European Centre for Medium Range Weather Forecasts

ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England. Tel: U.K. (0734) 876000, Int. (44 734) 876000, Telex: 847908

Number 27 - September 1984

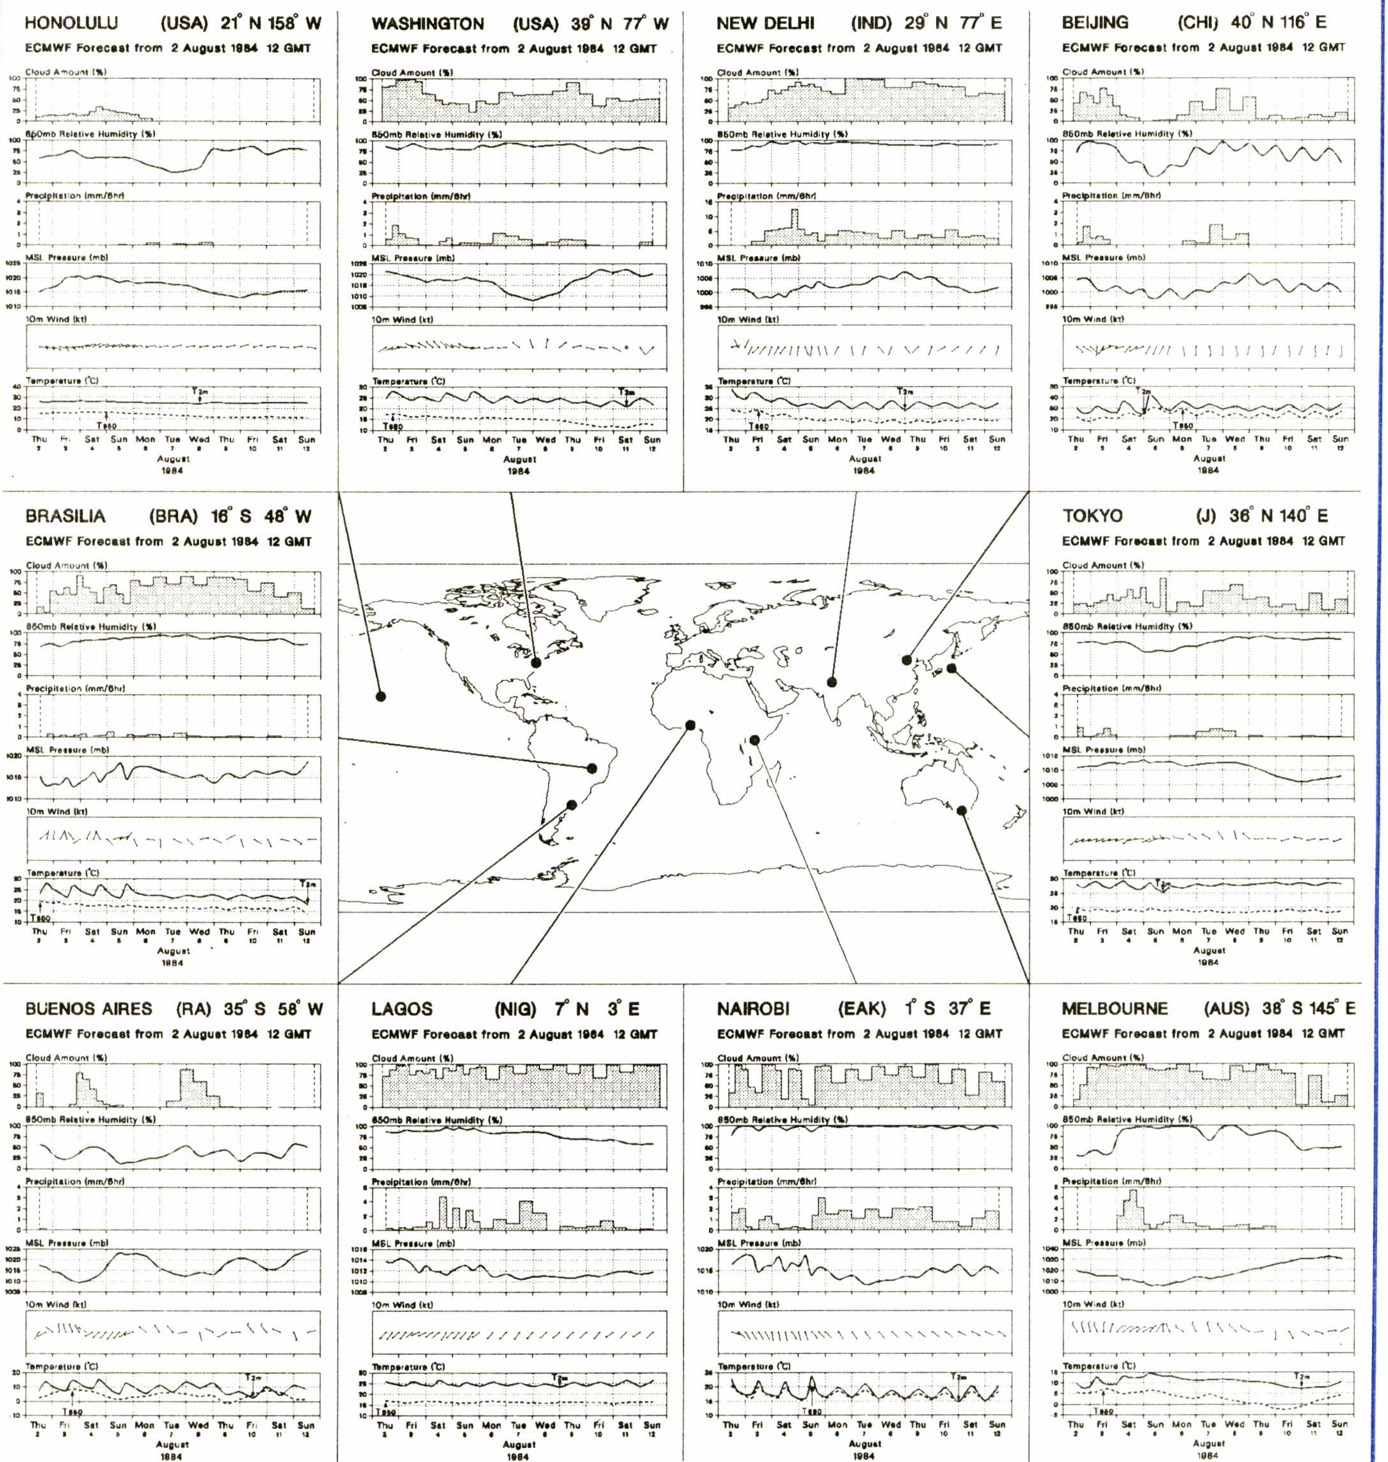

Page

IN THIS ISSUE

COVER: A set of 10 meteograms showing the evolution of model predicted meteorological variables (cloud, humidity, precipitation, pressure, wind and temperature) with forecast time. These values are interpolated from the 4 nearest grid points. Precipitation is accumulated in time, the other parameters are instantaneous values.

This Newsletter is edited and produced by User Support.

The next issue will appear in December 1984.

ECMWF'S METEOROLOGICAL OPERATIONAL SYSTEM

ECMWF started medium-range operational forecasting in the second half of 1979. The analysis and forecasting systems are at the heart of the operation and there are supporting programs involving file handling, data processing and facilities for monitoring and scheduling the whole suite. The complete Operational Suite is a highly complex system of interacting and simultaneously running programs with asynchronous data flows.

The basic flow of meteorological information in ECMWF is similar to the flow of meteorological information which any meteorological centre has to manage as regards the computer system, though there are constraints imposed on the system by the scale of the forecasts that ECMWF carries out (i.e. 10 days over a global domain).

ECMWF'S Meteorological Operational System (EMOS) can be broken down into a number of logical sub-processes: data acquisition, pre-processing, analysis, forecast, post-processing and dissemination of results. Additional to these are such aspects as the real-time operational supervision, control and scheduling of all the programs in the system, the graphical display of results and systematic archiving of data. The analysis and the forecast need to take advantage of the computing power of the CRAY X-MP. The other sub-processes involved with data manipulation, file handling and maintaining data integrity are implemented on the CDC CYBERs. The CYBER-CRAY link handles data transfers for input to the analysis on the CRAY, and also the products generated during the analysis and forecast from the CRAY to the CYBER for postprocessing and dissemination. The various sub-processes and data flows involved in EMOS are illustrated in the diagram (Figure 1 overleaf).

Data Acquisition

This program is invoked every time a group of GTS (Global Telecommunications System) bulletins from Bracknell or Offenbach arrives on the CYBER via the Telecommunications Subsystem. A back-up copy of the bulletins is made and updated as each new file of data is received. Checks are made for missing files and requests for repeat transmissions can be made by the Operators.

Preprocessing

The main tasks of preprocessing are :

- controlling the input GTS data.
- creating a file of input data for the analysis.
- providing statistical information on coverage and quality of data received.

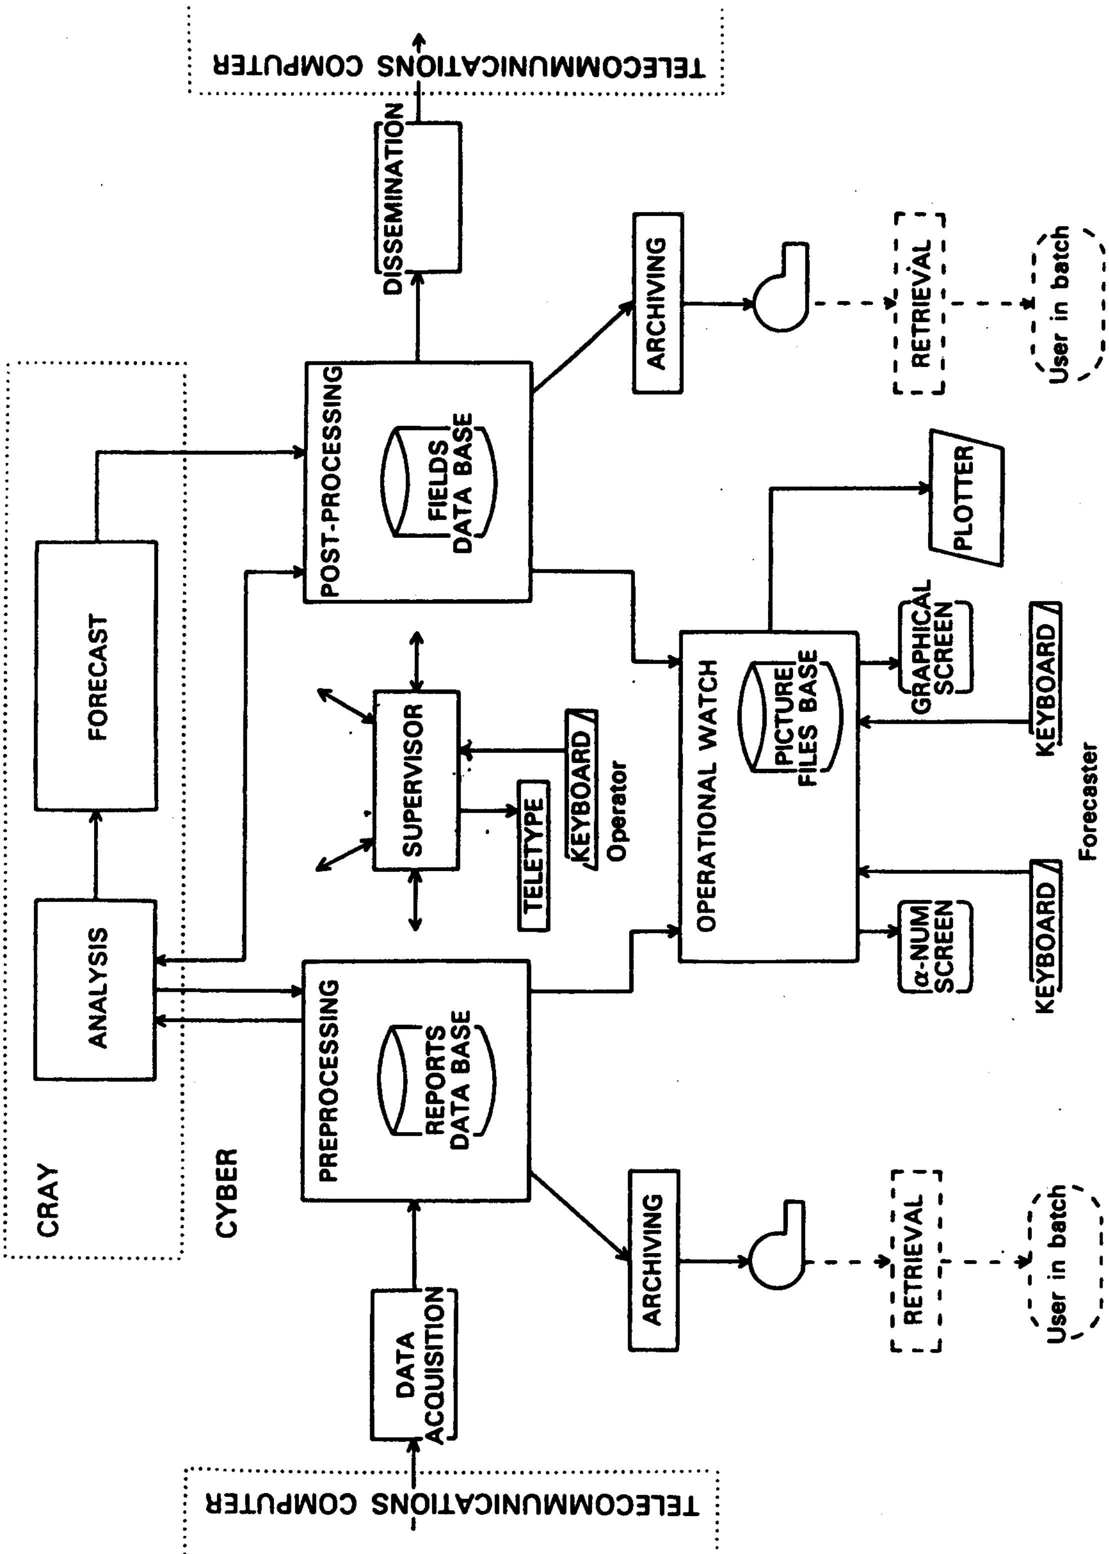

Figure 1 Diagrammatic illustration of the sub-processes and data flows in the Meteorological Operations System at ECMWF.

These three flows of information are asynchronous and are handled by independent systems centred around a data base of meteorological reports. When the data acquisition phase has accumulated a sufficient number of bulletins, the decoding program starts. The variable quality of GTS data implies a complicated program to recognise and decode the information from bulletins. Even with the most sophisticated software, a small percentage of the bulletins need to be rejected and written into a file for manual examination and correction by a meteorological analyst using an alpha-numeric VDU. A decoding program can be called to reprocess any corrected bulletins.

The decoding program keeps a record of input data in terms of statistics of the time of arrival, the presence and the quality of bulletins and reports. The statistics are saved for monitoring data receipt, and to influence the scheduling of data acquisition and preprocessing.

The last step before entering a report into the data base is the control of its quality. The internal consistency of each report is checked, and each meteorological parameter is verified according to climatological limits and, where possible, physical laws. The quality control program interrogates the data base to check for the presence of a matching observation. The program flags each suspicious meteorological parameter and can propose a substitute value; both the original and the substitution are kept in the data base. As the program adds or updates a report in the data base, it registers the time of arrival and the quality of the report to provide statistics for the meteorological analyst.

The reports data base is at the heart of the preprocessing and its efficiency is crucial since the execution of all the other subsystems depends on the ease and speed of access to its information. Data must also be read-accessible from any user's program.

The organisation and physical structure of the data base provides

- independence of a user's program from the physical location and file organisation of the data.
- control of access to data.
- centralised and easy back-up procedures.
- continuous availability of data on-line.
- possible asynchronous access to data.

<u>Analysis-Forecast</u>

The analysis scheme needs the reports organised in intervals of 6 hours. The following steps are necessary to provide a file of input data for the analysis for each 6 hour time interval:

- extraction from the data base of all the reports for the given interval.
- sorting to suit analysis requirements.
- formatting of variables for the analysis.

This input file is generated on the CYBER and passed to the CRAY for the analysis program. The forecast to 10 days follows the 12Z analysis starting around 2130Z and running for several hours. An important operational constraint for the analysis and forecast systems is their ability to be restarted after interruption or loss of data files. They are designed as a set of "black boxes" able to be restarted, able to know the status of their permanent files, able to reposition files and able to regenerate local files. Programs can receive or send permanent files from or to the CYBERS, as well as send messages to the central supervisor program. Some permanent files are sent to the CYBER for back-up on high density magnetic tapes in case of restart due to data loss on the CRAY. The checkpoints and copies of back-up files on tape are carefully organised to provide safety without cumbersome back-up or restart procedures in terms of overheads and operator actions.

Post-Processing

The post-processing handles, organises and distributes analysis and forecast results. Fields of data arrive on the CYBER from the CRAY computer at the end of designated analysis or forecast time-steps at 6- or 12-hour intervals. Large amounts of data are involved - each step passes over $3x10^6$ numbers from the CRAY and there are currently over 30 steps in a complete analysis/forecast cycle. The system has to transform and re-structure the data for visualisation on a screen, for automatic plotting of charts, and for transmission to Member States of about ten thousand fields of analysis and forecast every day. The GFDB (Global Fields Data Base) contains the output files from the analysis and forecast. The DDB (Dissemination Data Base) forms the base for staging fields for transfer to the telecommunications computer for transmission to Member States.

The history files from the analysis and forecast are retained as back-up of the GFDB.

Dissemination

A program is activated at the end of each post-processing sub-step to collect all the products for a timestep from the DDB to build up a file to be sent to each Member State. The file is forwarded to the telecommunications computer, with its destination indicated in a header.

Archiving

The archiving system runs every day to archive:

- the 3 day old reports data base files.
- the operational analysis files.
- the GFDB files containing the forecast and analysis results.

About 10 high density magnetic tapes are used every day to archive this data in duplicate 'A' and 'B' Tapes. The identification and the location of the files on the tapes is kept in a directory on a disk.

Supervisor

Within EMOS, many basic tasks are asynchronous (e.g. continued arrival of observational data, analysis and forecast fields produced at regular steps, dissemination of products following a predefined schedule). In addition to this, it must be possible to follow in detail the progress of the Operational Suite. A scheduler/monitor/supervisor (SMS) program which runs at a dedicated control-point on one of the CYBER computers is used to:

- schedule and initialise tasks.
- synchronise tasks.
- allow man-machine interaction.
- help restart procedures.

Files containing directives and job control cards scanned by the supervisor program control the progress of the Operational Suite. A basic principle used in SMS is that jobs within the Operational Suite can be grouped in "families". A "family" contains a set of jobs for which the executions are connected. Jobs belonging to different families are not necessarily linked. This concept of "family" assists in the management of the approximately 450 jobs which make up the Operational Suite. Families can be run independently, separately or in parallel and can be restarted. The supervisor program collects Operators' requests and program messages to centralise the management of the suite and to display supervisor messages.

Operational Watch

This is independent from the Operational Suite in terms of scheduling. This subsystem can provide information to the forecaster on duty while the Operational Suite runs, in reply to requests from an alpha-numeric or graphics terminal. An interactive utility analyses requests and calls overlays to display information on a screen or to plot it. Thus, the operational watch allows the data in the reports and post-processing data bases to be monitored.

- John Chambers

* * * * * * * * * * *

ECMWF GRID CODE PRODUCTS DISTRIBUTED VIA THE WMO GTS

The European Centre for Medium Range Weather Forecasts (ECMWF) has been distributing a selection of its output products via the Global Telecommunication System of WMO since 1 August 1981. The products currently disseminated comprise:

- mean sea level pressure and 500 hPa geopotential height analysis for 1200 GMT and derived forecasts with verifying times of +24, +48, +72, +96 and +120 hours. These fields are presented on a 5° x 5° latitude-longitude grid covering the Northern and Southern Hemispheres from the poles to latitude 20°.

- 850 hPa and 200 hPa wind analyses for 1200 GMT and derived forecasts with verifying times of +24 and +48 hours. These fields are presented on a 5° x 5° latitude-longitude grid covering a tropical belt from 35°N to 35°S.

The products are transmitted as a series of messages injected onto the GTS via the Regional Telecommunication Hubs (RTHs) at Bracknell and Offenbach. The products are issued from ECMWF during the course of the daily operational forecast run at the appropriate model timesteps, so that normally the analyses will be transmitted to the RTHs by about 2130 GMT and the 120 hour forecasts around midnight GMT.

The areas used for the dissemination products are shown schematically in the diagram (Figure 1 overleaf). The area limits are:

area	latitude	longitude	
1	20°N-90°N	90°W- 0°	Northern Hemisphere
2	20°N-90°N	180° - 90°W	Northern Hemisphere
3	20°N-90°N	90°E-180°	Northern Hemisphere
4	20°N-90°N	0° - 90°E	Northern Hemisphere
5	35°N-35°S	90°W- 0°	Tropical Belt
6	35°N-35°S	180° - 90°W	Tropical Belt
7	35°N-35°S	90°E-180°	Tropical Belt
8	35°N-35°S	0° - 90°E	Tropical Belt
9	20°S-90°S	90°W- 0°	Southern Hemisphere
10	20°S-90°S	180° - 90°W	Southern Hemisphere
11	20°S-90°S	90°E-180°	Southern Hemisphere
12	20°S-90°S	0° - 90°E	Southern Hemisphere

Each area is represented by grid point values given at the points of a regular latitude-longitude mesh with a 5°x5° interval. Values are given within a latitude row from left to right within the area, i.e. travelling eastwards along the row. In areas 01 to 04, latitude rows are given in order from south to north; in areas 05 to 12, rows are given in order from north to south. North and South pole values are represented by a single point and not a latitude row.

A fuller description of the ECMWF GTS products is given in ECMWF Meteorological Bulletin M1.8/1 - ECMWF grid code products distributed via the WMO GTS, J.D. Chambers, March 1984 (copies available on request).

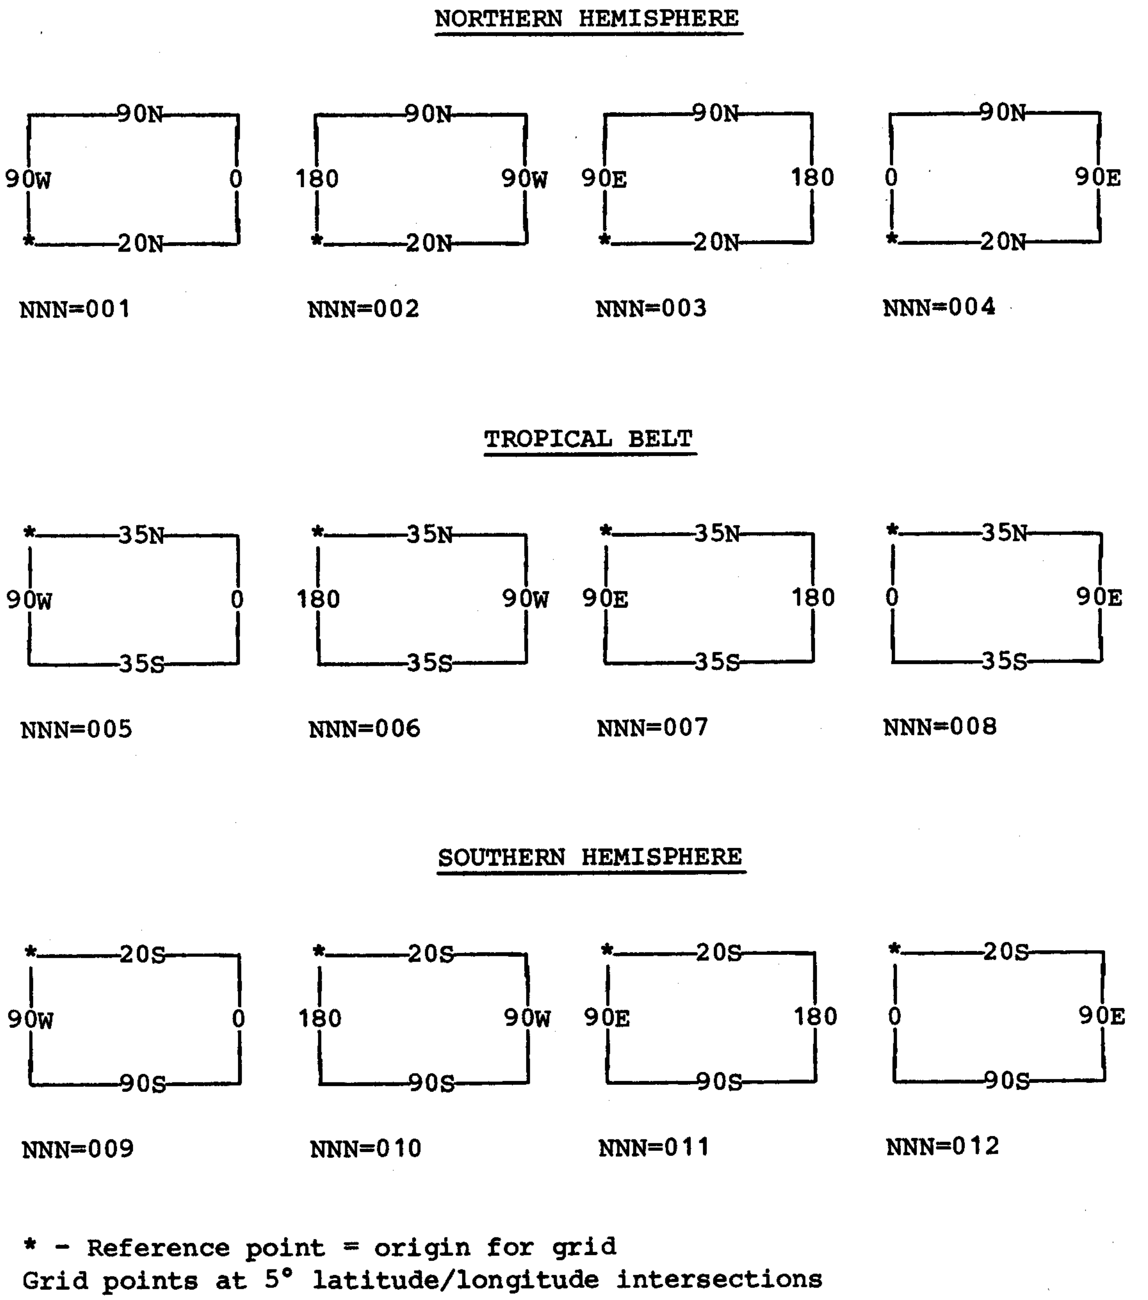

Figure 1: Product areas

- John Chambers

* * * * * * * * * *

THE IMPACT OF ANALYSIS DIFFERENCES ON MEDIUM-RANGE FORECASTS

A collaborative study has been carried out between the ECMWF, the U.K. Meteorological Office and the National Meteorological Centre (Washington) to examine the response of three advanced analysis systems to identical inputs of observational data, namely the FGGE IIb data. Amongst the many interesting results was a case which showed how the accuracy of the analysis over the northeast Pacific can be important for the success of medium range forecasts for Western Europe.

On 19 February 1979 a short upper trough approached the east coast of North America. At the surface, an extremely cold anticyclone began to move from the continent to the Atlantic. As the upper trough neared the coast, it triggered the development of a small but very intense system; this has become known as the President's Day storm.

Forecasts using the Centre's model were run from 00GMT 18 February with analyses from the ECMWF (EC) and the National Meteorological Centre (NMC). Fig. 1 shows the day 4 and day 6 forecasts along with the verifying analyses. Up to day 4, the behaviour of the storm was similar for both cases (compare Figs. 1c and 1e). However, by day 6, dramatic differences occurred (compare Figs. 1d and 1f). In one case, the low deepened in situ, while in the other, it moved rapidly northeast with much less deepening.

Examination of the differences between the forecasts indicated that they could be caused by differences in the EC and NMC analyses in the northeast Pacific. In this area, the main data sources are SATEM data (mainly microwave) along with surface ships, cloud wind data, AIREPS and ship PAPA (50°N, 145°W). Fig. 2 shows the EC and NMC analyses of the 500-700mb thickness at 00GMT 18 February; the observations are also shown. Clearly the NMC analysis has been faithful to the SATEM microwave reports, whereas the EC system has largely ignored them and drawn closely from the ship PAPA report. The differences between the analyses at the surface and near 250mb were less marked because of the use of the remaining types of data.

To test whether the differences in the analysis in this area were the causes of the differences in the behaviour of the President's Day storm, a transplant experiment was carried out: the NMC analysis in the northeast Pacific (covering 30° of latitude and 60° of longitude) was transplanted into the EC analysis and a forecast run from the resulting analysis. Fig. 3 shows the 300mb differences between the original EC based forecast and the forecast from the "transplant" dataset, at two day intervals (small differences outside the north Pacific have been excluded). The initial differences in the Pacific propagate downstream and amplify rapidly, reaching the mid-Atlantic by day 5. Clearly the differences between the forecasts from the EC and NMC analyses stem from analysis differences in the northeast Pacific.

These results illustrate some practical limits to predictability arising from analysis uncertainties. Analysis errors in baroclinically unstable regions are particularly important and their influence depends both on their magnitude and on the strength and direction of the flow in which they occur. The rapidity with which analysis errors can propagate downstream demonstrates the importance of having adequate accurate observations in baroclinically active regions for the success of medium range forecasts.

A fuller description of this case study and further results from the collaborative experiments can be found in the proceedings of the workshop on "Current problems in data assimilation" which have been published by the Centre.

\- Tony Hollingsworth

* * * * * * * * * * *

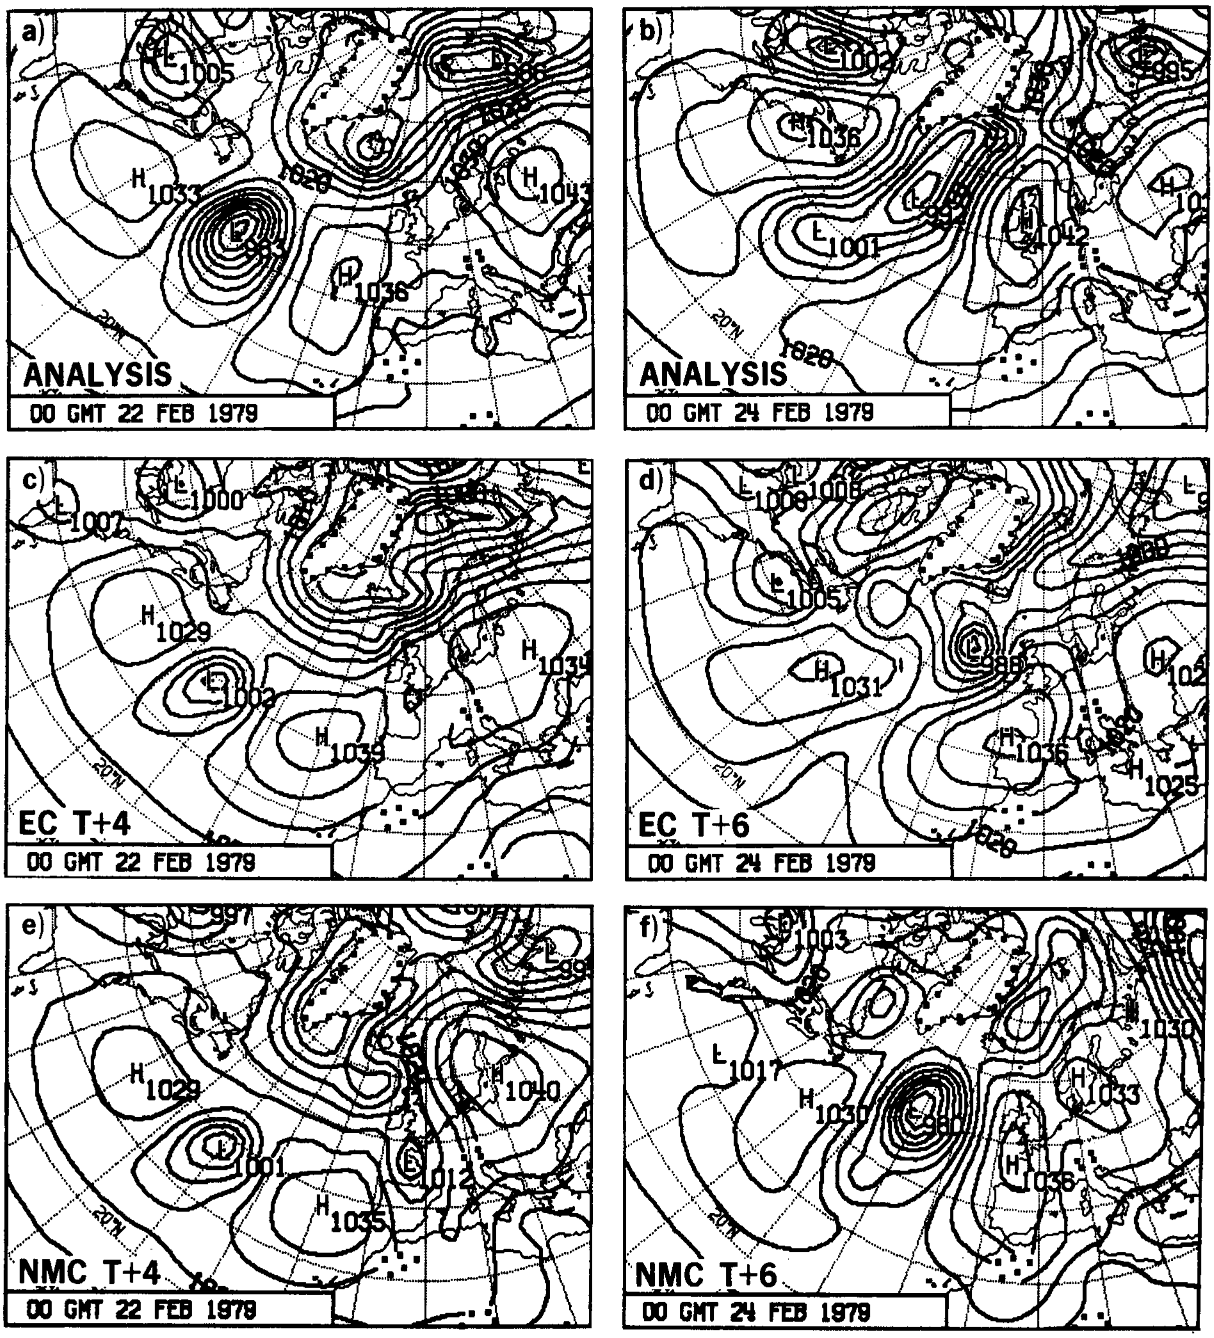

Fig. 1 **Analyses and forecasts for the mean sea level pressure field in the North Atlantic at 00Z on February 22 and 24 1979 ; a) and b) Verifying analyses for the 22nd and 24th (contour interval 5 mb) ; c) and d) Four and six day forecast with the EC model from the EC analysis valid at 00Z on February 18 1979; e) and f) Four and six day forecasts with the EC model from the US analysis valid at 00Z on February 18 1979.**

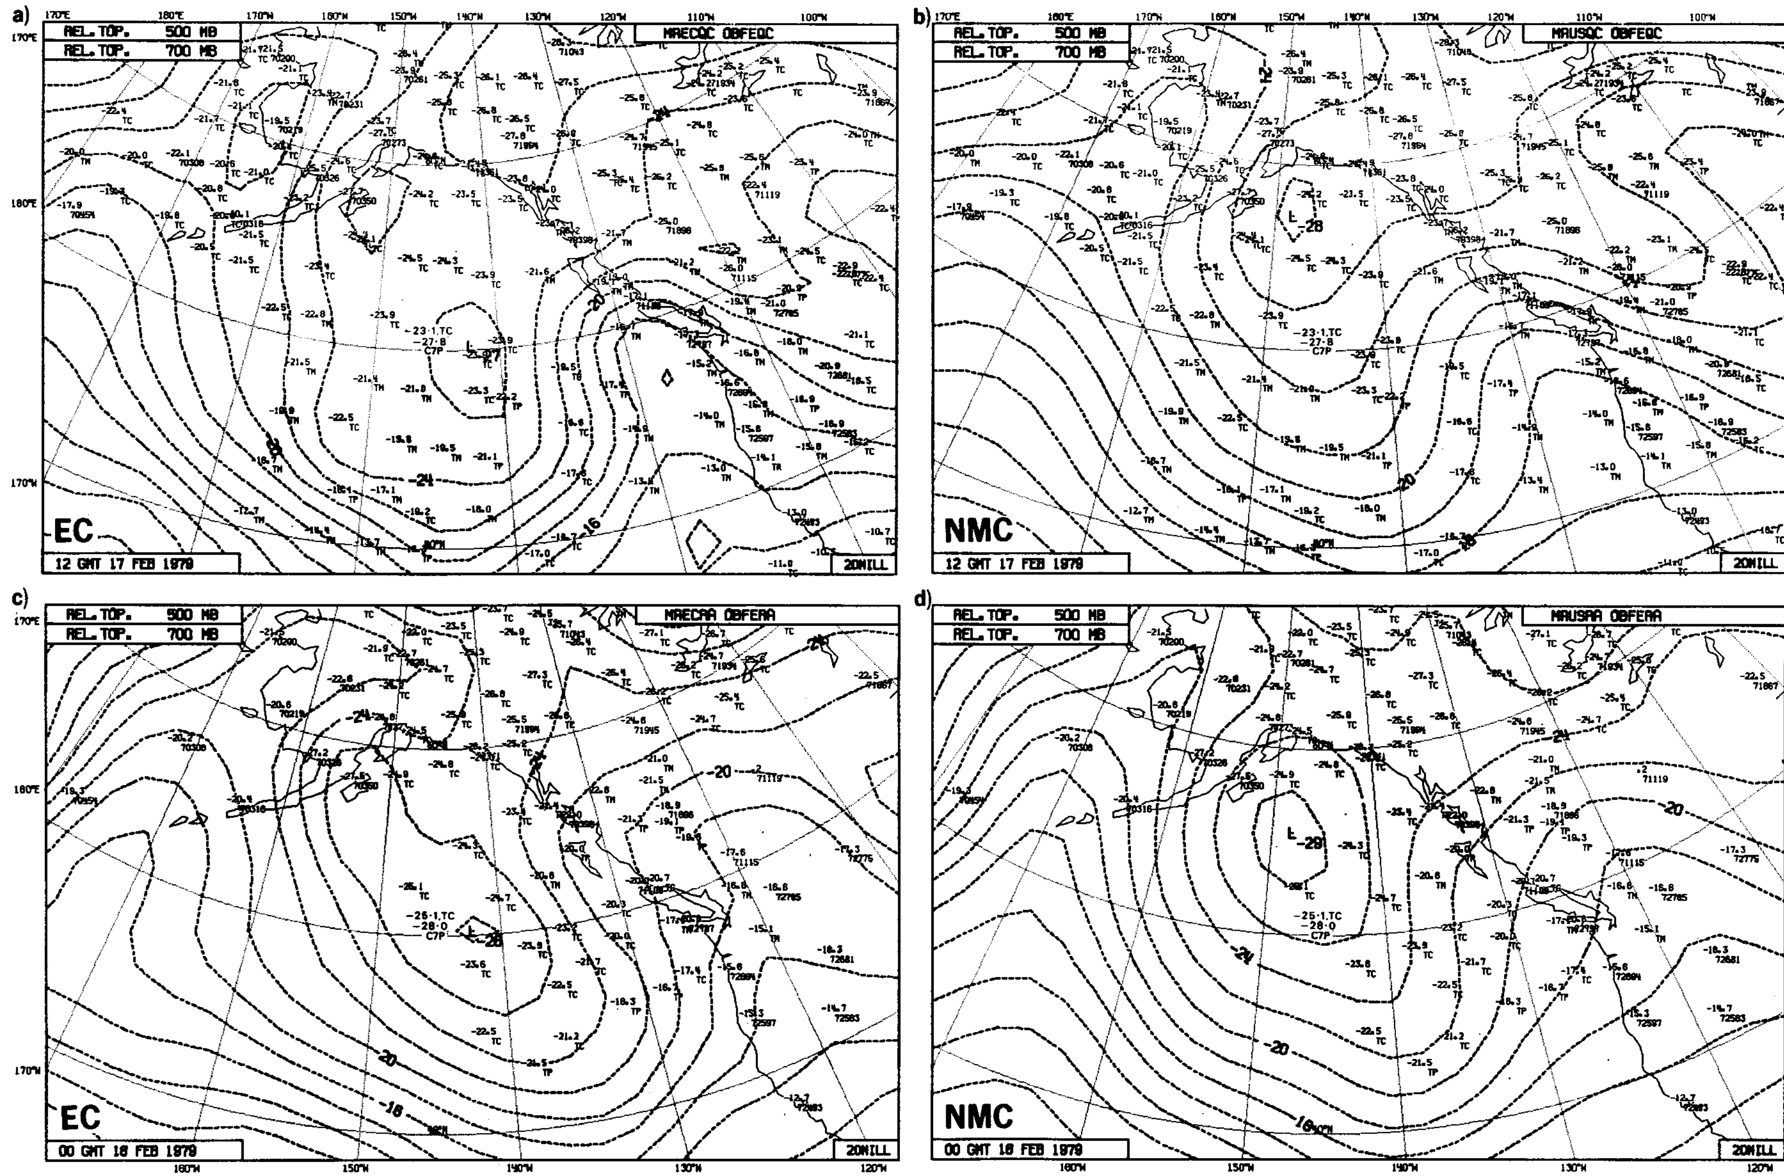

Fig. 2 Analysis of the 500-700mb thickness (converted to virtual temperature, contour interval 2K) over the North East Pacific at 1200GMT 17 February 1979 (top panels), and 0000GMT 18 February 1979 (bottom panels) as produced by the EC system (panels on the left) and the US system (panels on right). Also shown are the radiosonde and Satem reports. The different types of satellite retrievals are indicated as 'TC' for clear path, 'TP' for partly clear and 'TM' for microwave retrievals.

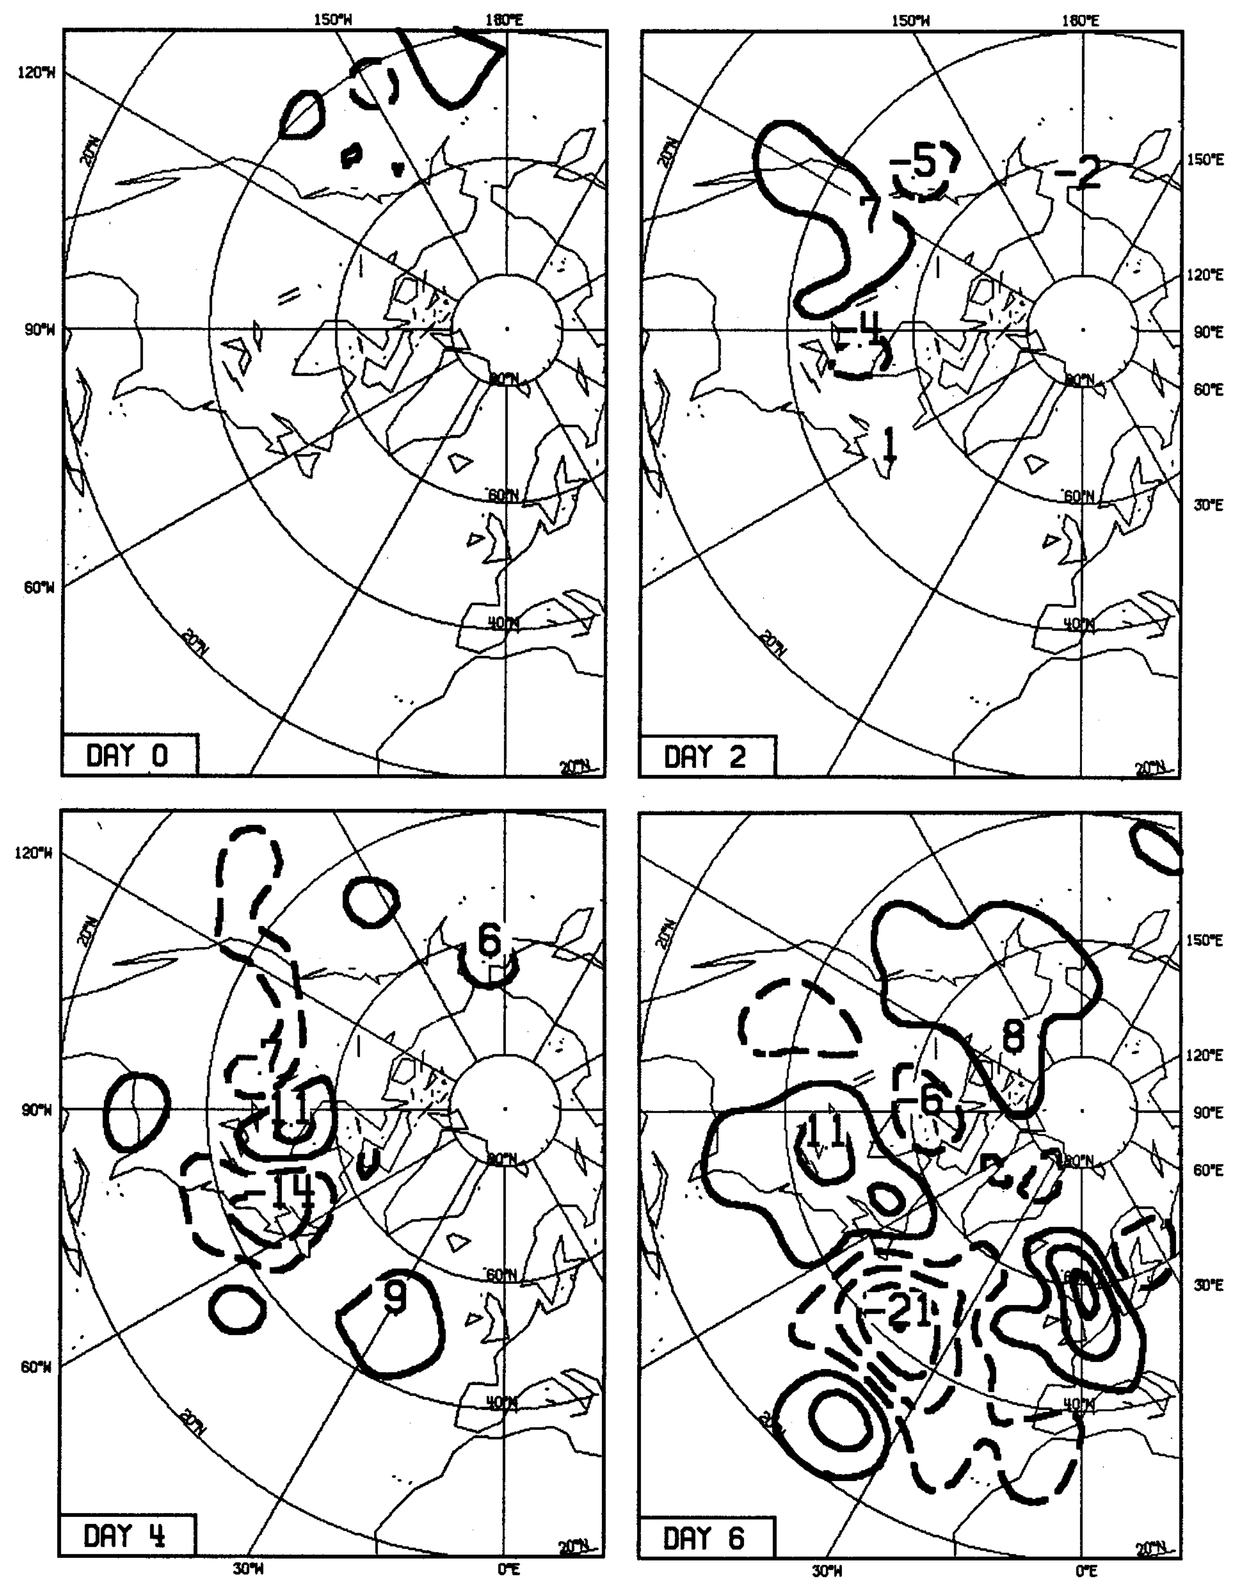

Fig. 3 Differences in the 300 mb height between forecasts with the ECMWF model using the original ECMWF analysis and the transport analysis, (differences outside the north Pacific have been excluded) valid 00.00 GMT 18 February 1979.
Contour interval: 5 dam (negative full, positive dashed)

RESULTS OF THE TRIALS OF THE REVISED DATA ASSIMILATION SYSTEM

In the June 1984 Newsletter (Number 26) there was a description of the substantial revisions made to various aspects of the analysis scheme on 22 May 1984. Here we describe the results of the trials of the new system which were carried out before operational implementation. It is convenient to consider separately the impact on the analyses and the forecasts.

Impact on the analyses

The impact of the changes has been evaluated in a great many ways, by using case studies, inspections of rejected data, evaluation of the relative roles of forecast, analysis and initialisation, and by evaluation of the impact on forecast skill.

Two main periods were chosen for the evaluation: 10 day periods in May and November/December 1983. Many parallel assimilations were carried out to test the effects of each of the main features of the changes in isolation, and in conjunction with the other changes. Numerous forecasts were run from these assimilations, and used in the case studies. Finally, when the experimental work was complete, there was a pre-implementation parallel run in the operational environment during May 1984. Thus, in effect, we had three periods, each of about 10 days, in which the final configuration was tested both in assimilation experiments and in forecast experiments.

The quality control modifications, the data selection modifications and the structure function modifications were found to work well. They led to distinct improvements in the quality of the analysis and the 6-hour forecasts. This could be seen in a variety of ways. The fit of the analyses and the 6-hour forecasts to the observational data was much closer. The data rejection rates were somewhat different, and provided a much improved discrimination capability. The overall magnitude of the analysis changes was reduced because of the more accurate short range forecasts. The initialisation changes were also smaller, and this indicated a better internal balance in the analyses. The analyses also showed a better analysis of the thickness and thermal wind fields, though there is still room for improvement in this area. This will come with incremental increases in resolution.

Impact on the forecasts

The standard means of evaluating a forecast at ECMWF is to examine the correlation of the anomaly (from climatology) of the predicted state and of the verifying analysed state. An example from the May trial is shown in Figure 1 (overleaf), for the extratropical northern hemisphere troposphere, for two forecasts based on the old operational analysis and an analysis produced using the new analysis system, as indicated. The limit of predictive skill is usually judged to be reached when the anomaly correlation falls below 0.6, as indicated by the horizontal dotted line in Figure 1 (overleaf). The new system is seen to produce an improvement of approximately 30 hours in the limit of predictability. However, this is an extreme case, in that it was the one in which the changes produced the largest improvement in forecast skill. Naturally there was a large range of gains in predictability from the 12 May cases, as illustrated in Table 1 overleaf.

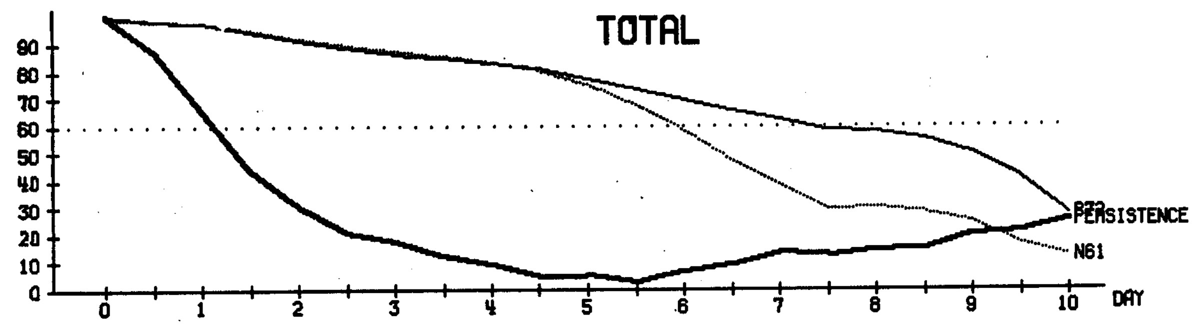

Fig. 1 Anomaly (from climatology) correlation scores measured against days of prediction for forecasts from 12GMT 6.5.83, for the northern hemisphere extra tropical (20°N - 82.5°N) troposphere (1000-200 mb). The continuous line is for the new analysis system ; the dotted line is for the old analysis system, and the bold continuous line is a persistence forecast. The horizontal line of dots mask the 60 % correlation threshold.

Table 1:

Gain in predictability in hours based on NH anomaly correlations											
6/5	8/5	10/5	26/11	27/11	28/11	29/11	30/11	1/12	2/12	3/12	4/12
+30	+9	-7	-11	-3	+9	-1	+11	+2	+3	+7	-15
Average = +2.8											

Based on the 12 cases, an average improvement of 2.8 hours is found; however this sample is very small, and exclusion of individual cases produces a quite different impression. The 6 hour forecasts are known to be better, at least in the rms sense, in their fit to verifying observations. Indeed, detailed investigation of the forecasts for 12Z 26.11.83, when the new system performed badly, suggests that the new analysis system produced a better analysis in the critical areas from where the errors in the forecast are seen to grow. No convincing explanation of the two most seriously degraded forecasts (26 November and 4 December) has been found.

A reassuring signal in the November sequence of forecasts is in the consistency of forecasts from successive days. Figure 2 (overleaf) is a plot of the time in the forecast when the 60% anomaly correlation is reached, for forecasts made on successive days, for the new and old analysis systems. The new thresholds are seen to be generally advanced in time compared to the old, as suggested by the above table. In addition, there is less variability in the quality of forecasts with the new system. This is perhaps to be expected, in that the improved quality control measures of the new system eliminate many of the random errors present in the initial states arising in the old system.

Table 2 below shows the 200 mb vector wind error for the 24 hour forecasts in the tropics for the May 1984 sequence of cases. The 24 hour forecasts are invariably better. However, the D+2 day forecasts (at the limit of our tropical predictive skill) show some gains and some losses.

Table 2:

Tropical scores - 200 mb rms vector wind error

	T + 24	
	Old	New
7 May	6.6	5.9
10 May	7.2	6.8
11 May	7.0	6.3
12 May	6.5	6.1
13 May	6.5	6.0
14 May	7.0	6.5
15 May	6.6	5.6
16 May	6.4	6.2
17 May	7.5	6.6
18 May	7.2	6.4
Average	6.85	6.27

Overall, the impact on medium range forecast skill was not clearcut, but nonetheless positive. However, it is encouraging that important improvements during the parallel run were noted in the tropical wind forecasts at 200 mb. There were also indications that the more accurate fit to tropical data resulted in a more active ITCZ over the tropical oceans, with presumably beneficial effects on the model's spin-up.

As a result of the extensive theoretical, experimental, and statistical work over the last 18 months, we have implemented a radical overhaul of the Centre's assimilation system. This has had a clear beneficial effect on the quality of the analyses and the short range forecasts, and a more modest but still positive impact on forecast skill in the medium range in mid latitudes. There are also clear improvements in the tropical forecasts.

The modifications are important not only for their immediate benefit, but also for the potential for future improvements that are implied.

- Dave Shaw, Peter Lönnberg,
Tony Hollingsworth and Rauno Nieminen

* * * * * * * * * * *

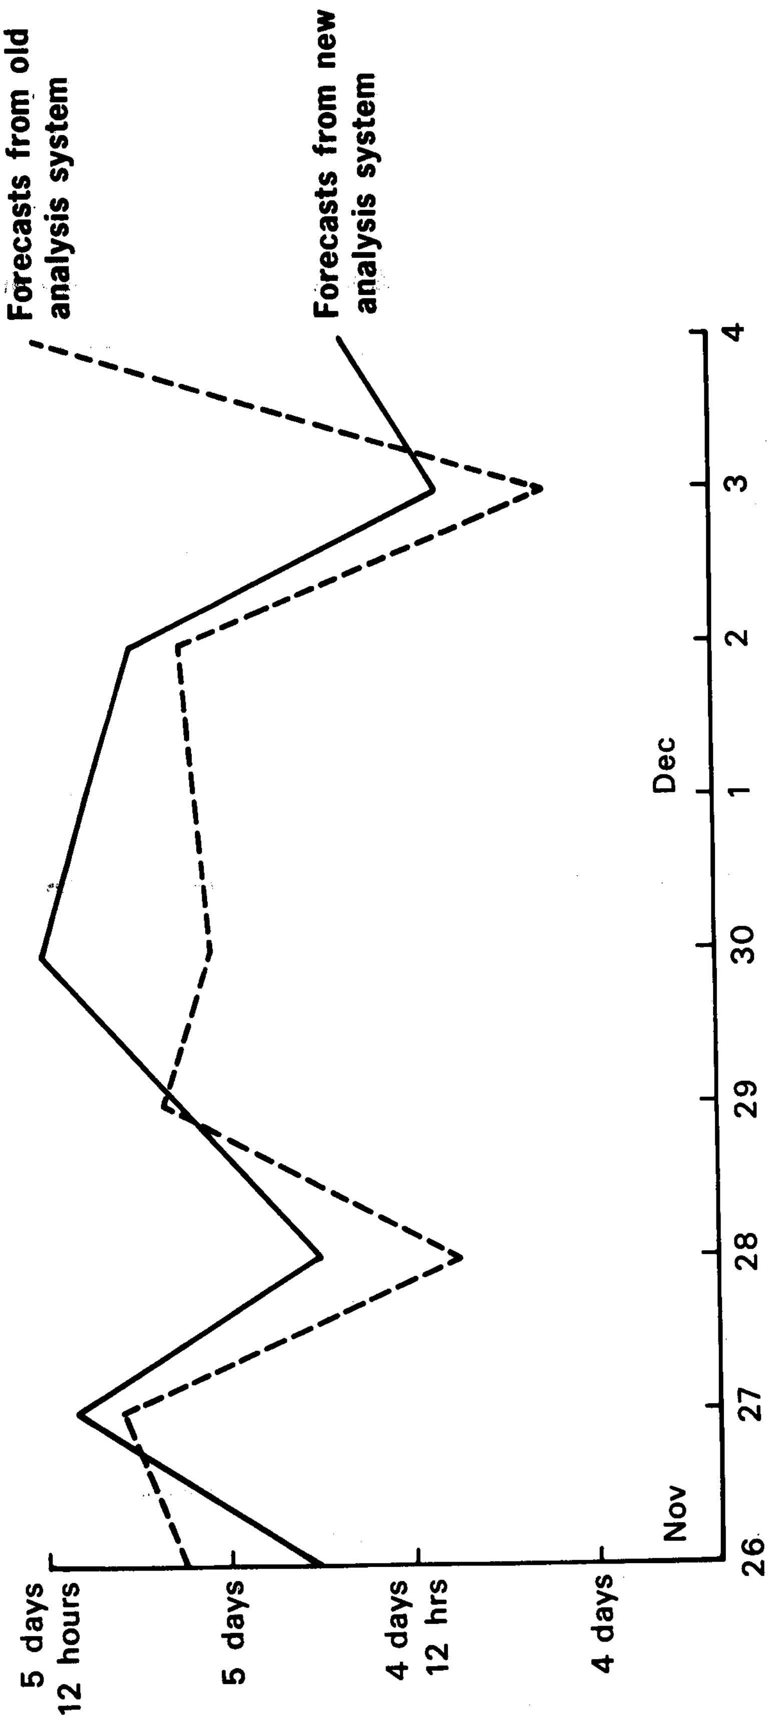

Fig. 2 Time in forecast at which 60% anomaly threshold is reached for forecasts run from successive days (12GMT start time) for 26.11.83-4.12.83

RECENT CHANGES TO THE ECMWF OPERATIONAL FORECASTING SYSTEM

Under this heading, changes to the ECMWF operational forecasting system which occurred during the quarter will be summarised in this and future editions of the Newsletter. Attention will be drawn to other articles in the Newsletter, or separate ECMWF publications, dealing with the reported changes in more detail. The most recent operational changes were:

i. The diurnal cycle of radiation was implemented in operations on 1 May 1984 (see ECMWF Newsletter No. 26, June 1984, article by H. Pümpel).

ii. Extensive changes to the ECMWF analysis system were implemented on 22 May 1984. The optimum interpolation statistics, i.e. forecast and observation error covariances, have been revised based on the performance of the ECMWF assimilation system. The horizontal error structure of the forecast error is now modelled by a series of Bessel functions rather than a Gaussian type model. The rejection limits for observations have been tightened. A modified data selection algorithm produces a more comprehensive and homogeneous usage of observations in data dense regions than before.

The modifications produced a closer agreement between observations, analysis and the first guess. A better discrimination between correct and incorrect data is now possible. The analysis resolution has improved (see ECMWF Newsletter No. 26, June 1984, article by P. Lönnberg).

- Horst Böttger

* * * * * * * * * * *

THE ECMWF METEOGRAM SYSTEM

Meteograms based on on-line data or data from the ECMWF archives (after the introduction of the spectral model on 21 April 1983) can be plotted with this system. Meteograms can be generated for any position on the globe and will display, for a 10 day forecast in black/white or colour, the cloud amount, 850mb relative humidity, precipitation, mean sea level pressure (MSL), 10m wind and temperature (2m and 850 mb).

With the interactive command driven processor, METGRAM, the user can generate a batch job sequence to retrieve the data for the meteograms and plot them. The data retrieval, data manipulation and post processing of the graphics output is performed on the Cybers whereas the generation of the meteograms is done by a Cray job.

To create a meteogram using METGRAM, one or more station definitions have to be entered. A station definition consists of a name, country code, latitude and longitude. However, a list of 30 stations is known by the METGRAM processor and for these stations, only the station name has to be entered.

Using the default for the other parameters, a batch job sequence will be generated and automatically entered into the Cyber input queue. The output will be a set of meteograms, based on on-line data, plotted in A4 size on the raster plotter.

Other commands for METGRAM can control:

- retrieval of data from archives
- size of the meteograms (A1-A6)
- graphics output quality ranging from reproduction quality to a simple version where the meteogram, after completion of the batch jobs, can be transferred to a graphics terminal (running at 9600 baud) in 25 sec
- a map layout, which, if specified, will group 10 stations at a time around a European or World map (see cover illustration)
- generation of colour output
- post processing to generate output for the raster plotter and/or store the metafile containing the graphics information.

A HELP facility is provided with the Metgram command processor and the LIST command will display the current settings of the parameters.

- Jens Daabeck

* * * * * * * * * *

COLOUR RASTER GRAPHICS AT ECMWF

Introduction

Following a competitive Invitation to Tender, the Centre installed an Aydin 5216 Colour Raster graphics terminal in November 1980. This was an experimental system to gain experience with colour graphics and with an interactive presentation system for meteorological data.

The system installed was equipped with a 19 inch monitor, capable of displaying a 1024 * 1024 resolution image at 30 Hz interlaced. It was able to display simultaneously 16 colours from a pallette of 64, it was equipped with 10 million bytes of local disk storage, and was connected to the Cyber mainframes by a 9.6 Kilobits/second serial interface.

A Matrix Camera hardcopy device was installed in 1981.

Applications Implemented

The Aydin had an Intel 8086 16 bit microprocessor built in to the system which was programmed in a language called FORTH. In order to make any use of the Aydin, FORTH programs had to be developed locally and stored on the Aydin disks. They could then be called into execution to display stored data or to store data from the serial link.

Initially, the programming effort spent on the Aydin was used to demonstrate possible applications which might be implemented later in an operational mode. These demonstrations were exhibited to various visitors and external committees during 1981 and showed that, with some FORTH programming effort, viable applications were possible on the Aydin. These demonstrations were also used as the basis for display material produced by the Centre.

During 1982, effort was spent on implementing an operational application on the Aydin. Basically, the Aydin was used to display observations and analysed fields produced daily by the operational suite.

In order to achieve this, 2 FORTH programs were produced, one to transfer vector information from the Cyber to the Aydin disk, and the second to display this data. The display program allowed zooming, displaying from an entire hemisphere to a single observation, and allowed contours, annotation, etc. to be selectively added and deleted from the display.

The Cyber contouring package was amended to produce suitable data in Aydin format and the data transmission between Cyber and Aydin disk took place early each morning before the operational analysts arrived for work. Upon their arrival, they were able to examine observations and analyse output as necessary to help identify problems.

This application demonstrated the usefulness of colour in the interpretation of meteorological results, since the colour of the observation was changed according to the observation quality flag, and the usefulness of being able to zoom in and out of a specific area of the hemisphere.

Recent developments in colour raster graphics

During 1983, it became apparent that improvements in display technology had overtaken the Aydin project. In particular, terminals with better resolution and a non interlaced display were appearing. In addition, they were provided with a high degree of local intelligence which did not require local programming effort to utilise. In short, the newer products were much easier to use and, due to their higher display quality, had a high user acceptance.

In part, these improvements occurred due to changes in technology such as higher speed microprocessors and much lower memory prices, but a significant factor was the rapidly increasing market for colour raster graphics terminals which occurred partly as a result of standardised software interfaces.

In 1984, it became clear that the cost of maintaining the Aydin hardware and the cost of supporting a non standard software system based on the FORTH language could no longer be justified by the Centre. Therefore, early in 1984, the Aydin service ended.

Since then, the Matrix camera has been connected to one of the Centre's other terminals, where it will continue to provide hard copy of colour images, on either 35mm colour slides or 8" x 10" colour prints.

The Future of Colour Raster Graphics

Despite the termination of the experimental service, it is clear that high quality colour raster graphics terminals can be extremely useful in interpreting the Centre's results. In particular, the ability to interface to the Centre's DISSPLA software will make programming the devices much easier, and a high speed hardware interface to the VAX will considerably reduce the data transfer time.

The Centre also benefited from the experience of using the many features of the Aydin. This experience will allow the Centre to select appropriate facilities to be incorporated in future terminals, and will allow a better assessment of the limitations of the equipment.

In order to progress in this field, the Centre is about to issue an Invitation To Tender for a high resolution graphical terminal to provide much improved facilities for operational usage.

- Peter Gray

* * * * * * * * * * *

COS AND CFT 1.13 NEW FEATURES

Systems Section are currently in the early stages of testing COS and CFT 1.13. It is hoped that it will be possible to make CFT 1.13 and the 1.13 libraries available for user testing in September. It is then proposed to implement COS 1.13 during October.

Since the current version ,COS X.13, is a pre-release version of COS 1.13, it already contains most of the new features in COS 1.13, in particular support for multitasking.

The main reason for moving to COS 1.13 is to use standard released software which is fully supported by Cray Research. This should make the correction of bugs in the system simpler, and also make the generation of new versions of COS simpler.

There are a few new features in COS 1.13 and its products:

- SSD is now a separate generic resource.
- ASSIGN will have a NOF (no overflow) parameter, which can be used in conjuction with SSD to prevent SSD datasets overflowing to disk.
- Permanent Dataset messages from ACCESS, ACQUIRE, DELETE, MODIFY and SAVE can be suppressed by use of the MSG parameter.
- ACCESS has an IR parameter which allows immediate return if the dataset is busy.
- The new product COPYU may be used to copy unblocked datasets.
- UPDATE has many enhancements, in particular DEFINE, IF, ELSEIF, ELSE and ENDIF directives which allow a Program Library to contain conditional text.

There are many improvements to CFT 1.13, which should result in faster code. The main ones of interest are listed below:

Code Generation

Code generation in CFT has been improved to generate significantly more efficient code. In particular, CFT 1.13

- takes functional unit and memory busy times into account
- resolves register path conflicts
- improves scheduling of scalar instructions
- bottom loads scalar loops (may be disabled if necessary)

B and T Register Optimisation

Use of the control statement option OPT=BTREG will cause CFT 1.13 to allocate local scalar variables to T registers rather than memory. This should give a significant speedup. Dummy arguments and variables named in SAVE, DATA, COMMON or NAMELIST are allocated to memory. If a program contains subroutines which depend on local variables retaining their value across calls, such variables will require SAVE statements to force them to be allocated to memory.

Re-entrant code

Use of the control statement option ALLOC=STACK will cause CFT 1.13 to generate code which is re-entrant, i.e. can be multitasked. This will cause program variables to be assigned to a runtime stack.

Gather and Scatter

Gather loads and scatter stores within loops no longer inhibit vectorisation, which should speed up such codes.

Informative messages for DO loops with dependencies

CFT 1.13 produces informative messages for DO loops with dependencies, indicating the type of dependency that is inhibiting vectorisation.

Special case handling of one-line DO loops

Single-line DO loops performing specialised functions (such as first-order linear recurrences) are automatically recognised, and are replaced by calls to appropriate $SCILIB routines. There are compiler options (SAFEDOREP, FULLDOREP and NODOREP) to enable or disable this replacement.

Align code on instruction buffer boundary

The new compiler directive CDIR$ ALIGN directs CFT 1.13 to align the next DO loop, label or entry point on an instruction buffer boundary. This can significantly improve the speed of certain DO loops.

- Richard Fisker

* * * * * * * * * *

THE USE AND MISUSE OF FLOWTRACE

FLOWTRACE is a timing facility built into the Cray Fortran compiler (CFT). It will report down to subroutine level on such things as:

time spent in each routine,
the number of occasions each routine is called,
listing of routines which call,a given subroutine, and listing
the routines a given subroutine calls,
subroutine linkage overheads.

It can be switched on using ON=F on the CFT control card. It can be switched on and off for specific parts of a program via CDIR$ FLOW and CDIR$ NOFLOW directives. The full details of how to use it are found in the CFT Reference Manual (SR-0009).

The information it provides can be very helpful in isolating the parts of a program which consume most resources and, hence, where efforts should be made to optimise the program. Unfortunately, FLOWTRACE has the undesirable effect of using considerable operating system resources to gather this information, so seriously affecting the service to all other users. Thus please, when using FLOWTRACE:

do not use it unnecessarily, remove the ON=F parameter as soon as you have completed your timing tests.

Only use it on short runs of your program, e.g. for 2 or 3 iterations only. (In a recent example, a user ran a 10 minute job with FLOWTRACE on, and tied up operating system resources for several hours).

Always use FLODUMP as well. If you use FLOWTRACE by itself and your program aborts, then you lose all the timing information which has been accumulated up to that point (maybe at considerable cost). By using the FLOWDUMP utility, you receive the timing information collected. This utility is described in the COS Reference Manual (SR=0011). An example of its use is:

```
CFT(ON=F)
LDR.
EXIT.
DUMPJOB.          needed to dump FLOWTRACE tables
FLODUMP.
   .
   .
   .
```

There are alternatives to FLOWTRACE. For example, there is the ECMWF timing package, called TIMING, fully described in the ECMWF Computer Bulletin B4.9/3. This provides timing information down to DO loop level (with a few restrictions as listed in the bulletin). It does not cause the same excessive overheads in the operating system as FLOWTRACE.

- Andrew Lea

* * * * * * * * * * *

SCHEDULED INTERRUPTIONS TO THE COMPUTER SERVICE

The present scheduled interruptions are listed below. They consist of preventative maintenance sessions, and "systems" sessions. Preventative maintenance sessions are regularly taken by the engineers to carry out the on-going programme of hardware maintenance. "Systems" sessions are booked as required by systems analysts to test new systems, or engineers to trace and fix faults. Also if changes to cable connections, power supplies or similar have to be done, they are scheduled in these periods if possible. These periods are valid irrespective of whether there is a U.K. public holiday or not. Times are local.

	CYBER 835	CYBER 855	CRAY X-MP
MONDAY	7 am - 9 am	7 am - 9 am	7 am - 9 am (2)
TUESDAY	7 am - 9 am (1)	7 am - 9 am	7 am - 9 am (5)
WEDNESDAY	7 am - 10 am (3)	7 am - 10 am	7 am - 10 am (2)
THURSDAY	6 pm - 8 pm	7 am - 9 am	7 am - 9 am
FRIDAY	7 am - 9 am (1)	7 am - 9 am	7 am - 9 am
SATURDAY	8 am - 2 pm (5)	8 am - 2 pm (5)	8 am - Noon (5)
SUNDAY	7 am - 9 am (4)	7 am - 9 am (4)	7 am - 9 am (4)

(1) On Tuesday and Friday, either the 855 or the 835 is available for 'systems' bookings, leaving the other system to support Intercom.

(2) The Cray analysts normally take one or both of these sessions each week but rarely take the Thursday or Friday period.

(3) The period 9 am - 10 am is not always taken because Intercom should be available, if possible, from 9 am.

(4) Very occasionally, all three machines may be taken simultaneously.

(5) Preventative maintenance.

If more time than can be provided during these periods is needed, extra time is made available where possible outside the working week (i.e. Monday - Friday 0830 - 1730), usually after 2 pm on Saturday or during the day on Sunday. As much notice as possible of such additional sessions is given.

Member States should note, however, that the NFEP is normally available at all times, and is able to store incoming jobs, if the Cybers are not available.

- Eric Walton

* * * * * * * * * * *

THE NAG LIBRARY

The NAG library is now available at Mark 10. The Mark 11 implementations should soon be released by NAG Ltd., the CDC implementation being due in August and that for the Cray in October.

The situation at ECMWF is that until the beginning of this year, the Mark 6 version of the library was still being used, in order that the Cyber and Cray versions remained compatible for reasons of documentation, etc. There was no Cray version at Mark 7, and problems were experienced with the Cray versions of Mark 8 and 9. The library has remained at Mark 6 at ECMWF for 4 years.

A large number of changes have taken place between the Mark 6 and Mark 10 versions. Some 80 out of 351 user-callable routines from the original Mark 6 have been deleted and 221 new user-callable ones have been added. There is therefore a great likelihood that someone may be using a routine which will no longer be available when Mark 10 is implemented. In all cases, a replacement routine is available, so there should be no major problems, although some program changes may be required. On the positive side, there are many more routines, including two completely new chapters. Amongst the major changes are:

New chapters:

Non-parametric statistics	(GØ8)
Time series analysis	(G13)

Major additions to chapters:

Roots of one or more transcendental equations	(CØ5)
Summation of series including Fourier transforms and inverse transforms	(CØ6)
Quadrature	(DØ1)
Ordinary differential equations	(DØ2)
Partial differential equations	(DØ3)
Curve and surface fitting	(EØ2)
Minimising or maximising a function	(EØ4)
Matrix operations including inversion	(FØ1)
Eigenvalues and eigenvectors	(FØ2)
Simultaneous linear equations	(FØ4)
Simple calculations on statistical data	(GØ1)
Special functions	(S)
Machine constants	(XØ2)

When we tranferred to the X-MP, we took the opportunity of changing the Cray implementation to Mark 10, without any adverse effect. We have not converted the Cyber version yet, but intend to do so as soon as possible, after users have had an opportunity to test their jobs with it. To test the Mark 10 version on the Cyber, use the control card:

```
NEXT(NAGLIB)
```

anywhere preceding the LOAD sequence.

Please try this version and, if you experience any problems, contact User Support.

- John Greenaway

* * * * * * * * * *

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 165). All other News Sheets are redundant and can be thrown away. The following News Sheets can be discarded since this list was last published: 54, 154, 156, 163.

No.	Still Valid Article
16	Checkpointing and program termination
19	CRAY UPDATE (temporary datasets used)
56	DISP
67	Attention Cyber BUFFER IN users
73	Minimum Cyber field length
89	Minimum field length for Cray jobs
93	Stranger tapes
118	Terminal timeout
120	Non-permanent ACQUIRE to the Cray
121	Cyber job class structure
122	Mixing FTN4 and FTN5 compiled routines
127	(25.1.82) IMSL Library
130	Contouring package: addition of highs and lows
135	Local print file size limitations
136	Care of terminals in offices
140	PURGE policy change
141	AUTOLOGOUT - time limit increases
144	DISSPLA FTN5 version
147	(20.7.83) NOS/BE level 577
152	Job information card
158	Change of behaviour of EDIT features SAVE, SAVEX.
	Reduction in maximum print size for AB and AC
162	DISPOSE problem on the Cray X-MP
164	CFT New Calling Sequence on the Cray X-MP
165	Maximum memory size for Cray jobs

* * * * * * * * *

CALENDAR OF EVENTS AT ECMWF

3-7 September 1984	Annual ECMWF seminar: "Data assimilation systems and observing system experiments, with particular emphasis on FGGE."
12-14 September 1984	12th session of Scientific Advisory Committee
18-20 September 1984	7th session of Technical Advisory Committee
25-27 September 1984	32nd session of Finance Committee
6-8 November 1984	Workshop: "The use and quality control of meteorological observations"
20-22 November 1984	20th session of Council
26-28 November 1984	Workshop: "Cloud cover and radiative fluxes in large scale numerical models - intercomparison and verification"
3-6 December 1984	Workshop: "Using multiprocessors in meteorological models"

* * * * * * * * * *

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in the ECMWF Newsletter plus those in the original ECMWF Technical 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.

*T indicates the original Technical Newsletter series

USEFUL NAMES AND 'PHONE NUMBERS WITHIN ECMWF

			Room*	Ext.**
Head of Operations Department - Daniel Söderman			OB 010A	373
ADVISORY OFFICE - Open 9-12, 14-17 daily			CB Hall	309
Other methods of quick contact:				
- telex (No. 847908)				
- COMFILE (See Bulletin B1.5/1)				
Computer Division Head	-	Geerd Hoffmann	OB 009A	340/342
Communications & Graphics Section Head	-	Peter Gray	OB 101	369
COMPUTER OPERATIONS				
Console	-	Shift Leaders	CB Hall	334
Reception Counter) Tape Requests)	-	Jane Robinson	CB Hall	332
Terminal Queries	-	Norman Wiggins	CB 035	209
Operations Section Head	-	Eric Walton	CB 023	351
Deputy Ops. Section Head	-	Graham Holt	CB 024	306
DOCUMENTATION	-	Pam Prior	OB 016	355
Distribution	-	Els Kooij-Connally	Library	430
Libraries (ECMWF, NAG, CERN, etc.)	-	John Greenaway	OB 017	354
METEOROLOGICAL DIVISION				
Division Head	-	Frédéric Delsol	OB 008	343
Applications Section Head	-	Rex Gibson	OB 227	448
Operations Section Head (Acting)	-	Horst Böttger	OB 004	347
Meteorological Analysts	-	Veli Akyildiz	OB 005	346
	-	Herbert Pümpel	OB 006	345
	-	Liam Campbell	OB 003	348
Meteorological Operations Room	-		CB Hall	328/4
REGISTRATION				
Project Identifiers	-	Pam Prior	OB 016	355
Intercom & Section Identifiers	-	Jane Robinson	CB Hall	332
Operating Systems Section Head	-	Claus Hilberg	CB 133	323
Telecommunications Fault Reporting	-	Stuart Andell	CB 035	209
User Support Section Head	-	Andrew Lea	OB 018	353
RESEARCH DEPARTMENT				
Computer Coordinator	-	David Dent	OB 123	387
Meteorological Education and Training	-	Bob Riddaway	OB 319	478

---

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (0734) 876000
international +44 734 876000