European Centre
for Medium Range Weather Forecasts

ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England
Reading (0734) 85411 Telex 847908

Number 14 - April 1982

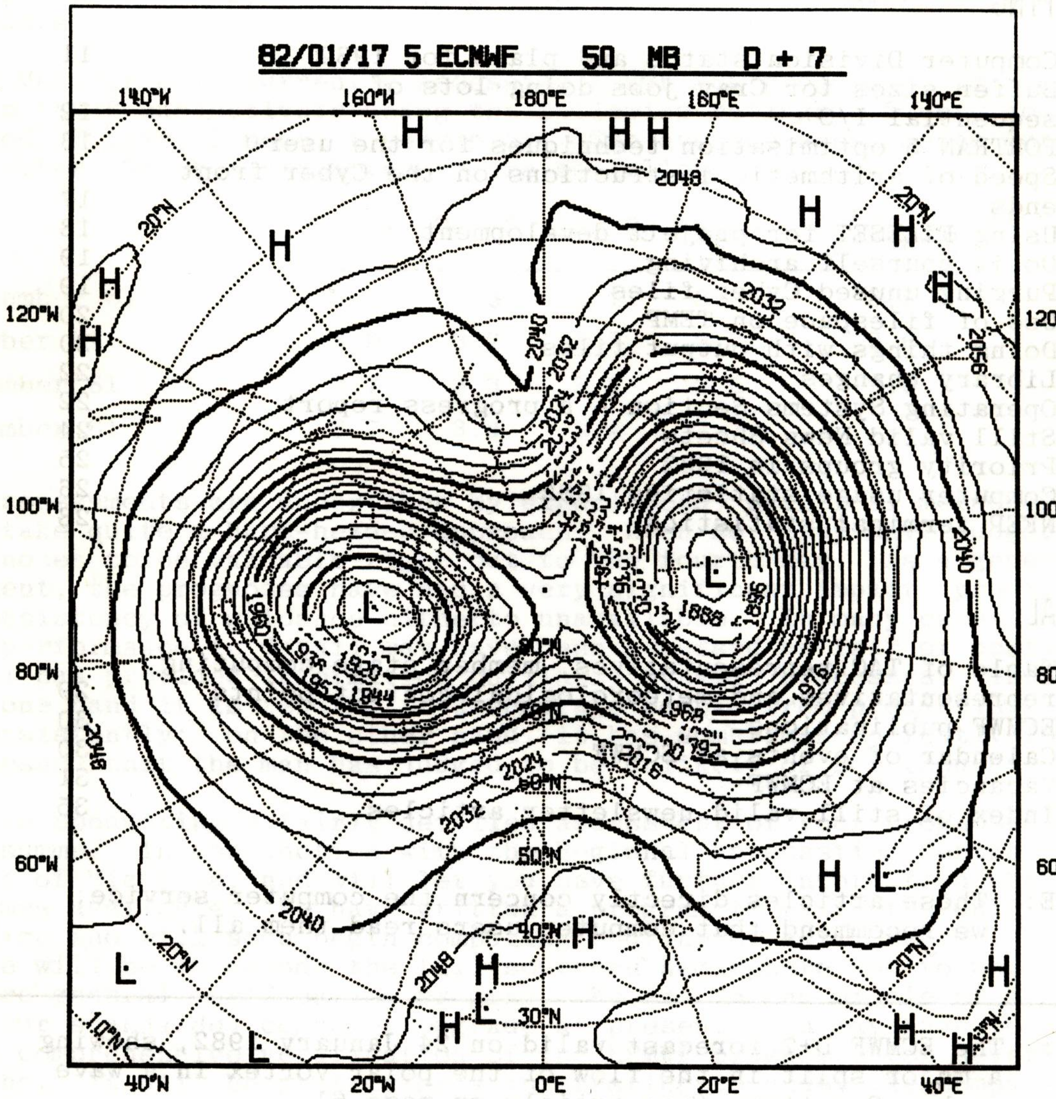

IN THIS ISSUE:

* NOTE: These articles directly concern the computer service; we recommend that computer users read them all.

COVER: The ECMWF D+7 forecast valid on 24 January 1982, showing a major split in the flow of the polar vortex in a wave number 2 pattern (see article on page 6).

This Newsletter is edited and produced by User Support.

The next issue will appear in June.

RECEPTION OF ECMWF FORECASTS IN AUSTRALIA

"The major operational event of 1981 at WMC Melbourne".

Dissemination of a very limited sub-set of ECMWF products, including MSL pressures and 500mb heights, to the world meteorological community began on 1 August 1981. The products include analyses and forecasts to D+4 for the Southern Hemisphere; they are transmitted in grid-code format on the Global Telecommunications System of WMO and can readily be decoded by Meteorological Services throughout the world. Feedback from non-Member State users has begun.

The piece below is extracted from a letter received from P. Noar, Head of the World Meteorological Centre, Melbourne, concerning the ECMWF products:

"WMC Melbourne is currently assessing ECMWF MSL prognoses for the area 90 to 180^{O}E subjectively using the following ratings: 5 very good, 4 good, 3 fair, 2 poor, 1 very poor. The average assessments for September, October and November are as follows:

	Day 1	Day 2	Day 3	Day 4	Overall 4 day guidance
September 81	3.9	3.3	3.1	3.0	3.9
October 81	3.9	3.7	3.3	2.7	3.6
November 81	3.7	3.3	3.1	2.5	3.0
December 81	3.7	3.4	3.4	2.7	

The assessments are undertaken by experienced, practical meteorologists who take quite a hard-headed approach. Even so, they have judged the prognoses to be valuable right out to the fourth day. As a general comment, the prognoses have had a very significant impact, providing a consistency of guidance hitherto unavailable. In many cases, where the performance in detail may have been assessed as fair or poor, the model has provided information on likely dynamic processes. Thus, on the one hand it is tending to discourage people from predicting inconsistently; on the other hand, it has sometimes correctly forecast processes that the man has judged to be unlikely or not thought of.

We are conducting a fairly detailed assessment of your prognoses during our summer, in conjunction with the Regional Forecasting Centre in the state of Victoria, and will let you have further information when it becomes available. We have written a program to write the GRID data to disc and will soon begin computing routine S1 skill scores. The next stage will be to decode the full southern hemisphere and to undertake a more general verification program. We are in the middle of a major (by our standards) computer upgrade at present, and may have to delay more comprehensive verification of the ECMWF prognoses for several months. It will be particularly interesting to undertake comparative verification studies for the southern hemisphere, and to identify systematic intensity or location errors.

All in all, I would describe the beginning of our reception of your product as being the major operational event of 1981."

- Austin Woods

* * * * * * * * * * *

ALPEX: THE ALPINE EXPERIMENT OF THE GARP MOUNTAIN SUB-PROGRAMME

Introduction

Recently many ECMWF Newsletter readers will have heard that 1 March 1982 marked the start of the two monthly Special Observing Period of an international meteorological experiment called ALPEX, and that ECMWF is playing a comparatively small but crucial part in this experiment by hosting the international data management effort associated with it.

The purpose of this note is to explain what ALPEX is and what its general goals and scientific objectives are. The material presented here was derived, with very few alterations, from the ALPEX Experiment Design document recently published in its final form by WMO/ICSU (Geneva, January 1982).

Readers are also reminded that a more detailed description of the meteorological data collection and of the related data management effort is contained in J. Martellet's article "The ALPEX Data Management and the International Alpex Data Centre" published in the ECMWF Newsletter No. 11, p.1 (September 1981).

The Mountain Problem

Mountain complexes, even of moderate size, are known to influence the atmosphere on practically all scales of motion. Orographically influenced phenomena on the planetary scale, such as blocking, are, at present, a subject of renewed scientific interest. They affect the weather worldwide. Planetary waves may interact with smaller scales to produce powerful synoptic disturbances, e.g. lee cyclogeneses. In the European countries surrounding the Alps, the prediction of the highly weather-affecting "Genoa Cyclones" in the lee of the Alps is a notorious forecasting problem. There are indications that subsynoptic processes may play an important role in the initial stages of lee cyclogenesis.

There is also insufficient knowledge for the design of adequate methods to treat mountain effects in numerical weather forecasting models. Computational grids are generally too coarse, and similar problems also exist in analysis, initialisation and data assimilation.

On the synoptic and subsynoptic scales, mountains affect the weather in a variety of ways. The drag of an entire mountain range and its various components have never been fully determined. While these components are fairly well known in principle (form-drag , momentum transport by lee waves, turbulent dissipation due to flow separation, etc.), their simultaneous quantitative measurement is lacking. Closely related to these phenomena and to leewaves are the severe downslope winds known by different names in various parts of the world (the chinook winds of the Rocky Mountains, the foehn and mistral of the Alps, and the bora of the eastern Alps and Yugoslavian coastal mountains). While there is general agreement that nonlinear effects are important, the question of the hydrostatic or nonhydrostatic character of some of these phenomena is still being debated.

The fact that the Alps are associated with well-known examples of practically all the aforementioned phenomena, that the Mediterranean area south of the Alps has the highest cyclogenetic frequency of the Northern Hemisphere and that the Alpine area is already fairly well covered by routine meteorological observations, have all contributed to the selection of the Alps as the field site for this programme. Other factors have been the relatively long climatological records available for this area and the scientific and operational interest of surrounding nations.

The Scientific Objectives of ALPEX

The primary goal of the GARP Mountain sub-programme in general, and of ALPEX in particular, is to provide a basis for designing improved procedures for quantitatively representing in numerical models the dynamical and physical effects of orography on various scales. It is clear that ECMWF shares great interest in such a scientific goal.

The detailed scientific tasks of ALPEX, therefore, include the study of:

a. the air flow, mass and moisture fields over and around the mountain complex under various synoptic conditions;

b. the formation of cyclones in the lee of the mountain barrier with special emphasis on the subsynoptic nature of the associated processes;

c. the drag of the mountain complex upon the atmosphere;

   the vertical transport of horizontal momentum as a function of height near a mountain range;

   the dissipation of gravity-inertial wave energy over and downwind of the mountain range;

d. local mountain winds, such as foehn, bora, mistral;

e. the role of sensible and latent heat flux over the Mediterranean Sea in connection with lee cyclogenesis;

f. the effects of differential radiative heating introduced by the mountain range;

g. the effect of mountain complexes on precipitation;

h. the physical processes responsible for the development of severe floods, wind storms, and storm surges.

The Experiment Area and Timing

The Experiment Design envisages an "inner" and an "outer" experiment area. The inner area surrounds the Alpine massif and contains nearly all special observing systems dedicated to ALPEX. The outer experiment area defines the large-scale flow in which the orographically disturbed air is embedded and encompasses, therefore, Western Europe, the Mediterranean and part of the Eastern Atlantic, see Fig. 1, p. 1 of Newsletter No. 11).

There are two experiment periods: a low-intensity ALPEX Observational Period (AOP) of about one year (September 1981 - October 1982) and a high intensity Special Observing Period (SOP) scheduled for March and April 1982. According to climatological statistics, the SOP should encounter at least two cases of lee cyclogenesis (three have already occurred in the first fifteen days!).

The Composite Observing System

The World Weather Watch is enhanced in the inner experiment area and many special platforms, such as ships, aircraft, balloons, radars and microbarographic arrays are being deployed during the SOP. Four or more ascents per day are made for many of the upper-air sounding stations in the inner experiment area during the SOP. Several of the ALPEX ships and research aircraft have upper-air capability giving wind and thermodynamic parameters. Meteosat II is making "rapid scans" for high resolution cloud winds on days of special interest. The NAOS (North Atlantic Ocean Station) ships in the east Atlantic also operate with increased sounding frequency.

About 30 microbarographs are at present measuring pressure profiles across the Alps and the Yugoslavian coastal mountains. Up to fifteen aircraft are participating in ALPEX, among them three highly instrumented long and medium-range research aircraft from the USA and Federal Republic of Germany. The flight tracks are tailored to the specific scientific tasks listed previously. They contain upstream missions during lee cyclogenesis, dropwindsonde missions over the Mediterranean Sea and vertically stacked multi-aircraft missions flown along the sections of surface microbarographic stations placed across the mountain ranges.

The Oceanographic Programme

The oceanic component of ALPEX is called MEDALPEX for "Mediterranean ALPEX". It is directed toward understanding the effects of wind forcing on the dynamics of the Western Mediterranean and the Adriatic Sea, especially under conditions of severe weather. Among the phenomena being studied are meso-scale eddies and storm surges. All studies are intimately tied to the enhanced meteorological observations during ALPEX. This programme, in which six nations participate,with eight to ten ships, numerous buoys and tide gauges, is coordinated by the Intergovernmental Oceanic Commission (IOC) following a resolution by its Assembly to participate in ALPEX. Some of the ships are making 6-hourly upper-air soundings during the SOP to fill the existing gaps in the essentially land- and island-based atmospheric network.

The Field Operations

The Swiss airport of Genève-Cointrin is the primary operations base for ALPEX. The Italian Marco-Polo Airport in Venice serves as secondary base. In addition, a number of national operations centres have been activated by the participating countries. The Operations Centre at Geneva houses about 100 technicians and scientists.

The Participants

The following nations participate actively in ALPEX and have made special observing systems or other resources available to the field phase and to the Data Management phase:

Austria	Belgium
Canada	Czechoslovakia
France	Germany (Federal Republic)
Greece	Hungary
Italy	Netherlands
Poland	Romania
Spain	Switzerland
U.S.A.	U.S.S.R.
Yugoslavia	

Additional support is being given by Israel, Portugal and Turkey.

The Data Management Effort

Experience gained in earlier GARP field experiments such as GATE, FGGE and MONEX has suggested that more emphasis should be given to an early assembly of a "quick-look" data set to be released within a few months of the termination of the field phase. The final validated and quality controlled (Level II-b) data set is to be released 12-18 months later. It is hoped that the quick-look data will allow preliminary research to begin almost immediately. The Level II-b data will consist of a comprehensive collection of validated conventional data as well as processed data from the special ALPEX platforms.

To accomplish these tasks is the goal of the International ALPEX Data Centre hosted by ECMWF. It is supported by 12 National ALPEX Data Centres in the participating countries and 12 Special ALPEX Data Centres, the latter having an international role in the compilation of special platform data, such as NAVAID radiosondes, microbarographs, research aircraft, etc. Nearly all data sets will be archived by the World Data Centres and will be available to all users.

Organisational Aspects

ALPEX is the last of the major international field projects of the Global Atmospheric Research Programme (GARP). As with earlier GARP Projects, it has been planned by a number of international working groups established jointly by the International Council of Scientific Unions (ICSU) and the World Meteorological Organisation (WMO). Their Joint Scientific Committee (JSC) and its "Joint Planning Staff" located at the WMO Secretariat in Geneva are responsibile for the implementation of ALPEX.

Most participating countries have their own national organisations and programmes, but all of them have appointed so-called "Focal Points", i.e. individuals who represent their country in all international scientific and technical matters. The field operations are directed by the international Scientific Director (J.P. Kuettner) and the Operations Director (T. Gutermann) stationed at the ALPEX Operations Centre in Geneva.

- Stefano Tibaldi

* * * * * * * * * * *

THE POLAR VORTEX SPLIT, 17 - 24 JANUARY 1982

At 1400 GMT on 18 January 1982, the following STRATALERT was issued by the Freie Universität, Berlin:

A% 10-MBAR ANALYSIS OF 17 JAN

LOW 281 77N 90W, COLD MINUS 77 77N 13W
HIGH 316 51N 153E, WARM MINUS 28 53N 136E
HIGH 312 46N 52W, WARM MINUS 31 50N 80W

B% NO SATELLITE DATA AVAILABLE TODAY. ACCORDING TO FORECASTS OF THE EUROPEAN MEDIUM RANGE CENTER A SPLIT OF THE POLAR VORTEX IS EXPECTED DURING THE NEXT FEW DAYS. WEAK HEAT FLUX NORTHWARDS THROUGH 60N at 30 MBAR ON 17 JAN. INCREASING HEIGHT WAVE NUMBER TWO.

C% STRATALERT EXISTS. WARM AIR OVER EASTERN CANADA INTENSIFYING. MAJOR WARMING IN PROGRESS. SPLIT OF VORTEX EXPECTED WITHIN THE NEXT FEW DAYS.

Figure 1 shows the initial 50 mb analysis of 17 January 1982. It was the forecast from this analysis which made it possible to mention the expected split of the polar vortex in the STRATALERT, which is issued daily. Figure 2, on the cover of this Newsletter, shows the D+7 forecast, valid on 24 January; the major split of the flow in a wavenumber 2 pattern is evident. The verifying analysis, Figure 3, shows that the split was indeed very well advanced at this stage. In forecasting these stratospheric flow patterns during such warming events as this, there is a tendency for the model to deepen the lows too much and to misplace the centres too far to the east; this can be seen especially in the position and intensity of the low over Siberia. From a time sequence of stratospheric maps (not shown) for the warming period, it can be seen that the individual lows are travelling faster westward in the analysis than in the forecast. A wavenumber-frequency analysis carried out for the winter 1980/81 forecasts has clearly shown that this behaviour of the model is a systematic error, as it could be shown that, in the forecast, the ultralong planetary waves travelling westward are overdeveloped and too slow compared with the analysis.

- Ernst Klinker

* * * * * * * * * * *

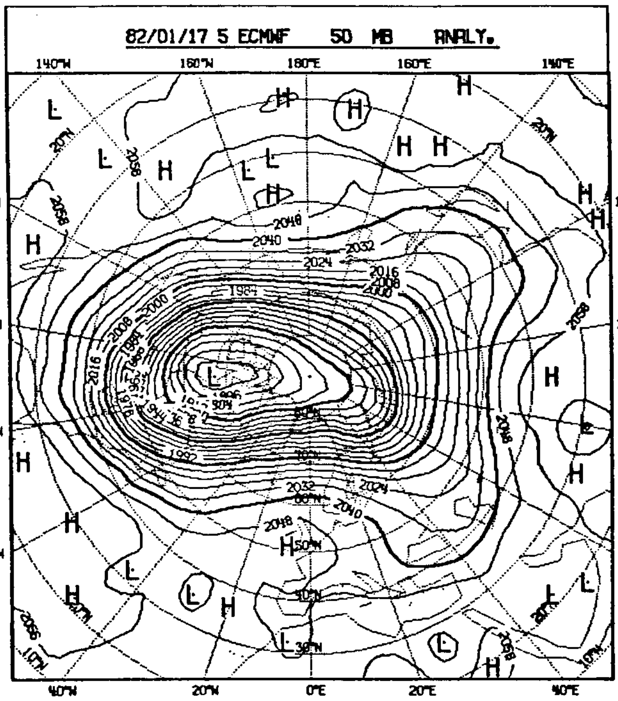

Figure 1

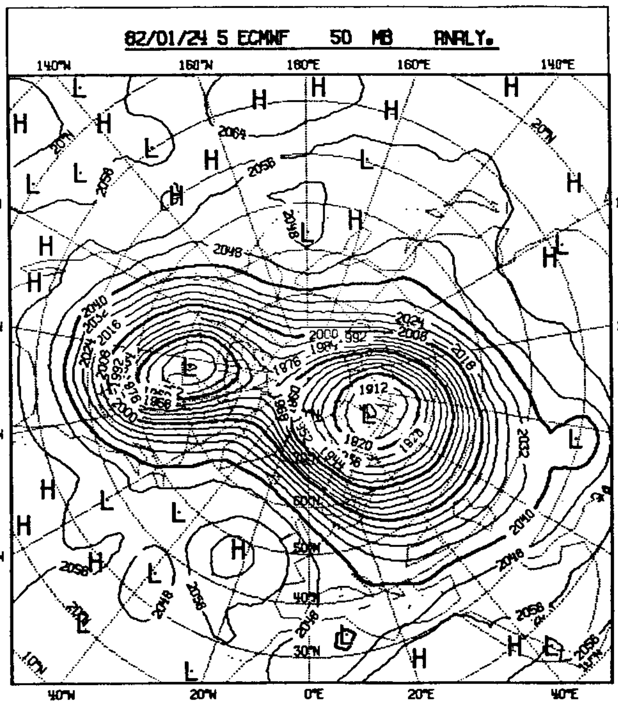

Figure 3

THE USE OF TEMPERATURE AT 2M ABOVE THE MODEL SURFACE:

ONE OF THE ECMWF EXPERIMENTAL PRODUCTS

A recent ECMWF Newsletter article by Austin Woods (Number 13, February 1982) described some of the ECMWF experimental products including such parameters as total precipitation, wind at the 10m level (above the model surface), temperature at the 2m level and the model predicted cloud amount. It was stressed that "extreme caution should be exercised in making use of these parameters." This is true especially when attempting to use the direct model output from these parameters as guidance in local forecasting of weather conditions. The model orography and the coastlines are not in agreement with reality; with a model resolution of 1.875 degrees in latitude and longitude, local details in the wind and temperature pattern will not be reflected in the gridpoint values.

Temperature at 2m level as a predictor

In recent verification studies of model predicted precipitation (Åkesson, Technical Memorandum No. 24; Johannessen, Technical Memorandum No. 51), it was shown that, for such a discontinuous parameter, area averaging has to be applied in order to achieve meaningful verification results and to make the best use of the forecast. In a further study of direct model output verification of near surface weather parameters at 17 locations in Europe (Åkesson et al, Technical Memorandum No. 47), it was found that over homogeneous terrain, the 2m temperature can give useful guidance in predicting local temperature conditions up to 120 hours of forecast time. For stations in mountainous areas or under maritime influence, the mean error becomes large and the bias has to be taken into account when using the 2m temperature as guidance in local forecasting.

There are two ways to overcome the problem:

1. Instead of using the absolute value of the 2m temperature, the predicted change in temperature (ΔT) is computed, that is

   $$\Delta T = T_F - T_A$$

   where T_F and T_A are the forecast and analysed values respectively.

   By taking the difference from the analysed temperature, the bias ($T_O - T_A$ where T_O is the observed temperature) is implicitly corrected but the mean error depending on the forecast time remains. This method does not need any further statistical investigation.

2. Stratify the forecasts into summer and winter and apply a simple linear regression between T_O and T_F on historical data.

   $$\overset{*}{T}_F = a + b\,T_F$$

   The intercept a on the T_O-axis will allow for the bias and the slope b depends on the ratio between the standard deviations of T_O and T_F. Thus, $\overset{*}{T}_F$ is the computed temperature that can be used instead of the T_F direct model output. The advantage of this procedure is that model errors depending on the forecast time will be statistically corrected. The disadvantage is that the prognostic equation has to be developed from historical data which is normally not easy to obtain. Further, in order to derive stable equations, a minimum number of events is needed; normally two years of data are regarded as being sufficient. If the periods are too short, climatological anomalies will strongly influence the result. On the other hand, if the linear regression technique is limited to the one predictor case, the statistically interpreted forecasts will still be very close to the direct model forecast and there is no danger of the equation becoming unstable when applied to independent data.

Examples of 3-day mean temperature forecasts

Even though temperature as a continuous parameter can be used to give useful guidance in local forecasting right into the medium range at many locations, it is still affected by model errors in phase and amplitude. One way to cancel out at least minor phase errors is to average in time and space. Figure 1 (heavy line) gives an example of a 3-day mean temperature timegraph observed at Hannover in Germany, January to March 1981. The terrain around the observing site is homogeneous and therefore suitable for application of the model forecast, which is shown in the same figure (light line). The forecast time is 120 hours and the direct model output of the 2m temperature is averaged over the 4 gridpoints surrounding the station. The absolute correlation between both curves is .87. Note that the period of ten days or so in February when the averaged temperature stayed below $0^{o}C$ is captured by the model 5-day forecast as well as the subsequent warming. Large model errors can be seen, however, towards the end of January, at the beginning and end of February and in mid-March. At least during this 3 month period the model exhibits a tendency to be one or even two days too slow to capture extreme changes.

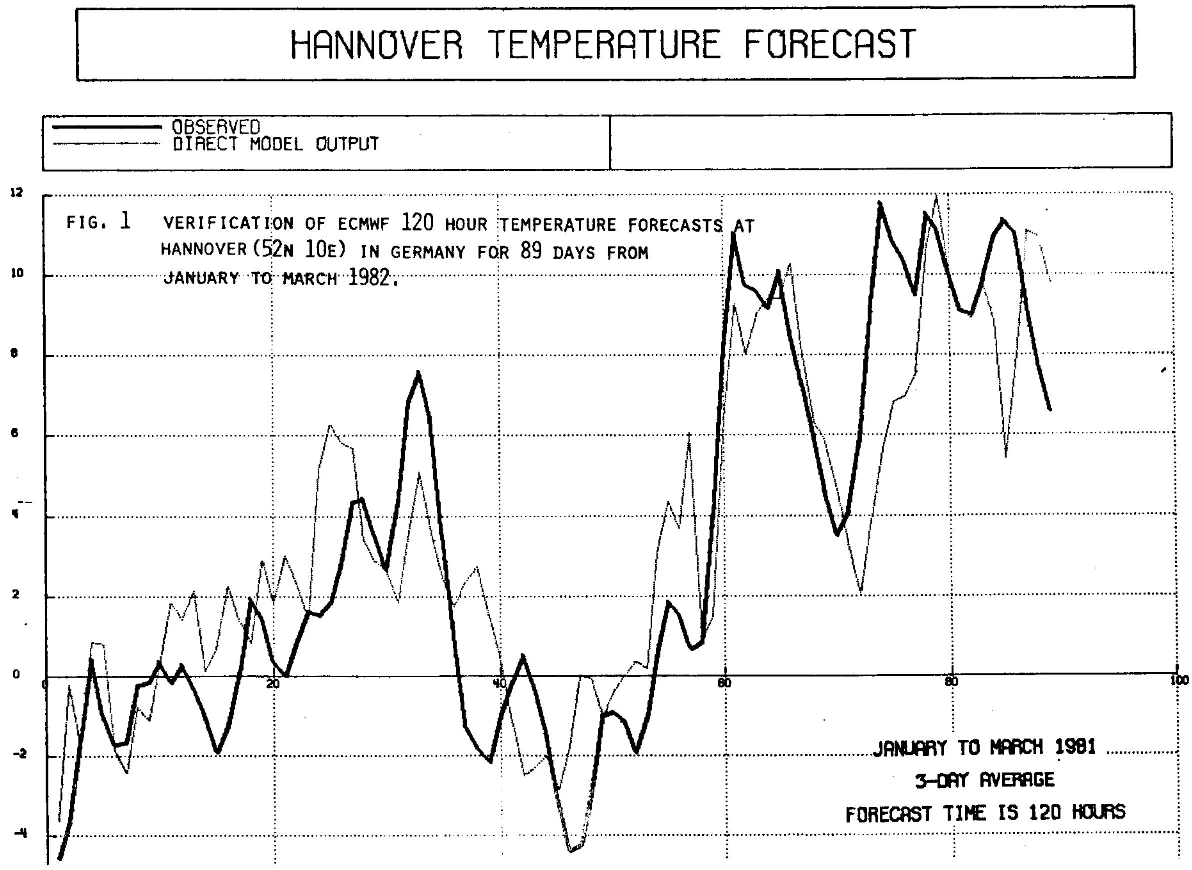

FIG. 1 VERIFICATION OF ECMWF 120 HOUR TEMPERATURE FORECASTS AT HANNOVER (52N 10E) IN GERMANY FOR 89 DAYS FROM JANUARY TO MARCH 1982.

Figure 2 shows the observed and model predicted 2m temperature for the same period at Braunlage (52N 11E) in Germany. Braunlage is situated in a small mountainous area in northern Germany at an elevation of about 600m. These mountains are not reflected in the model orography. Thus a strong bias becomes apparent between both curves. After a linear regression, where the 1000mb temperatures averaged over 16 gridpoints and forecast step 108 hours, was chosen as a predictor (heavy dashed line), the observations can be fitted satisfactorily. The 16 point average resulted in a smoothing avoiding extreme forecast errors but the main effect is the correction for the bias.

With caution, such a prognostic equation with one predictor derived from one season only may be applied to independent data in the following year. Climatological anomalies of each of the seasons have to be taken into account. Especially over snow covered ground in winter, the direct model output gives misleading results, as the change in snowcover is only poorly reflected in the ECMWF assimilation cycle.

Careful testing of experimental products in such statistical applications as these will be required by users in Member States; with experience, it is expected that output products from ECMWF will be applied increasingly to design user-oriented medium range forecasts for the end customers of the Meteorological Services of the Member States.

ECMWF is currently building up a European archive of its model output. This will, together with user-friendly retrieval facilities, enable the Member States in future to carry out their own local verification and interpretation studies based on ECMWF products. This system will be introduced during the autumn ECMWF Seminar/Workshop on the "Interpretation of numerical weather prediction products" and participants of the Workshop will get the opportunity to have practical instruction and to gather their own first experience with the ECMWF European Data Base and the statistical verification and interpretation package that will be provided with it.

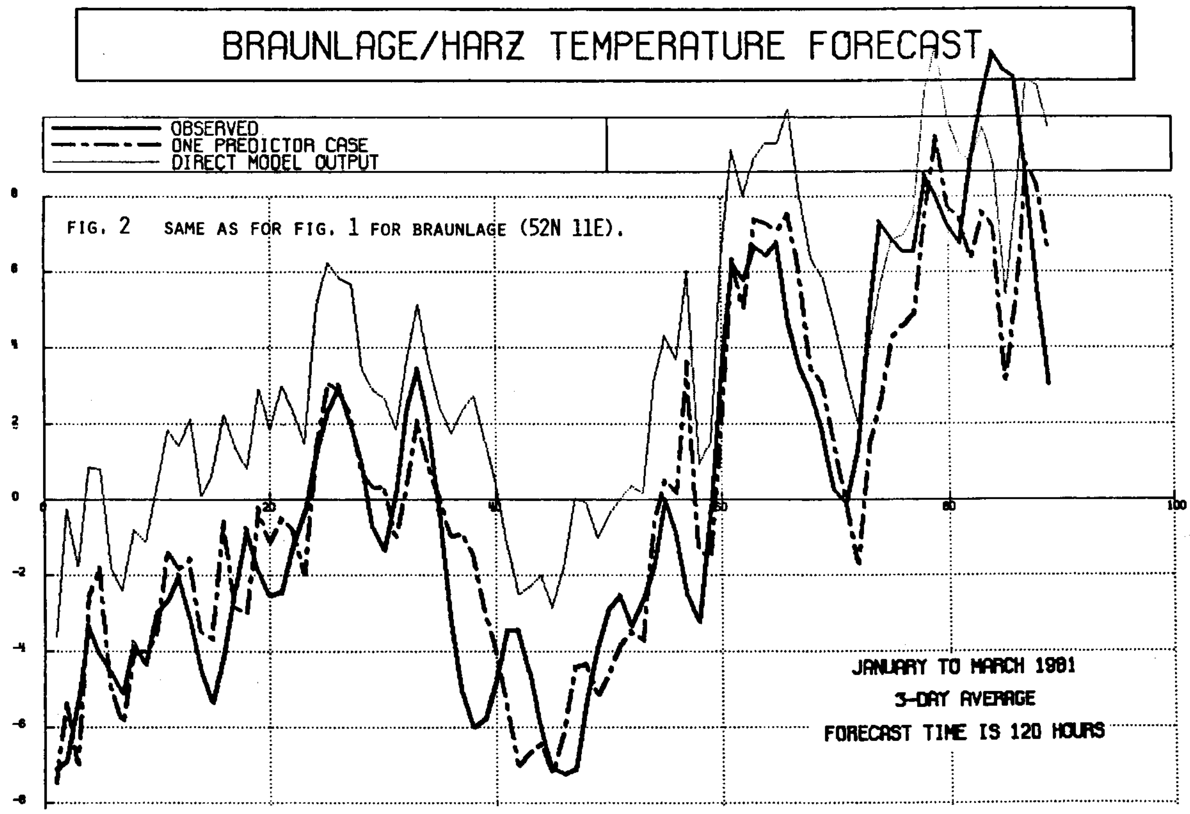

FIG. 2 SAME AS FOR FIG. 1 FOR BRAUNLAGE (52N 11E).

- Horst Böttger

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

COMPUTER DIVISION STATUS AND PLANS FOR 1982

In the following, I try briefly to outline the current status and plans of the Computer Division of ECMWF:

Personnel

(i) On 1 February 1982, the former Telecommunications and Graphics Section was renamed "Communications and Graphics Section" (CGS). At the same time, the section was increased to a total of 7 staff by transferring two members of staff from the Operating System Section (OSS). The new section comprises the following people:

Peter Gray, Head of Section	(transferred from OSS)
Jens Daabeck	(since 1 February 1982)
Richard Dixon	(transferred from OSS)
Klaus Gradischnig	(since 18 January 1982)
Emil Hinrichs	(to start on 1 July 1982)
Philippe Quoilin	
Howard Watkins	

The section is responsible for the software maintenance and development of the RC systems and the Cray-Cyber link. Furthermore, it coordinates all graphics activities, including the Aydin colour raster display. In addition, the section will be in charge of the planning, acquisition and implementation of the planned high speed data highway and data handling processor.

(ii) Following Peter Gray's transfer to head the new section, CGS, his successor, Claus Hilberg, has been appointed Head of Operating Systems Section, with effect from 1 June 1982.

(iii) After the departure of Pierre-Pascal Regnault, the workload of the Computer Operations Section (COS) was distributed differently and a successor is currently being sought, with a slightly modified description of duties, as an "Operational Analyst".

Current Status

The second Cyber, a Cyber 170-730E, is currently undergoing its final acceptance. As no serious errors have occurred so far, it can be expected to pass its trials successfully on 17 March 1982.

On 12 March 1982, the Centre took delivery of a Datagraphix alphanumeric COM device. It will go into off-line operation as soon as possible and will provide the Centre with the capability to produce microfiche on site.

Plans for 1982

Central Computer Facilities

Besides the acquisitions already mentioned, there are no plans for enhancing the in-house computer facilities this year. However, from 1 July 1982, for 18 months, the Centre has acquired 2,250 hours' elapsed time on the Cray 1 machine at Harwell. This time will provide additional resources for Centre research usage.

- Geerd-R. Hoffmann

* * * * * * * * * *

*BUFFER SIZES FOR CRAY JOBS DOING LOTS OF SEQUENTIAL I/O

If you have a Cray job which does a large amount of sequential blocked I/O, then please ensure that it uses a buffer size which is larger than the default. The default buffer size is 4 blocks long (4 x 512 words). If more than 4000 blocks, which corresponds to about 5% of a DD-19 drive, are to be transferred, then you should include an 'ASSIGN' control card in your JCL (or call the ASSIGN subroutine) to increase the buffer size to 18 blocks (22 octal),

e.g. ASSIGN(DN=DATA,BS=22)

If you are processing 50% or more of a DD-19, i.e. 40000 or more blocks, then use a buffer size of 36 blocks,

e.g. ASSIGN(DN=DATA,BS=44)

Note that if the dataset is to be ACQUIRED, the ACQUIRE must be made before the ASSIGN.

The amount of I/O which a job performs is reported in the logfile at job termination,

```
e.g.    DISK BLOCKS MOVED ---- 62890
        PHYSICAL I/O REQUESTS-  4193
```

The memory used for disc buffers is part of the job's field length, so if you only need the dataset for part of the time, remember to RELEASE it as soon as it is no longer required. The memory used for the buffers is <u>not</u> released but can be re-used as buffer space for other datasets. To actually release the memory you need to use an RFL statement.

- Neil Storer

* * * * * * * * * *

FORTRAN 5 OPTIMISATION TECHNIQUES FOR THE USER

Part 1, from a talk presented by Paul L. Derby, CDC, at the VIM 35 Language Processors Committee.

The least expensive optimisation technique is the use of the FTN5 OPT parameter. Since the source remains unchanged, the program is still portable. It is just as maintainable as it was before the OPT parameter. Very little programmer time is required. An OPT value of 2 or 3 will cause the FORTRAN 5 compiler to spend extra time during the compilation process that will result in reduced execution time. The FTN5 compiler does the following to a program compiled with OPT values of 2 or 3:

1. Uses registers instead of memory for intermediate results. There are cases in which variables will be optimised out of existence.

2. Moves invariant code out of DO loops. Invariant code is code that can be evaluated once before the DO loop and not reevaluated each trip through the DO loop.

3. Replaces time consuming operations with faster operations. A hardware multiply is faster than a divide. A multiply of a 'reciprocal' is the same as a divide except that it takes less processor time.

4. Eliminates unnecessary operations by removing any 'dead code'.

The actual optimisation performed at each OPT level is described in Chapter 11 of the FORTRAN 5 reference manual.

Source Optimisation Overview

All of the following need to be determined before source code optimisation begins:

1. that the program generates the correct results, and

2. that the program truly does need to be optimised, and

3. that the FTN5 OPT parameter has been exploited.

If all of the above are true, it might be necessary to make changes in the actual source code. The first task is to determine where the central processor is spending most of its time.

It is not productive to optimise a program code sequence that is only executed once or twice in the program. Determining where the CPU time is spent can be accomplished by using the Fortran library routine SECOND. SECOND returns the CPU time since the beginning of the job. If the result of the SECOND function is saved before a code sequence and compared with the result from SECOND after the code sequence, the amount of CPU time used in the code sequence can be determined.

Source Changes to Reduce Central Processor Time

SECOND should be used particularly to monitor DO loops and array manipulations since these repetitive operations generally use the largest amount of CPU time. Optimisation should be concentrated on code sequence in DO loops. Straight line code optimisation is not likely to be beneficial. The following source code optimisation concentrates on code in DO loops and array manipulations.

"Unrolling" DO loops - Each pass through a DO loop consists of testing the loop control variable, performing a unit of "work", incrementing the loop control variable and transferring to the loop control variable test. The amount of "work" performed in each pass of the DO loop can be increased by unrolling the DO loop.

```
          Normal                      Unrolled 2

          DO 100 I = 1, 300           DO 100 I = 1, 300, 2
      100 X (I) = 0.0                 X (I) = 0.0
                                  100 X (I + 1) = 0.0
```

This unrolled loop will probably execute twice as fast since the overhead of the test, increment and jump in the "normal" case, is larger than the actual work performed. Unrolled loops generally require more scratch registers so it is suggested that unrolling be limited to 3.

Combine loops - "Similar" back to back DO loops can be combined into one loop with one loop control variable. Again, the amount of work performed in each pass through the loop will increase in proportion to the overhead of the test, increment and jump.

Array subscripts - The first subscript (row) of a multi-dimensional array should be varied faster than the second (column) etc. The reason for this is that multiple dimension arrays are stored by rows then columns then planes etc. To get from one row to the next consists of incrementing by one memory location. To get from one column to the next requires incrementing by the number of rows. The process is even more involved if the plane subscript is incremented faster than the "lower" subscripts. There is a double price to pay for incrementing subscripts in the wrong order. The obvious price is the extra calculations to jump around in memory. The other price is due to the way memory is "banked" in the hardware. Typically, each 8th, 16th and 32nd word in memory lies in the same bank. Since only one word per bank can be accessed at a time, referencing a word in the same bank increases the CPU overhead. Incrementing to the next memory locations uses a different bank and thus minimal CPU overhead.

Array sizes - the first N-1 dimensions of a multi-dimensional array should be a power of 2. To index from one column to the next column or one plane to the next plane requires an expensive multiply instruction. The multiply is used to index (column) * (rows). If the number of rows is a power of 2, the multiply can be replaced by a faster executing shift instruction.

The first word of each back-to-back array should be offset by multiples of 4. Consider the following code fragment:

```
          DIMENSION A(128), B(128)
          DO 100 I = 1, 128
      100 A (I) = B (I) + 10.0
```

In this example, the elements of A and B that are required in each pass through the loop are in the same memory bank, thus hardware conflicts result. If the dimension of A is increased to 132, the required element in each pass through the loop will be in different banks. The loop will run twice as fast.

Avoid Mixed Mode Arithmetic - Mixed mode arithmetic should be avoided like the plague. When it cannot be avoided, there are some tricks to minimise the problem. Consider the following statement, keeping in mind that expressions are evaluated left to right.

```
I1 = R1 + I2 + I3 + I4
```

I2 must be converted to real to be added to R1, etc. The result is converted to integer when it is stored in I1. One technique to minimise the cost is to put R1 on the far right. The best technique is to change the statement to:

```
I1 = INT (R1) + I2 + I3 + I4
```

The problem is even more complex and expensive when double precision and complex variables are involved.

Avoid COMMON - COMMON, especially for scalar variables, should be avoided unless absolutely required to pass variables between modules. If variables (scalars, DO loop control variables, etc.) are in common, every subroutine or function reference requires the compiler to store the variable before the reference. The compiler does this because it cannot determine whether the variable is actually used in the referenced sub-program. Placing "scratch" variables in COMMON actually turns off many parts of the FTN5 optimiser.

Avoid EQUIVALENCE classes - Equivalence classes can inhibit optimisation in less obvious ways. Consider the following example:

```
DIMENSION X(100)
EQUIVALENCE (X(1), W)
W = Y
PRINT *, X(I)
```

Without the EQUIVALENCE statement, the compiler would eliminate the statement W=Y because the value of W is not used again in the program. However, since W is equivalent to X(1), and the PRINT statement might reference X(1), the assignment statement cannot be eliminated.

Conditional Branches - The fastest conditional branch is the assigned GO TO. The slowest is the computed GO TO. The IF statement is somewhere in the middle. A three branch IF should be used instead of a computed GO TO when there are only 3 possible cases.

Variables should be initialised with DATA statements - The use of DATA statements saves CPU time and memory since no executable code is generated. Obviously, a DATA statement cannot be used if a variable needs to be reinitialised in the program. The structure of a DATA statement for multiple arrays can have an impact upon the size of the relocatable modules on disc before it is processed by the loader. Consider the following examples:

```
DATA (A(I), B(I), I=1,250) /500*0/
DATA (A(I), I=1,250), (B(I),I=1,250) /500*0/
```

Both will accomplish the same task. The second will use considerably less disc space to hold the relocatable modules in which the DATA statement resides.

<u>Formatted I/O Optimisation</u>

An area of optimisation that is rarely discussed is the CPU time required to execute formatted I/O statements. Some of the CPU time is outside the user's control and a part of the normal overhead of preparing variables in internal format for representation on a line printer or terminal. The part that the user can control is described here. The formatted I/O processor actually "interprets" FORMAT statements. This is necessary since the processor must process compile-time formats created by the user at execution time.

Use explicit repeat counts - The format statement (10F10.5) will execute faster than the (F10.5,F10.5...) variant and even faster than the (10(F10.5)) variant.

Put "positioning" spaces in the specification field - Consider the following two formats:

```
901 FORMAT (10 (1X, F10.5))
902 FORMAT (10F11.5)
```

Both format statements will generate the same output (provided the data is well behaved). The format interpretation of statement 902 will execute nearly twice as fast as statement 901.

Avoid blanks in execution time formats - Embedded blanks in execution-time format statements should be avoided. The compiler will remove embedded blanks in format statements declared with a FORMAT statement at compile time. Execution time declared format specifications are the responsibility of the user. The execution time format processor will remove the extra blanks but it takes more CPU time.

Array I/O is faster than multiple scalar I/O - The I/O list that is processed by the formatted I/O processor can have different formats. The most efficient I/O list entry is the entire array or an array with an implied DO and successive array elements specified. Less efficient is a list of scalar variables. The formatted I/O processor can be persuaded into processing scalar variables as an array by equivalencing the scalar variables to an array and processing the array. Keep in mind that the above use of EQUIVALENCE might degrade the optimisation of DO loops.

Intermediate files - Files that hold intermediate results to be read in later but never printed,should use unformatted I/O. Unformatted I/O has a triple advantage over formatted I/O for intermediate files.

1. There is no overhead of interpreting the format statement.

2. There is no overhead in converting internal representation of variables to printer-readable display code.

3. There is no loss of precision since each bit of the internal representation of the variable is recorded on the file.

Internal files - Internal I/O is the same as formatted I/O except that the destination (or source) is a character variable instead of a file. The same format processing is required by internal files. All of the above guidelines should be used.

* * * * * * * * * * *

Part 2 will follow in the next issue of the Newsletter.

*SPEED OF ARITHMETIC INSTRUCTIONS ON THE CYBER FRONT ENDS

This note provides details of the speeds of the arithmetic instructions on the two Cybers. This is part of a full report on the differences between the machines currently being researched.

Some points to note are that the rounded floating add and difference are 2 and 2.5 times slower on the 730E than their non-rounded counterparts. Also double precision sum and difference are 8 to 11 times more costly on the 730E than on the 175.

Users are therefore advised:

1. not to use Floating Point Round (which will continue to be the default) on add and subtract, unless they really need to.

2. not to run on the Cyber 730E any programs which heavily use Double Precision Floating Point Arithmetic (see User Support to arrange for appropriate ST parameter to put on the jobcard).

Code	Description	Time 730E	Adjusted time 730E	Time 175
24xxx	Normalise	336	126	75
25xxx	Round and Normalise	672	252	75
30xxx	Floating sum	224	84	100
31xxx	Floating difference	224	84	100
32xxx	Floating D.P. sum	2296-3080	861-1155	100
33xxx	Floating D.P. difference	2296-3080	861-1155	100
34xxx	Round Floating add	448	168	100
35xxx	Round Floating difference	560	208	100
36xxx	Integer sum	112	42	50
37xxx	Integer difference	112	42	50
40xxx	Floating product	672	252	250
41xxx	Round Floating product	672	252	250
42xxx	Floating D.P. product	672	252	250
44xxx	Floating Divide	2968	1113	500
45xxx	Round Floating Divide	2968	1113	500

Notes:

1. All times are in nano-seconds (10**-9).

2. The adjusted 730E times are the real times multiplied by 3/8ths. This is the time that the user is charged for.

3. Some of the 730E execution times are data dependent.

4. D.P. means double precision.

- Gary Harding

* * * * * * * * * *

*USING FILESET FOR PROJECT DEVELOPMENT

1. Introduction

When designing and developing computer programs which are non trivial (i.e. more than a few pages of code) the value of a modular approach has long been recognised. This involves dividing the program into routines, each performing a logical task. The routines may then be written and tested, each one separately, before attempting to link them together into the final program. The normal method of storage for source code of this nature at ECMWF is the CDC UPDATE library. This article suggests how FILESET may be used as a tool in the development of such an UPDATE library.

2. The disadvantages of using UPDATE for development

An UPDATE library is ideal for maintaining source code requiring minor alteration from time to time. If major modifications, deck replacements, partial code extractions, etc. are frequently required, as in the development phase of a project, UPDATE can be quite expensive in terms of computer resources. The UPDATE process is an all-or-nothing system. The complete process must be pre-defined to UPDATE - it is not possible to request individual sub-processes. Thus FILESET, which is fast and interactive in operation, has certain attractions as a development aid.

3. A suggested methodology using FILESET

A fileset can be created for each project being developed. A large project may be split into sub-projects for this purpose. As routines are written, they are added to the fileset in the form of UPDATE decks, and given the name DECKS/routine name. Common decks may usefully be grouped together as a single fileset member called DECKS/AAACOM, enabling their extraction at the front of any file created by a multiple GF (see later) due to the alphabetic order of the fileset index. Job control, libraries, etc. associated with the project may be retained on the same fileset using different group names, e.g. JOBS/JOB1, LIBS/LIB1, etc.

The routines are coded, test compiled, tested, etc. one at a time. They are entered on the fileset as UPDATE decks, preceded by an appropriate *DECK card. An UPDATE input file containing all of the source code is created thus:

```
GF,UPDIN,DECKS/*,PO=NR
EDIT,UPDIN
D /*EOR/*
END
```

giving the required file on UPDIN. An UPDATE library, compile file, etc. may now be produced by:

```
UPDATE,I=UPDIN.
```

Testing of the complete program may then be carried out using the UPDATE generated COMPILE file. Further routine modifications, additions, etc. may be performed at the FILESET level until a reasonably stable UPDATE library can be produced. When this stage is reached, the UPDATE NEWPL may be catalogued and used as the basis for production code.

- Rex Gibson

* * * * * * * * * *

* DO IT YOURSELF ARCHIVING

An archiving package, ARCHSET, which allows users to perform their own archiving and retrieval of Cyber permanent files is now available.

The archiving and retrieval mechanism is based on FILESET. Initially, a fileset is created on temporary disc space. Normal fileset commands are used to manipulate this fileset which is then written to tape by means of a special command. This process does not purge the permanent files. Retrieval is accomplished by copying from tape to temporary disc when, once again, normal fileset commands may be used. Changes (i.e. additions, replacements or deletions) to an existing archived fileset must be carried out using 2 tapes for security reasons.

It is important to use ARCHSET only for archiving permanent files which are unlikely to be required for lengthy periods. Frequent use of the package will result in excessive tape mounts and temporary disc space usage.

To make the task of archiving as easy as possible, the interactive procedure SELPUR belonging to the procedure library (PROCIN) has been enhanced. It now allows the user to select files for archiving. An option allows selection of files unattached for N months. SELPUR generates a batch job which can be submitted to perform the archiving and purge the files.

A Bulletin describing ARCHSET is being printed.

- David Dent

* * * * * * * * * * *

*PURGING UNUSED CYBER FILES

With the introduction of a user file archive facility (ARCHSET - see preceding article) we can now deal more effectively with Cyber permanent file storage. The philosophy being adopted is that the <u>user</u> should be responsible for the control of his own files. For some time, users have been able to list their files (via AUDIT and AUDPF), and copy them to tape using standard copying utilities. The introduction of ARCHSET now gives users a much more comprehensive file management facility.

The user permanent file area (SYSSET) is always heavily used, and often overfull: even allowing a generous buffer area, this has caused problems in the past. A recent survey of the files held in SYSSET showed that some 20% had not been accessed for 6 months or more.

The Cyber permanent file area is meant for active files only; those which will not be required for a period should now be archived by users. All users are requested to check their file holdings, to purge unwanted files and archive files which are not being used now, but may be required later. After successful archiving, they should purge the file from SYSSET.

In order to ensure the file area SYSSET is kept for active files only, files unused for a specific period will be purged on a monthly basis without any prior warning to users. The first such purge will be done in 6 weeks' time, i.e. mid-May, when all files not accessed for 6 months or more will disappear. It is the user's responsibility to take an archive copy before those files disappear. The user can then use that copy to restore the file, if it is wanted sometime in the future. A similar purge will be carried out at every subsequent mid month. Should it be necessary to reduce the period for which files can remain unused in SYSSET, a warning will be given via the usual means.

The normal control of SYSSET space will continue to operate as described in ECMWF Computer Bulletin B8.1/2.

- Andrew Lea

* * * * * * * * * *

* USE OF FILESPACE ON TEMP

Short-lived permanent files are commonly put onto the TEMP set. This set now occupies 2 885 units, so that files may reside on either unit. The master unit has VSN=TEMP01; however, if this becomes full, files will be automatically placed on the second unit, providing that space is requested in the correct way:

```
MOUNT, VSN=TEMP01, SN=TEMP.
REQUEST, lfn, SN=TEMP.
```

Some user jobs request space using:

```
REQUEST, lfn, SN=TEMP, VSN=TEMP01.
```

This method does not allow the file to reside on the secondary unit and the job fails if TEMP01 becomes full. For this reason, the VSN parameter should NOT be put on a REQUEST statement for TEMP space. In general, the VSN for a disc pack should ONLY be specified on a MOUNT statement.

- David Dent

* * * * * * * * * *

* DOING THINGS WITH OUTPUT FILES

On both the Cyber and the Cray, the normal sequence of events at job termination includes the process of appending the dayfile to the end of the OUTPUT file before sending it to the output queue.

However, it is possible for a user to exercise some control over the destination of both the output file and the dayfile.

The basic mechanism returns the job output to the location from where it was submitted.

1. <u>In a Cray job</u>

   a. Use a DISPOSE statement to change the destination of the output file $OUT, e.g.

      DISPOSE(DN=$OUT,TID=C,DEFER)

      sends $OUT to the Cyber printer regardless of where the job came from. The statement may conveniently be placed near the beginning of the job, as disposition is deferred until the job terminates. The destination of the dayfile is unaffected so that the job output is now split into 2.

   b. DISPOSE(DN=$OUT,SDN=LISTING,ID=uid,DC=ST,DEFER)

      causes the output file to be retained on Cyber disc as a permanent file. Again the dayfile is routed to the location where the job originated.

   c. DISPOSE(DN=$OUT,DC=SC)
      ASSIGN (DN=$OUT,DC=SC)

      causes all output to be sent to a 'scratch device' (i.e. thrown away) <u>including the dayfile</u>. This sequence could usefully be positioned as the last control statements of the job when terminating normally. If it is not placed after EXIT, an abnormal termination will therefore print job output in the usual way.

2. <u>In a Cyber job</u>

   The ROUTE control statement serves a similar function to DISPOSE, e.g.

   a. ROUTE,OUTPUT,DEF,TID=C.

      However, in this case, the dayfile is duplicated and included with the rerouted output as well as being sent to the original destination.

   b. ROUTE,OUTPUT,DEF,DC=SC.

      Causes the output file to be scratched but the dayfile is retained and routed to the originating location.

   c. APR(11)

      A convenient trick for terminating the job and scratching ALL output, including the dayfile. Please note that this is an undocumented 'feature' of NOS/BE and hence cannot be guaranteed to work in future systems.

- David Dent

* * * * * * * * * *

*LIBRARY CHANGES

The changes to ECLIB which were announced in the February Newsletter (page 20) have now been implemented. These are:

New routines:

MAKEPL, ORDERPL (Cyber)	creation and organisation of UPDATE libraries
UPDLIST (Cyber)	indexed listing from UPDATE source file
SD2UV,UV2SD	conversion between wind components and speed/direction
POPCNT (Cyber)	obtains number of 1 bits in a word
XREF (Cray)	global cross reference mapping

Changes:

FFT99,FFT991	improved (30% faster) versions of existing routines
CONVAR	improved, easier to use package for map conversions

An additional change has been made to the routine INDEX. Since Fortran 77 compilers recognise INDEX as an intrinsic function, it is necessary to add an EXTERNAL statement to the Fortran source in order to use the ECLIB routine INDEX. Therefore, an alternative entry point INDX has been added so that the routine may be called with either name.

Programs to be compiled under FTN5 or, in the future, CFT level 1.10, should therefore be changed to either

a. use INDX, or

b. use INDEX with an EXTERNAL statement.

The procedure SELPUR has been enhanced to assist user archiving (see page 19).

- David Dent

* * * * * * * * * *

OPERATING SYSTEMS SECTION - A PROGRESS REPORT

Since the previous progress report in the February 1981 Newsletter, the section has achieved a lot. The items listed below should be mentioned in particular:

1. Preparing for the acquisition and installation of a second Cyber front end; this entailed

   i. extensive benchmarking to select the most suitable Cyber model and the optimum configuration;

   ii. upgrading to NOS/BE level 530 and introducing FORTRAN 77 Standard (FTN5) in May 1981;

iii. upgrading to NOS/BE level 538 with Multimainframe and ECS support in August 1981;

iv. enhancing the EMOS Supervisor to support the operational suite running on 2 interconnected Cybers as well as the Cray;

v. integrating the Cyber 730E such that the dual Cybers support the Cray and share the front end tasks; this is done in such a way as to enable either Cyber on its own to maintain a service by supporting

The Cray Station
The Operational Suite
The Interactive Service.

vi. The Cyber 730E was delivered in December 1981 and completed its Provisional Acceptance Test on 18 December. By 4 January, it was running production work during the prime shift and, by the end of January, it was fully integrated into the production service.

vii. After a disappointing period when hardware and software problems resulted in a rather poor Cyber service, by mid February, these problems had been overcome, and since that time, both Cybers have been giving a very good and reliable service.

2. Improving the performance and reliability of the Cray 1 system by

i. the installation of additional disc capacity in August 1981; the addition of 3 DD-29 double density discs doubled the available user disc space, and significantly reduced the amount of data being staged from the Cyber.

ii. The introduction of COS 1.10 in October 1981; because of the increased effort put into testing first COS 1.09 (which was not introduced) and then COS 1.10, this system was introduced without any problems.

iii. The development of a statistical performance analysis package to give day by day monitoring of the Cray performance, and to enable identification of jobs which are not using the Cray efficiently;

iv. the introduction of 'dynamic station buffers' allowed 140K of Cray memory, previously used by the station, to be made available to user jobs.

3. Making many additional user facilities available, including

i. the FORMAL text processor;

ii. the SPACE control card to prevent a job from running if there is insufficient space available on a private pack;

iii. a single user editor called EDIT;

iv. the COMMAND facility to issue Cyber control cards from within Fortran programs;

v. the ARCHSET utility to allow user archiving of Cyber Permanent Files to tape;

vi. installation of an offline COM microfiche production service.

During the remainder of 1982, the section will concentrate on:

1. maintaining and improving the level of service, which is by far the most important and most time consuming aspect of our work;

2. introducing the CFT 1.10 compiler on the Cray, which incorporates features of the Fortran 77 standard;

3. upgrading to NOS/BE level 552 on the Cybers, to improve the reliability and efficiency of the Cyber operating system;

4. enhancing the facilities for analysis and reporting of Cyber utilisation and performance, thereby enabling us to improve the level of service on the Cybers;

5. introduction of a 'direct spooling' of user generated FICHE files, without the user's needing to use a tape, for the COM device.

- Tony Stanford

* * * * * * * * * * *

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 129). All other News Sheets are redundant and can be thrown away.

No.	Still Valid Article
16	Checkpointing and program termination
19	CRAY UPDATE (temporary datasets used)
47	Libraries on the Cray-1
53	Writing 6250 bpi tapes (EEC parameter)
54	Things not to do to the Station
56	DISP
65	Data security on Cyber and Cray
67	Attention Cyber BUFFER IN users
73	Minimum Cyber field length
89	Minimum field length for Cray jobs
93	Stranger tapes
98	Cray symmetric multiply (rounding factors)
108	SUBMIT
114	Cray jobcard memory parameter
116	TEMP
118	Terminal timeout
119	New version of FICHE
120	Non-permanent ACQUIRE to the Cray
	Local terminal line speeds
121	NOS/BE Version 1.5, level 538 (including new Cyber job class structure)
122	Cyber FORTRAN News (level 538, FTN4 & FTN5)
126	Unnecessary waiting for permanent files
127	(25.1.82) Cyber 730E introduction - user interface advice
	IMSL Library
129	Optimisation problems with CFT 1.09
	Terminal fault reporting and testing

The following News Sheets can be thrown away since this list was last published: 127 (17.12.81).

* * * * * * * * * * *

PRIORITY GROUPS IN 1982

After initial experience of running the priority allocation scheme, it was decided to merge LOW and NORMAL priorities; this took place at the beginning of 1981. Now, after more experience, further simplification has taken place for 1982. The changes are outlined below.

i. For Member States, the numeric priority 5 will be allowed within the NORMAL group. Over-allocation at project level will not be allowed.

ii. For Centre staff, the BASE group will be abolished. There will be no upper limit to the amount of NORMAL priority that can be used, however, HIGH priority will still be restricted to 10% of what the machines are expected to deliver. This HIGH priority will be available to a few specific project names only.

In all other respects, the priority allocation scheme will be the same as that in operation now. For 1982, the priority structure will be as follows on both the Cyber and the Cray:

Priority Group	Octal Priority	Usage
Normal	1,2,3,4*5**	to be used for all normal priority work. 90% of the allocation will be available at NORMAL for Member States. No limit for Centre users.
High	6	to be used in special situations where a high priority has to be given to work. 10% of the allocation will be available at HIGH for Member States. For Centre users only specific projects, which cannot be used at any other priority, will have access to HIGH priority.
Met. Operations	10-13	to be used <u>only</u> for ECMWF's operational forecast.
Operator	14-17	used only for some housekeeping tasks and the Cray station.

* Default priority will remain at 4 for all users.
** Priority 5 only available to Member States.

Note that the NORMAL group now covers numeric priorities 1,2,3,4 (and 5 for Member States). Thus, users who wish to structure their own workload may still submit jobs at priorities 1,2,3,4 (or 5) to gain faster or slower turn around as required.

This scheme came into effect on Monday, 4 January 1982.

- Andrew Lea

* * * * * * * * * * *

Total = total usage less those jobs classed as systems overheads
OO = operational suite running
EC = Centre users
MS = Member State users, including Special Projects

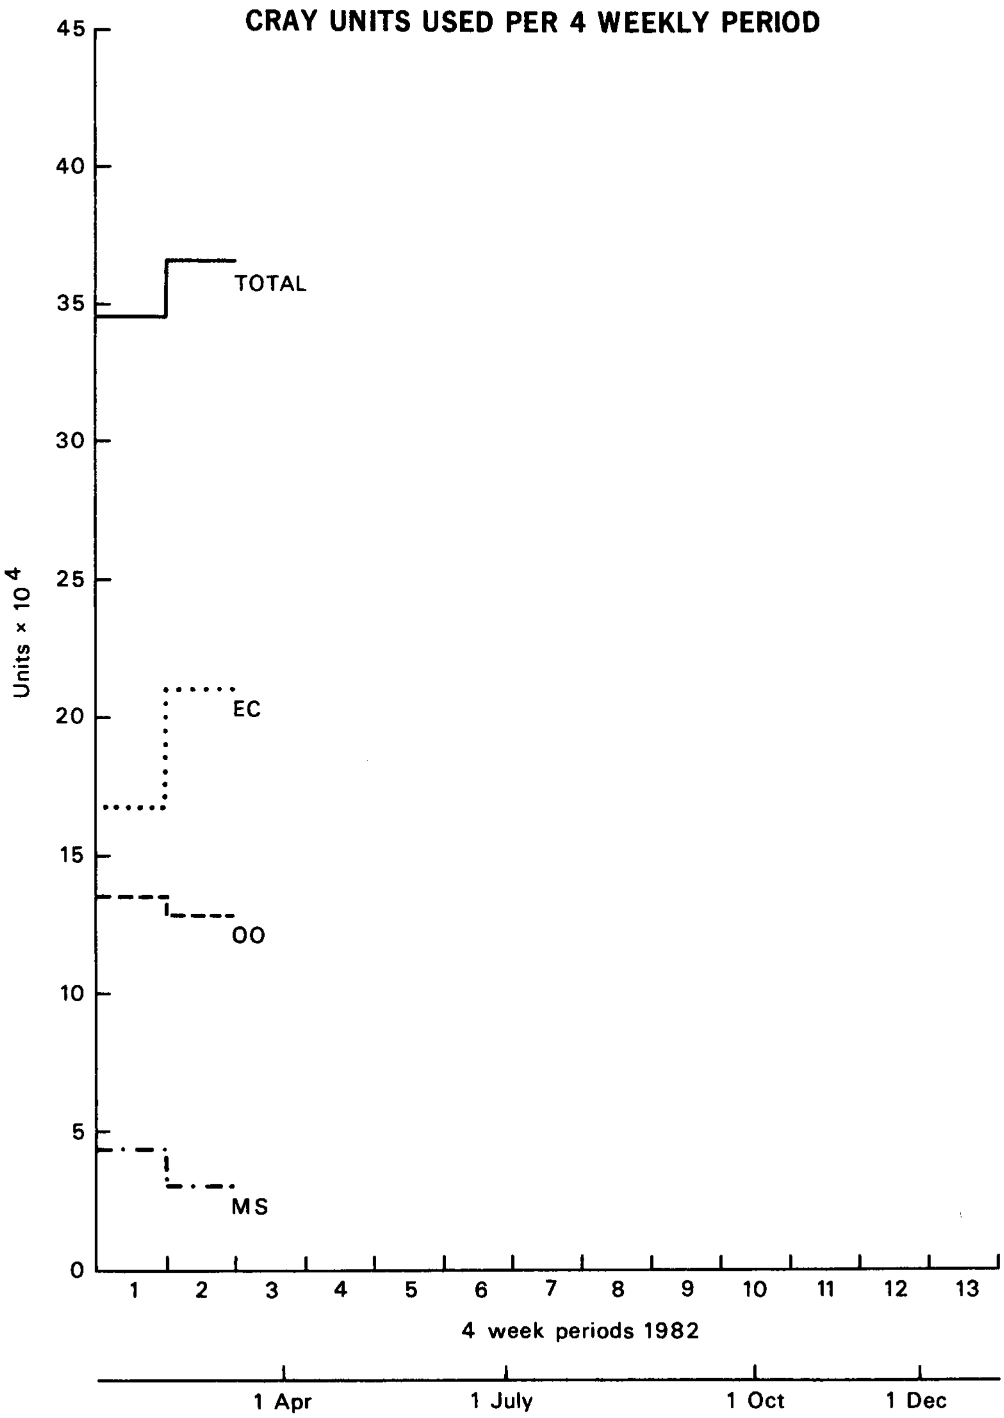

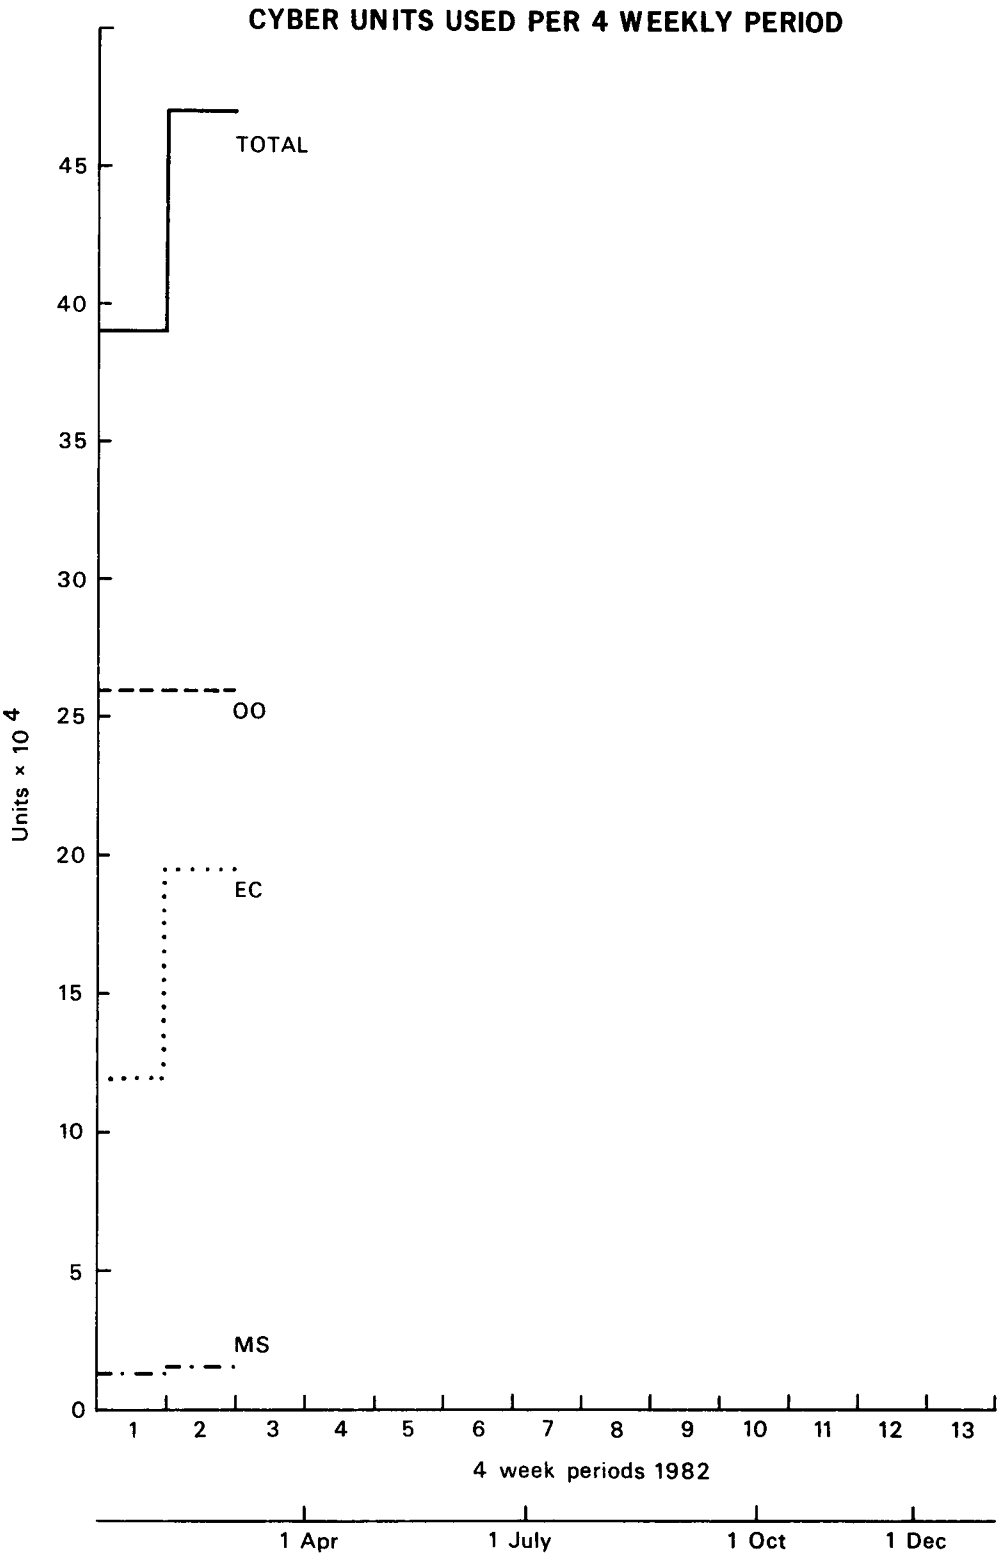
CYBER UNITS USED PER 4 WEEKLY PERIOD
TOTAL
OO
EC
MS
Units × 10 4
45
40
35
30
25
20
15
10
5
0
1
2
3
4
5
6
7
8
9
10
11
12
13
4 week periods 1982
1 Apr
1 July
1 Oct
1 Dec

NFEP TERMINAL STATISTICS

From 01/01/82 to 28/02/82

COUNTRY	TOTAL DATA (KCHAR/DAY) INPUT	TOTAL DATA (KCHAR/DAY) OUTPUT	AVERAGE DATA RATE (CH/SEC) INPUT	AVERAGE DATA RATE (CH/SEC) OUTPUT
Denmark	6.07	1425.5	115.0	229.3
F.R. Germany	4026.1	2324.9	57.4	331.5
Ireland	0.5	1816.4	79.4	214.1
Spain	0.0	261.6	0.0	13.6
France	70.3	1567.1	303.7	307.5
*Greece	0.0	422.3	0.0	8.3
*Italy	0.0	211.8	0.0	13.2
*Yugoslavia	0.0	39.07	0.0	6.3
Netherlands	2.0	488.9	65.0	69.2
Austria	0.006	890.5	61.4	209.8
Portugal	0.2	421.0	156.6	24.9
Finland	4.9	1917.8	168.1	333.7
Sweden	63.9	3985.9	316.3	303.8
*Turkey	0.0	96.3	0.0	6.4
United Kingdom	4354.3	940.4	53.3	198.2

*low speed line only

Explanations

TOTAL DATA — Input data is data acquisition (UK and Germany only) plus remote job entry (medium speed lines only). Output data is batch output (medium speed lines only) plus dissemination data.

DATA RATE — gives the average transmission speed in characters per second for an input or output file, respectively, including overheads at all levels of protocol.

Comments

DATA RATE — Input: Real-time data from Germany and the United Kingdom are sent when available and, therefore, the values for those two countries are, as expected, low.

Output: The data rate output for Portugal is abnormally low because of the changeover from low to medium speed line during the above period.

- Philippe Quoilin

* * * * * * * * *

TABLE OF TAC REPRESENTATIVES, MEMBER STATE COMPUTING REPRESENTATIVES AND METEOROLOGICAL CONTACT POINTS

Member State	TAC Representative	Member State Computing Representative	Meteorological Contact Point
Belgium	Dr. W. Struylaert	Dr. W. Struylaert	Mr. E. De Dyker
Denmark	Mrs. A.M. Jørgensen	Mr. P. Henning	Mr. H. Voldborg
Germany	Dr. W. Buschner	Dr. R. Lamp	Dr. K-O. Wegner
Spain	Mr. B. Orfila	Mr. M. Hortal	Mr. R. Font Blasco
France	Mr. J. Lepas	Mr. L. Jouaillec/ Mr. J.P. Bourdette	Mr. J. Lepas
Greece	Mr. G. Barbournakis/ Mr. D. Katsimardos	Mr. G. Sakellarides/ Mr. J. Iakovou	Mr. A. Kakouros/ Mrs. M. Refene
Ireland	Mr. W.H. Wann	Mr. D. Murphy	Mr. S. Miller
Italy	Dr. G. de Florio	Dr. G. de Florio	Dr. M. Conte
Yugoslavia	Mr. Z. Butigan	Mr. M. Gavrilov	Mr. S. Nickovic
Netherlands	Dr. A.P.M. Baede	Mr. G.D.G. Folkers	Mr. W.M. Reinten
Austria	Dr. G. Wihl	Dr. G. Wihl	Dr. H. Gmoser
Portugal	Mr. S. Cristina	Mr. M.J. Rodrigues de Almeida	Mr. O. Valente Rasquinho
Switzerland	Mr. M. Haug	Mr. G. Siegwart	Mr. M. Schönbächler
Finland	Mr. P. Nurminen	Mr. T. Hopeakoski	Mrs. P. Saarikivi
Sweden	Dr. L. Moen	Mr. S. Orrhagen	Mr. G. Salomonsson
Turkey	Major Gen. M.C. Özgül	Major Gen. M.C. Özgül	Major Gen. M.C. Özgül
United Kingdom	Mr. D.H. Johnson	Dr. A. Dickinson	Mr. C. Flood

ECMWF PUBLICATIONS

Technical Memorandum No. 44:	On the mass balance correction in diagnostics studies.
Technical Memorandum No. 45:	A survey of the performance of GCM's in the tropics - P. 1: other models.
Technical Memorandum No. 46:	Biases in the ECMWF data assimilation system.
Technical Memorandum No. 47:	First results of direct model output verification of near-surface weather parameters at 17 locations in Europe.
Technical Memorandum No. 48:	On the effect of model resolution on numerical simulation of blocking.
Technical Memorandum No. 49:	On the effect of energy/enstrophy conservation in the finite difference scheme of the ECMWF gridpoint model.
Technical Memorandum No. 50:	Report on second meeting of Member State Computing Representatives 14-16 October 1981.
Technical Memorandum No. 51:	Verification of ECMWF quantitative precipitation forecasts over Europe January 1980 to April 1981.
Technical Memorandum No. 52:	The DOCTOR system - A DOCumenTary ORiented programming system.
Technical Memorandum No. 53:	An evaluation of three ten-day forecasts for the Australian region during FGGE SOP-2.
Technical Report No. 29:	Orographic influences on Mediterranean lee cyclogenesis and European blocking in a global numerical model.
FGGE-III Daily Global Analysis Part 3:	June - August 1979.
ECMWF Forecast Report No. 14:	April - June 1981.
ECMWF Forecast and Verification Charts:	up to 28 February 1982.

* * * * * * * * * *

CALENDAR OF EVENTS AT ECMWF

28/29 April — Council 15th session

Please note that Research Department Courses A1 and A2, due to take place on 4-28 May and 1-11 June respectively, have been cancelled due to lack of response.

14 - 25 June	Research Department course B
15 - 18 June	TAC 4th session
18/19 November	Council 16th session

* * * * * * * * * *

VACANCIES AT ECMWF

Fluency in one of the working languages of the Centre (English, French, German) and a good knowledge of at least one of the others, is required for all posts.

POST: HEAD OF MODEL DIVISION, RESEARCH DEPARTMENT

MAIN DUTIES: To be Head of the Model Division in the Research Department and in that capacity be responsible for the development of the Centre's forecasting models.

To undertake, in collaboration with the Head of Research, the long-term planning of the research at ECMWF and in particular to formulate research programmes in the following areas:

- General formulation of forecasting models;
- Numerical methods for the integration of forecasting models;
- Parameterisation of sub-grid scale processes in numerical models.

Additionally, the Head of the Model Division has the responsibility for the implementation and management of the research programme, and it is expected that he/she will initiate research into unexplored areas.

QUALIFICATIONS: A university education (Ph.D.) or the equivalent in meteorology or in a related field, several years experience as a group leader for research into, and development of, numerical forecasting models and a good research record in the field of numerical weather prediction.

STARTING DATE: As soon as possible.

CLOSING DATE: 23 April 1982

POST: ASSISTANT TO THE HEAD OF RESEARCH IN THE FIELD OF EDUCATION AND TRAINING

MAIN DUTIES: Organising, in co-operation with representatives from the Member States, the Centre's educational programme. The present programme comprises training courses in dynamical meteorology and numerical weather prediction, seminars and workshops. The assistant is expected to give some of the training courses.

Acting as scientific editor, co-ordinator and supervisor of the technical production of ECMWF scientific publications. The scientific editing will involve the following publication series:

Technical Reports
Technical Memoranda
Seminar Proceedings
Workshop Proceedings
Lecture Notes

In addition, the Education Assistant is involved in other Centre publications which are edited in the Research and Operations departments.

Acting as Chairman of the Library User's Group and general supervision of the Library.

Providing general administrative assistance to the Head of Research.

In addition to the above duties, the Education Assistant is expected to either carry out independent research or take part in one of the Centre's research projects.

QUALIFICATIONS: A university education or equivalent in meteorology and several years of experience in arranging meteorological courses as well as experience in lecturing. Editorial experience will be regarded as an advantage. The candidate should have a demonstrated ability to work independently and must be able to formulate, in co-operation with the Head of the Research Department, educational and training programmes and take the responsibility for their implementation.

STARTING DATE: As soon as possible.

CLOSING DATE: 23 April 1982

POSSIBLE OPERATOR VACANCIES

COMPUTING ENVIRONMENT for all 3 possible posts:

The computer facility currently comprises a CRAY-1 and a CDC CYBER 175 linked by telecommunications computers (RC 8000) to a private network spanning Member States. A second CDC CYBER computer was added to the facility late in 1981.

POST: COMPUTER OPERATOR (CLERICAL ASSISTANT)

ENVIRONMENT: The post is in the Computer Operations Section, which is responsible for the operation of large-scale Control Data and Cray computers. The Tape Library consists of approximately 12,000 reels of tape. The clerical work is distributed between the Tape Library, Computer Operations, Electrical/Mechanical Engineering, Telecommunications and Management for the Section.

MAIN DUTIES: Updating, maintaining daily equipment and media performance records and graphs, and producing long-term statistics.

Keeping all maintenance and service sheets, and maintaining a ledger detailing history of all equipment.

Assisting in the Magnetic Tape Library, as required; duties include:

a. updating records for use and performance of tapes;

b. receiving and allocating tapes for other establishments;

c. general Tape Library duties (allocating tapes to internal users, maintaining microfiche service, removing and re-allocating tapes into fire-safe, updating various listings using a computer terminal).

Typing, copying and distributing of all Section papers, reports and procedures.

Maintaining records of shift staff, annual leave and other absences.

Maintaining stores records (incoming in conjunction with packing notes, outgoing as per weekly withdrawal forms), and periodic counting of stores.

Assisting with other duties as necessary.

These duties do not require regular shift work. However, circumstances may call for shift duties as peripheral operator.

QUALIFICATIONS: Secondary education and at least three years' experience or higher technical education. Specific training, however, will be given on the job. Good typing ability is essential.

The post demands practical work of a repetitive nature. Candidates should be methodical and have a good degree of common sense.

STARTING DATE: As soon as possible.

CLOSING DATE: 23 April 1982

For both the following posts, applicants must be prepared to work a shift system including weekends and public holidays and, in times of emergency, to stand in for operators in other shifts.

FUNCTION: CONSOLE OPERATOR

MAIN DUTIES: The Console Operator reports to a Shift Leader and is responsible for controlling the computer systems via central consoles to provide a computer service for local batch work and supporting remote batch stations and terminals.

Additionally, the Console Operator will perform the operation of off-line equipment, the control of input and output batch workloads, performing maintenance procedures on peripherals, reporting faults and monitoring performance of all equipment essential to the smooth running of Computer Operations.

QUALIFICATIONS: Candidates must have at least 2 years' operating experience in a large-scale scientific computing environment and be able to demonstrate their knowledge. CDC experience is an advantage but not essential. The post calls for a good standard of secondary education.

STARTING DATE: As soon as possible.

CLOSING DATE: 15 April 1982

FUNCTION: COMPUTER OPERATOR

MAIN DUTIES: The Computer Operator reports to a Shift Leader, and will perform the operation of peripherals, off-line equipment, the control of input and output batch workloads, performing maintenance procedures on peripherals, reporting faults and monitoring performance of all equipment and carrying out any other tasks which are essential to the smooth running of Computer Operations.

Additionally, the Computer Operator will perform duties for assisting in the control of the computer systems via central consoles to provide a computer service for local batch work and supporting remote batch stations and terminals.

QUALIFICATIONS: Candidates must have at least 18 months' operating experience in a large-scale scientific computing environment and be able to demonstrate their knowledge. CDC experience is an advantage but not essential. The post calls for a good standard of secondary education.

STARTING DATE: As soon as possible.

CLOSING DATE: 15 April 1982

* * * * * * * * *

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 Other methods of quick contact: - telex (No. 847908) - COMFILE (see Bulletin B1.5/1)		CB 037	308/309
Computer Division Head	- Geerd Hoffmann	OB 009A	340/342
COMPUTER OPERATIONS			
Console	- Shift Leaders	CB Hall	334
Reception Counter) Terminal Queries) Tape Requests)	- Judy Herring	CB Hall	332
Operations Section Head	- Eric Walton	CB 023	351
Deputy Ops. Section Head	- Graham Holt	CB 035	209
DOCUMENTATION	- Pam Prior	OB 016	355
Libraries (ECMWF,NAG,CERN,etc.)	- John Greenaway	OB 017	354
METEOROLOGICAL DIVISION			
Division Head	- Frédéric Delsol	OB 008	343
Applications Section Head	- John Chambers	OB 007	344
Operations Section Head	- Austin Woods	OB 107	406
Meteorological Analysts	- Ove Åkesson	OB 106	380
	- Veli Akyildiz	OB 104A	379
	- Horst Böttger	OB 130	310
	- Rauno Nieminen	OB 104A	378
	- Herbert Pümpel	OB 106	380
Meteorological Operations Room		CB Hall	328/443
REGISTRATION (User and Project Identifiers, INTERCOM)	- Pam Prior	OB 016	355
Research Department Computer Co-ordinator	- Rex Gibson	OB 126	384
Systems Software Section Head (Acting)	- Tony Stanford	CB 132	325
TELECOMMUNICATIONS			
Fault Reporting	- Stuart Andell	CB 024	306
User Support Section Head	- Andrew Lea	OB 003	348

* CB - Computer Block
OB - Office Block

**The ECMWF telephone number is READING (0734) 85411, international +44 734 85411

Please note that this number will be changed to Reading 876000 at 06.30 hours on 15 May 1982.