MB

European Centre
for Medium Range Weather Forecasts

ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England

Reading (0734) 85411 Telex 847908

Number 13 - February 1982

IN THIS ISSUE:

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

COVER: Professor Doctor E. Lingelbach, representing the President of the Council at a ceremony in December last year bidding farewell to Mr. Jean Labrousse, departing to become head of la Météorologie Nationale, France, and welcoming Doctor Lennart Bengtsson, the new Director of the Centre as from 1 January 1982.

This Newsletter is edited and produced by User Support for the Operations Department of ECMWF.

The next issue will appear in April.

NUMERICAL FORECASTS IN THE 1980's

One of the fundamental problems for the meteorological community in general, and for the Centre in particular, is to what extent the numerical forecasts will continue to improve. During the short time (about 30 years) that numerical forecasts have been around, we have seen a very substantial improvement, and today, 3 and 4 day forecasts by ECMWF are as good as 1 day forecasts in the 1950's. Will this development continue and what further improvements can possibly be achieved during this decade? I will try to give you some idea of what I believe is feasible.

The error of a numerical forecast is due to two different reasons: the first is incomplete knowledge of the initial conditions (the analysis has an error because of inaccurate observations with an unsatisfactory coverage). The second part of the error is related to the fact that the model is no exact replica of the real atmosphere, but just a model.

The initial error has, unfortunately, an inherent growth, which leads to a doubling of the error in 3 to 4 days, the error growth generally being faster for small errors. This error growth varies very much with the atmospheric flow pattern and is related to the time scale of individual weather systems and their interactions. The initial error can be reduced in two ways: (i) by better observations or a better coverage of observations, or both; (ii) by an improvement of the data assimilation, whereby observations are assimilated more frequently into the model. It can be shown by means of simple statistical calculation that there is scope for a continued reduction of the initial error by increasing the frequency of data assimilation from 6 to 3 hours and, thereby, making use of all available surface pressure reports. It also follows from such calculations that this improvement is increased as the model itself is improving. Assuming that we will be able to use surface pressure observations every third hour (this can be done by using the reported surface pressure tendencies) and that the satellite temperature and wind sounding will be rationally explored, there is possibly a potential improvement of one day, (perhaps for day 4 to day 5) to be gained.

The errors of the numerical model are due to errors in the numerical formulation (mainly resolution), to errors in the description of physical processes,and to the specification of the earth surface and the physical processes there. A further improvement of the horizontal resolution is justified. When the Centre's first model was planned, a horizontal grid distance of 100km was proposed. This was very well justified, both theoretically and synoptically, but could not, unfortunately, be implemented. We now believe that a corresponding resolution could be achieved with the new computer system at present being built. No firm figures can be given at this time as to what improvement this could give, however, recent numerical experiments at the Centre with high resolution spectral models suggest major improvements in the development of intense cyclones. These improvements of short range forecasts are bound to spread into the medium and extended range when other aspects of the model and of the data-assimilation have been brought into line.

Possibly the area where the greatest gains are to be achieved is in the improvement of the parameterisation of physical processes. The physical processes in the atmosphere take place on a scale smaller than the grid distance we can ever hope to implement and must therefore be described in a statistical way. There are no straightforward ways to do this, and present research has to be guided very much by physical intuition and by an heuristic approach, since observational information is insufficient. The progress will, therefore, in all likelihood be slow. The greatest deficiency is at present observed in the Tropics, and a careful study of the performance of the model in this area will, hopefully, reveal model weaknesses. The specification and modelling of the surface processes need much more attention than they are at present given. Previously, the major effort was devoted to ocean processes, but recent findings indicate that the land processes, mainly the description of the evapo-transpiration, are of fundamental importance. The modelling of the orography (mountains) is another area of prime importance, in particular for forecasts at middle and high latitudes. There are now already suggestions that the systematic error observed in the Centre's model may, to a large extent, be due to an incomplete description of the orography.

In conclusion, we may be confident that our numerical forecasts will continue to improve by a few days, by and large, given the level of economical resources that the Centre now has at its disposal. However, this calls for a very substantial scientific effort in years to come,and for the provision of successively better computer resources, with a capacity towards the second half of this decade which is an order of magnitude more powerful than our present one.

- Lennart Bengtsson

DIRECTIONS OF RESEARCH AND DEVELOPMENT AT ECMWF IN 1982

1. Introduction

It is useful to put the programme of research and development for 1982 within the context of the Centre's four year plan for the period 1982-85. The main goal for the four years is to develop and test a new forecasting system which will be available for operational forecasting in 1986. The requirements for a new system have arisen partly as a result of the research and development carried out at the Centre over the last few years, and partly as a result of demands for new facilities for the current and future research programme.

The development and subsequent evaluation of a spectral representation for the horizontal discretisation has shown that spectral models perform better than grid point models for approximately the same cost; this has certainly been the Centre's experience with the T63 (triangular truncation at total wavenumber 63) spectral model and the operational N48 grid point model. As a result of the various comparisons between the T63 spectral model and the N48 grid point model, the following general comments can be made:

i. T63 performs, objectively, better than N48, the improvement being largest near the surface (the geopotential height forecasts) and near the tropopause (temperature forecasts);

ii. The differences in the objective measures of skill used (mainly the anomaly correlation and the standard deviation between forecast and analysis), although not large, are shown to arise from small forecast differences and are, nevertheless, synoptically significant; in particular, most of the phase differences found in the early stages of the forecasts, are in favour of T63;

iii. The two models suffer from similar systematic deficiencies; in particular, they tend to over-develop low pressure systems over Western Europe and Northern Pacific;

iv. N48 was shown to cool the atmosphere more than T63 at 500mb in mid-latitudes; part of this extra cooling is explained by the use of a larger amount of horizontal diffusion in N48, inhibiting large scale latent heat release.

The implementation of the spectral model, whilst not particularly urgent, is necessary because it gives improved forecasts and addresses some of the systematic deficiencies mentioned in Meteorological Bulletin 2.2/2, namely phase speed errors and the tropospheric cooling.

A variety of terrainsfollowing vertical coordinate systems have been developed and tested in the N48 global grid point model. The most successful and appropriate for forecasting and data assimilation are those in which there is a smooth transition from a sigma-like system near the earth's surface to a pressure coordinate in the lower stratosphere. Rather than implement these improvements into the present operational system in an inefficient and clumsy manner, the Centre has decided to first make the spectral model with a hybrid coordinate system the basis of the next generation of forecasting systems and then to introduce the first version, after a short test period, into operations at the end of 1982. This new system will then be used as the main "test vehicle" for research and experimentation during 1983-1985.

1982 can be viewed as an important turning point for the Centre's research and development, and this can be best appreciated by listing the timetable for the development and validation of the second generation system, which is:

- 1981/82 development of a new version of the present system and its implementation into operations (fourth quarter of 1982)
- 1983/84 development of the second generation forecasting system
- 1984/85 validation and pre-operational testing of this second generation system.

The purpose of the remainder of this article is to highlight and discuss the main aspects of the Centre's plans for 1982, section 2, and to summarise the goals for 1982, section 3.

2. Research programme

2.1 Modelling (Numerical and dynamical)

a. Dynamical studies

The purpose of our theoretical work is to provide guidance to the programme of numerical experimentation, and to increase our understanding of predictability and the causes of the large scale systematic errors of our models.

The theoretical work will be a continuation of our present studies on the interaction between the tropics and middle and high latitudes, forcing and maintenance of the long waves and the study of bifurcation as a mechanism for the development of blocking in the atmosphere. The work on tropical/extra-tropical interactions is particularly important because

i. our models perform less well in the tropics even on the largest scales, and

ii. there is a great deal of evidence showing that model errors in the tropics affect the extra-tropical forecasts considerably.

Our study of the inherent predictability of our models as a function of the model (parameterisation, resolution, etc.) and weather type and its geographical and seasonal variation will receive a considerable boost from our present visiting scientist programme (Professor E. Lorenz - MIT, USA, and Professor M. Wallace - University of Washington, Seattle).

b. Numerical studies

The main work is to investigate the effect of further increases in the horizontal and vertical resolution, to review our present numerical procedures for the vertical description, and to develop finite element schemes for the vertical discretisation.

Resolution experiments carried out at the Centre, some of which have been reported in Technical Memorandum No. 23, show clearly the improvement in the forecasts brought about by increasing horizontal resolution (T21 → T40 → T63) in the spectral model. In addition, more modest improvements have been obtained by increasing the vertical resolution from the standard 15 levels to 20 levels. More experimentation with higher horizontal and vertical resolution is necessary in order to assess the impact on the accuracy of our short to medium range forecasts, and to investigate the reasons for the apparent increase in the amplitude of our systematic errors with increasing resolution.

In addition to these resolution studies with the present models, a number of different finite element techniques for the vertical discretisation will be evaluated, the aim being to find, by the end of 1982, a more accurate alternative to the present vertical differencing scheme which is a simple second order divided difference scheme.

2.2 Modelling (Parameterisation)

a. Convection

Diagnostic evaluation of the model's large-scale errors indicates that insufficient latent heat release is a main source of error and, in particular, the convection scheme used in the operational system (a "Kuo" scheme) performs less well than other parts of the parameterisation scheme.

The Centre now has five convection schemes (Kuo, an Arakawa-Schubert scheme, a parameterisation based on some recent ideas by Lindzen, a scheme designed by two of our consultants, Miller and Moncrieff (Imperial College, London) and an adjustment scheme designed by Kurihara (GFDL)). The main task for 1982 is a systematic evaluation of each of these schemes in our global models.

b. Boundary layer parameterisation

There is evidence that improvements obtained in the Centre's medium range forecasts over the last two years have, in the main, been a result of changes in the boundary layer formulation. Early in 1980, the boundary layer parameterisation in the operational system was changed. This improved significantly the performance of our model during winter, see Met. Bulletin 2.2/1. In this revised formulation, the exchange coefficient for momentum was chosen to be larger than that for heat and moisture for stable conditions, whereas for unstable conditions, it was smaller. Prior to this change, the exchange coefficient for momentum had been smaller for all stabilities. The revised model had less of a tendency to produce over large features and had less intense cyclones.

In addition, a recent series of experiments have provided clear illustration of the model's sensitivity to apparently small changes in the boundary layer formulation, and obviously more experimentation and careful documentation of this sensitivity is an important line of research, and could lead to an improvement in the model's handling of the decay of mature cyclones, namely the occlusion process.

The new work for 1982 will be the development and evaluation of one of the Mellor-Yamada higher order closure schemes for the planetary boundary layer; it is expected that a tested version will be made available by the end of 1982. The main reason for developing such a complex, and computationally expensive, scheme is the success achieved by Miyakoda (GFDL, Princeton, USA) for extended range integrations.

2.3 Data assimilation

The data assimilation scheme used for both operations and research is an intermittent 4-dimensional procedure in which observations are assimilated every 6 hours. Apart from the necessary data checking procedures, this assimilation scheme consists of a 3-dimensional statistical optimum interpolation technique and a non-linear normal mode initialisation procedure. Although this system, as at present programmed for the CRAY-1 computer, can accommodate a large class of different formulations and experiments, it is nevertheless restricted in several aspects and it is necessary to develop a more flexible scheme in order to accommodate formulations and approaches being suggested by the Research Department; of particular interest are:

i. continuous data assimilation and/or higher temporal resolution; a 3 hour assimilation frequency is likely to be the next development;

ii. direct interaction between the analysis and initialisation;

iii. higher vertical and horizontal resolution; the latter will require theoretical study of such questions as the possibility of using structure functions with divergent wind components.

The importance of a higher resolution data assimilation which matches closely that of the model cannot be stressed too strongly. In order not to create a long "spin up" time for some physical processes such as latent heat release, observations must be assimilated by a system which has basically the same resolution as the prediction model, and which does not smooth out excessively, realistic model structures which are not resolved by the observing system. If this is not the case, model improvements will not necessarily imply better forecasts.

Much of this new development is bound up with the development of the new forecasting system and it will lead eventually to a complete revision of the data assimilation scheme.

In addition to this new work, development of the present system, analysis and initialisation, will continue. This work includes:

Analysis of surface fields;
Estimation of first guess biases (the biases in the short range forecasts) and structure functions, estimation of observational biases and errors;
Incorporation of non-adiabatic effects in the non-linear normal mode initialisation and the treatment of tropical modes in the initialisation.

2.4 Diagnostic studies

Extensive diagnostic packages are now in a mature stage of development and are providing more extensive climatological information than has been previously available. This data is essential for describing the growth of the mean errors of the model including, of course, areas such as the southern hemisphere and the tropics, where such information has been scanty.

The verification group has started a project to define the mechanisms (changes in mountain torque, standing eddy fluxes, etc.) which lead to the growth and evolution of the model's systematic errors, from the six hour forecast through to the end of the forecast. This will exploit the diagnostic packages referred to above, which are now available for the forecasts as well as the analyses. Additional tools for this work include suites of programs to carry out wave number frequency analyses and diagnosis, in the model's sigma coordinates, of diabatic heating.

3. Summary of goals and requirements for 1982

We summarise the Centre's main goals as:

3.1 Implementation, into operation, of the research and development carried out over the last few years

This involves the construction of a new forecasting system with the spectral model as the main model; grid point and finite element models will also be available for high resolution experiments. The models will use a smoothly varying hybrid coordinate system.

3.2 Provision of a high resolution system, so that very high resolution integrations can be carried out easily, albeit inefficiently, on the CRAY-1. The present versions of the programs will allow a limited assessment of the impact of higher resolution, though the full programme of work is not expected to begin until the new spectral model becomes available for research, end of the second quarter of 1982.

3.3 Re-design of the data assimilation schemes in order to provide higher resolution analyses and to allow the evaluation of other approaches.

3.4 Extensive diagnosis of the model's systematic errors

Our diagnostic scheme has taken a considerable effort to build up and, with the availability of the FGGE analyses, it should be possible to obtain a clearer understanding of the model's large scale errors. It has been estimated that the removal of these errors would improve, on average, the skill of our present models by two days.

- Dave Burridge

* * * * * * * * * * *

THE RELEASE OF EXPERIMENTAL METEOROLOGICAL PRODUCTS TO MEMBER STATES

1. Introduction

Member States have, by now, had considerable operational experience of ECMWF products and are becoming more aware of the quality and limitations of the forecasts, especially in the later stages of the forecast period. Following requests from some Member States for products up to H+240 for test purposes, it has been decided that products up to this time should be released. However, in order to make clear the experimental (as opposed to operational) nature of the later forecast products, the list of products is divided into two categories, called "operational products" and "experimental products".

The category of "experimental products" also includes such parameters as total precipitation, wind at the 10m level, temperature at the 2m level and model-predicted cloud. The computation of total precipitation is documented in the ECMWF Forecast Model Documentation Manual. The interpolation of the wind and temperature to the near-surface levels is described below, as is also the formulation used in deriving the model-predicted cloud.

It must be stressed that extreme caution should be exercised in making use of these parameters. Before they are to be seriously considered for operational use, it is essential that an extended period of testing be carried out so that a thorough appreciation of their limitations can be gained. They are model-dependent; the resolution of the model is 1.875° in latitude and longitude, so that coasts, irregular terrain and water-to-ice boundaries, for example, are not resolved in detail. Thus the direct model output values of 10m wind and 2m temperature cannot be expected to agree with locally observed values near the coasts or over high ground. Even over relatively homogeneous terrain, the model output value at a point must be regarded as a representative value for a large area rather than a point value for a single location.

For reasons of uniformity, all verifying times of such parameters as total precipitation, wind at 10m, temperature at 2m and model-predicted cloud are included in the category of "experimental products" although some Member States have now had some experience of using some of these parameters at the earlier verifying times and are starting to use them operationally.

Thus, the category of operational products includes standard parameters (geopotential height, temperature, wind, humidity) up to H+180 on standard pressure levels, and mean sea level pressure. The category of experimental products includes those which the Centre feels are suitable for testing or experimental use (e.g. in the form of mean fields) in the later stages of the forecast but which are in general not suitable for immediate operational use without a period of testing and evaluation for their quality and usefulness.

2. A description of some of the experimental products

The surface temperature is the temperature of the model surface and, thus, is the sea surface temperature over the oceans, for example.

The 2m temperature is the temperature calculated from the surface values and those at the lowest sigma-level of the model. The temperature is interpolated using the model's drag coefficient for temperature C_{DT} where:

$$C_{DT} = f\left(Ri, \frac{z_{15}}{z_o}\right)$$

where Ri is the Richardson number, z_{15} is the height of the lowest model level and z_o is the roughness length. The roughness length has been computed from a high resolution orography over land, and takes into account the variance exhibited by the high resolution orography. Over sea z_o is a diagnostic quantity and is computed every time step from the wave height using the Charnock formula. Over land, z_o varies from 1.5mm to 8.5m; the global average value is around 10 to 20cm.

The function f is highly complicated and thus no simple inversion of the temperature profile is possible. However, the drag coefficient can formally be obtained from another relationship:

$$C_{DT} = \left(\frac{k}{\ln \frac{z_{15} + z'_o}{z'_o}}\right)^2$$

where k is the Karman constant and z'_o is a modified roughness length which varies with the stability. From this relationship, the value of z'_o can be obtained and a logarithmic profile of the temperature between the surface and the lowest model level can be derived, thus allowing the calculation of the temperature at a level between the model surface and the lowest model sigma level:

$$T(z) = T_{surf} + (T_{15} - T_{surf}) \frac{\ln \frac{z+z'_o}{z'_o}}{\ln \frac{z_{15} + z'_o}{z'_o}}$$

It is again stressed (see Introduction above) that the value of the 2m temperature is strongly influenced by location; in particular, near ocean boundaries and over high ground, it cannot be expected to agree with locally observed temperatures. Furthermore, at this time, there is no diurnal cycle in the model and, therefore, the 2m temperature should be regarded as an indication of the daily mean temperature.

The wind at 10m is the wind calculated in a similar way to the 2m temperature (see above), using a drag coefficient for wind C_{DW}. Like the 2m temperature, it is strongly influenced by location. It is again stressed (see Introduction above) that over locally rough or irregular terrain and over high ground, it cannot be expected to agree with locally observed winds.

The model-predicted cloud is the field that is used in the radiation computation of the model or, more precisely, is a combination "as seen from above" of the partial cloud cover computed for all the 15 levels of the model. The computation at a given level uses the humidity and temperature predicted by the model at that particular time; the vertical combination is done by assuming maximum overlapping of adjacent cloudy points. It is found that this parameter verifies best against the infra-red rather than the visual satellite imagery, i.e. frontal cloud bands are well represented, while extensive areas of low stratus or stratocumulus cloud or sea fog do not appear. For further information and an example of this field, see: J-F. Geleyn. Pseudo-satellite picture presentation of model results. ECMWF Newsletter Number 1, February 1980.

Member States taking experimental products for research or development work are expected to send a report on this work to the Centre.

For further details on these products, including instructions on ordering them, see Met. Bulletin M3.4/1(3) which is currently being distributed.

- Austin Woods

* * * * * * * * * * *

ECMWF METEOROLOGICAL TRAINING COURSE
4 MAY - 25 JUNE 1982

The 1982 training course will be very similar to the 1981 course; only slight re-arrangements in some of the lectures have been made. The 1982 course is divided into three parts:

A1: Lectures on basic dynamic meteorology and NWP
A2: Applications of material covered in A1 to the ECMWF model
B : Operational aspects of the ECMWF forecasting system

Part A1 will last four weeks and is intended as an introductory course on dynamic meteorology and numerical weather prediction. Included in this part will be a two-day computer course, giving the basics of how to run simple jobs on the CYBER and the CRAY. The course literature for part A1 will be the Lecture Notes published in 1979 by ECMWF. The student who so wishes can undertake a special project and will be assisted by a tutor selected from the scientific staff of the Centre. In this case, however, it is recommended that the student attends parts A1 and A2, to allow more time to complete the project.

Part A2 is a two week theoretical course, with lectures oriented towards the application to the ECMWF forecasting model of material covered in A1. Therefore, a background knowledge corresponding to the level reached on course A1 is assumed for A2. For those not attending A1, but participating in A2, it is recommended that they read the Lecture Notes beforehand. The course literature for A2 will consist of ECMWF Technical Reports and the ECMWF Forecast Model Documentation Manual.

Part B of the course, which will run for two weeks, is quite different from A1 and A2. It will cover operational aspects of the ECMWF forecasting system, and is directed towards those who are, or will be, using the ECMWF products actively. Others, for example university scientists, are also welcome to attend the course, but it will be explicitly aimed at those who will be at the receiving end of the communication lines and who will be interfacing between the Centre and the eventual end user of the forecasts. The B course will be oriented around laboratory sessions, in which the students will examine both current forecasts and data in the Centre's archives.

In order to make it easier for students to obtain the necessary funding to attend the 1982 training course (or part of it) the ECMWF will reimburse travel expenses for one delegate per member country.

Application forms have been sent out to all Member State Meteorological Services and should be returned to ECMWF before 19 March 1982. This early closing date for applications is necessary in order to find accommodation for the delegates.

- Erland Källén

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

*CRAY DATASET STORAGE

1. Choosing which computer

Recent requirements for the processing of data originating on the Cray have led to a review of our policy on the storage of data created by the Cray. We have already experienced situations where jobs, better run on the Cyber, have been run on the Cray with a severe impact on Cray effectiveness.

The following factors should be taken into account when you are organising your work:

> The Cray is a large and powerful computing machine, which is fully loaded; there is little prospect of acquiring extra Cray computing facilities until at least mid-1982.
>
> The Cybers are not completely loaded, since the power increased by 30-50% when the installation of the Cyber 730E took place at the end of 1981.
>
> The Cray CPU is about 5 to 20 times the Cyber 175 CPU power, but the Cray discs have a slightly slower access time and only about 4 times the data rate of the Cyber discs. This means that jobs requiring a relatively small amount of CPU power in relation to their I/O requirements should be run on the Cyber whenever possible, since the CPU and I/O power is better balanced. Random access I/O will almost always be slower on the Cray than on the Cyber as the seek time is slower on the DD-19 or DD-29 Cray discs (50ms average) than on the 885 Cyber discs (25ms average).
>
> Transferring a file to or from the Cyber in binary blocked format is as fast as transferring in transparent mode and uses a similar amount of Cyber and Cray resources.
>
> Movement of jobs, particularly I/O limited jobs, from the Cray to the Cyber has the effect of releasing scarce Cray power for the "number crunching" workload.

2. Processing existing Cray datasets

Users who have stored Cray datasets on the Cyber in transparent format (DF=TR) will have some difficulty in using this data on the Cyber since the file is in Cray format and the data is also in Cray format. However, a utility BBCONTR (see later) has been produced to alleviate this situation by converting transparent format files to binary blocked files. Binary blocked files (DF=BB) are stored in Cyber format, though the data within each record is still in Cray format. However, useful data reduction can take place on the Cyber by manipulating binary blocked files.

So our advice is NEVER to dispose datasets to the Cyber in DF=TR format unless you are sure you will never want to use those datasets on the Cyber. Probably the only datasets which should be disposed in DF=TR format would be those containing Cray binary programs, or Cray unblocked datasets, which cannot, in any case, be disposed in other than transparent format.

Format of DF=BB Files on the Cyber

Each Cray dataset can consist of a number of Cray files. A Cray file can consist of a number of Cray records. A Cray record consists of a number of data entities, usually each entity being a multiple of 64 bits in length.

When stored in binary blocked format on the Cyber

A Cray dataset	=	Cyber file
	=	Record manager data between BOI and EOI in RT=W, BT=I format
Cray file	=	Cyber binary record
	=	Record manager section in RT=W, BT=I format
Cray record	=	Record manager record in RT=W, BT=I format.

The bit string comprising the Cray record is exactly equal to the bit string comprising the Cyber record manager record.

Selecting Files

If you have a program on the Cray which needs to select only a part of a dataset (e.g. every sixth file on a 30 file dataset) it is in everybody's interest for you to create on the Cyber a dataset containing only those files required. The SKIPF and COPYBR Cyber utilities can be used to skip or copy Cray files in a binary blocked dataset and the resulting dataset will be transferred to the Cray in, for instance, one sixth of the time taken to transfer the original dataset, in the example above.

If you need to select data at record level within a Cray dataset, you will have to write your own FTN program on the Cyber. Binary blocked files are stored in RT=W, BT=I file format on the Cyber, so a program to select records is not difficult to produce, writing the selected records to a new dataset for transmission to the Cray. Note, however, that the record length on the Cyber will be longer than on the Cray, due to the difference in word length.

Temporary storage space for the output of a selection job, prior to its transfer to the Cray, is available on setname TEMP. Such files should be purged following their transfer to the Cray.

Converting multifile tapes to single file tapes

Following a scientific experiment, you may wish to ACQUIRE all the files in a multifile set to the Cray. If you have the kind of job which acquires one file to the Cray, processes it, deletes it, acquires the next file, etc., continue as you do now, but remember to delete unwanted space as soon as possible.

If, however, you need to ACQUIRE all files before doing any processing, convert your multifile set to a single file dataset containing several files.

The Cyber procedure MULTSIN will perform this conversion. It is called as follows:

```
PROCIN.....
MULTSIN,(VSNIN=vsn1, VSNOUT=vsn2, NF=number, TDEN=GE)

where VSNIN  is the vsn of the input multifile set
      VSNOUT is the vsn of the output tape
      NF     is the number of files in the set
      TDEN   is tape density (GE by default)
```

Tapes are requested and unloaded within the procedure.

<u>Converting Cray datasets on the Cyber</u>

Datasets on the Cyber (on tape or disc) which have been transferred from the Cray with DF=TR may now be converted on the Cyber as though DF=BB had been used. This may be useful where:

1. Any of the data is required for use on the Cyber.

2. Only part of the data is subsequently required on the Cray.

The utility may be accessed by:

```
ATTACH,ECLIB.
LIBRARY,ECLIB.
BBCONTR.
```

There are 3 possible forms for the BBCONTR control card:

```
BBCONTR.
BBCONTR{,ifile}{,ofile}.
BBCONTR{,I=ifile}{,O=ofile}.
```

where ifile is the name of the Cyber file containing Cray format data (default is not specified = CRAY) ofile is the name of the Cyber file to receive Cyber format data (default is not specified=BBFILE). This file is rewound before processing begins.

Either or both of ifile and ofile may be a magnetic tape. Converting a full 2400' 6250bpi tape will take around 5½ minutes' elapsed time and 20 c.p. seconds.

To transfer part of a file to the Cray, notice that Cray files may be skipped or copied by the Cyber skip or copy record control cards,

e.g. To build a composite file containing files 7 and 13 from a transparent format tape:

```
:
:
LABEL,TAPE,......
ATTACH,ECLIB.
LIBRARY,ECLIB.
BBCONTR,TAPE,DISK.
REWIND,DISK.
RETURN,TAPE.
REQUEST,RESULT,*PF.
SKIPF,DISK,6.
COPYBR,DISK,RESULT.
SKIPF,DISK,5.
COPYBR,DISK,RESULT.
RETURN,DISK.
CATALOG,RESULT.....
```

The resulting file can be transferred subsequently to the Cray with DF=BB, if required. Notice that a LIMIT card may be needed, if you are converting very large files,

e.g. To convert a multifile transparent format tape to a multifile binary blocked tape:

```
ATTACH,ECLIB.
LIBRARY,ECLIB.
REQUEST(ITAPE,MF,GE,NORING,E,VSN=vsn1)
REQUEST(OTAPE,MF,GE,RING,N,VSN=vsn2)
LABEL(IN,R,M=ITAPE,L=label1).....first file
LABEL(OUT,W,M=OTAPE,L=label1)
BBCONTR(IN,OUT)
LABEL(IN,R,M=ITAPE,L=label2).....second file
LABEL(OUT,W,M=OTAPE,L=label2)
BBCONTR(IN,OUT)
etc.
```

Although file selection using BBCONTR is possible within a Cray SPOT job (i.e. as specified in an ACQUIRE TEXT field) , this practice is not recommended,as the conversion has to be repeated for each execution of ACQUIRE. Also, the performance of the Cray is severely degraded,since a logical path between Cray and Cyber is blocked for substantial periods while the SPOT job executes.

For the same reason, although a loophole currently allows other operations to be specified in a TEXT field, this possibility will be disallowed in the future.

Finally, it may be desirable to do some real processing of Cray data stored on the Cyber. To do this, the dataset needs to have been disposed not only in binary blocked format, but also to have been passed through the CONVERT utility to convert it to Cyber number format. Since CONVERT on the Cray runs almost as fast as COPY, you should review whether it is feasible always to convert data to Cyber format.

So, our advice would be extended to include a request to store datasets in Cyber number format, using CONVERT to handle reformatting.

3. Summary

Disposed datasets should never be stored on the Cyber in 'transparent' format (DF=TR).

Data selection/reduction should be done by a Cyber job, and only the useful data passed to the Cray.

Whenever possible, data should be converted to Cyber format before disposition, using CONVERT.

Data processing requiring relatively little CPU time should be done on the Cyber.

- Peter Gray

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

ABOUT THE COST OF CALLING SUBROUTINES

Los Alamos Scientific Laboratory has measured the CP times required to call a subroutine during an exercise aimed at evaluating two super-computers, the Cray 1 and Cyber 205.

Their findings were:

Subroutine call takes: 720ns + 60ns/parameter on the Cray 1

Some measurements were similarly made at ECMWF with two particular subroutines recently included in ECLIB (SBYTE and GBYTE designed to manipulate bit fields, one at a time) and which take, on their own, in the region of a few microseconds.

We compared the cost of calling the subroutines as against executing the same code 'in line'. By 'in line' we mean inserting the code of the subroutines in place of the call to the subroutine. The results are presented in the table below:

	CRAY 1			CYBER 175		
	Subroutine*	In line* code	Call overhead	Subroutine*	In line* code	Call overhead
GBYTE	2.97	1.80	65%	9.93	4.17	138%
SBYTE	3.64	2.45	48.6%	11.28	5.40	108.9%

*times are given in microseconds and concern one single call.

Comments

1. The cost of one call, calculated using the Los Alamos formula, would amount for each subroutine on the Cray 1 to 960ns (720 + 4(parameters) * 60). In our tests, we counted 1170ns for GBYTE and 1190ns for SBYTE. The results are in very much the same order of magnitude; the difference could be due to many factors, amongst which, the "nature" of our test, which cannot give an 'average' figure similar to the one from Los Alamos.

2. Among the various tools which can be used to optimise code, moving part of it one level higher (i.e. using fewer subroutines) is worth considering when this is not prohibited by memory occupation. The two subroutines involved in our tests occupy about 20 CM words each and treat the most general cases (bytes spanning words).

3. The next Fortran standard (Fortran 8X) could ease the replacement of a subroutine call by in-line code by means of a MACRO facility. This would also make it easy to evaluate the performances of in-line code versus subroutine calls.

4. Finally, the figures given in the table cannot be taken as 'general case figures' but are meant to show that in these circumstances (i.e. with short code taking a few microseconds) programs making extensive use of "short" subroutines could be improved.

- Michel Miqueu

* * * * * * * * * * *

*CYBER 730E PROGRESS

The Cyber 730E mainframe was delivered to ECMWF on Saturday, 28 November 1981. It was assembled rapidly and power was applied during the same weekend. Local CDC engineers, assisted by an engineer from the Toronto factory, then continued installation testing and the software was loaded for the first time on Thursday, 3 December.

From then until the start of the provisional acceptance test on Thursday, 17 December, installation testing of all the hardware and software continued. This took longer than normal, due to the necessity of maintaining the normal Cyber 175 service and due to the non CDC peripherals (e.g. the Cray, the NFEP and the plotters) which needed special testing.

Eventually, all was ready and the 48 hour provisional acceptance test began at 0730 on 17 December. Other than a small disc problem, causing 7 minutes' downtime, the machine ran without major incident. Various problems occurred, mostly minor in nature, which are being resolved during the 90 day final acceptance period. However, during both nights of the test, there were delays to the operational suite, totalling 4 hours on the first night and 1½ hours on the second. During, and since, the provisional acceptance trial, minor changes to the disc and software configuration appear to have eliminated the problems which caused the delays.

During the Christmas and New Year period, the Cyber 730E was used by Operations Department staff to perform further software installation testing and to produce and check operating procedures.

The machine then entered service on 4 January, processing only batch jobs during the prime shift. Since that time, the service has gradually built up. From 13 January, the Cyber 730E began to handle all Cray transfers. This is an interim measure, pending the transfer of Intercom to the Cyber 730E, expected on 25 January. It is also worth mentioning that there have been no unexpected incompatibilities between the Cyber 175 and the Cyber 730E.

Experiments are currently underway to confirm that the Cyber 730E can process the operational suite. Though the suite runs more slowly, we have already shown that the Cyber 730E is capable of handling the operational forecast if the Cyber 175 were to be unavailable. Further tests will establish the full extent of the delay possible.

Looking ahead, it is expected that the Cyber 730E will be used in its final role from the beginning of February. At that time, the two Cybers will be used as follows, under normal circumstances:

Cyber 175	Cyber 730E
Large batch jobs	Intercom jobs
Majority of operational suite	Printer/reader/plotter driving
	Small operational suite jobs
Cray station	Small batch jobs (larger jobs at night)
	Housekeeping jobs
	NFEP handling

The additional disc drives have had a positive effect on the configuration. They have allowed TEMP space to be doubled, providing more space for short-lived permanent files, and the use of removable disc packs has almost been eliminated. This has reduced the operator workload sufficiently to allow two Cybers and a Cray to be operated by the same number of staff as previously, and should also reduce problems experienced due to lower reliability when disc packs are frequently exchanged.

There are also a number of other, minor, configuration changes taking place. The third meeting of the Technical Advisory Committee decided that punched cards should be phased out at the Centre. Hence, late in 1981, one of the two central site card readers and the online card punch were returned to CDC. A limited online card punching service is now provided by a modified data preparation punch. The money thus released was used to rent two further CDC 679-6 magnetic tape decks, which read 1600 or 6250 bits per inch tape at 150 inches per second; they were installed early in January 1982 and will improve the turnaround of magnetic tape jobs.

In addition, a 679-3 tape deck, used to read 800 or 1600 bits per inch tapes has been exchanged for a 679-6 deck. The few 800 bits per inch tapes processed by the Centre can be handled easily by the one remaining 679-3 deck.

During the spring of 1982, an enhancement to the CDC 2551 telecommunications processor will provide an extra 24 ports to the Cyber, thus reducing present queues for interactive ports, and providing facilities for visitors, to replace the punched cards previously used.

- Peter Gray

* * * * * * * * * * *

*SOME CHANGES RESULTING FROM THE INTRODUCTION OF THE CYBER 730E

This article tries to explain a few differences to the user interface which will result from having two Cyber front ends.

1. After 25 January, Intercom will normally be run on the Cyber 730E. All outputs specifying a user terminal identifier (TID) will be kept on the same Cyber as that which Intercom is running on. It is expected that when Intercom is running on the Cyber 730E, the recent problem of bad response time should be much reduced.

2. Dependency Jobs will only be run on one front end at any time. However, it may be that when the dependency jobs are transferred from one Cyber to the other, a dependency job may be left running and may not be able to transfer dependency, in which case, the system will request the operator to manually transfer the dependency. This action should reduce or eliminate dependency problems, but our advice not to use the dependency features remains valid.

3. Tape Jobs may be run on either or both Cybers.

4. Batch Jobs which do not have dependency or require tapes, may run on either front end. Such jobs will be 'load levelled' between the two Cybers, so as to maximise throughput of these jobs. Note that any jobs with STPAK on the job card can expect considerable delay in execution, and users are requested to stop using STPAK as soon as possible.

5. The Cray Station will normally be up on both Cybers, with most transfers going via the Cyber 175. It will be transparent to CRAY jobs whether one or both Cybers are in production and which Cyber handles the transfers.

6. The Printers and Plotters can be driven from either front end. Normally they will be driven from the Cyber 730E.

7. The discs are shared between the two front ends, so that permanent files and private packs can be accessed from either front end.

8. Job Visibility (with a few limitations) allows jobs in any mainframe to be seen from either front end.

i. Q command by default shows the queue on the Cyber to which the user is connected. However, it is possible to enter:

Q,q,MCY to see Cyber 175 queues
Q,q,MCZ to see Cyber 730E queues
Q,q,ALL to see queues on both Cybers

where q is queue type, i.e. I,O,P,E

It is intended to make the ALL parameter the default during the next few weeks.

ii. MYQ command shows the queues in both Cybers.
Note that MYQ may miss a file if it is transferring from one queue to another at that moment.

iii. CQ command shows the CRAY queues from either Cyber.

9. Job Control does not allow jobs or queue files on one Cyber to be controlled from the other Cyber. Thus, if a user enters the commands:

DROP, DIVERT, EVICT, KILL, PRIOR, REVERT

these will only be effective upon the Cyber to which the user is logged in. To control jobs in the other Cyber, it is necessary to send a request to the operator.

However, the CRAY job control commands CDROP and CKILL can be entered from either Cyber and will work as normal.

10. CP Time is adjusted on the Cyber 730E so that the CP time accrued is actually 3/8 of the CP time used. This is done to ensure that the CP time parameter on the job card is valid for either mainframe. However, depending upon the nature of a given job, it may use slightly more or less CP time on one Cyber than on the other.

Unfortunately, there are a few jobs which use considerably more CP time on the Cyber 730E than on the Cyber 175. The reasons for this are being investigated, and meanwhile, such jobs should be run on the Cyber 175.

Please note that double precision arithmetic is very much slower on the Cyber 730E, and any jobs doing a lot of double precision calculations should ideally be run on the Cyber 175.

11. Old Absolute Binaries which were created prior to the introduction of NOS/BE Level 538 on 24 August 1981 should be recompiled as they may not run correctly on the Cyber 730E.

Our plan is that almost all the time at least one Cyber is available, supporting the above functions. Soon we hope to restrict the need for taking both Cybers for systems development to two mornings per week.

- Tony Stanford

* * * * * * * * * * *

COMMAND

A Fortran callable subroutine 'COMMAND' is now available on the Cybers to allow job control statements to be executed from a user program. Each control statement is supplied as a subroutine parameter, allowing several statements to be executed by one subroutine call. During execution, the running program is swapped out, and is restarted when all submitted statements have been executed, or when an error has been found. For FTN4 programs, the parameters must be Hollerith string, L format. For FTN5 programs, the parameters should be character data type, e.g. (using FTN5):

```
INTEGER ERROR
CHARACTER *50  C1,C2
..
C1 = 'ATTACH,FILESET,ID=XYZ.'
C2 = 'GF,TAPE1,GROUP/ELEMENT.'
..
CALL COMMAND (ERROR,C1,C2)
IF(ERROR.NE.0) WRITE(*,'("ERROR IN FILESET")')
```

A Computer Bulletin with a full description is in preparation. (Before trying to use COMMAND, please ask for a copy of the draft bulletin.)

- Walter Zwieflhofer

* * * * * * * * * * *

USE OF CONTROL CARDS IN EDIT

A new version of EDIT, which allows the use of most control cards from within EDIT, has been released. To use this facility, prefix the control card with an *.

```
E.g. *ATTACH,FRED,ID=XXX
     *FILES
     *SUB,FRED
```

The control cards which cannot be used are those which are processed specially by INTERCOM, such as ETL, KILL or EVICT and those which would call EDIT recursively, such as EDIT, EF or the proclib procedure PRINT, when used with both the ASCII and S parameters. If you do call EDIT recursively, you will get one of the following error messages:

```
     INTERNAL ERROR DETECTED
or   EDIT FILES SCRAMBLED
```

However, you may be able to recover your session by using the EDLOG file.

Note that the CPU time used by all commands is deducted from your job step allocation, so that if you expect to use more than the default CPU allocation (10 seconds), you should use ETL to increase your time limit, before you enter EDIT.

- Gary Harding

* * * * * * * * * * *

LIBRARY CHANGES

1. ECLIB

Recent changes to ECLIB have included the following:

RIO (Cray)	bug fixes
GPGRAPH	additional options
SDSP (Cray)	new routine to locate address of DSP area
UCOPY2	new routine to copy vectors

The following changes and additions will be introduced shortly:

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 (Cray)	obtains number of 1 bits in a word
XREF (Cray)	global cross reference mapping

Changes:

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

Note that XREF is an addition to Cray ECLIB. It has been available on Cyber (FTN) for some time. Future work will include making XREF available for FTN5 programs.

2. Procedure Library

Recent PROCLIB changes were as follows:

PROCIN	more speed improvements so that, if a private version is kept by specifying parameter 'KEEP', subsequent calls using PROCIN,ID=uid,AC=ac. will be very fast as the mechanism to find 'uid' and 'ac' from the LOGIN command is not necessary.
TYP & DAYF	can be executed from within EDIT
FICHE	now accepts a local file as input
SUB	now issues MYQ instead of FILES command
EVICTAL	evicts all remote output files
MULTSIN	combines files from a multifile tape into a single file (see article entitled 'Cray Dataset Storage', Page 11).

3. Even faster Fourier Transforms

As mentioned under 1., a new version of the Fast Fourier Transform package will soon be available. It is intended to be compatible with the present version of FFT99 and FFT991 with the following differences:

a. faster (reduces cp time by 30-35%)

b. N no longer has to be even

c. the setup mechanism is simplified to one subroutine call (previous setup routines FAX and FFTRIG will still work)

d. the dimension of the WORK area is reduced.

- David Dent

* * * * * * * * * * *

ALPHANUMERIC COM SYSTEM

The Centre has completed its tender evaluation for an alphanumeric COM system. An order has been placed with DatagraphiX Ltd., for delivery of a model from their Autocom II range on 13 March 1982. This is an offline system,which has its own micro computer with floppy discs, and which will receive its data via magnetic tape. This system could be upgraded in the future to work on line.

Initially, the only changes,from the users' point of view, will be a dramatic decrease in turnaround time, from the present 48 hours minimum to between 1 and 2 hours, and that copy fiche will be produced on the same type of film as the masters.

Plans for the future use of the system include the ability to put files for the COM directly into the NOS/BE film queue. This will remove the need for users on INTERCOM to run a batch job. It is also intended to use the intelligence of the COM system to produce its own indexes. Will any users who have problems with the current fiche index, or who would prefer a different type of index for their application, please contact me, so that the necessary arrangements can be made.

- Gary Harding

* * * * * * * * * * *

COMPUTER USAGE STATISTICS 1981

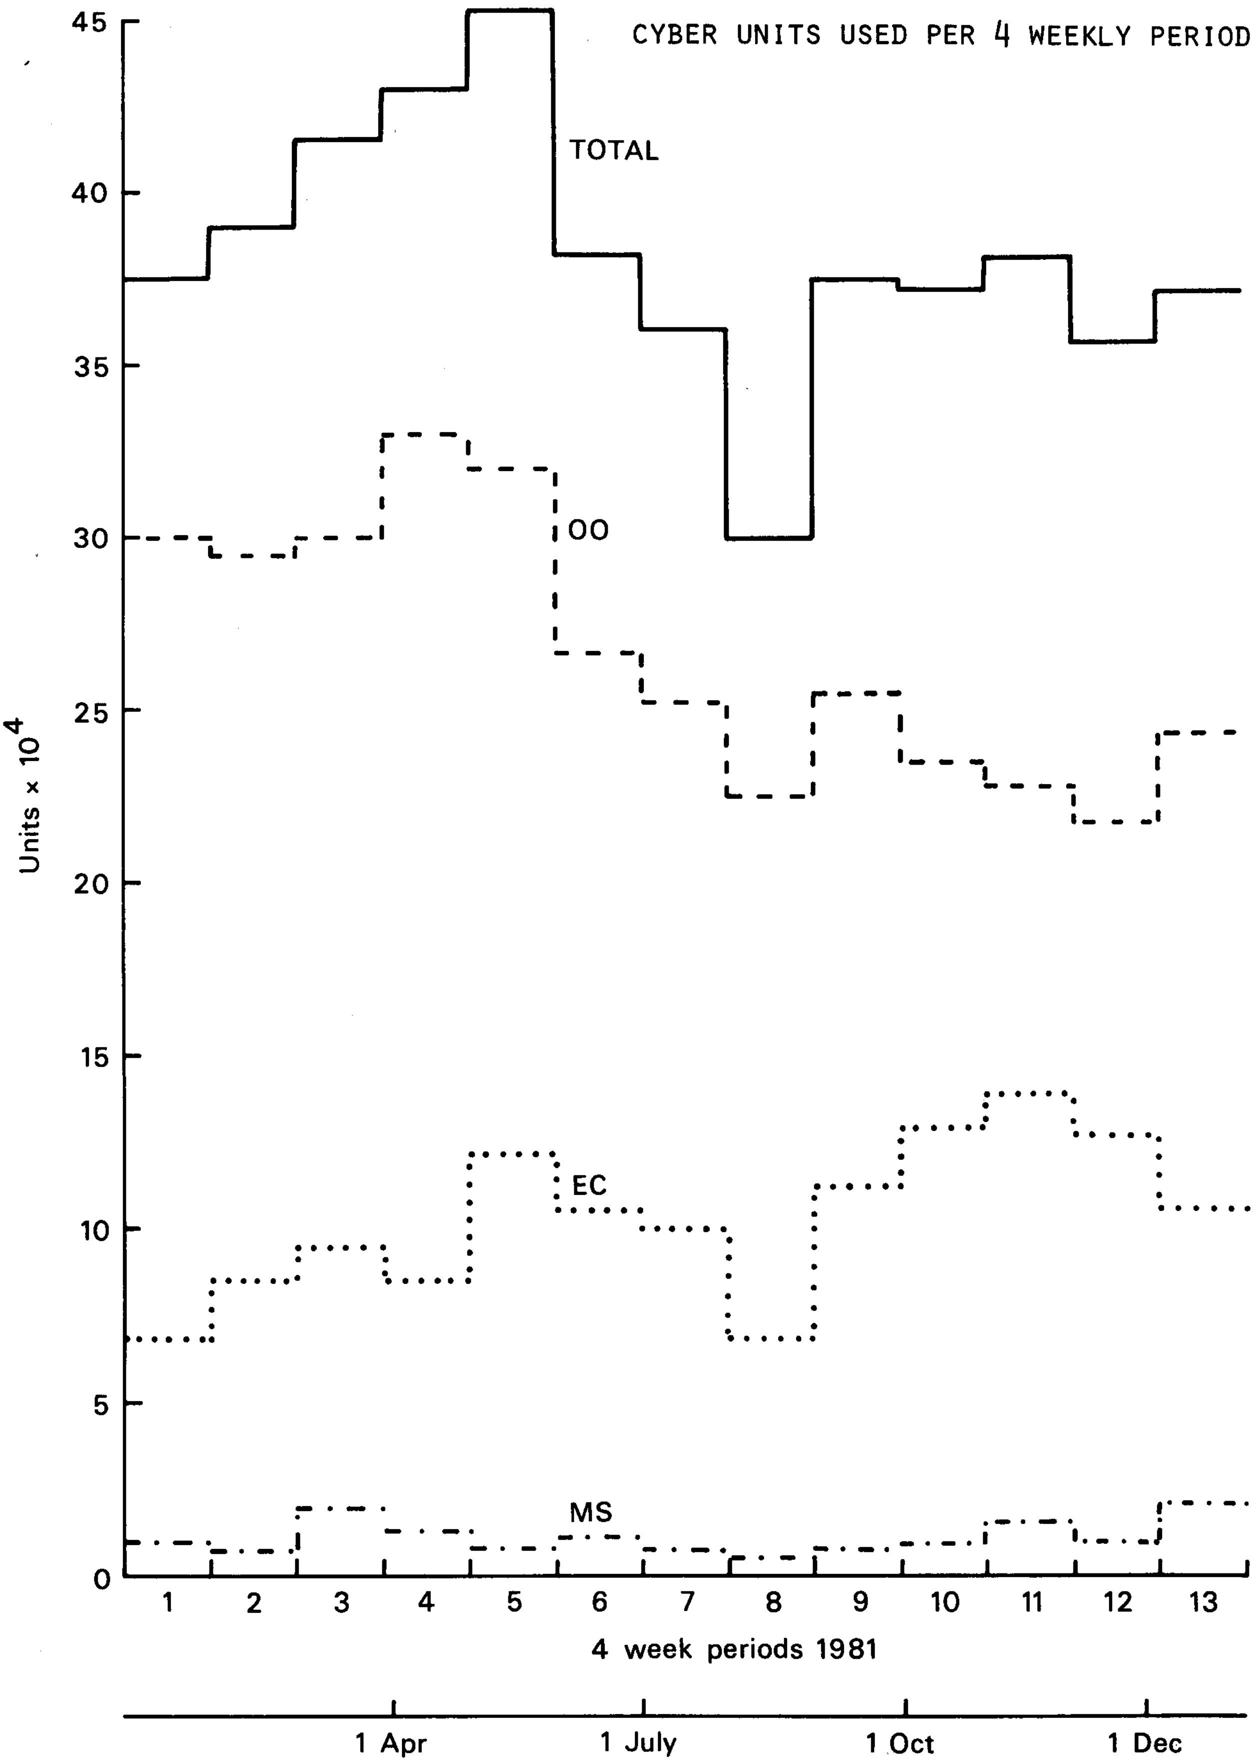

Total = total usage less those jobs classed as systems overheads
OO = operational suite running
EC = Centre users
MS = Member State users, including Special Projects

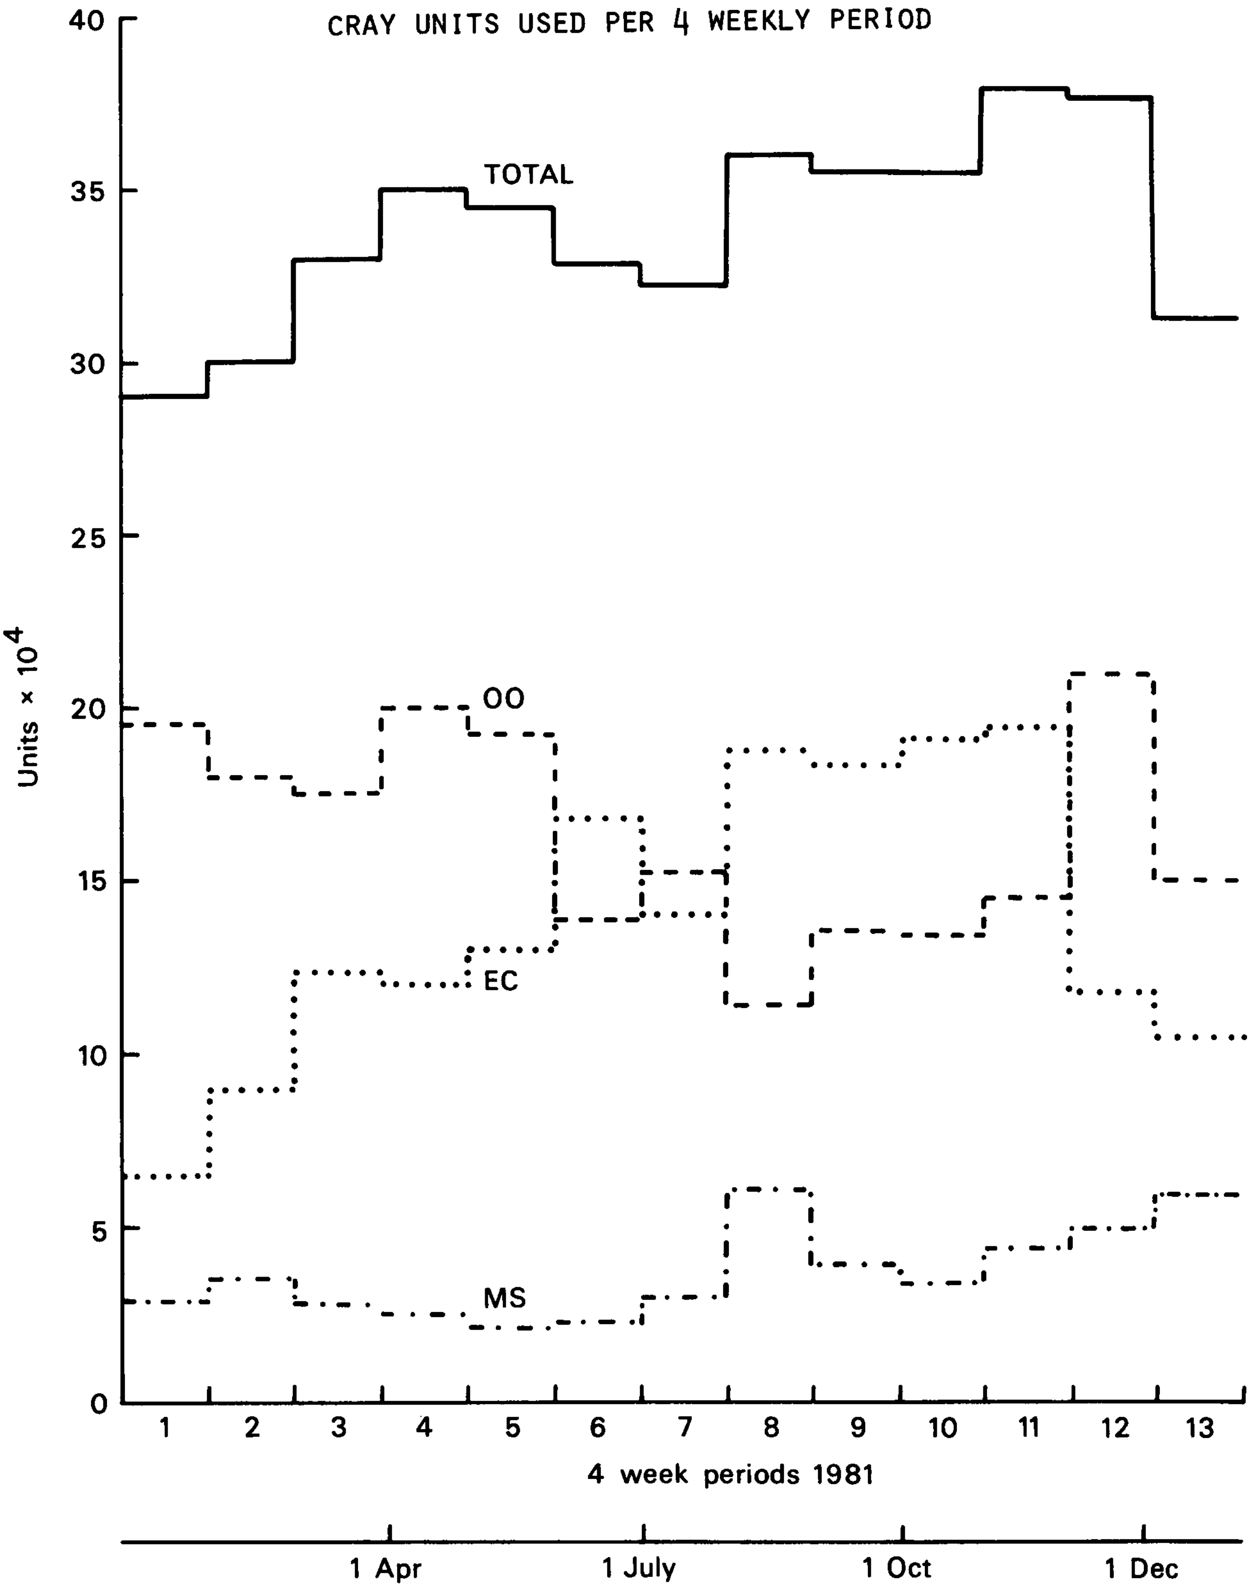

NEWS SHEETS

Computer News Sheets are now sent as remote output files to those member states who have RJE facilities. Additionally, the texts of news sheets are held as elements of a fileset and may be examined, either interactively or in a batch job, as follows:

```
ATTACH,FILESET,NEWSET.
LD,FM=FD.                    lists all section headings
```

Each news sheet is identified by:

a. group name = month in current year, e.g. JANUARY.

   A news sheet for December 1981 has a group name = DECEMBER81

b. element name = 'N' followed by a sequence number.

Therefore, to examine a particular document, type

e.g. GF,JANUARY/N128

Then use EDIT or PAGE to scan it. Note that the text is in ASCII and contains text processor commands. If using PAGE, the command

```
L,=/NUMBER/
```

will display from the start of the text. Documentation entries as displayed by the LD Command will be removed when an item becomes out of date.

- David Dent

* * * * * * * * * * *

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 127). 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	(17 Dec. '81) Priority groups in 1982
127	(25 Jan. '82) Cyber 730E introduction - user interface advice
	IMSL Library

The following News Sheets can be thrown away since this list was last published: 101, 106, 122 (12 Oct.'81), 124 and 125.

* * * * * * * * * * *

ECMWF PUBLICATIONS

Technical Report No. 27:	The energy budgets in North America, North Atlantic and Europe based on ECMWF analysis and forecasts.
Technical Report No. 28:	An energy and angular momentum conserving finite-difference scheme, hybrid coordinates and medium-range weather prediction.
ECMWF Forecast Report No. 13:	January - March 1981
ECMWF Forecast Charts:	Up to December 1981
ECMWF Verification Charts:	Up to December 1981

* * * * * * * * * * *

CALENDAR OF EVENTS AT ECMWF

15 - 19 March	Computer course B	see details in Newsletter No. 12 (Dec. 1981)
22 - 26 March	Computer course C	
29 March - 2 April	Computer course D	
28/29 April	Council 15th session	
4 - 28 May	Research Department course A1	see details p. 10
1 - 11 June	Research Department course A2	
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 both posts.

POST: OPERATIONAL ANALYST

COMPUTING ENVIRONMENT: The computer facility currently comprises a CRAY-1 and a CDC CYBER 175 linked by telecommunications computers (RC8000) to a private network spanning Member States. A second CDC CYBER computer has been added to the facility late in 1981.

The computer facility provides a continuous service, 24 hours per day, 365 days a year. Therefore, it might be required of the analyst to work outside the Centre's normal working hours from time to time.

MAIN DUTIES:

- Controlling the operational implementation of applications and basic software.
- Liaison with applications and basic software teams, including communications, to develop the technical operational interface to these teams.
- Involvement with final testing and quality assurance for software enhancements.
- Ensuring that operations procedures are accurate and up-to-date, for all software and hardware changes.
- Assisting with the training of computer operators.

The analyst will work in the Computer Operations Section and report to the Head of that Section. There is a need for a dynamic involvement, working closely with other staff providing the operational and meteorological service.

QUALIFICATIONS: University education or equivalent, completed by training in programming and analysis. A solid background in software engineering, communications and Computer Operations is highly desirable.

STARTING DATE: As soon as possible.

CLOSING DATE: 11 February 1982

POST: SENIOR SCIENTIST

MAIN DUTIES: The responsibilities of this post, in the Diagnostic Section of the Research Department, will include diagnostic research on atmospheric behaviour and on the performance of both the data assimilation system and the forecast model. The intent is to identify the sources of errors in the Centre's forecasts.

QUALIFICATIONS: A university education (Ph.D.) or equivalent and several years post-graduate experience in dynamical meteorology, diagnostic studies and related areas of meteorology. Familiarity with numerical weather prediction or general circulation modelling in the areas of modelling or diagnostics would be an advantage.

STARTING DATE: As soon as possible.

CLOSING DATE: 5 March 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	- Peter Gray	CB 133	323
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