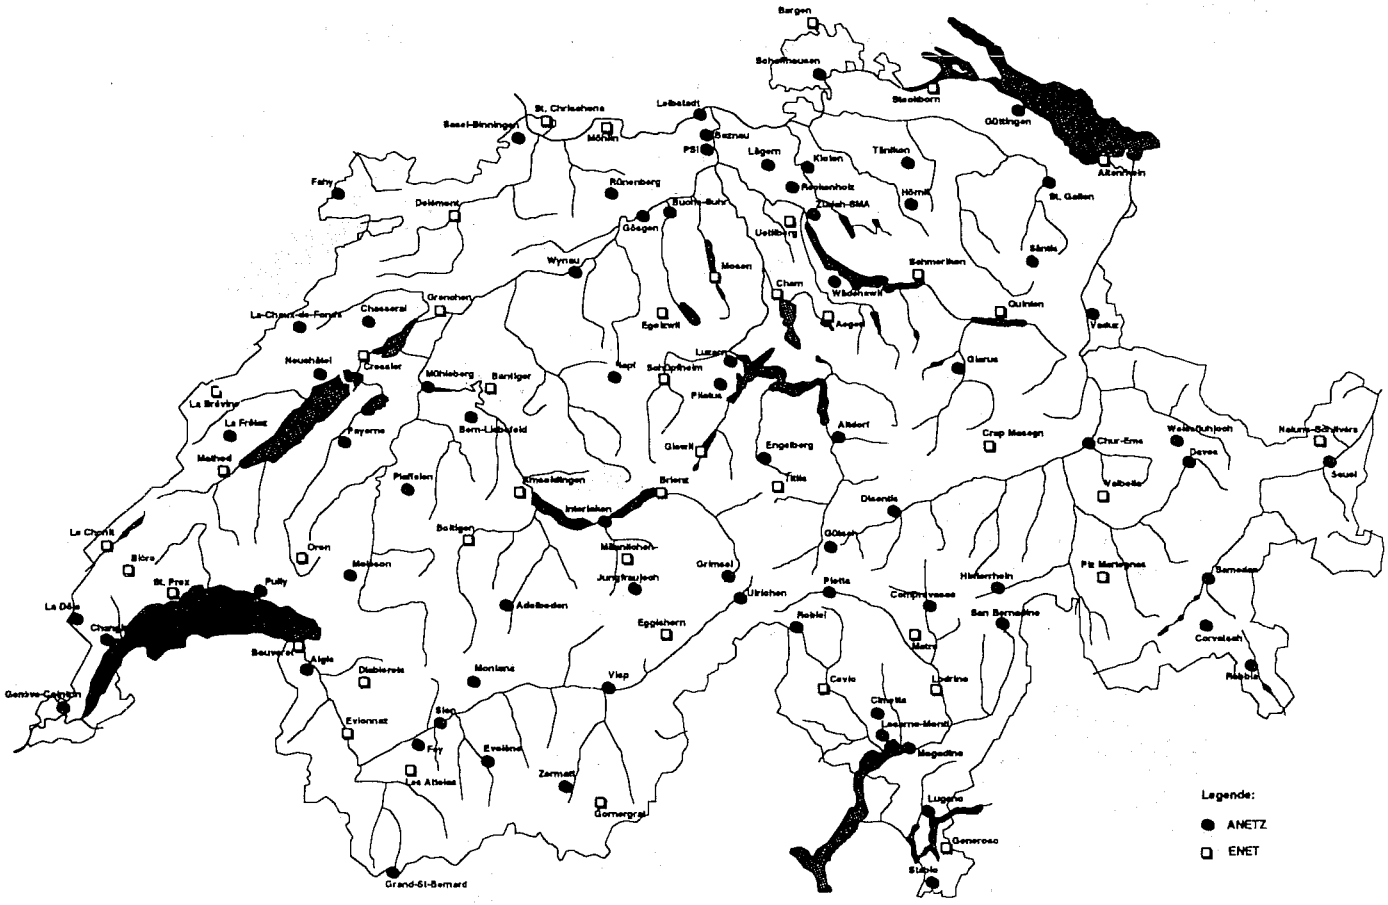

3. Connection to ECMW

The SMI-Network is connected over a 64 kbps line to ECMWF. The dissemination files will be transferd periodically via TCP/IP filetransfer to the GTS-Servers. From there on they can be visualized on SUN-workstations with Micromagics and Magics. The interactive access to the ECMWF-Servers is made possible by telnet.

4. Projects run at ECMWF

a) Operational Weatherforecast
- Dissemination products
- MARS data retrievals
- Precipitation probabilities (based on MARS)
- Precipitation forecast
- Surface temperature forecast
- Trajectory model of KMNI

b) Research division
- MARS data retrieval for historical case studies

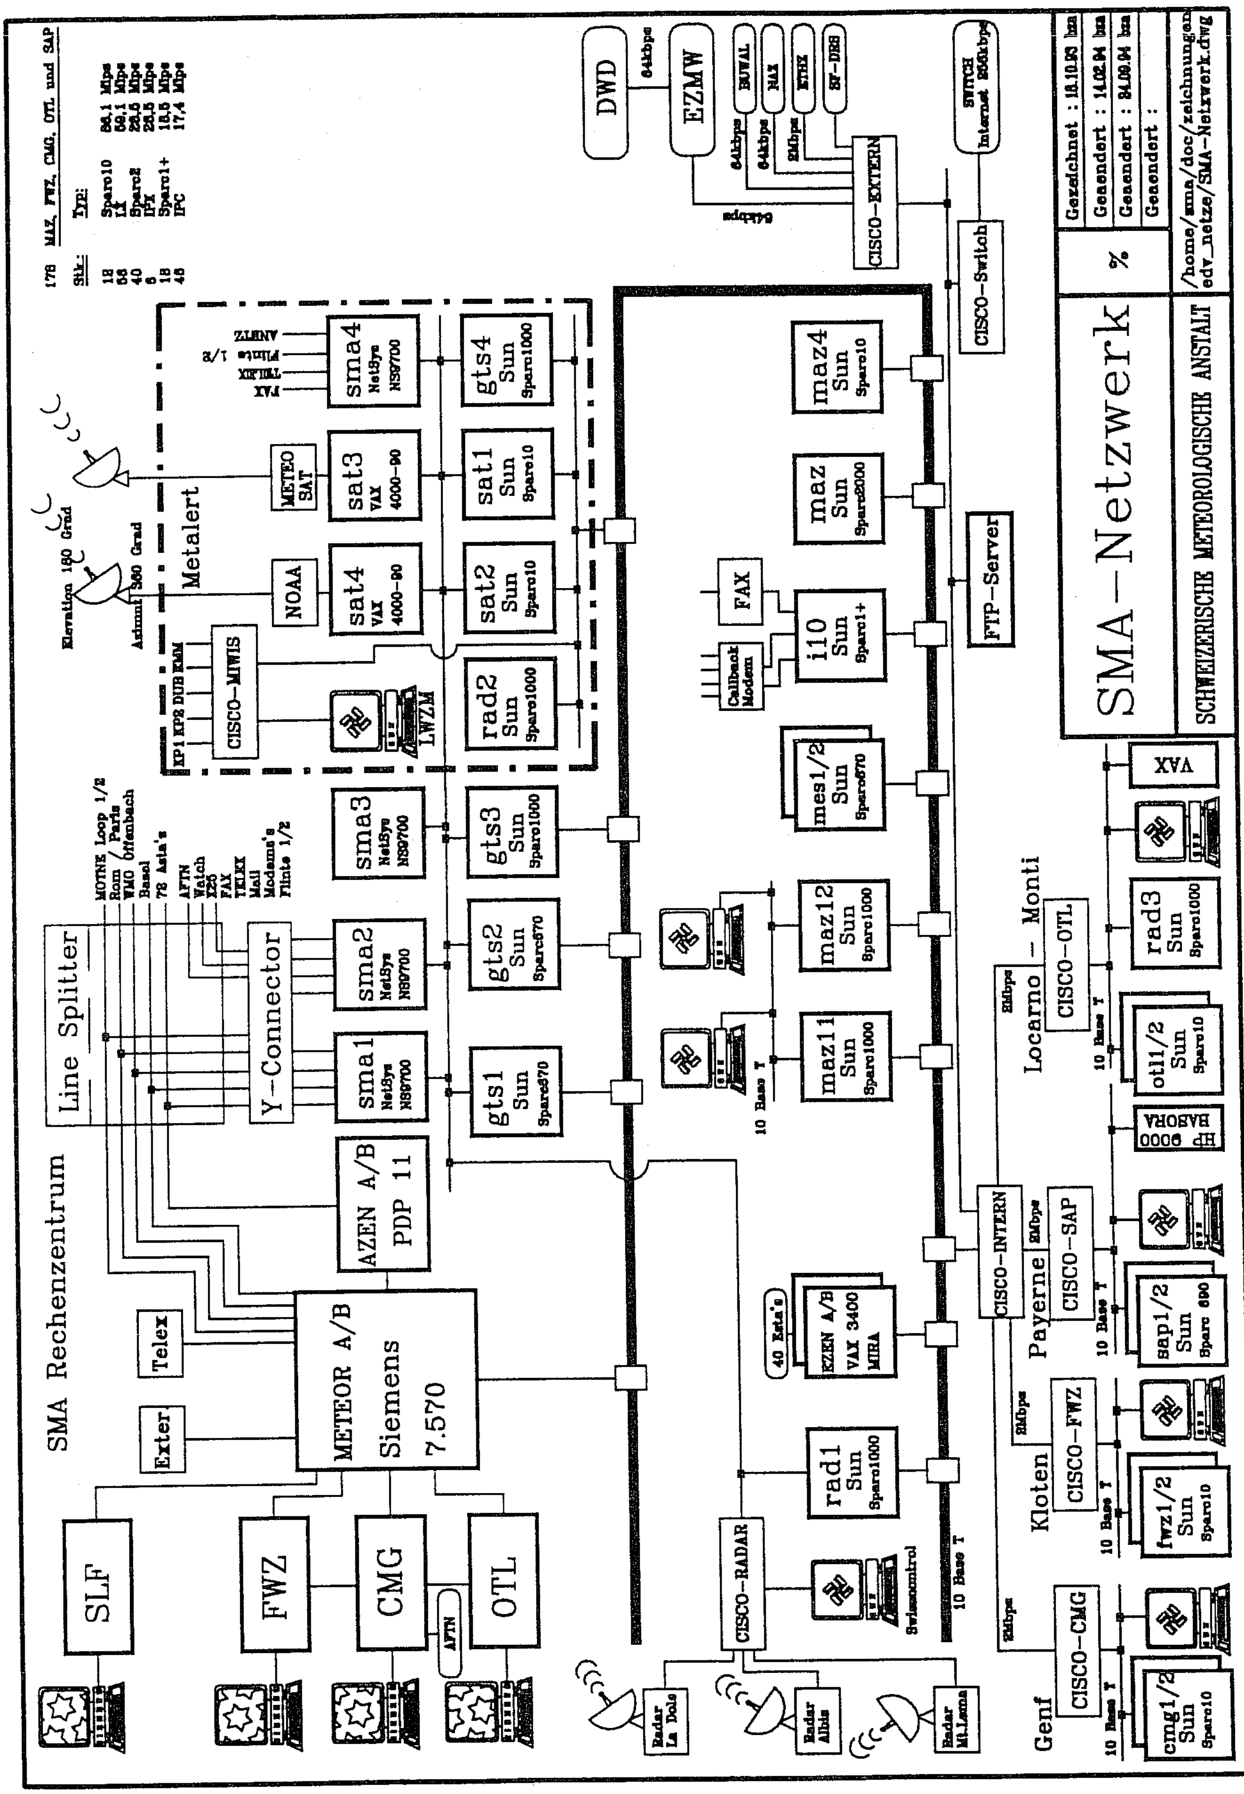
SMA-Netzwerk
SCHWEIZERISCHE METEOROLOGISCHE ANSTALT
SMA Rechenzentrum
Metalert
Line Splitter
Y-Connector
METEOR A/B Siemens 7.570
AZEN A/B PDP 11
Telex
Exter
SLF
FWZ
CMG
OTL
AFTN
sma1
sma2
sma3
sma4
gts1
gts2
gts3
gts4
sat1
sat2
sat3
sat4
rad1
rad2
rad3
maz
maz4
maz11
maz12
mes1/2
i10
NOAA
METEO SAT
CISCO-MIWIS
LWZM
CISCO-RADAR
CISCO-INTERN
CISCO-EXTERN
CISCO-Switch
FTP-Server
FAX
DWD
EZMW
VAX
HP 9000
Radar La Dôle
Radar Albis
Radar Mt.Lema
Genf
CISCO-CMG
cmg1/2
Kloten
CISCO-FWZ
fwz1/2
Payerne
CISCO-SAP
sap1/2
Locarno - Monti
CISCO-OTL
otl1/2
10 Base T

COMMENTS AND QUESTIONS

B. Barg — Are you, or have you, switched from the Ingres to the Empress database systems?

B. Bachofner — We have moved from Ingres to Empress. Empress is used for all new projects.

UK METEOROLOGICAL OFFICE

A. Dickinson

Member States Computing Representatives' Meeting, 27-28th September, 1994

1. Computer equipment and connections to ECMWF

The computing environment at the UK Met Office is shown in Figure 1. The main computer systems consist of a 16 processor Cray C90, a 4 processor Cray EL and a HDS EX100 (IBM compatible) mainframe. The role of the EX100 as a front-end system is gradually being phased out and it is now acting more and more in the role of a file server. All research users are connected to a workstation network based primarily on HP 9000-735 servers and X-terminals. The Cray C90 resources are split 70:30 between climate modelling and numerical weather prediction.

The telecommunications link to ECMWF is shown in Figure 2. Links to ECMWF go via a MicroVax gateway which is connected to the Office-wide ethernet network known as CDN (Central Data Network). A second MicroVax acts as backup. The line to ECMWF is rated at 64 Kbits/sec and uses X25 protocols. The connection supports the receipt of observational data and the dissemination of ECMWF products. Along with Offenbach, the UK provides one of the connections between ECMWF and the GTS.

The telecommunications link will be upgraded early in 1995. The new configuration is shown in Figure 3. The MicroVax gateway will be replaced by a more powerful VAX 4000/500 system. The MicroVax only supports 100Mbytes of disk space, which severely limits the staging space available for file transfers. The new system will be installed with 8 Gbytes of disk space so that much larger file transfers can be supported. A CISCO MGS/A router will be used to support both DECNET and TCP/IP communications. This will allow the capacity of the link to be upgraded beyond 64Kbits/sec up to a maximum of 2Mbits/sec. DECNET will continue to be used for connection to the ECMWF computer facilities, while TCP/IP will be used for connection to ECNET.

2. Projects, experience and plans

It is expected that our Cray C90 allocation of resources will be distributed as follows over the new few years:

European Climate Network - 15%
Ensemble forecasting - 70%
Diagnostic studies - 15%

The heading "Diagnostic studies" covers a number of projects whose main use of the computer is to access MARS data. The ensemble forecasting work is being carried out in

conjunction with the ECMWF. Two 33 member ensembles are run each week at a total cost of 5000 units. The elapsed time from first job submission to receipt of the final forecast is around 4 days - a usage that is consistent with a *pro rata* consumption of our annual allocation. However, this turn around is dependent on more or less sole use by the UK of the Member States' large memory queue. If other Member States begin to use this queue, it will be difficult to maintain this level of turn around and we will have problems in realising our annual allocation of units. Further adjustments to the priorities of the C90 queuing system may be required.

Once the upgrade to the telecommunications link is established early in 1995, we will review our requirements for data transfers and implement a upgrade to the band width of the connection.

The UK plans to use the Centre's T3D in order to further develop and optimise the message passing version of the Unified Model . The Unified Model is the main modelling code used for climate prediction and NWP by the Met Office. The message passing layer has been developed as part of the GP-MIMD 2 contract funded by the European Union. The main areas of work will involve the identification of inefficiencies on the T3D, hopefully their resolution, and the performance of scalability tests. We will be using PVM as the message passing language accessed through the Parallel Utilities Library from the Edinburgh Parallel Computing Centre.

UK METEOROLOGICAL OFFICE MAINFRAME COMPUTERS

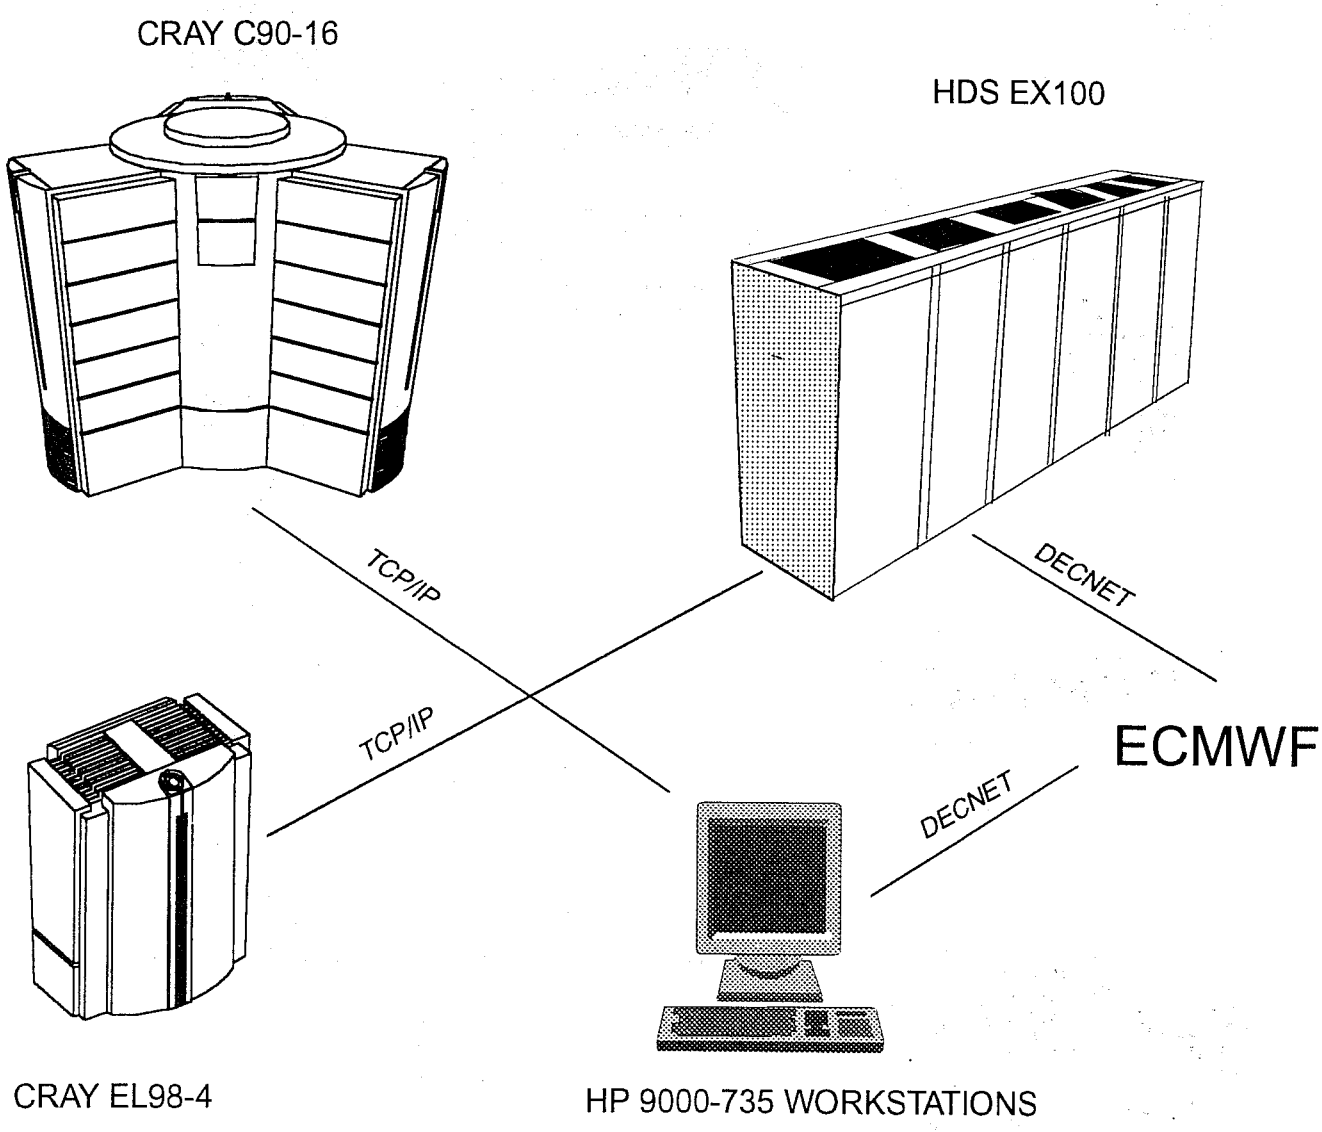

Figure 1

BRACKNELL-ECMWF LINK
September 1994

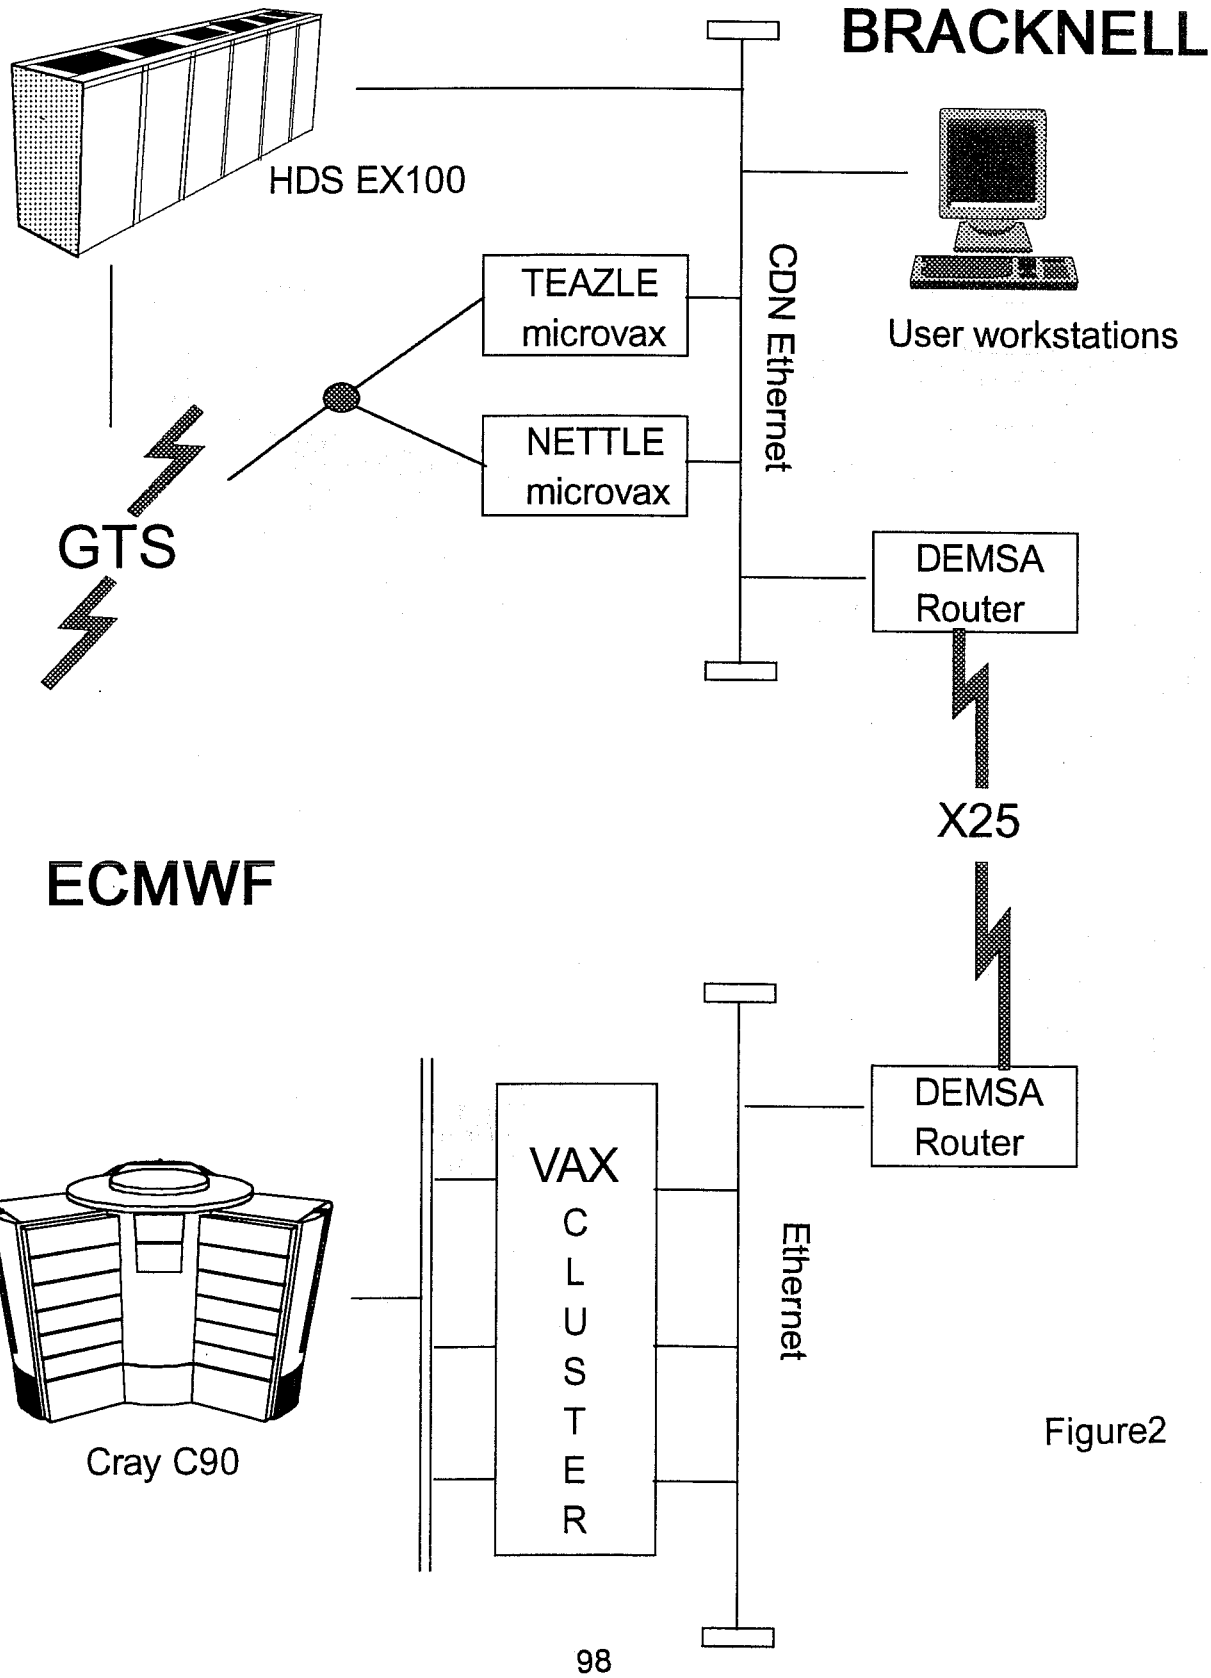

Figure2

BRACKNELL-ECMWF LINK
Early 1995

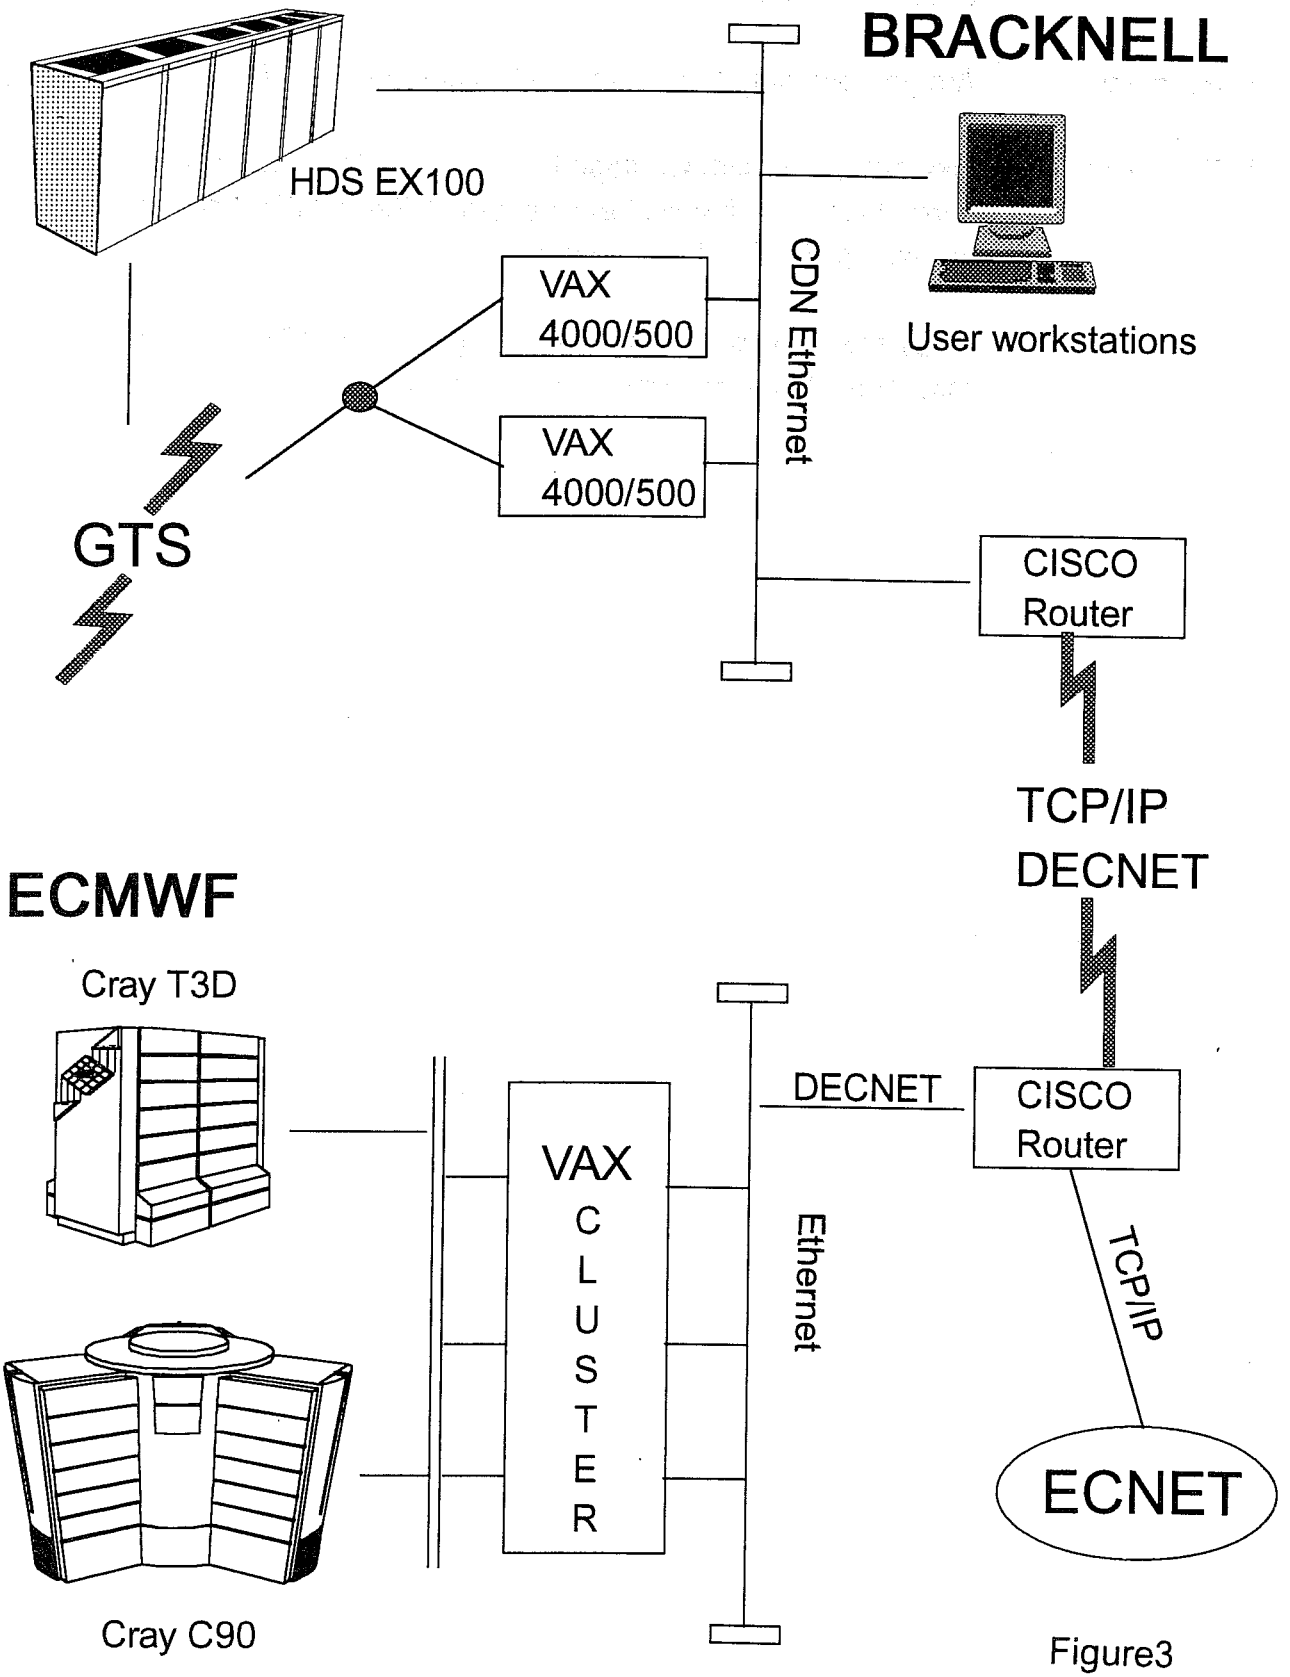

Figure3

COMMENTS AND QUESTIONS

G-R. Hoffmann — If you do begin to experience problems of turnaround, please let us know early enough so that we can, hopefully, do something about it in time.

S. Pasquini — Are you planning to move to Unix systems on your networks?

A. Dickinson — Like everyone else, we hope to move to "open" systems, but the current reality is that we have several different operating systems in use e.g. MVS, VMS, Tandem.

S. Orrhagen — In Sweden we are experimenting with Open VMS, one of the most "open" standard systems available today.