ECMWF Newsletter

Number 59 - September 1992

ECMWF Computer Networks - 1992

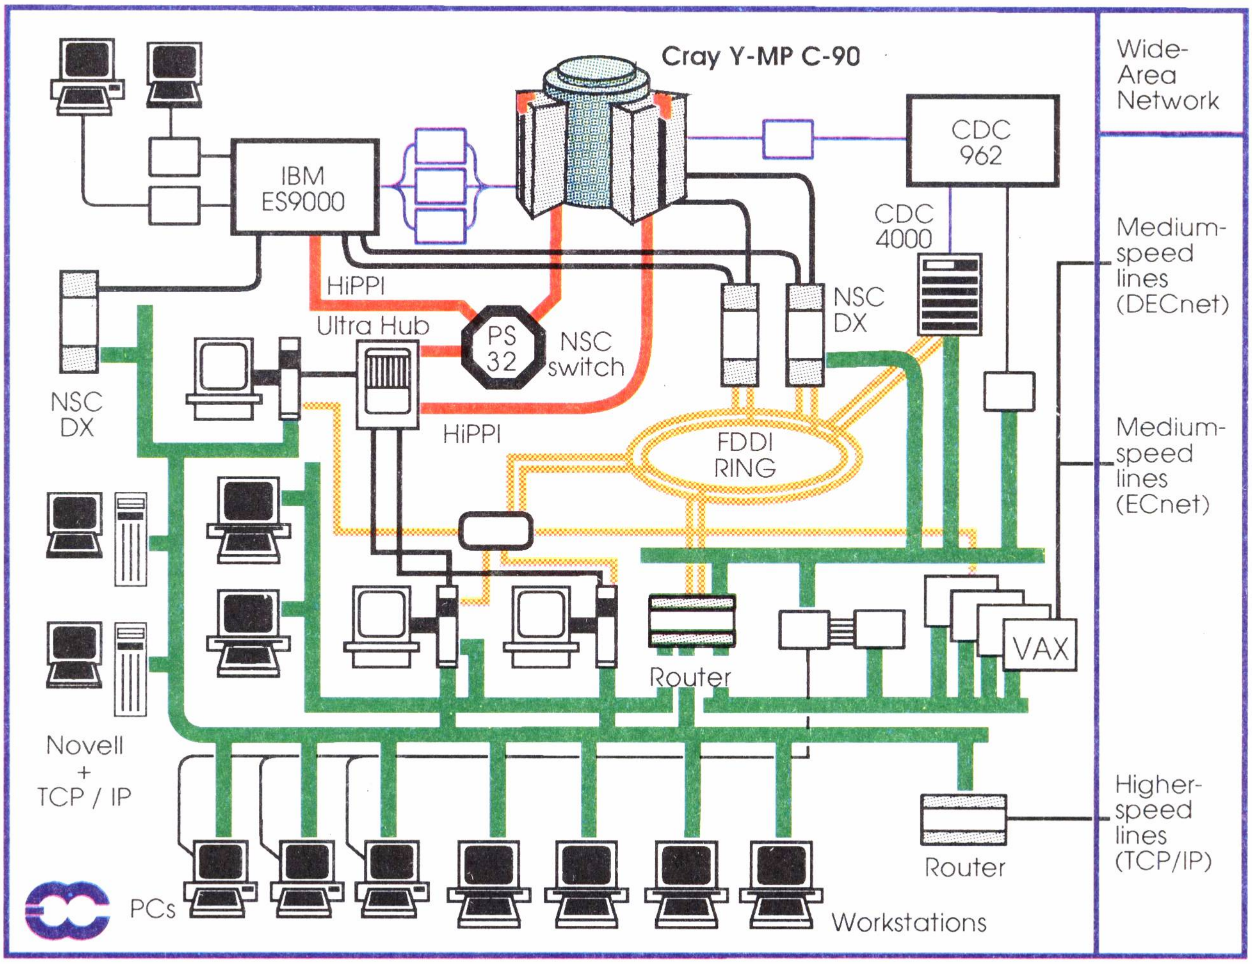

Shinfield Park, Reading, Berkshire RG2 9AX, England. Telephone: U.K. (0734) 499000
International (+44 734) 499000, Telex: 847908 ECMWF G, Telefax (0734) 869450

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

IN THIS ISSUE

METEOROLOGICAL

COMPUTING

COMPUTER USER INFORMATION

GENERAL

COVER: ECMWF Computer Networks, 1992 (see article page 20)

This Newsletter is edited and produced by User Support.

The next issue will appear in December 1992.

Two of the articles in this issue of the Newsletter illustrate the need, and concern, to maximise the efficiency of operations on the Centre's computer system. These relate to the Supervisor Monitor Scheduler (SMS) and the Common File Server (CFS).

The meteorological article this time describes pre-operational trials of a scheme for assimilating satellite data, in the form of multi-spectral radiances, more directly into the numerical weather prediction system at ECMWF. This scheme has proved beneficial, and has been implemented operationally since June this year.

Also in this edition is an article giving up-to-date information on the ECMWF local computer network.

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

CHANGES TO THE OPERATIONAL FORECASTING SYSTEM

Recent changes

On 9 June 1992, two changes were introduced in the analysis:

- humidity data from SYNOP observations were excluded from the analysis, to reduce excessive convective precipitation in short and early medium-range forecasts;
- a redundancy check on aircraft data was introduced so that new high density data (ACARS) can be correctly used in the OI scheme.

Since 23 June 1992, the temperature data from the NOAA satellites used in the analysis of the Northern Hemisphere have been derived from cloud cleared radiance data received from NESDIS, using a 1d-variational inversion technique. A significant improvement in the quality of the analysis and forecast is expected from this change, particularly over the Pacific and North America. A positive impact should also be seen in the medium-range over Europe.

On 17 August 1992, a change was made to the forecast model, to include:

- introduction of vertically non-interpolating semi-Lagrangian scheme and smaller time filter;
- suppression of inversion clouds at the lowest three model levels rather than just the lowest level;
- introduction of a prognostic equation for sea-ice temperature to replace the use of climatological values;
- other minor parametrization changes.

Experimentation on this change suggested that the RMS error of the forecast should be reduced and the anomaly correlation of height improved in the medium-range, particularly in winter. It also suggested that the new model version will improve the day-to-day consistency of forecasts, given its lower and more realistic level of eddy activity. Initial operational results have indeed shown a distinct improvement when compared to other NWP centres.

Planned changes

Improved representation of cloud/radiation interaction and of surface and planetary boundary layer processes are under development, but likely dates for operational implementation cannot yet be given.

- Bernard Strauss

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

A NEW APPROACH TO THE USE OF TOVS SATELLITE DATA AT ECMWF

Introduction

The TIROS Operational Vertical Sounder (TOVS) on the NOAA series of polar-orbiting satellites measures multi-spectral radiances which are related to the atmosphere's temperature and humidity structure (see *Smith et al.*, 1979). The scheme which has been developed for assimilating information from these radiances more directly into the numerical weather prediction (NWP) system at ECMWF is known as one-dimensional variational analysis or 1DVAR. Extensive testing of the scheme through forecast impact experiments has demonstrated clear and consistent benefits for forecast skill in the northern hemisphere extra-tropics. As a result, the 1DVAR scheme was implemented operationally at ECMWF in June 1992. The scheme and the results of the forecast impact experiments are described in more detail by *Eyre et al.* (1992).

Background

It is widely recognized that satellite sounding products have a beneficial effect on NWP analyses and forecasts in the southern hemisphere. Indeed they are essential for maintaining the level of forecast skill currently achieved. In the early 1980s, it was also demonstrated that they had, on average, a positive impact when used in data-sparse areas in the northern hemisphere (for example, see *Halem et al.*, 1982; *Bengtsson*, 1985; *Kashiwagi*, 1987). More recently, however, as NWP systems have improved, ECMWF and other major NWP centres have experienced increasing difficulty in demonstrating consistent positive impact on northern hemisphere forecasts from temperature profiles retrieved from satellite sounding data by conventional techniques (for example, see *Andersson et al.*, 1991; *Kelly et al.*, 1991; *Thoss*, 1991).

The main reason for these problems is well understood. It is implicit in the analysis of the inversion problem for satellite sounding data given by *Rodgers* (1976) and has been discussed more recently by *Eyre and Lorenc* (1989). In brief, the problem arises from the rather poor vertical resolution of satellite sounding radiometers. Although the current operational instruments of the TOVS system contain more than twenty spectral channels, their weighting functions are broad and overlapping (see Fig. 1) and they supply fewer than twenty independent pieces of information (perhaps six or seven). There are components of the atmospheric profiles, namely those representing high-order vertical structure, to which the satellite sounder is blind. The information on these components can be supplied only by prior or "background" information.

In a conventional retrieval system the prior information is independent of a NWP model. In the system used operationally by NOAA/NESDIS, it comes from a large library of recent pairs of satellite measurements and collocated radiosonde profiles. Although the methods for using such information are now quite sophisticated and complex, they cannot really compensate for the

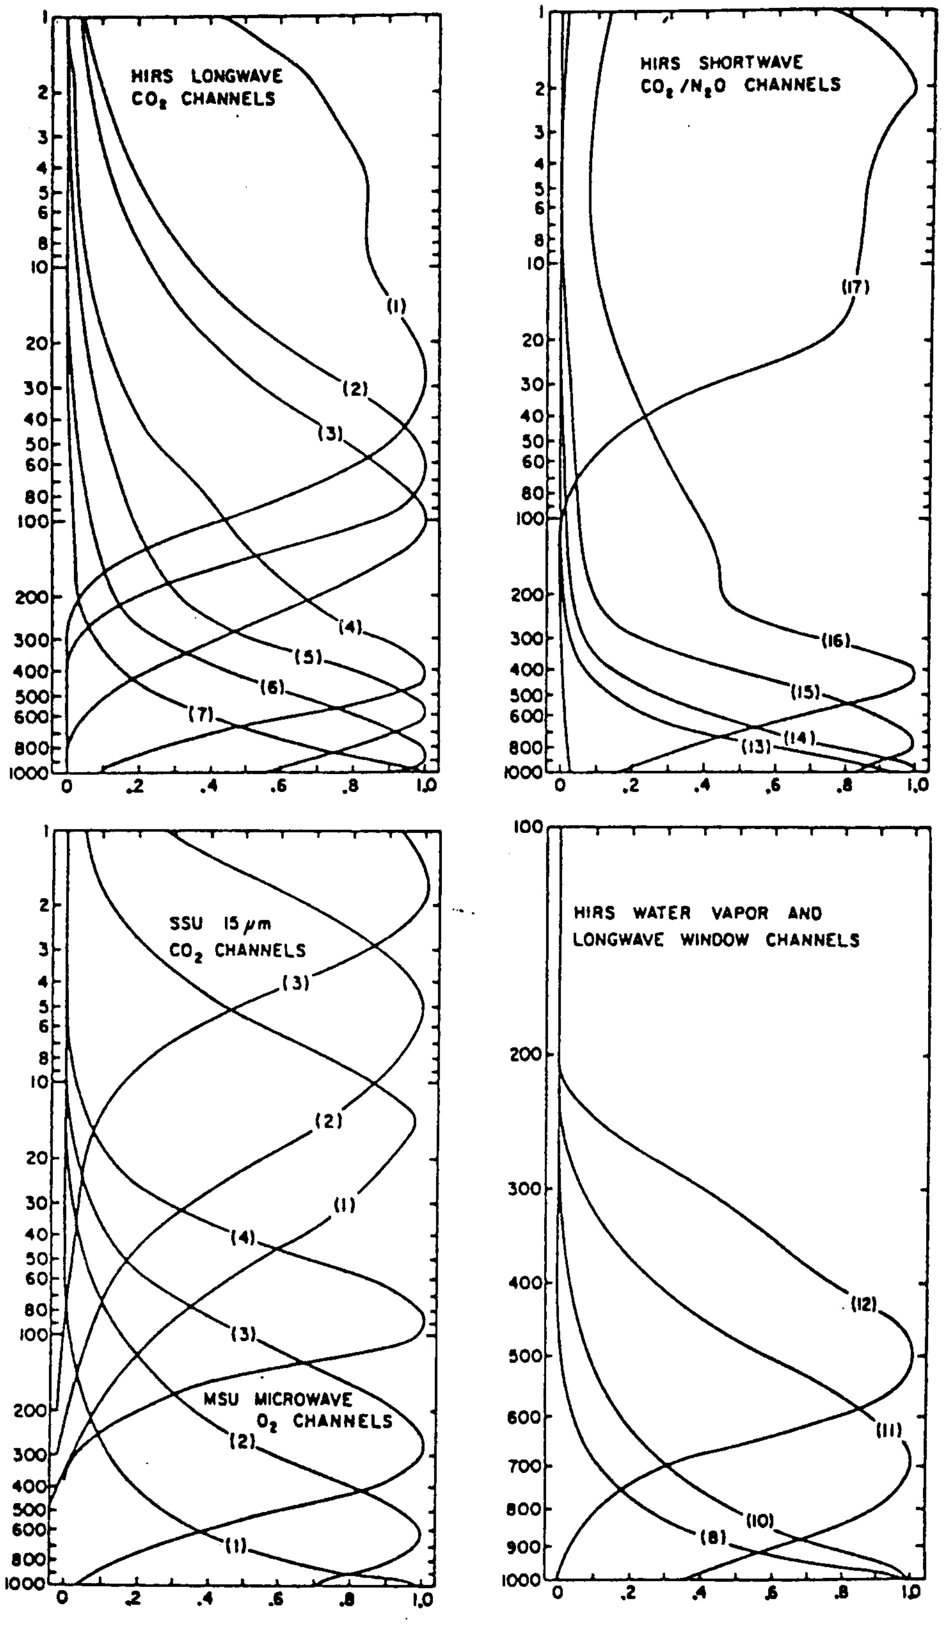

Fig. 1 TOVS weighting functions (taken from *Smith et al.*, 1979).
HIRS - High-resolution Infra Red Sounder
MSU - Microwave Sounding Unit
SSU - Stratospheric Sounding Unit

fundamental "blindness" of the measurements to certain components of the profile, and they often supply information on these components which is less accurate than can be supplied by a NWP system. In such cases, the assimilation of independently retrieved temperature/humidity profiles will tend to degrade rather than improve the accuracy of the NWP model fields.

The solution to this problem is to stop using satellite sounding data in the form of retrieved profiles, as though they were poor-quality radiosonde data, and to start treating them for what they really are — radiance measurements. Météo-France and the UK Meteorological Office (UKMO) adopted these ideas some time ago and have implemented them operationally to a limited extent. At Météo-France a scheme was developed for assimilating TOVS brightness temperatures into a three-dimensional OI (optimal interpolation) analysis (*Durand and Juvanon du Vachat*, 1985; *Durand*, 1985). Its main limitation is that it approximates the problem as linear and cannot treat properly the nonlinear link between brightness temperatures and NWP model variables. At the UKMO, an equivalent one-dimensional (vertical only) approach was developed in which a forecast profile and its expected error covariance provide the constraints on the inversion (*Eyre et al.*, 1985; *Turner et al.*, 1985). It makes similar approximations of linearity. Experiments with a fully nonlinear one-dimensional scheme have also been performed (*Eyre*, 1989).

At ECMWF, a number of initiatives have been undertaken in order to make more direct use of satellite radiance data within the assimilation system. The UKMO's one-dimensional nonlinear optimal estimation scheme has been imported and an interface to ECMWF's NWP model developed. The scheme has been adapted to run on cloud-cleared TOVS brightness temperatures (rather than on raw TOVS radiances for which it was originally developed). It has been applied for the first time to global data, and this has involved the development of a new radiance monitoring and tuning scheme. In parallel with these activities, the development of a three-dimensional variational analysis (3DVAR) has been a major project at ECMWF (*Pailleux et al.*, 1991). One component of this scheme is the assimilation of global TOVS brightness temperature data (see *Pailleux*, 1990). Much of the development involved for TOVS is common between the 3DVAR and the one-dimensional scheme, including the radiative transfer model and the radiance monitoring/tuning system. Moreover, both approaches are based on the solution of the same variational equation. For these reasons, and to emphasize the link between the two projects, the one-dimensional optimal estimation scheme implemented at ECMWF is known as "one-dimensional variational analysis" or "1DVAR".

Theory

The variational approach to the assimilation of data into an NWP system has been described by a number of authors (e.g. *Lorenc*, 1986). When applied specifically to the assimilation of satellite radiance data we find the state of the atmosphere (which in one dimension reduces to a profile of temperature and humidity) which minimises a pre-defined "cost" or "penalty" function. In its simplest form this function comprises two terms. The first is a background term which is a measure of how well a particular profile agrees with our prior or background estimate of the atmospheric state, inversely weighted by the expected error in the background profile. The second is an observation term and measures how well radiances evaluated from a particular profile agree with those actually measured by the satellite. This observation term also has an inverse weighting based on expected errors in the measurements and the expected error in the profile to radiance mapping (i.e. the radiative transfer model).

The combined background and observation costs are simultaneously minimized with respect to every element of the atmospheric profile. A number of methods exist for finding the minimum of such a cost function and the TOVS 1DVAR scheme uses that of Newtonian iteration given by *Rodgers* (1976) which is practicable due to the relatively small order of the profile vector (about 60 elements). The profile that minimizes the cost function is "optimal" in that it is the most probable on the basis of both the background information and the measured radiances. However, to ensure this, an accurate specification of the inverse weighting factors (i.e. the error covariances of the background profile, measured radiances and forward model) is crucial (see *Watts and McNally* 1988; *McNally* 1990). These ensure that the solution does not draw too closely to the radiance data when it is noisy or difficult to simulate, and that the solution is not over-constrained by its agreement with a background profile in which we have little confidence.

Implementation

The TOVS brightness temperatures used are (at present) the global, cloud-cleared data generated by NOAA/NESDIS and available in Europe in near-real time as part of the "120 km BUFR TOVS" data set. These data have already undergone substantial pre-processing at NESDIS (see *Smith et al.*, 1979) followed by cloud-clearing (*McMillin and Dean*, 1982). The cloud-clearing route is identified with the data and can be either "clear", "partly cloudy" or "cloudy". The following TOVS channels are used at present (see fig.1 for weighting functions): HIRS channels 1-7 and 10-15 and MSU channels 2-4 are used for "clear" and "partly cloudy" soundings; HIRS channels 1-3 and MSU channels 2-4 are used for "cloudy" soundings.

The vertical atmospheric profile is represented by the temperature at 40 pressure levels (1000 to 0.1 hPa), the humidity at 15 pressure levels (1000 to 300 hPa), surface air and skin temperatures, and surface pressure. The background estimate of this profile is obtained from the NWP model. Nominally a 6-hour forecast is used: in fact forecast fields at 3-hour intervals are interpolated quadratically in time and bilinearly in space to the location of each TOVS sounding within the "analysis window", which contains all the data between +3 and +9 hours from the previous analysis.

Systematic errors in the simulated brightness temperatures and those measured by the satellite can be significant and great care is taken to remove them. A spatially-varying bias correction is currently used which is based on air-mass prediction by the microwave channel measurements. A full description of the bias correction scheme is contained in *Eyre* (1992).

From the 1DVAR "retrievals" which pass the internal quality control, a further selection is made to reduce the data density to a spacing of about 250 km. Priority is given to soundings over sea before those over land. Clear, partly cloudy and cloudy soundings are selected preferentially in this order. Further technical details are given in *Eyre et al.* (1992).

Interface to the OI analysis

Previous work at ECMWF (*Kelly and Pailleux*, 1988) has shown that NESDIS temperature/humidity retrievals could be assimilated to best effect if used on a limited number of thick layers compatible with the information content of the data. Layer-mean virtual temperatures are presented to the OI analysis on 7 layers: 1000-700, 700-500, 500-300, 300-100, 100-50, 50-30, 30-10 hPa. Layer water vapour contents are presented on the 3 lowest layers. For technical reasons, it has been easiest to use the 1DVAR products within the OI analysis system by integrating the profiles to the same layers and replacing the NESDIS-derived quantity with the 1DVAR equivalent. Since at present the 1DVAR does not use any SSU data (see fig. 1), it does not provide significant information in the top layer, 30-10 hPa. Consequently the NESDIS product for this layer is not replaced.

The 1DVAR may be considered as the first stage in the assimilation of radiance data into the NWP fields, in which the radiance information is projected on to a set of vertical levels at each observation point. The second stage, performed by the OI analysis scheme, involves the projection of information on to the analysis variables. A problem which arises in such a two-stage analysis is the correlation of error between the output of the first stage and the background for the second. This should be allowed for during the second stage. The problem is analogous in principle to the "super-ob" problem described by *Lorenc* (1981). The solution proposed by *Lorenc et al.* (1986) has been adopted which in practice involves amplifying the 1DVAR first-guess increments, but giving them less weight in the analysis.

On 23 June 1992 the 1DVAR products were introduced into the operational OI analysis at ECMWF in the following configuration (on the basis of results from a number of forecast impact experiments):

a) Northern Hemisphere extra-tropics:
NESDIS products have been removed except for the layer 30-10 hPa. 1DVAR products have replaced NESDIS products in the layers between 100 and 30 hPa. In addition, 1DVAR products have been introduced over the sea for the layers between 1000 and 100 hPa (from which NESDIS products were removed in May 1991);

b) Tropics (defined by 20N-20S for clear and N*, 30N-30S for cloudy soundings):
No sounding data are used here, i.e. no change;

c) Southern Hemisphere extra-tropics:
NESDIS products continue to be used above 100 hPa and over the sea below 100 hPa, i.e. no change.

1DVAR Performance

The development and testing of the 1DVAR scheme has been performed mainly through forecast impact studies. In each experiment, 1DVAR products have been assimilated into the NWP system for several days. After each day of assimilation, a 10-day forecast has been run and compared with a verifying analysis. Corresponding control forecasts have been run from assimilations without 1DVAR products. In all, over 100 days of assimilation have been completed from four different periods.

Two pre-operational trials were conducted, both using the model at T213L31 for assimilation followed by forecasts at T106L19. Forecasts run at T106L19 from the operational analyses (T213L31) were used as controls (labelled "OPS 106"), and operational analyses provided verification. The first trial was run for the period 7-12 April using model cycle 41, and the second for 15-21 May using cycle 42. Six days of assimilation followed by six 10-day forecasts were run in each case.

Results are shown in Fig. 2 for the two periods combined (12 cases). The impact of the 1DVAR products is strong and favourable. For the northern hemisphere as a whole (Fig. 2a), the mean impact is clearly positive from day three. Over North America (Fig. 2b), the impact is strong and appears at day two in the mean scores with some large improvements, even at day three. Over Europe (Fig. 2c), the signal is also positive, although the impact tends to emerge later in the forecast, around day four or five. This is consistent with the fact that the largest analysis changes are observed over the Pacific, and so their effect is first felt over N America and over Europe only later in the forecast period.

Some experimental forecasts were also run from the same analyses at T213L31 and compared with operational forecasts. The impact of the 1DVAR on the forecast skill was not significantly affected by the change in forecast resolutions for the cases investigated.

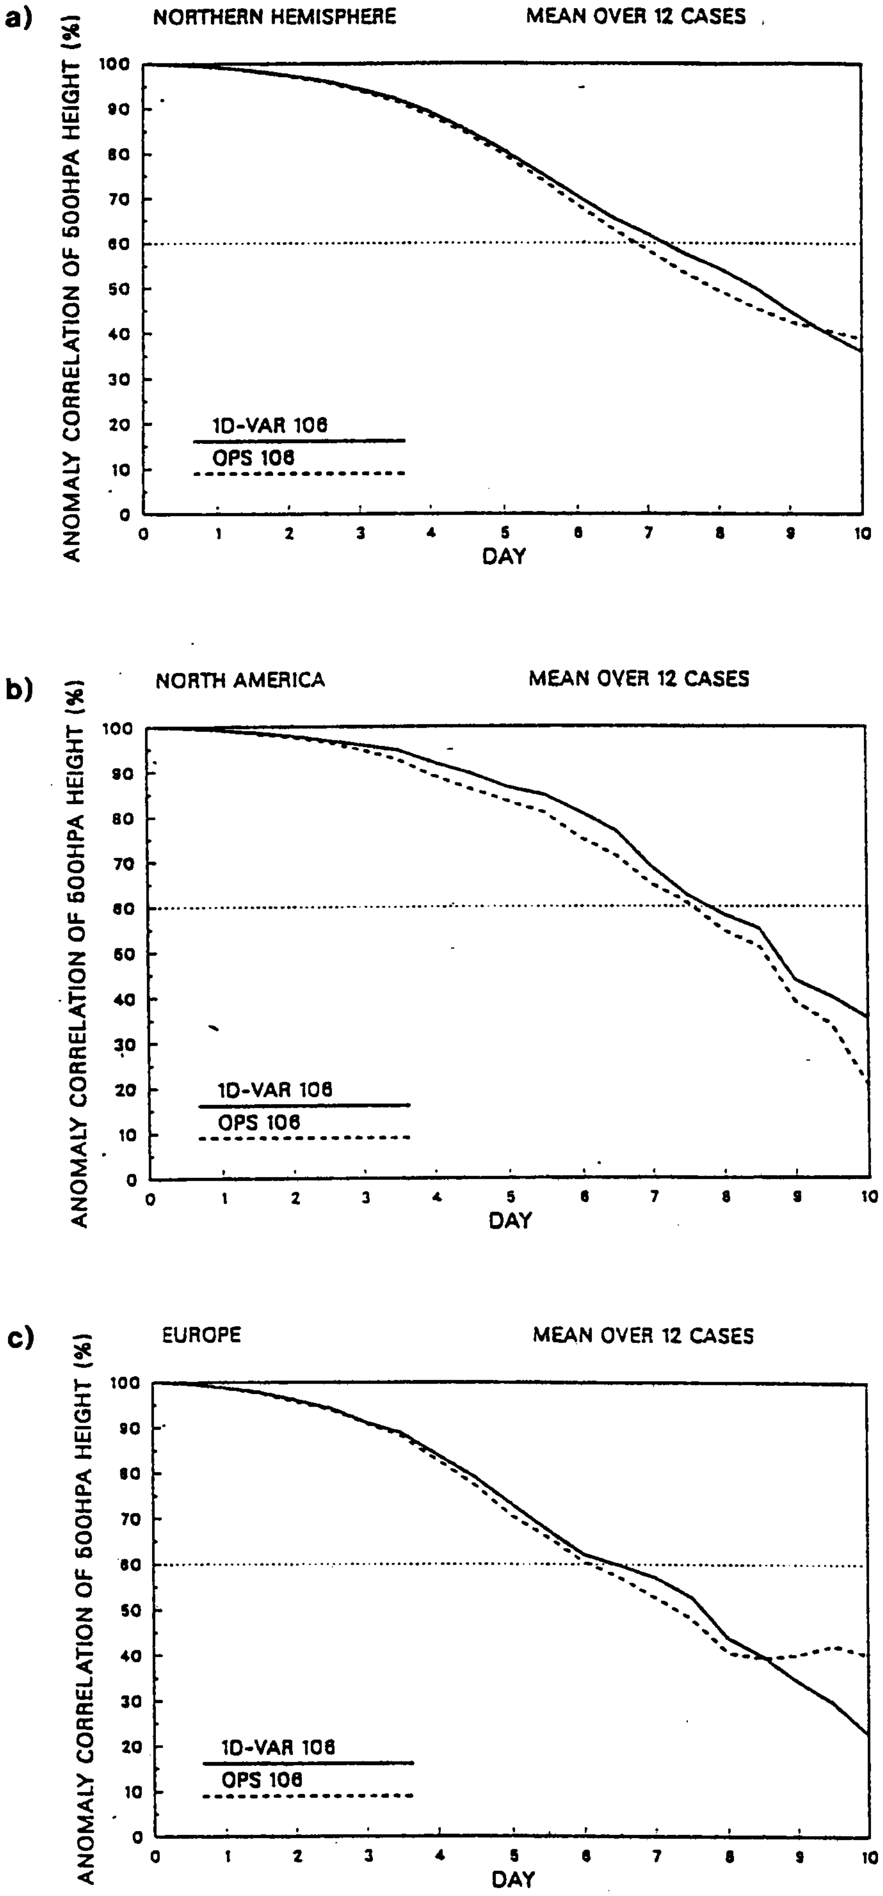

Fig. 2 Results of 1DVAR pre-operational trials in April/May 1992.
Mean anomaly correlations for 500 hPa height for (a) northern hemisphere, (b) North America and (c) Europe and the Mediterranean.

Future Developments

The impact in winter and spring on northern hemisphere forecasts was considered sufficiently clear and consistent to merit operational implementation of the 1DVAR. However, it is recognized that aspects of the scheme's performance have not been investigated and there are several areas in which the scheme can be extended and improved.

When real-time availability of data improves, implementation of 1DVAR in the southern hemisphere will be considered. Technical changes at NOAA, Washington DC, should make this possible by the end of 1992. In order to demonstrate positive impact of 1DVAR in the southern hemisphere, it may be necessary to improve the performance over sea-ice, particularly in the winter. A simple heat balance model has recently been implemented for sea-ice which should improve this aspect of the model "background" and thus assist the performance of the 1DVAR.

There is clear evidence that the humidity information in the 1DVAR products should be beneficial in the tropics, at least in improving the mean analyzed fields (see *Eyre*, 1992). However, experiments so far suggest that the inclusion of 1DVAR products in the tropics would have a small adverse effect on forecasts in the extra-tropics. To improve this, a number of aspects deserve attention including the weights given to 1DVAR products in the OI and the specification of the vertical correlation of forecast error.

There is no reason in principle why 1DVAR products should not also be assimilated over land. Positive impact might be expected in data sparse areas.

Now that 1DVAR is running routinely, the associated monitoring systems are providing a wealth of information on its performance. This will be used to re-tune all the statistical parameters and those in the interface to the OI analysis. When the re-tuning is complete, extension of the scheme to areas other than the northern hemisphere extra-tropical oceans will be considered.

- Tony McNally

REFERENCES

Andersson E, A Hollingsworth, G Kelly, P Lönnberg, J Pailleux and Z Zhang, 1991. Global observing system experiments on operational statistical retrievals of satellite sounding data. Mon Wea Rev, 119, 1851-1864.

Bengtsson L, 1985. Medium-range weather forecasting at ECMWF. Adv Geophys, 28B, 3-54.

Durand Y, 1985. The use of satellite data in the French high resolution analysis. Proc ECMWF Workshop on "High Resolution Analysis"; Reading; 24-26 June 1985; ECMWF Report, pp.89-127.

Durand Y and R Juvanon du Vachat, 1985. Developments in mesoscale analysis using raw satellite data. Tech Proc 2nd Int TOVS Study Conf; Igls, Austria; 18-22 February 1985; Ed. W P Menzel; Report of CIMSS, University of Wisconsin-Madison, pp.39-57.

Eyre J R, 1989. Inversion of cloudy satellite sounding radiances by nonlinear optimal estimation. Q J R Meteorol Soc, 115, 1001-1037.

Eyre J R, 1992. A bias correction scheme for simulated TOVS brightness temperatures. ECMWF Tech Memo 186.

Eyre J R, G A Kelly, A P McNally and E Andersson, 1992. Assimilation of TOVS radiances information through one-dimensional variational analysis. ECMWF Tech. Mem. 187.

Eyre J R and A C Lorenc, 1989. Direct use of satellite sounding radiances in numerical weather prediction. Meteorol Mag, 118, 13-16.

Eyre J R, P D Watts, J Turner and A C Lorenc, 1985. Research and development on TOVS retrievals in the U.K. Tech Proc 2nd Int TOVS Study Conf; Igls, Austria; 18-22 February 1985; Ed. W P Menzel; Report of CIMSS, University of Wisconsin-Madison, pp.94-100.

Halem J, E Kalnay, W E Baker and R Atlas, 1982. An assessment of the FGGE satellite observing system during SOP-1. Bull Am Meteorol Soc, 63, 407-429.

Kashiwagi K, 1987. On the impact of space-based observing systems in the JMA global forecast/analysis system. J Meteorol Soc Jap, 65, 189-219.

Kelly G, E Andersson, A Hollingsworth, P Lönnberg, J Pailleux and Z Zhang, 1991. Quality control of operational physical retrievals of satellite sounding data. Mon Wea Rev, 119, 1866-1880.

Kelly G and J Pailleux, 1988. Use of satellite sounder data in the ECMWF analysis system. ECMWF Tech Memo 143.

Lorenc A C, 1981. A global three-dimensional multivariate statistical analysis scheme. Mon Weather Rev, 109, 701-721.

Lorenc A C, 1986. Analysis methods for numerical weather prediction. Q J R Meteorol Soc, 112, 1177-1194.

Lorenc A C, W Adams and J R Eyre, 1986. The analysis of high resolution satellite data by the Meteorological Office. Proc ECMWF Workshop on "High Resolution Analysis"; Reading; 24-26 June 1985; ECMWF Report, pp.69-87.

McMillin L M and C Dean, 1982. Evaluation of a new operational technique for producing clear radiances. J Appl Meteorol, 21, 1005-1014.

McNally A P, 1990. Satellite sounding of tropospheric temperature and humidity. Oxford University D.Phil Thesis (Department of Atmospheric Physics).

Pailleux J, 1990. A global variational assimilation scheme and its application for using TOVS radiances. Preprints WMO Int Symp on "Assimilation of observations in meteorology and oceanography; Clermont-Ferrand; 9-13 July 1990; WMO Report, pp.325-328.

Pailleux J, W Heckley, D Vasiljevic, J-N Thépaut, F Rabier, C Cardinali and E Andersson, 1990. Development of a variational assimilation system. ECMWF Tech Memo 179.

Rodgers C D, 1976. Retrieval of atmospheric temperature and composition from remote measurements of thermal radiation. Rev Geophys Space Phys, 14, 609-624.

Smith W L, H M Woolf, C M Hayden, D Q Wark and L M McMillin, 1979. The TIROS-N Operational Vertical Sounder. Bull Am Meteorol Soc, 60, 1177-1187.

Thoss A, 1991. Cloud motion winds, validation and impact on numerical weather forecasts. Proc Workshop on "Wind extraction for operational meteorological satellite data"; Washington D C; 17-19 Sept 1991; Eumetsat Report.

Turner J, J R Eyre, D Jerrett and E McCallum, 1985. The HERMES system. Meteorol Mag, 114, 161-173.

Watts P D and A P McNally, 1988. The sensitivity of a minimum variance retrieval scheme to the values of its principal parameters. Tech Proc 4th Int TOVS Study Conf; Igls, Austria; 16-22 March 1988; Ed. W P Menzel; Report of CIMSS, University of Wisconsin-Madison, pp.399-412.

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

SUPERVISOR MONITOR SCHEDULER (SMS)

Introduction

To use the Centre's available resources to best effect requires efficient scheduling of the operational work load. This, in turn, necessitates that the combination of batch jobs at any time forms a good mix, and that a complex set of dependencies between the jobs be respected. In other words, it is necessary throughout the course of the daily operations to start the right set of jobs at the right time; care must also be taken to ensure that all previous jobs on which the current set of jobs depend have been successfully completed.

The ECMWF Meteorological Operational System (EMOS) requires that over 650 batch jobs be run each day. These are made up of more than 250 separate jobs, some of which are run only once, and others many times. In addition, there are a number of jobs which must be run monthly, and some which must be run on specific days.

The need for automated facilities to schedule the submission of the operational tasks was recognised early in the Centre's development. The first implementation of a programmed "Supervisor Monitor Scheduler", or SMS, was based on a French system, then in use at DMN in Paris. Versions 1 and 2 of SMS were written in Cybil, and could run only within the CDC NOS/BE environment. When NOS/BE was phased out, assistance from Cray Research enabled a replacement SMS version 3 to be written in PASCAL, to be capable of operating on Cray or VAX/VMS systems.

In June 1991 a project was initiated to produce SMS version 4, written using the C language. The new system is now complete, has been tested, and was introduced into daily operations on 5 May 1992.

What is SMS?

Essentially, SMS is a scheduler. Using a simple definition language, it enables complex sets of tasks (computer jobs) to be defined, and the dependencies between them to be stated. It encourages the user to view the set of tasks to be performed as a tree-like structure, similar to the structure of a UNIX directory for a file system. All leaf nodes are tasks - computer jobs which, preferably, perform a single well-defined function. The root node is SMS itself. Nodes adjacent to the root are called "suites", while all remaining nodes are "families". This enables sets of related tasks to be identified as a family, and sets of related families to be identified as a suite; dependencies can then be defined at the level of suite, family, and task. Additional dependencies can be defined with respect to date, time, day of the week, periodic frequency, specific events set by tasks, and status of task termination (abort or normal end).

However, SMS is not just a scheduler. It retains information about the suites, families and tasks under its control which allows their progress to be monitored. This is especially important within an operational environment, as a full appreciation of the progress being made and of any problems as they arise is essential for reliability.

Having a good scheduler, supported by high quality monitoring facilities, is not sufficient. If the monitoring information indicates a need for action, then there is a basic requirement that such action should be supported. Thus SMS also enables supervisory intervention. Operators and authorised administrators are given facilities to alter status information, re-queue tasks, and to re-define (partially or fully) the workload being handled by SMS.

Why develop SMS version 4?

Version 3 of SMS was written in PASCAL; it handled inter-process communications through a file system; it worked reliably and well, but had been designed and implemented originally for the Cray Operating System (COS), and for VAX/VMS. Although it had been possible to port version 3 to UNICOS, it was felt that use of the C language, and use of Remote Procedure Calls (RPC) for inter-process communications would produce a more portable and more efficient product for future use within UNIX and VMS environments.

There were also a number of features which experience at ECMWF had shown would be very desirable additions to SMS. In particular, sufficient flexibility within a new version would allow the use of SMS to be extended to support Research Department experiments. The use of UNIX style commands, simple edit facilities, and environment variables would make the system easier to use, more standard in its notation, and less dependent on separate job scripts for each task. There seemed little point and much danger in maintaining many copies of a specific script, each with slight variations from the others. The provision of improved visualisation had also been an important long-standing requirement.

The SMS project

The principle objective of the project was to rewrite SMS in C and improve the inter-process communications. Additional resources, in the form of one man-year of consultancy, were committed. A set of requirements was drawn up, together with a timetable. The aim was to generate the new version, together with both a dumb terminal based "Command and Display Program" (CDP), and an X-based visualisation facility, XCdp. The development was scheduled to include as many of the proposed enhancements as possible, and to be ready for operational support in May 1992.

Otto Pesonen, a consultant from Finland, was responsible for SMS and CDP; Baudouin Raoult from the Meteorological Applications Section developed XCdp. By December 1991 a prototype system had been put together, and demonstrated internally. By May 1992 objectives had been met in full and on time; in addition, many features were invented and added during the development phase. It had also been possible to demonstrate the pre-operational system to the computer operators, obtain valuable feed-back, and incorporate many of the resulting ideas.

Version 4 of SMS, together with CDP and XCdp have been used operationally since May 1992; they have performed well, and have been extremely reliable. Follow-up work has led to the incorporation of a set of post-operational changes to correct problems detected in use; these were of such a minor nature that they had not affected the reliability of the operational service. The initial documentation has been prepared as a pre-print; it will now be reviewed, revised, and corrected, to be issued later in the form of Meteorological Bulletins. The unusual step has been taken of printing the pre-print versions because of the considerable demand for information, especially from Member States' meteorological services, experienced since the product went operational.

The visualisation of information

XCdp enables many aspects of the tasks controlled by SMS to be visualised. Features exist to display graphs of dependencies, to use colour to convey status (queued, active, aborted, completed, etc.), to trace and follow the timeliness of the tasks, and to access both dynamic and pre-prepared information. Such information can include the documentation of the tasks, giving easy access to suggestions for corrective action if problems arise.

It is beyond the scope of this article to present all of the features supported, but sample diagrams on the following two pages give the flavour of the product. These include:

- the Tree Window - this illustrates the composition of suites and families with respect to their constituent tasks, with colour being used dynamically to indicate status;
- the Triggers Window - indicating family and task dependencies, and enabling the user to visualise what needs to happen before a task is triggered;
- the Time Table Window - using timing information from previous runs this window shows which tasks normally run when, and for how long; the current time is shown as a vertical line, and colour is used to indicate status - it is thus possible to see which tasks are running early or late;

TREE WINDOW

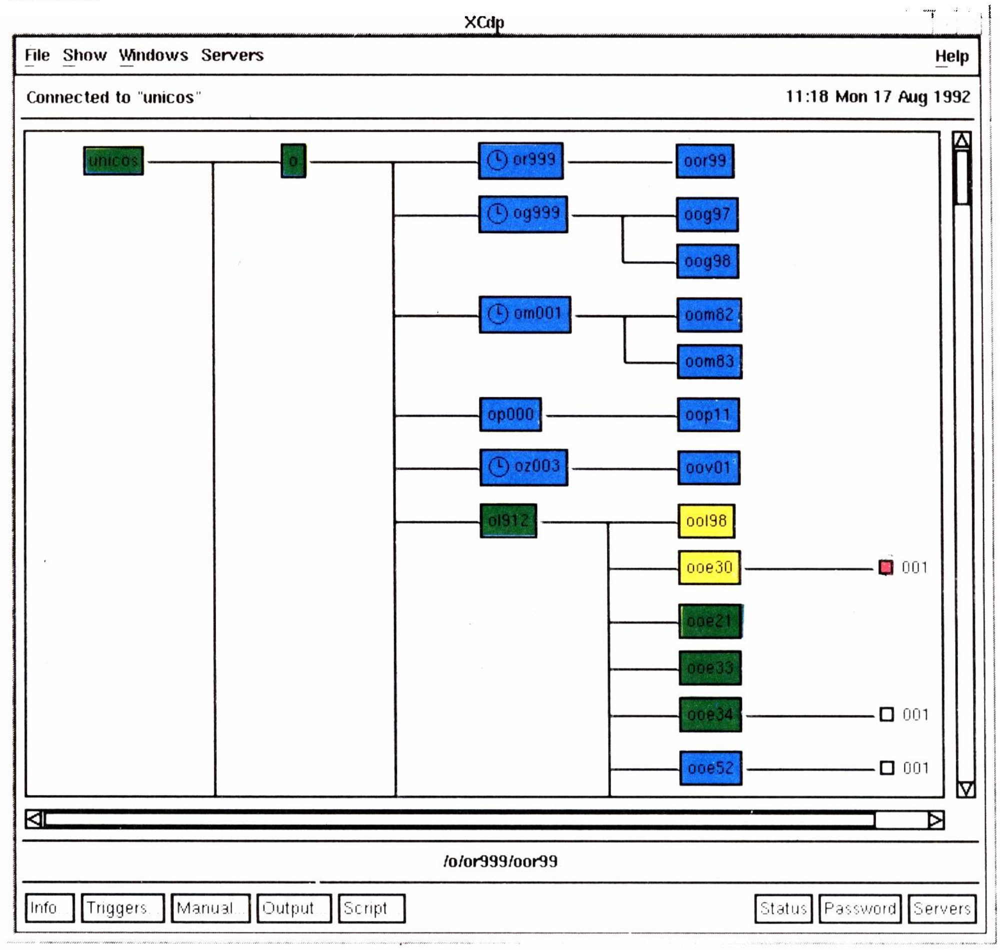

WHY WINDOW

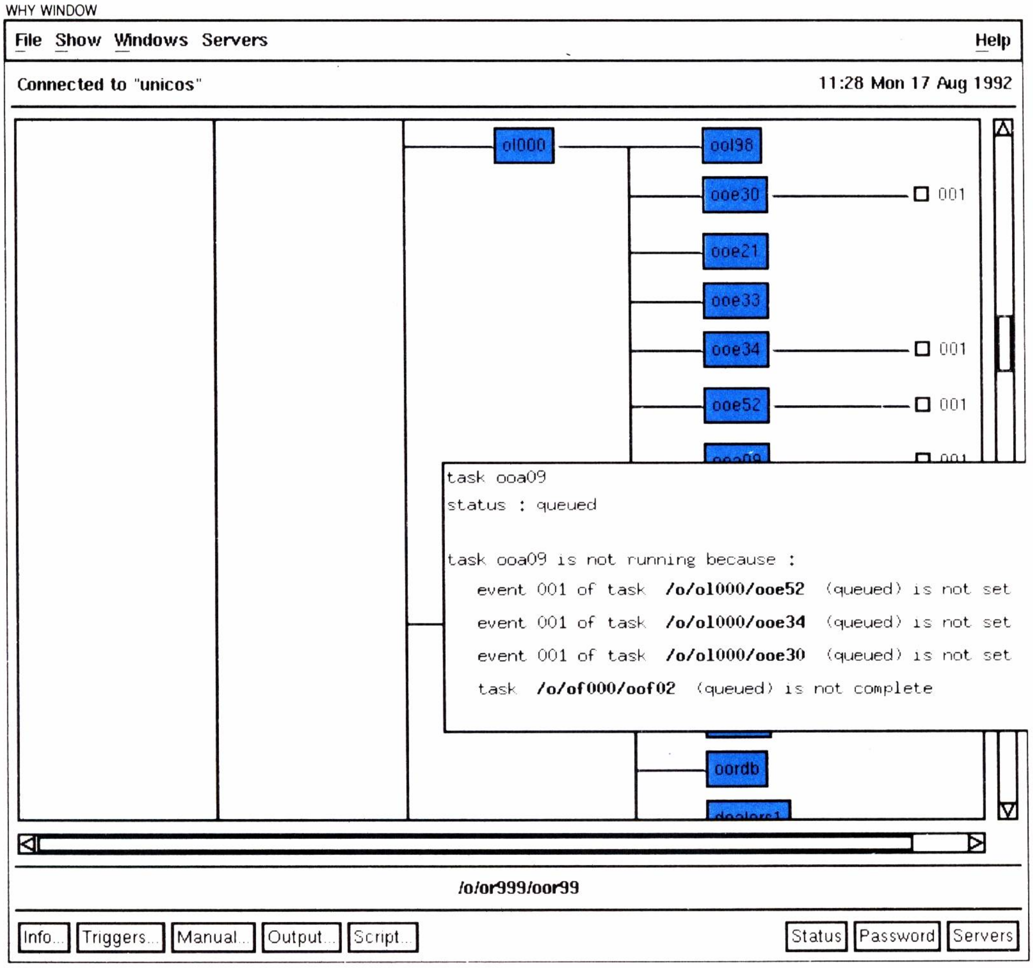

TRIGGER WINDOW

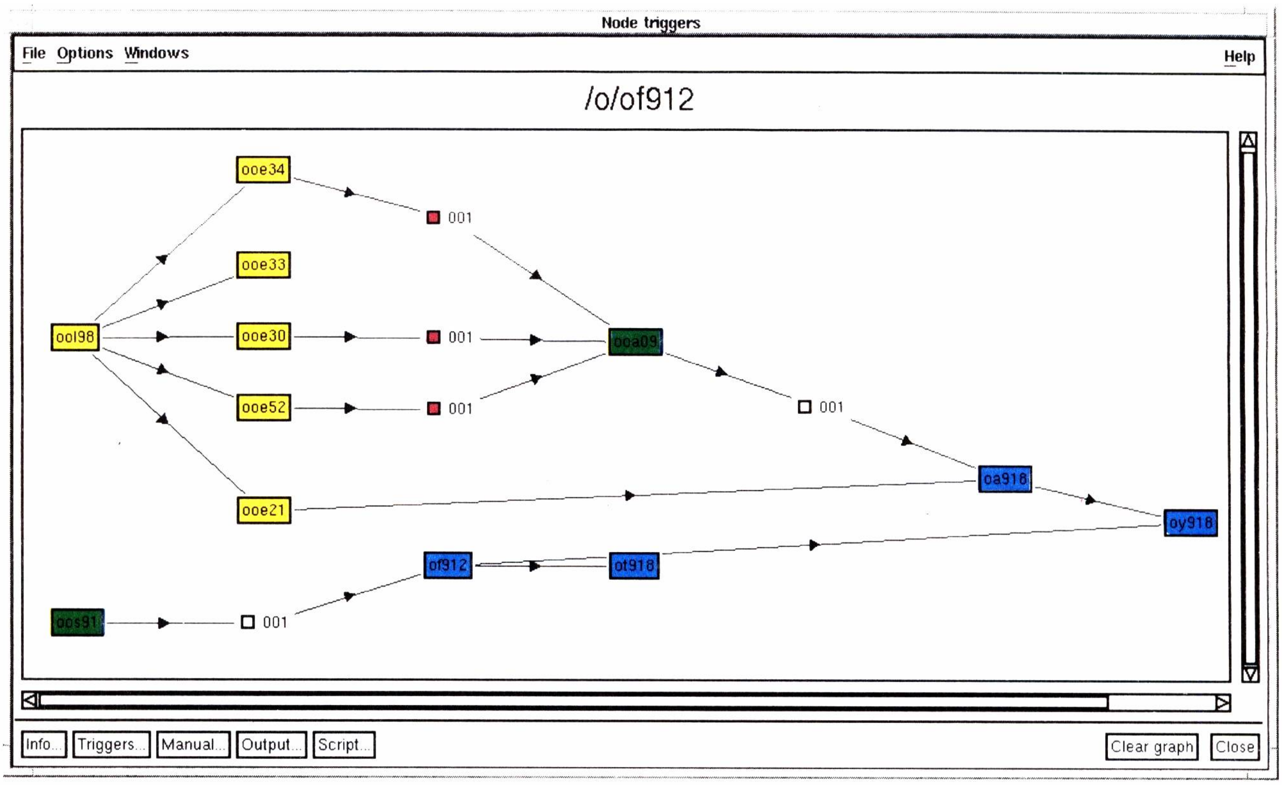

TIME TABLE WINDOW

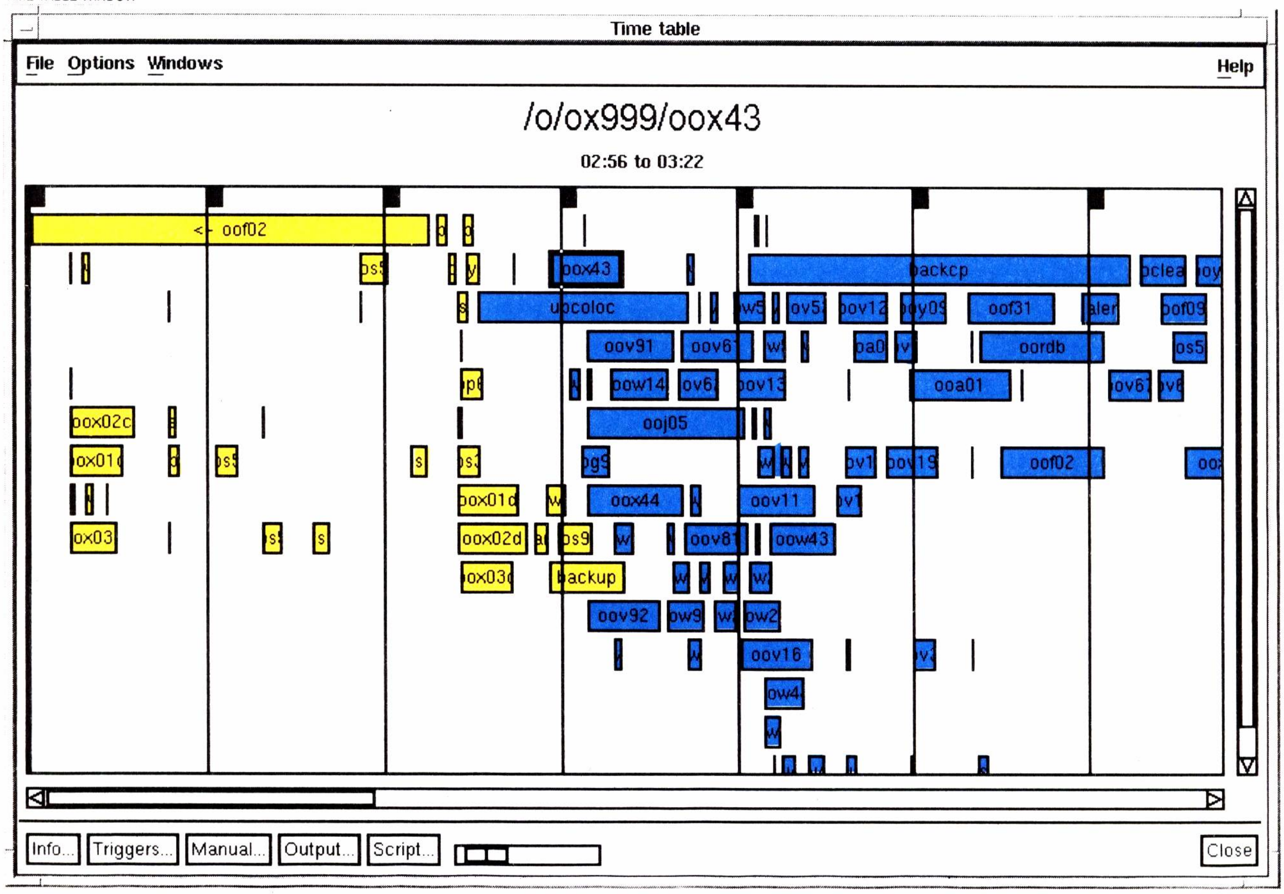

the Why Menu - by using the mouse to position the cursor over specific tasks on selected windows (e.g. the Tree Window) it is possible to click into the Why Menu; this gives access to information explaining why a queued task has not yet been started.

Characteristics of SMS

SMS solves scheduling problems of a general nature. It enables the user to run large numbers of programs which have mutual dependencies, together with dependencies on date, day, time, etc., in a controlled environment, and across a network of machines. It is thus of interest to Member States, and to external organisations.

The following list of technical characteristics of the new system may be of interest:

a) SMS V4 is written in standard C. It has been compiled on a number of systems to verify its portability (CDC, DEC, SUN, Cray, Convex, IBM, Silicon Graphics, etc.). Although intended primarily to support process control in a UNIX environment, facilities for VAX VMS have been included;

b) Inter-process communications are based on SUN Remote Procedure Calls (RPC), which currently use TCP/IP and UDP/IP. UDP is used only for broadcast facilities;

c) Batch jobs to be controlled by SMS are called "tasks". Traditionally, ECMWF has used the concept of a "suite" comprising a set of "families", each of which consists of a set of related "tasks". The purpose of SMS is to enable such structures to be described, scheduled using a rich diversity of dependency information, monitored, and supervised dynamically; all processing is automatic by default, but operator supervision and re-scheduling is available where necessary;

d) Many suites can be controlled under one SMS; also several versions of SMS may run at the same time; SMS V4 allows more structure than previous versions - it is possible to have "super families" between suite and family - the philosophy of the UNIX file directory structure has been followed;

e) A suite does not have to complete each day; it can use a wall clock (ie actual date/time) or a hybrid clock (real clock but fixed date until changed). No suite begins until commands have been sent to SMS for the suite to be defined and started - such commands may come from a batch job (e.g. of a previous suite) or from interactive operator commands;

f) SMS has an edit facility; at ECMWF a standard set of environment variables is used for operational tasks. Values related to date, time, operations, test operations, etc. are set. Using the edit facility any environment variable can be altered to a specific value for the duration of a suite, a family, or a task;

g) SMS helps in terms of enabling mixtures of jobs to be designed with the load capacities of specific machines in mind. This is a function of suite design and implementation. SMS does not, however, communicate with the queuing systems in order to take specific avoidance action in this respect;

h) There are extensive checkpointing and recovery procedures within SMS; these can be configured in several ways. Thus, checkpoints can be taken at selected intervals, or after every change of state. Checkpoints are additionally backed up, providing considerable recoverability after computer outages;

i) A command and display program (CDP) enables the visualisation of systems controlled by SMS, supports access to information concerning possible recovery actions, and provides the interface by which commands are passed to SMS. CDP is text based, and can be invoked from a terminal or terminal emulator running on a PC, not capable of X window support;

j) An X-based visualisation program (XCdp) is also available; so far it has been implemented on various platforms with MOTIF including SUN, Convex, Silicon Graphics, VAX Ultrix, and IBM AIX. It enhances the visualisation of systems controlled by SMS, illustrates aspects of systems being controlled in a graphical manner, enables icon and mouse based interaction, and provides a superior interface to the operators and users using SMS. The functionality of CDP is available from within XCdp.

- Rex Gibson

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

NEW LOCAL COMPUTER NETWORKS AT ECMWF

ECMWF has used computer networks in its operations for many years: its very first computer installation at Reading included inter-computer communications between the Cray, the CDC Cyber, and the Regnecentralen telecommunications machine. In 1983 Control Data's LCN/RHF network was added, which eventually connected the IBM, CDC, VAX and Cray; and in 1985 Ethernet was added, at first for the VAXes, and progressively for other machines thereafter.

Times and data rates move on, and what seemed a generous bandwidth in 1985 became at first merely adequate, and then a bottleneck. The 10 megabits-per-second capacity of the Ethernet can today, in principle, be consumed by a single Unix workstation. Even a PC network server can generate enough traffic to flood the Ethernet.

Fortunately, the technology of local area networks is moving almost as fast as that of computer systems. In recent years, standard networks have been developed which have ten or a hundred times the capacity of the Ethernet.

ECMWF is installing networks based on two of these newer technologies - FDDI and HiPPI.

FDDI

The most popular of the new networks is known as FDDI (Fibre Distributed Data Interface). It can transmit a maximum of 100 megabits per second between one network station and another, using light in optical fibres instead of electrical signals in copper wires.

The optical fibres used are made of glass, with a diameter of 125µ - about the same as a human hair. Within the fibre is an inner core of half this diameter, down which the light travels. The fibres are quite delicate, and need protection in several layers of plastic sheathing.

A piece of equipment connected to FDDI needs to have two fibres brought to it - one bringing incoming light (and therefore data), and one carrying outgoing light and data. The optical connections are made through a special plug and socket arrangement, which takes the polished, cut ends of the fibres and holds them in close alignment with the receiving or transmitting device.

The light used by FDDI is invisible - a wavelength of 1.3µ in the long infra-red (visible light has a range of wavelengths around 0.6µ). This wavelength is chosen because the transmitting and receiving devices can be made to operate efficiently, and because the loss of light during its passage through the glass fibre is lower than for shorter wavelengths.

An FDDI network is always constructed as a ring, in which each unit receives from its 'upstream' neighbour, and transmits to its 'downstream' neighbour. A packet of data may have to be passed many times from hand to hand around the ring before it reaches its destination. The ring must be complete in order to work, since the destination also continues to pass the data on around the ring, until it finally returns to the transmitter. The transmitter will see, from a tag placed in the data by the receiver, whether it was correctly received.

This means that the FDDI is critically dependent on every unit working correctly - if even one of them stops handing on packets, the ring is completely useless.

To overcome this, equipment with only a single pair of FDDI fibres is connected to an FDDI concentrator - a tiny FDDI ring in a single box, with electronic switching to include or exclude loops of fibre running out to separate devices. If the device is functioning correctly, the concentrator will switch it into the ring; if it fails, the concentrator will remove it and re-connect the ring around the failing device.

Some devices attached to FDDI have two sets of fibre connections, rather than one. This allows them to be chained together directly into a ring which will continue to work even if one of the devices, or connections, fails. The spare pair of fibres is used to complete the ring in this case.

ECMWF has just accepted equipment from Network Systems Corporation which connects the Cray C90 and the IBM ES9000 to the FDDI. Also on FDDI is the CDC 4360 server. During 1992, all the workstation servers and the VAX cluster will be attached to the FDDI network, providing a high-performance network service between the key machines in the machine hall.

Over the next few years, there will be a need for FDDI levels of performance in other areas of the network, even down to office workstations. ECMWF is not planning to install optical fibres to offices, however; fibre and its connections are still quite expensive. Much development work is being done in the networking industry which should permit FDDI data rates on ordinary office-grade network cables. By the time the Centre needs a capacity of 100 megabits per second to the offices, it should be possible to use what is effectively high-grade telephone wire to carry it.

HiPPI

Although FDDI properly addresses the needs of computers in the middle range, it is unable to meet the demands for bandwidth between the highest performance machines. Standards and equipment for very high performance networks are still under development, but the connection standard known as HiPPI is quite stable, and is available on most top-end machines. In particular, the Centre's C90 and IBM ES9000 machines both have HiPPI connections.

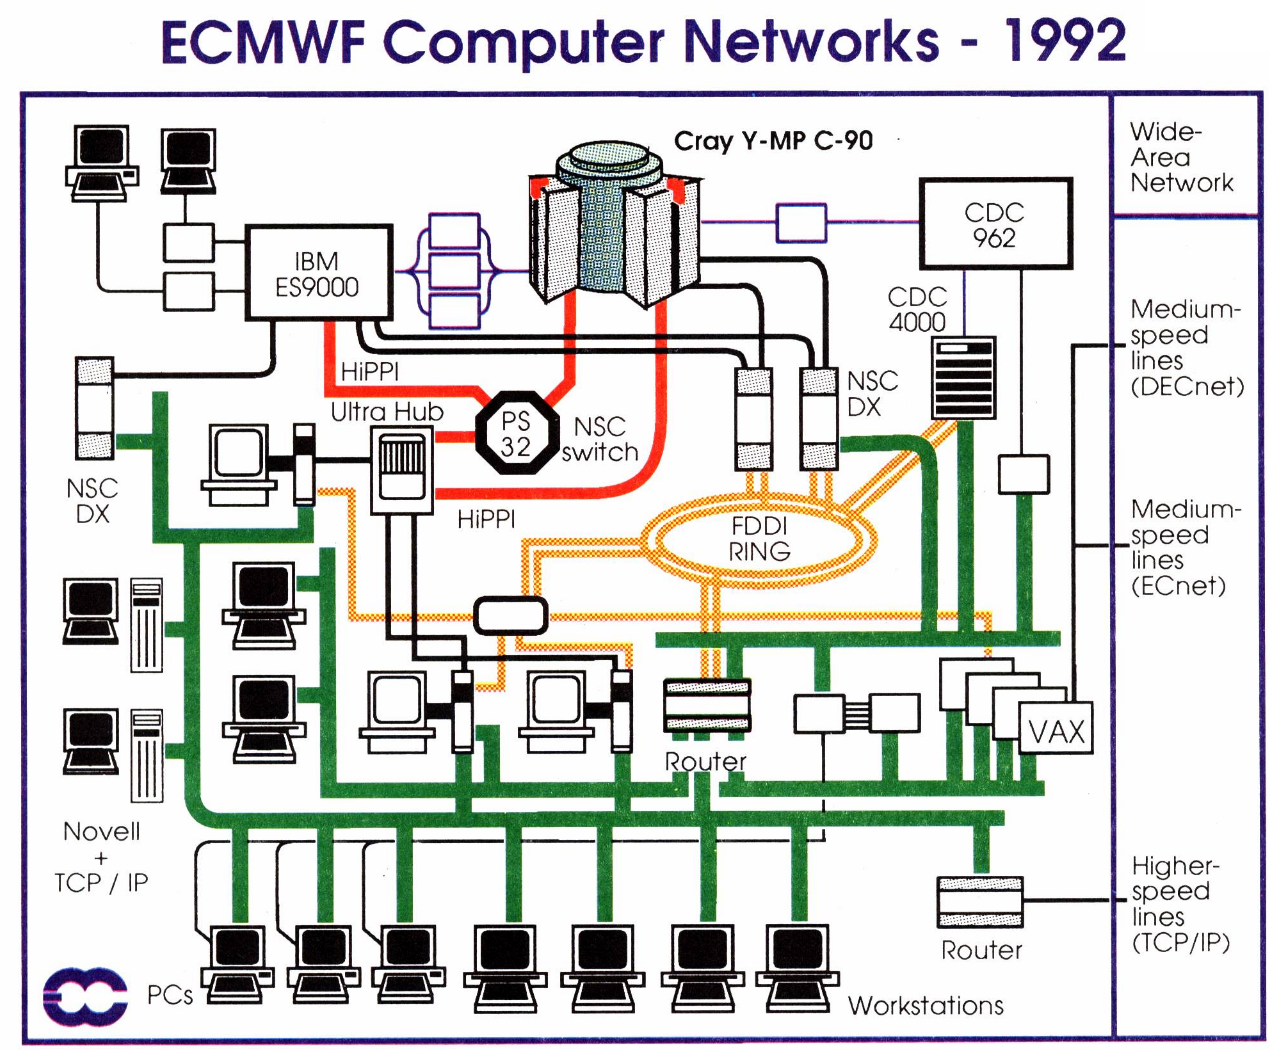

Fig. 1: The ECMWF Computer Networks, 1992

HiPPI at present is not really a network, but simply a connection standard. It was developed at Los Alamos National Laboratories in the USA, starting in 1987. There it had been found that connections at super-high data rates between a number of machines were needed, and at first it was suggested to Cray Corporation that the HSX channel (which Cray had developed for their own machines) should be standardised and published. As Cray were unwilling to allow this, Los Alamos went ahead and developed their own specification (which nevertheless retains a strong resemblance to the Cray HSX).

HiPPI specifies a channel which can carry up to 800 megabits per second, or up to 1600 megabits per second in the double-width version. It is now an approved ANSI standard, and is being adopted not only for high-performance computers, but also for peripherals: many I/O devices (especially disks) are now available with HiPPI interfaces.

The physical connection is carried on twisted-pair copper wiring, 32 bits wide. Each connection is unidirectional (SIMPLEX): a pair of connections is needed before machines can exchange data. The connection can only span 25 metres, which restricts its use to the computer room environment.

The HiPPI configuration at ECMWF is quite small. The IBM ES9000 has one HiPPI channel, and the Cray C90 has two. It is intended that one of the C90 channels should be connected directly to the IBM machine, and that the other should be connected via the Ultra Hub 1000.

In order to provide the necessary flexibility in this configuration, and to overcome some restrictions of the IBM HiPPI connection (which is not quite standard), a Network Systems HiPPI switch will also be installed. This provides a circuit-switched ('telephone exchange') interconnection between the different HiPPI channels. It could be expanded in future to allow the large machines access to big, high-performance disk arrays and other mass-storage technology.

Early tests with the HiPPI equipment so far installed at ECMWF indicate that an actual machine-to-machine data rate in excess of 300 megabits per second can be achieved. This is a good match for the requirements of the current generation supercomputer and file storage systems.

- Dick Dixon

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

STILL VALID NEWS SHEETS

Below is a list of News Sheets that still contain some valid information which has not been incorporated into the Bulletin set (up to News Sheet 286). All other News Sheets are redundant and can be thrown away.

No.	Still Valid Article
204	VAX disk space control
205(8/7)	Mispositioned cursor under NOS/VE full screen editor
207	FORMAL changes under NOS/VE
212	MFICHE command from NOS/VE
214	NAG Fortran Library Mark 12 News Sheets on-line
224	Job information cards
235	VAX public directory - how to create
236	Alternative VAX graphics service for in house users
248	Changes to the Meteogram system
253	Copying/archiving NOS/VE catalogs to ECFILE Copying complete UNICOS directories to ECFILE
254	UNICOS carriage control
260	Changes to PUBLIC directories for VAX users
261	Meteogram system on UNICOS

No.	**Still Valid Article**
265	Lost UNICOS outputs submitted via RJE or VAX Microfiche changes
266	Reminders on how to import/export magnetic tapes
267	Checking on your UNICOS account usage
268	Changes to WMO FM 92 GRIB
270	Changes to the Meteogram system
271	New ECFILE features on UNICOS
276	Periodic deletion of all Cray /tmp files
280	UNICOS on-line documentation: docview Further UNICOS differences
281	File transfer via FTP (possible problems)
283	New features for Member State batch users (RQS 1.1)
284	UNICOS 7 features & differences
286	C90 - Y-MP/8 compatibility

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

ECMWF CALENDAR 1992

28 - 30 Sep	Scientific Advisory Committee - 20th session
28 - 30 Sep	Member States' Computer Representatives - 7th meeting
30 Sep - 2 Oct	Technical Advisory Committee - 17th session
6 - 8 Oct	Finance Committee - 49th session
9 - 12 Nov	Workshop: Variational assimilation with emphasis on 3-dimensional aspects
23 - 27 Nov	Workshop: Parallel processing
2 - 3 Dec	Council - 37th session
24 - 28 Dec	***ECMWF holiday***

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

ECMWF PUBLICATIONS

ECMWF Annual Report 1991

Workshop Proceedings: Fourth Workshop on Use of Parallel Processors in Meteorology, 26-30 November 1990

Workshop Proceedings: Third Workshop on Meteorological Operational Systems, 18-22 November 1991

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

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in the ECMWF Newsletter series. As one goes back in time, some points in these articles may have been superseded. When in doubt, contact the author or User Support.

METEOROLOGY

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

USEFUL NAMES AND 'PHONE NUMBERS WITHIN ECMWF

		Room*	Ext.**
DIRECTOR	- David Burridge	OB 202	2001
HEAD OF OPERATIONS DEPARTMENT	- Michel Jarraud	OB 010A	2003
ADVISORY: Available 9-12, 14-17 Monday to Friday			2801
Other methods of quick contact:	- Telex (847908 ECMWF G)		
	- Telefax (+44 734 869450)		
	- VMS MAIL addressed to ADVISORY		
	- Internet mail addressed to Advisory@ecmwf.co.uk		
REGISTRATION			
Project Identifiers	- Pam Prior	OB 225	2384
User Identifiers	- Tape Librarian	CB Hall	2315
COMPUTER OPERATIONS			
Console	- Shift Leaders	CB Hall	2803
Reception Counter	- Tape Librarian	CB Hall	2315
Tape Requests	- Tape Librarian	CB Hall	2315
Terminal Queries	- Norman Wiggins	CB 026	2308
Telecoms Fault Reporting	- Michael O'Brien	CB 028	2306
ECMWF LIBRARY & DOCUMENTATION - DISTRIBUTION	- Els Kooij-Connally	Library	2751
LIBRARIES (ECLIB, NAG, etc.)	- John Greenaway	OB 226	2385
METEOROLOGICAL DIVISION			
Division Head	- Horst Böttger	OB 007	2060
Applications Section Head	- Rex Gibson	OB 101	2400
Operations Section Head	- Bernard Strauss	OB 004	2420
Meteorological Analysts	- Andreas Lanzinger	OB 003	2425
	- Ray McGrath	OB 005	2424
	- Anders Persson	OB 002	2421
Meteorological Operations Room	-	CB Hall	2426

		Room*	Ext.**	Bleeper
COMPUTER DIVISION				
Division Head	- Geerd-R. Hoffmann	OB 009A	2050	150
Systems Software Sect.Head	- Claus Hilberg	CB 133	2350	115
User Support Section Head	- Andrew Lea	OB 227	2380	138
User Support Staff	- Antoinette Alias	OB 224	2382	154
	- John Greenaway	OB 226	2385	155
	- Norbert Kreitz	OB 207	2381	156
	- Dominique Lucas	OB 206	2386	139
	- Pam Prior	OB 225	2384	158
Computer Operations Section Head	- Peter Gray	CB 023	2300	114
Security, Internal Networks and Workstation Section Head	- Walter Zwieflhofer	OB 142	2352	145
GRAPHICS GROUP				
Group Leader	- Jens Daabeck	OB 016	2375	159
RESEARCH DEPARTMENT				
Head of Research Department	- Anthony Hollingsworth	OB 119A	2005	
Computer Co-ordinator	- David Dent	OB 123	2702	

---

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (0734) 499000, international +44 734 499000, or direct dial to (0734) 499 + last three digits of individual extension number, e.g. the Director's direct number is (0734) 499001.

DEC MAIL: Contact scientific and technical staff via VMS MAIL, addressed to surname.

Internet: The ECMWF address on Internet is **ecmwf.co.uk**
Individual staff addresses are firstname.lastname, e.g. the Director's address is David.Burridge@ecmwf.co.uk