



European Centre
for Medium Range Weather Forecasts

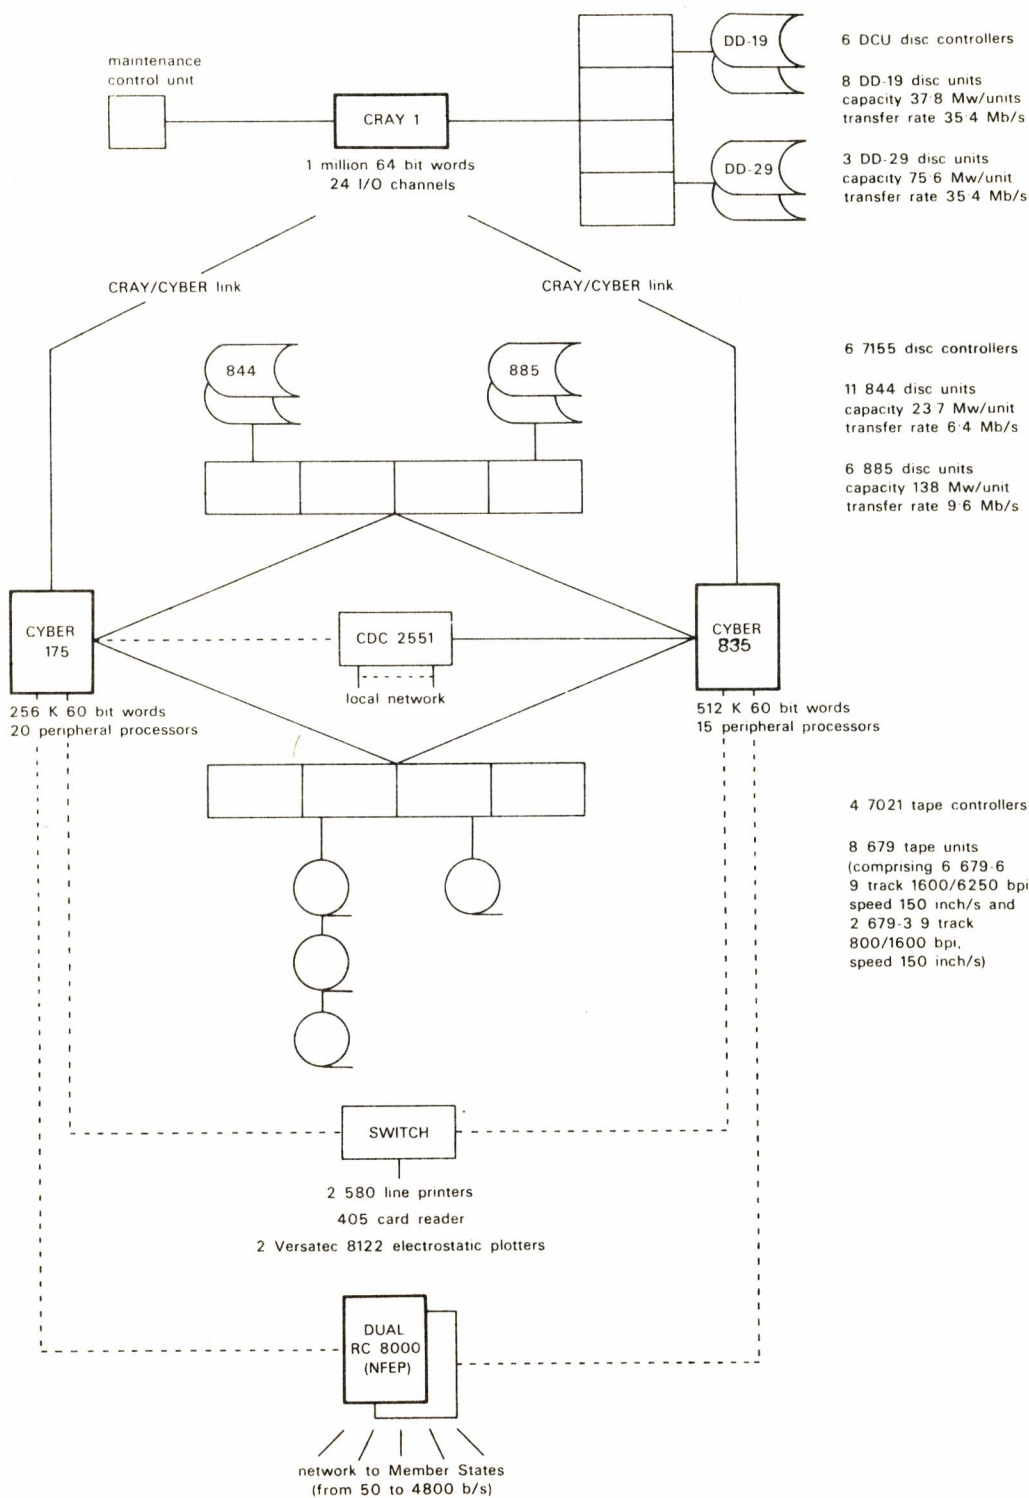
ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England

Reading (0734) 85411 Telex 847908

**NOT TO BE
TAKEN AWAY**

Number 15 - June 1982



IN THIS ISSUE:

| METEOROLOGICAL | Page |
|---|------|
| Development of a new forecasting system at ECMWF | 1 |
| Verification of ECMWF forecasts at the National Meteorological Center, Washington, USA | 4 |
| | |
| COMPUTING | |
| Progress in the Operations Department | 5 |
| Cray User Group - Dallas 27-29 April 1982 | 6 |
| Cray X-MP and SSD | 7 |
| *The use of the CFT flowtrace facility | 7 |
| Highlights from the VIM/ECODU Conference, St. Paul, USA | 8 |
| *Execution time differences on the Cyber front ends | 9 |
| *Buffering or non-buffering on Cyber? | 10 |
| Fortran 5 optimisation techniques for the user | 10 |
| NFEP terminal statistics | 14 |
| *FORMAL - a text processor | 15 |
| In-house alphanumeric COM system | 15 |
| Punched cards at ECMWF | 15 |
| ECLIB documentation | 16 |
| Still valid News Sheets | 17 |
| Computer usage statistics 1982 | 18 |
| | |
| GENERAL | |
| Computer user training courses | 20 |
| ECMWF publications | 20 |
| Calendar of events at ECMWF | 20 |
| Technical Advisory Committee - participants | 21 |
| Index of still valid newsletter articles | 22 |

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

COVER: Computer configuration at ECMWF as on 31 March 1982,
showing the second Cyber front end.

This Newsletter is edited and produced by User Support.

The next issue will appear in August.

DEVELOPMENT OF A NEW FORECASTING SYSTEM AT ECMWF1. Introduction

As mentioned in D. Burridge's description of ECMWF research work for 1982 (ECMWF Newsletter No. 13), the main goal is the development of a new forecasting system, the basis of which will be a model using a spectral representation in the horizontal, and a hybrid coordinate system in the vertical. The first version will soon be available for research and development projects, after a short test period, and the target date for its introduction into operation is January 1983.

The purpose of this short paper is to describe the main features of this first version and the status of the development work.

2. Description of the model2.1 General design and coding principles

The main goal during the designing and coding has been to achieve the greatest flexibility and simplicity of use for any potential "customer". In order to increase the simplicity, a memory manager package has been written, adding to the clarity of the code (meaningful array names are used rather than pointer displacements in a large all-purpose array), and making easier the requesting of new working space without having to know the internal structure of the model.

The documentation of the model and its updating have been made easier and more automatic by the adoption of a new coding standard (DOCTOR, Technical Memorandum No. 52) developed at ECMWF.

2.2 Adiabatic part

As stated in the introduction, the technique chosen for the horizontal discretisation is the spectral one: the prognostic fields are expanded in terms of spherical harmonics, and only a subset of the total expansion is used (a "truncation" procedure). All the non-linear terms (including the physical parameterisations) are evaluated on an associated grid, called the Gaussian grid, which is an almost regular latitude longitude grid. One of the main advantages of the spectral technique is that it allows an exact computation of the horizontal derivatives (within the chosen truncation) and thus, a greater accuracy in the treatment of the horizontal advection. An immediate consequence is a reduction in the phase error for fast moving systems.

The vertical coordinate is of a hybrid type, following the terrain near the surface (sigma-like) and changing smoothly to pressure in the lower stratosphere. Amongst many others, one advantage is the possibility of having the same stratospheric levels in the analysis and in the model. The technique used for the vertical discretisation is a classical finite differences one, but alternative techniques (using finite elements) are being investigated.

A particular emphasis has been placed on achieving the maximum flexibility for research without serious loss of efficiency in the operational version.

Various scanning structures (2, 3, 5 or 9 scans through the data) will make experiments with very high resolution feasible.

There are almost no constraints in the choice of the truncation. In particular, this can be used to reduce the horizontal resolution used in the stratosphere (in connection with the hybrid coordinate).

It will also be possible to modify the number of points on each latitude line (and in practice to reduce it towards the poles).

The introduction of new advective variables will be relatively straightforward.

2.3 Physical parameterisations

Contrary to the adiabatic part, it was decided not to change drastically the physical parameterisations for the first version of the new forecasting system. It was, however, decided to have a complete recoding with three main objectives:

- i. "cleaning" and optimisation of the various routines;
- ii. automatic documentation extraction;
- iii. introduction of small changes, either made necessary by the hybrid coordinate or in order to get a better physical description of atmospheric processes.

Three of these changes can be mentioned as being particularly important:

Firstly, in stability computations, the potential temperature is replaced by the dry static energy, in order to avoid expensive p^{R/C_p} calculations. Secondly, in condensation processes, a distinction is made between liquid water and ice/snow phases. Lastly, a moisture dependence of the specific heat of the atmosphere is introduced (this is also done in the energy conversion term of the adiabatic equations). Preliminary results suggest that the second and third above-mentioned modifications could go a long way towards curing the mid tropospheric cooling problem of the ECMWF model.

3. Present development status

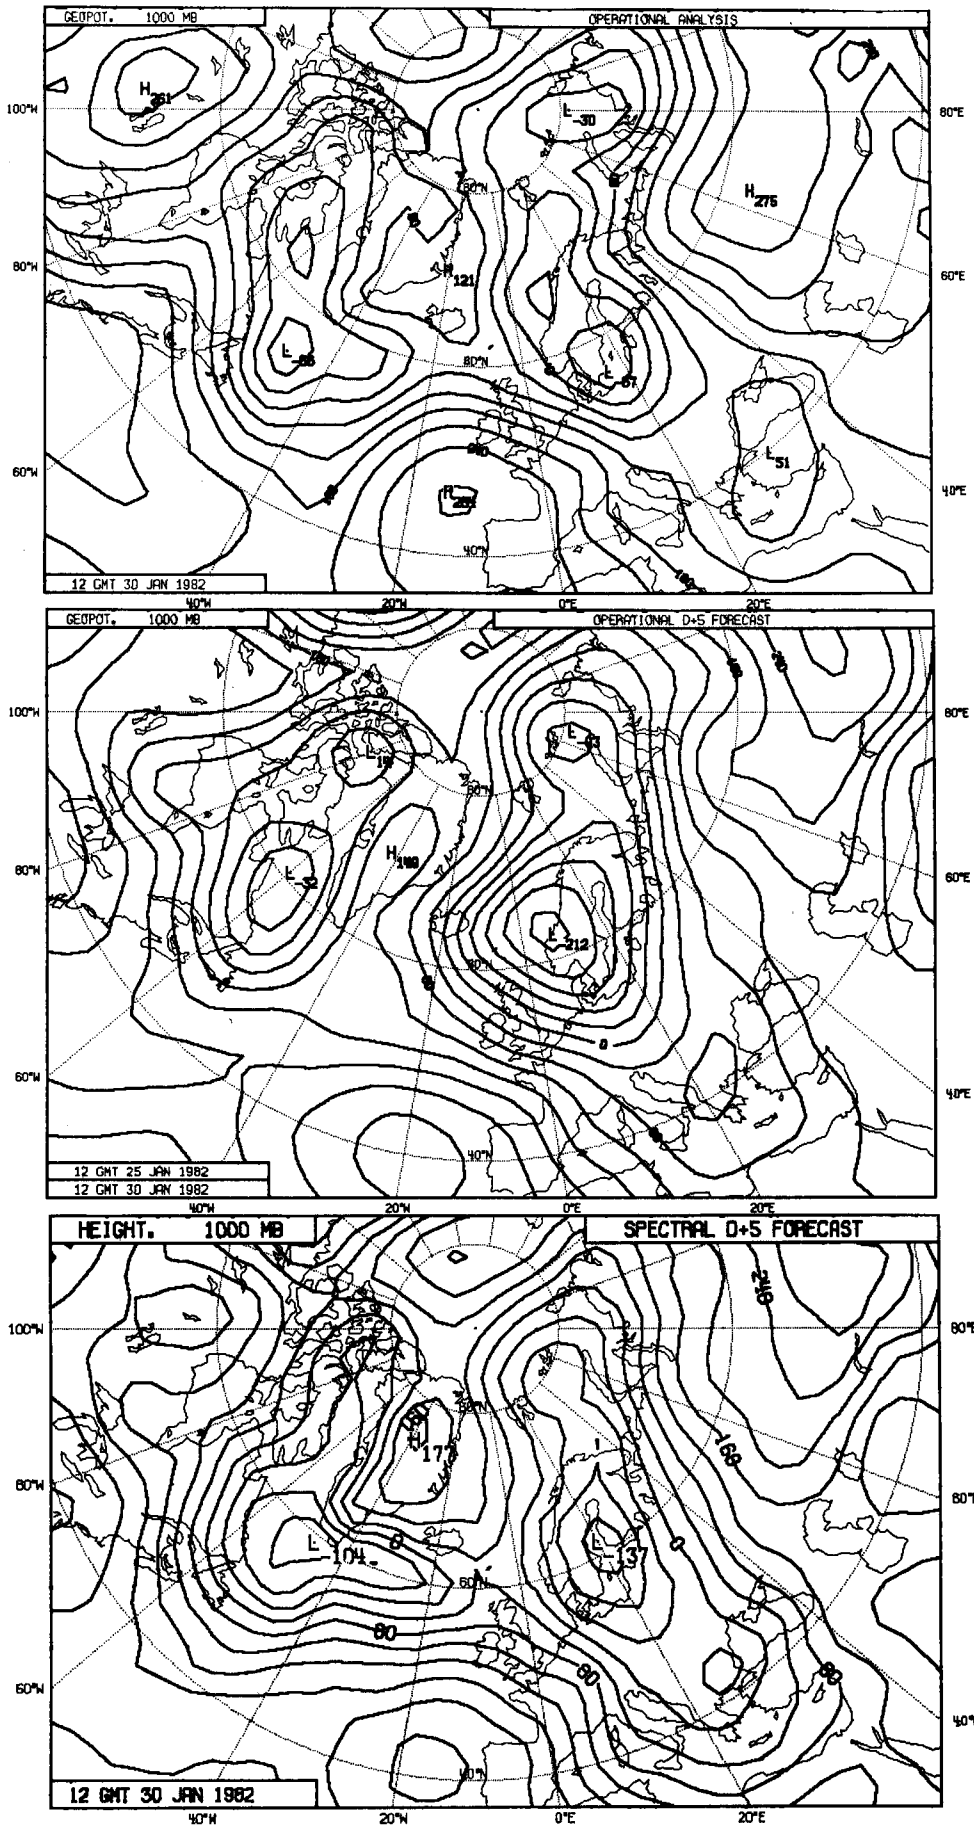
The adiabatic part of the model has been successfully tested and temporary interfaces have been written, making the new model fully compatible with all the software available for the present spectral model. It can, in particular, already be used for much research development work (for example, the testing of finite elements technique). The recoding of the physical subroutines is well advanced and should be completed by the time this newsletter is published. The new model will thus be fully available from this summer onwards as the main Research "test vehicle".

For operational implementation, the normal mode initialisation will be recoded in order to work on the spectral components rather than on grid point values, as at present. Some modifications are also to be brought to the interfaces between the model and the analysis in the data assimilation.

In conclusion, we would like to stress that, so far, development is progressing according to schedule (as presented in Newsletter No. 13) and we are confident that, from the beginning of next year, users will benefit from forecasts which are slightly better on average, and much better several times a month, than the present forecasts. A recent example of striking differences over Europe and the North Atlantic is shown in Figure 1.

- Jean-François Geleyn
Rex Gibson
Michel Jarraud
Adrian Simmons

* * * * *



Maps of 1000 mb height over Europe and the North Atlantic. The upper map shows the analysis for 30 January 1982, the middle map the 5-day operational forecast verifying on this date and the lower map the corresponding spectral forecast. The operational forecast exhibits a spurious ridge along 40°W and a much too deep low centred over the Norwegian coast, with an erroneous cyclonic circulation from Iceland to Western Europe. The improvement in all these features, brought about by use of the spectral model, is evident.

VERIFICATION OF ECMWF FORECASTS AT THE NATIONAL
METEOROLOGICAL CENTER, WASHINGTON, USA

ECMWF GRID code products, including 500 mb geopotential height forecasts from D+1 to D+5 are being decoded and verified at NMC Washington. For verification, radiosonde data from over 100 radiosonde stations over the northern hemisphere are used. The scores for February 1982 and March 1982 have been communicated to ECMWF by J.A. Brown, Chief of the Development Division at NMC, and are shown in Figure 1. The ECMWF forecast fields are from grid point values at 5° intervals. Also plotted in Figure 1 are the equivalent scores from the NMC operational spectral model. The lower errors of the ECMWF forecasts are obvious; the 5-day ECMWF March forecast scores were equivalent to those of the 4-day NMC forecasts.

- Austin Woods

* * * * *




Fig. 1 Standard deviation of forecast error northern hemisphere scores for the NMC spectral model (solid line) and for the ECMWF GRID code 5° resolution fields for February 1982 (top) and March 1982 (bottom).

PROGRESS IN THE OPERATIONS DEPARTMENT

The routine production and dissemination of operational medium range forecasts has been maintained reliably. No forecasts have been lost since September 1980. Operational delays have occurred, but have, in general, been short.

The range of products available to Member States continues to grow, and now includes some experimental material, such as the temperature at 2m, and forecasts for days 8-10. The number of forecast products disseminated to Member States now exceeds 7000 per day. The substantial increase in the number of products disseminated to Member States has led to dissemination delays of the order of 30-90 minutes for some products, notably for 48-72 hour forecasts. This is primarily due to the limited capacity of the current dissemination link between the Cybers and the NFEP; it is hoped that this problem will be resolved in the near future. The dissemination of a limited range of the Centre's products to non-Member States via the Global Telecommunications System (GTS) began on 1 August 1981 and products have been received with enthusiasm. From 1 July 1982, tropical 24 and 48 hour forecasts will be added to the subset of data distributed via the GTS.

Access to the Centre's archives has now been facilitated for Member State users by the distribution of full documentation on the GETDATA utility. A product development group has provisionally been set up to establish further access facilities which will support Member States' local verification and interpretation projects. The final aim of the group is the availability of a system which allows selective retrieval and use of data for local verification and interpretation. The test system which is at present being prepared will be demonstrated to the participants at the seminar and workshop in September this year.

The telecommunications network is progressing well: eleven Member States are now connected to the Centre by medium speed lines, the latest additions being Finland, Austria, the Netherlands, Portugal and Spain. Only four Member States remain dependent on low speed lines, as Switzerland and Belgium receive an agreed selection of products via Offenbach and Bracknell respectively. One new development has been the installation of a second low speed line to Italy, enabling direct Cyber access. This is an experimental arrangement for two years.

Finally, the computing and technical facilities on which the forecast depends have been developing steadily. The second Cyber front-end, the Cyber 835, was installed at the end of November and completed the Final Acceptance period on 17 March 1982. It has continued to run reliably since then and is fully integrated into the system, normally providing valuable additional resources to support the Cray, and back-up on those occasions when the Cyber 175 is unavailable. The Cray service has been improved by the installation of two additional DCU-2 disc units and three DD-29 disc drives, in July 1981. This enhancement resulted in the doubling of the available on-line storage space. Output media available have now also been extended by the acquisition of in-house alphanumeric COM (Computer Output on Microfiche) facilities, a Datagraphix Autocom II device now replaces the former bureau service. The faster, more acceptable turnaround time has already increased user demand considerably.

The next significant enhancement of the Centre's computing resources will take place on 1 July 1982, when the Centre will begin to lease time on the Cray-1 machine at Harwell. This will provide a considerable amount of additional "number-crunching" resources for the Research Department's work.

- Daniel Söderman

* * * * *

CRAY USER GROUP - DALLAS 27-29 APRIL 1982

The hosts for the Cray User Group meeting were Atlantic Richfield Oil and Gas Company (ARCO), based in Dallas, Texas.

Cray reported that their 1981 revenues increased by 67% over 1980, to \$101.6 million. During the six months since the last CUG meeting, 8 new machines were delivered. Cray now employs 1200 staff worldwide.

The Cray X-MP and SSD announcements were also made, and are more fully covered below.

During the software presentations, the most burning issue was the state of CFT 1.10. Margaret Loftus, the Vice President of Software, has put in a great deal of effort to resolve the CFT problems. However, at the time of the CUG, the best estimate for a release date was "not before July 1982."

Cray have contracted with the University of Manchester, England, to produce a PASCAL compiler. This should be running later in 1982 and Cray intend to use PASCAL as a systems implementation language in the future. Work has begun on a new Fortran compiler, written in PASCAL.

COS 1.11 should be released in June or July, at about the same time as CFT 1.10.

Various other software technical issues were covered, and schedules discussed.

Cray also discussed their plans for their own corporate datacentre. By the end of 1982, it should contain:

- 2 Cray 1S systems
- 1 Cray X-MP
- 1 SSD system
- 1 Amdahl V7
- 1 IBM 4341
- 1 Cyber 73

When compared with facilities 5, or even 3 years ago, it is clear that Cray now realise the importance of providing powerful facilities for both marketing and development purposes.

- Peter Gray

* * * * *

CRAY X-MP AND SSD

On 26 April 1982, Cray Research Inc. announced two new products, the Cray X-MP, a multiprocessor version of the Cray 1, and the Solid state Storage Device (SSD) which is a large semiconductor memory, accessed rather like a disc system.

The Cray X-MP contains 2 Cray CPUs, software compatible with the Cray 1 CPU. Each CPU has a 9.5 nanosecond clock period giving 105 MFLOPs per processor, and accesses a shared 2 or 4 million word memory. The dual processor system allows either loosely coupled multiprogramming (1 job using each CPU) or single program multiprocessing (1 job using both CPUs). This latter facility will be supported by both COS and CFT, and common versions of the systems software will run on both the Cray 1 and the Cray X-MP.

The memory bandwidth has been increased eightfold and vector register to memory operations now chain with other operations. On the X-MP, chaining always occurs, so there are no problems due to missed chain-slot time. These features mean that the Cray will provide good throughput on long vector processing. The Cray X-MP comes complete with an Input/Output Subsystem, containing 2 to 4 I/O processors.

The Cray SSD contains 64, 128 or 256 million bytes of MOS semiconductor memory (8, 16, or 32 million words). The SSD is driven rather like a disc subsystem by the user, except that seek time is zero (compared with 50 milliseconds average for a DD-29 disc), and except for a transfer rate of up to 100 million bytes/second to the Cray 1S or up to 1250 million bytes/second to the Cray X-MP (compared with 4.5 million bytes/second for a DD-29 disc).

Clearly the SSD will improve the performance of any input/output limited jobs and early comparisons of the SSD linked to a Cray 1S, indicate an I/O throughput improvement of 25 to 30 when compared with a DD-29 disc-based system. Greater improvements should be seen when the SSD is linked to the Cray X-MP.

- Peter Gray

* * * * *

* THE USE OF THE CFT FLOWTRACE FACILITY

By using the parameter 'ON=F' on the CFT control card it is possible to obtain a detailed timing summary for the execution of the program and subprograms being compiled. This information is useful in determining which subroutines are worth optimising. Unfortunately, this facility has the undesirable effect of causing the operating system to use a lot of CPU time. It would therefore be helpful if users would remove the 'ON=F' parameter from the CFT control card as soon as they finish their timing tests. Do not leave the parameter on the control card if you do not intend to use the timings produced.

Another point with Flowtrace concerns those programs which abort, without producing the desired timings. There is a utility called 'FLODUMP', which can be used to give the timing details in this case.

An example of how to use FLODUMP is:

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

FLODUMP is described in the COS Reference Manual (SR-0011).

- Neil Storer

* * * * *

HIGHLIGHTS FROM THE VIM/ECODU CONFERENCE, ST. PAUL, USA

The new CYBER Front-End

The new CYBER, so far referred to as CYBER 170-730E, passed its final acceptance in mid March and was formally announced by CDC at the worldwide VIM/ECODU conference (26-29 April) in St. Paul, USA. The machine is part of a complete range of computers called the CYBER 170 Series 800 (825,835,855,865,875). Our model is a CYBER 170-835 and will henceforth be referred to as the CYBER 835.

A CYBER 835 contains a microcoded central processor. MOS central memory can be expanded from 0.5 Mwords to 2 Mwords. Memory can be partitioned into central memory and unified extended memory. Unified extended memory (UEM) is functionally similar to ECS (Extended Core Storage) or ESM (Extended Semi-conductor Memory) but cannot be shared with other CPU's. The model 835 also contains a 2K or 4Kword cache memory which holds the most recently accessed central memory words and words expected to be used next, resulting in a shorter access time for data and program instructions. There is also an instruction lookahead register which can hold up to three words (with up to 4 instructions per word) ahead of the current instruction in the regular program sequence, and two instruction words in a branch sequence. As with all 170 series machines, up to 20 peripheral processors and 24 hardware channels are available. The whole system is organised in a three-section cabinet (central processor, central memory, I/O unit) and is refrigerant-cooled.

Our model has a 512 Kword central memory (half of it being used as UEM since the NOS/BE operating system is restricted to using 256 Kwords), a 2 Kword cache memory and 15 peripheral processor units (PPU's). A further 5 PPU's will be installed in the near future, to allow better support of a full service when the CYBER 175 is not available. The machine runs NOS/BE and is connected with the CYBER 175 via shared discs. All tape units, the unit record equipment (printers, plotters, card reader) and the INTERCOM service can be connected to the 835 or the 175 to guarantee full backup.

The power of a CYBER 835 is approximately 3/8 of a CYBER 175. NOS/BE on the 835 has been modified to reflect this in resource reporting, thus showing 175 units. However, units shown on the 175 and the 835 can differ considerably for certain jobs, due to the different architecture of the 835.

During prime shift the CYBER 835 mainly services the interactive users, while at all other times batch work is admitted as well.

Graphics at VIM/ECODU

- Frank Hollmann

CDC announced a number of graphics applications including the Environmental Image System CDC-McIDAS. Included within McIDAS are algorithms to extract wind, moisture and temperature information from satellite-derived imagery and from surface and balloon sounding reports. Other features are image enhancements (false colour techniques), examination of observations and determination of cloud motions. The McIDAS workstation connects to a Cyber and consists of a standard alphanumeric terminal and a RAMTEK 9400 graphics controller equipped with a high resolution monitor and with a tracker ball. The alphanumeric terminal is used for the normal login procedure and to control the execution of the McIDAS application program running on the Cyber. The graphics results (1024 x 1180 resolution) are displayed immediately on the colour monitor via high speed interface (2.8M bits/sec.).

In a demonstration centre at the conference many of the CDC graphics applications including McIDAS were shown running on a Cyber 170-825 mainframe.

- Jens Daabeck

* * * * *

* EXECUTION TIME DIFFERENCES ON THE CYBER FRONT ENDS

As shown in the article on the speed of arithmetic instructions in the April Newsletter, all floating point addition and subtraction instructions take 100 n.s. on the Cyber 175. However, on the Cyber 835 they vary from 224 n.s. to over 3000 n.s. Similar results also apply to the non-arithmetic instructions.

The result of the above differences can be seen in the execution times for various jobs. Although we make an adjustment of 3/8ths to the execution times on the Cyber 835, this figure is only an average. Some jobs will appear to use less CPU time on the Cyber 835 and some will appear to use more. One particular type of job that falls into the latter category is some large I/O jobs. Investigations into the cause of the problem are still progressing, but initial results suggest that it may be caused by the particular mix of instructions used by Cyber Record Manager. Users who have jobs which hit this problem will need either to increase their job's time limit so that it will not run out of time if run on the Cyber 835, or to force their jobs to run on the Cyber 175; User Support can advise how to do this.

- Gary Harding

* * * * *

* BUFFERING OR NON-BUFFERING ON CYBER?

1. Some figures

WRITE 200 records of 15139 word length costs 30.8 CP secs.
 " 2000 " of 15168 " " " 1.3 CP secs.

These figures were obtained on the Cyber 175 for 'DIRECT ACCESS' files created by using the Fortran 77 statement WRITE(1,RECORD=I)ARRAY; the figures are identical if CALL WRITMS is used.

Notes: i) 15168 is a multiple of 64, the nearest to 15139
 ii) I/O is unformatted and the iolist consists of only one item, the array name.

2. Record size = n*64 words(one PRU)

The file organisation which is used by the Cyber Record Manager (CRM) to create such files is 'Word Addressable'. If the user chooses for his record size an exact number of PRU's, CRM will transfer to disc from the Working Storage Area (user's array) without using the buffer associated to the file, hence saving all the memory transfers needed to move the user's array to the file buffer.

3. Moral

Whenever there is to be heavy usage of Direct Access files (Fortran 77) or READMS/WRITMS on the Cyber, making the record size a multiple of PRU's (by padding records out, for instance), can save quite a substantial amount of resources. Do not forget that the figures listed in 1. apply only to 'unformatted' WRITE with a single entire array as iolist.

- Michel Miqueu

* * * * *

FORTRAN 5 OPTIMISATION TECHNIQUES FOR THE USER

Part 2 from a talk presented by Paul Derby, CDC, at the VIM 35 Language Processors Committee (Part 1 appeared in Newsletter No. 14)

Central Memory Optimisation

Central memory is a resource that can be reduced. Care must be exercised since reducing the amount of central memory generally results in an increase in other resources, especially elapsed time. "Legitimate" reasons for optimising central memory include:

1. The hardware/software limit has been exceeded.
2. The user validation limit has been exceeded and cannot be increased.
3. The operating system swapping algorithm and scheduler thresholds "favour" smaller programs.
4. Large portions (initialisation, termination, etc.) of the program are not executed for long periods of the run.

There are many techniques to reduce memory. Some of the techniques are described below.

I/O Buffer Reduction

Each file that is accessed by Fortran 5 requires that a CIO buffer be allocated in central memory. The default size of each buffer selected by Fortran 5 is 2000 octal words. Optimisation of the CIO buffer size consists of reducing the buffer size when a large buffer is not required and totally releasing the buffer when the file is no longer required in the run.

A relatively inactive file or a file which contains a small amount of data does not need a 2000 word buffer. The CIO buffer size can be changed with one of three techniques:

1. Specify the file name and the buffer size on the PROGRAM statement, e.g. "...TAPE10=65..." (Note that 65 words is the smallest CIO buffer for files on disc).
2. Specify BFS on a FILE control statement before the program is executed.
3. Specify BUFL on an OPEN statement in the FORTRAN program. This option allows the programmer to determine the file "activity" and select the buffer size at execution time.

The minimum buffer size for a file depends on the device on which the file resides. Disc files require 65 words. Tape files require enough words to hold one block from the tape. Terminal files that do not use graphics-type devices require 0 words since CRM "funnels" all terminal output through one internal CIO buffer.

There are two differences in buffer allocation that the FORTRAN 5 programmer should be aware of that are different from FORTRAN 4. The first is that CIO buffers are allocated dynamically by CMM in the dynamic area when the file is first referenced. The second change is that CIO buffers can be released (memory given back to CMM and field length reduced) by means of the FORTRAN 5 CLOSE statement. As soon as a file is no longer needed, a CLOSE should be performed. The CLOSE will release the CIO buffer and associated CRM modules, if no other files of the same type are using the modules.

Loading Techniques

There are many cases in a large, long running program where entire modules are not referenced. Typical modules include initialisation and termination modules. Overlays, OVCAPs and segmentation provide techniques to store the unneeded modules on disc and bring them quickly into memory when needed.

Overlaying is a very old loading technique to reduce net memory requirements. The technique is well described in the FORTRAN 5 reference manual and the LOADER reference manual. Current users should be aware of a change that was made in the loading efficiency of overlays. The change is called the Fast Overlay Loader (FOL). FOL generates an index pointing to each overlay on disc. This allows the loader to randomly access the disc to read in the overlay. This eliminates many sequential disc accesses during execution of the typical overlay program. The only two things that the programmer needs to do is to count the primary and secondary overlays and place the count in the main overlay statement "OVERLAY(ABS,0,0,OV=n)".

"Overlay capsules" (OVCAPs) is a new loading technique for the FORTRAN 5 programmer. Overlays have several disadvantages that OVCAPs have eliminated.

1. An overlay program requires a minimum field length allocation sufficient to hold the largest overlay. This "space" is non-existent for OVCAPs.
2. Overlays are restricted to 3 levels. OVCAPs can have up to 4095 levels.
3. Overlays cannot be "unloaded". OVCAPs can be unloaded when no longer needed and the memory is given back to CMM (and the field length normally decreases if the block is at the end of the field length).
4. Overlays must be loaded at one and only one memory location. OVCAPs can be loaded at any memory location.

The FORTRAN 5 programmer's interface to OVCAPs consists of 3 sub-routines - XOVCAP, LOVCAP and UOVCAP. The subroutine calls are described in the FORTRAN 5 reference manual. Throughout the discussion of OVCAPs, one should keep in mind that all OVCAPs are a logical extension of the main overlay, i.e. each OVCAP can reference modules (entry points) and common blocks in the main overlay. OVCAPs can also reference (cross-link) entry points in other OVCAPs. Common blocks cannot be shared between OVCAPs except via the main overlay.

Segmentation is a loading technique that is similar to overlays but with most of the restrictions removed. A segment structure can consist of up to 4095 "levels" compared with 3 for overlays. Segmentation does not require that the program be recompiled to change the structure. Source changes are not required to convert a program for a basic load to segmentation. Segmentation is described in the Loader reference manual.

Data Storage

In many applications, the amount of memory required for storing data is larger than the memory for storing the instructions to operate on the data. The amount of memory required for storing the data structure can be reduced by some of the techniques listed below.

EQUIVALENCE - The Equivalence statement can be used to "share" memory between large arrays that are not needed at the same time.

Matrix Storage - Many scientific applications involve the use of large arrays to store matrices. In many of these matrices, only a small percentage of the elements have non-zero (interesting) values. This is called a sparse matrix. A sparse matrix can be stored as an (n, 3) array. Element (n,1) could be used to store the row number, element (n,2) to store the column number and (n,3) to store the actual data value.

Data Packing - Data packing is a technique of storing data in which more than one small integer value is stored in a single word. If the value of integer variables in an array is small, it is not necessary to use an entire 60 bit word to represent the value. With the use of the MASK and SHIFT functions, multiple integer values can be inserted into a single word. Obviously, this technique requires more CPU time.

Double Precision - Single precision on CYBER machines maintains accuracy to 14 decimal places. Other machines will use double precision to get the desired accuracy. Programs that are converted from other machines probably do not need double precision on CYBER machines. The change to single precision will reduce memory requirements and also reduce CPU time, since the use of double precision has a large premium on CPU time.

Programmer Time

The time to develop, debug and maintain a program is classified as programmer time in the following paragraphs. There are many new features in FORTRAN 5 that make it easier and thus less expensive to develop, debug and maintain programs.

IF statement. The FORTRAN 5 IF statement has been enhanced to include a block structure construct. The IF/THEN/ELSE construct allows the programmer to eliminate statement labels, eliminate GO TOs and generate more readable and more maintainable code. Consider the readability of the following FORTRAN 5 code fragment:

```
IF (X .GE. 0) THEN
  Y = SQRT (X)
ELSE
  PRINT *, 'SQRT OF NEGATIVE VALUE', X
ENDIF
```

The ELSE IF construct is very powerful for testing different values of a variable that can have a wide range of values.

```
IF (AGE .GT. 65) THEN
  SALARY = WAGES * .7
ELSE IF (AGE .GT. 60) THEN
  SALARY = WAGES * .6
ELSE IF (AGE .LT. 18) THEN
  SALARY = WAGES * .52
ELSE
  SALARY = WAGES
END IF
```

Note that the code has been indented. This adds significantly to program readability and hence maintainability. In order to decrease CPU time the condition most likely to occur could be placed first in the sequence of IFs.

Character data type - Another major feature of FORTRAN 5 is the character data type. It allows much simpler character handling than the Hollerith character strings and the masking and shifting required to process them. In addition, character data variables are not dependent on the number of characters per word, so character handling programs are portable. Standard constructs exist to change individual characters in a large character variable, concatenate separate character variables, reference substrings of character variables and move character variables.

Array bounds - FORTRAN 5 allows both upper and lower bounds to be specified. For example, a program fragment dealing with an array of years from 1964 to 1981 might be easier to comprehend if the array is declared:

```
INTEGER YEARS (1964:1981)
```

Summary

The purpose of this paper has been to outline various ways in which FORTRAN 5 programs can be constructed to perform better on Control Data computers. Program optimisation, program style, compiler optimisation parameters, and programmer time considerations have all been mentioned. Hopefully, these ideas will help the Fortran programmer to construct and evaluate the needs of each particular application.

* * * * *

NFEP TERMINAL STATISTICS

From 15.3.82 - 2.5.82

| COUNTRY | AVERAGE | | | |
|----------------|------------------------|--------|--------------------|--------|
| | TOTAL DATA (KCHAR/DAY) | | DATA RATE (CH/SEC) | |
| | INPUT | OUTPUT | INPUT | OUTPUT |
| Denmark | 18.4 | 1552.9 | 127.3 | 221.9 |
| F.R. Germany | 5311.7 | 2952.0 | 70.8 | 305.3 |
| Ireland | 28.2 | 2313.2 | 186.2 | 208.1 |
| Spain | 25.3 | 578.2 | 234.2 | 114.0 |
| France | 196.6 | 1618.3 | 138.7 | 307.5 |
| *Greece | 0.0 | 537.5 | 0.0 | 13.2 |
| *Italy | 0.0 | 307.1 | 0.0 | 13.3 |
| *Yugoslavia | 0.0 | 49.5 | 0.0 | 6.5 |
| Netherlands | 45.8 | 1077.9 | 143.4 | 184.3 |
| Austria | 61.6 | 1190.8 | 230.5 | 201.7 |
| Portugal | 5.7 | 1983.6 | 204.5 | 199.9 |
| Finland | 18.1 | 2042.7 | 34.3 | 294.2 |
| Sweden | 9.9 | 4542.0 | 263.8 | 277.9 |
| *Turkey | 0.0 | 119.0 | 0.0 | 6.6 |
| United Kingdom | 5927.4 | 1174.9 | 74.3 | 184.9 |

*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.

* * * * *

* FORMAL - A TEXT PROCESSOR

A text formatting program called FORMAL has been available on the Cybers for several months on an experimental basis. It is now generally available to the computer user community and in conjunction with EDIT and the QUME daisy wheel printer may be used to generate high quality, well-formatted documents.

The raw text must be a Cyber ASCII file so it is natural to use the editor EDIT to create and modify the text. The text formatter is available in the procedure library and can be invoked by typing

 FORMAL.....
or RUNFORM.....

Documentation in the form of a reference bulletin and a user guide is in preparation.

- David Dent

* * * * *

IN-HOUSE ALPHANUMERIC COM SYSTEM

Our DatagraphiX COM (microfiche) system completed its provisional acceptance trial successfully on the 7th April 1982 and went directly into service. The turnaround for a fiche job is now only a few hours compared with the minimum of 2 days with the previous bureau service. In the first four weeks of service 1200 fiche were produced, this is the equivalent of approximately 200,000 pages of paper output.

A new fiche system is currently being developed. This will allow users to produce fiche directly from any job, including an INTERCOM session, as it will no longer be necessary to mount a magnetic tape. The new systems will also produce a new type of index on the fiche, be able to produce fiche with full ASCII upper and lower case files, and allow users to have more control over the indexing performed. If you wish to try this system on an experimental basis before its release, due sometime in late June, please contact me.

- Gary Harding

* * * * *

PUNCHED CARDS AT ECMWF

As mentioned in the February ECMWF Newsletter (Cyber 730E progress, pages 16/17) the third meeting of the Technical Advisory Committee decided that punched cards should be phased out at the Centre. Now that the additional 2551 ports into Intercom are installed, card users can move to Intercom use. We urge all remaining card users to transfer to VDU use as soon as possible. So far the following changes in equipment have taken place:

One central site fast card reader and the on-line card punch have been returned to CDC (the money released has been used to rent two extra tape drives).

A simple on-line card punch service has been provided (via the NFEP) for cases when cards are still required, eg to send to another site. However, this service is rather limited, for example:

- it cannot punch binary characters;
- there are problems in handling decks of more than 400 cards;
- there are problems in handling the punching of several files in one session (contact the Advisory Office if you wish to do either of these latter functions).

In the future, the following will take place:

The two UNIVAC card punches will be disposed of as soon as possible.

The two IBM card punches will be returned to IBM at the end of the year.

The two DECISION DATA punches will be kept, one to provide the on-line card punch capability via the NFEP (see above), the other for emergency card deck preparation.

- Andrew Lea

* * * * *

ECLIB DOCUMENTATION

The complete set of ECLIB writeups has now been rewritten and extended to include all routines currently in ECLIB. The documents are being printed and will be distributed shortly. The originals are available in a Cyber fileset and additional copies may be obtained with the following command sequence:

```
ATTACH,DOCSET.  
GF,ECLIB/name,DOCSET.  
FORMAL,name,PRINTER.  
or FORMAL,name,CONNECT.
```

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

| <u>No.</u> | <u>Still Valid Article</u> |
|------------|--|
| 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 |
| 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 |
| 130 | Cyber software: PACKS; SPACE Contouring package: addition of highs and lows |
| 131 | File storage on TEMP |


The following News Sheet can be thrown away since this list was last published: 116

* * * * *

- 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



COMPUTER USER TRAINING COURSES

The Centre is offering another series of training courses for Member States personnel and ECMWF staff. Information has been sent to the Member States and nominations to attend are now invited. Nominations from ECMWF staff will shortly be invited via Section Heads.

Course B BASIC USAGE 11-14/15 October

This is intended for anyone who will be actually programming the Cray-1, to give them sufficient experience to run simple work. It will also introduce them to some of the Cyber facilities they may need to complement their Cray activity. Prior knowledge of another computing system, plus a knowledge of Fortran is required. An optional 5th day (15 October) is devoted to explaining how to use ECMWF's meteorological database and archive system.

Course C CRAY USER COURSE 18-22 October

An in-depth course for those who will make heavy use of the Cray-1 and its many unique facilities. Intending participants will be expected to know how to run simple jobs on the Cray-1.

Course D CYBER USER COURSE 25-29 October

Again, an in-depth course, this time for those who will make detailed use of the Cybers in conjunction with the Cray-1. Among other things, it will cover Cyber file handling (including private tapes and discs), and program preparation facilities (UPDATE, debugging etc.). Intending participants will be expected to know the basics of a Cyber NOS/BE system.

- Luigi Bertuzzi

* * * * *

ECMWF PUBLICATIONS

Technical Memorandum No. 54: The effects at ECMWF of the change of codes used for surface observations on 1 January 1982.

Technical Memorandum No. 55: A program for optimal regression analysis.

ECMWF Forecast Charts: up to 30 April 1982

ECMWF Verification Charts: up to 25 April 1982.

* * * * *

CALENDAR OF EVENTS AT ECMWF

| | |
|---------------------|---|
| 15-18 June 1982 | 4th Session of the Technical Advisory Committee |
| 11-15 October 1982 | Course B) |
| 18-22 October 1982 | Course C) see details above |
| 25-29 October 1982 | Course D) |
| 18-19 November 1982 | Council 16th session |

TECHNICAL ADVISORY COMMITTEE

15 - 18 June 1982

LIST OF PARTICIPANTS

| | |
|--|--|
| Austria | Dr. G. Wihl |
| Denmark | Mrs. A.M. Jørgensen |
| Finland | Mr. P. Nurminen |
| France | Mr. J. Lepas Mr. M. Trochu |
| Fed. Rep. of Germany | Dr. W. Buschner |
| Greece | Mr. D. Katsimardos |
| Ireland | Mr. W.H. Wann |
| Italy | Dr. G. de Florio |
| Netherlands | Dr. A.P.M. Baede Mr. R. Brinkhuysen |
| Portugal | Mr. S. Cristina |
| Spain | Mr. B. Orfila |
| Sweden | Dr. L. Moen |
| Turkey | Mr. N. Yaman |
| United Kingdom | Mr. D.H. Johnson Mr. R. Wiley |
| Finance Committee representative | Professor Dr. K. Cihak |
| Scientific Advisory Committee representative | Professor C.J. Schuurmans |

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.

| | <u>Newsletter</u> | | |
|---|-------------------|-------------|-------------|
| | <u>No.*</u> | <u>Date</u> | <u>Page</u> |
| <u>CRAY-1</u> | | | |
| Buffer sizes for jobs doing much sequential I/O | 14 | Apr. 82 | 12 |
| Computer Architecture | T2 | Apr. 79 | 10 |
| | T3 | June 79 | 10 |
| | T4 | Aug. 79 | 8 |
| Control of Cray jobs | 7 | Feb. 81 | 16 |
| COS 1.10 - new features | 12 | Dec. 81 | 5 |
| Dataset storage | 13 | Feb. 82 | 11 |
| File transfer to a named device | T2 | Apr. 79 | 14 |
| Public Libraries | T5 | Oct. 79 | 6 |
| Submit and Dispose | 8 | Apr. 81 | 6 |
| <u>CYBER 175</u> | | | |
| Arithmetic instructions - comparative speeds of execution on the Cyber front ends | 14 | Apr. 82 | 17 |
| Cyber front ends - execution time differences | 15 | June 82 | 9 |
| Buffering or non-buffering on Cyber? | 15 | June 82 | 10 |
| CMM-Fortran interface | 10 | Aug. 81 | 11 |
| Cyber 175 processor described | 1 | Feb. 80 | 6 |
| Cyber 730E - progress | 13 | Feb. 82 | 16 |
| - changes caused by the introduction of | 13 | Feb. 82 | 17 |
| Dynamic file buffers for standard formatted/unformatted data | 3 | June 80 | 17 |
| ECLIB - additions | 11 | Sept. 81 | 13 |
| - changes | 14 | Apr. 82 | 22 |
| Formatted I/O - some efficiency hints | 4 | Aug. 80 | 9 |
| FTN4 to FTN5 conversion | 6 | Dec. 80 | 15 |
| FTN5-effective programming | 9 & | June 81 | 13 |
| | 10 | Aug. 81 | 13 |
| -optimisation techniques | 14 & | Apr. 82 | 13 |
| | 15 | June 82 | 10 |
| Graphics - COM | 7 | Feb. 81 | 13 |
| - hints on memory and time saving | T6 | Dec. 79 | 20 |
| - libraries | T5 | Oct. 79 | 8 |
| Libraries - NAG developments | T5 | Oct. 79 | 7 |
| - public libraries | T5 | Oct. 79 | 6 |
| Jobs - hints on processing | T2 | Apr. 79 | 23 |
| - input queue delays | 4 | Aug. 80 | 12 |
| Magnetic tapes - hints on use | T2 | Apr. 79 | 17 |
| - LOOK9 analysis program | T3 | June 79 | 18 |
| - EEC parameter | T4 | Aug. 79 | 14 |
| - making back-up copies | 1 | Feb. 80 | 9 |
| - Stranger tapes | 5 | Oct. 80 | 10 |
| Permanent files - automatic purging policy | 14 | Apr. 82 | 19 |
| - dumping | T6 | Dec. 79 | 15 |
| - RT=W, BT=I file structure | T1 | Feb. 79 | 18 |

GENERAL

| | | | |
|--|----|----------|----|
| COMFILE | 11 | Sept. 81 | 14 |
| Cyber-Cray link software | 2 | Apr. 80 | 13 |
| Cyber-Cray speed comparison | T3 | June 79 | 19 |
| Cyber-Cray I/O efficiency comparison | 1 | Feb. 80 | 11 |
| Fortran 77 | 5 | Oct. 80 | 6 |
| Mass Storage Systems (MSS) | 5 | Oct. 80 | 8 |
| Member State Technical and Computing Representatives and Meteorological Contact Points | 14 | Apr. 82 | 29 |
| News Sheets still valid | 15 | June 82 | 17 |
| Output files - controlling destination of, in Cray and Cyber jobs | 14 | Apr. 82 | 20 |
| Priority - parameter on the JOB card | 7 | Feb. 81 | 8 |
| - groups in 1982 | 14 | Apr. 82 | 25 |
| Resource allocation for 1982 | 12 | Dec. 81 | 8 |
| Resource allocation - Council rules for | 6 | Dec. 80 | 10 |
| SMHI Computer Links | 9 | June 81 | 6 |
| Technical Advisory Committee - 3rd session | 10 | Aug. 81 | 1 |
| Telecommunications schedule | 12 | Dec. 81 | 10 |
| Upper and lower case text files | 11 | Sept. 81 | 15 |

METEOROLOGY

| | | | |
|---|----|----------|---|
| ALPEX: the alpine experiment of the GARP mountain sub-programme | 14 | Apr. 82 | 2 |
| Alpex data management and the international Alpex data centre | 11 | Sept. 81 | 1 |
| Baltic Storm of October 1980 | 6 | Dec. 80 | 2 |
| ECMWF Analysis and Data Assimilation System | T3 | June 79 | 2 |
| ECMWF Operational Forecasting Model | 5 | Oct. 80 | 2 |
| " " " " | 6 | Dec. 80 | 7 |
| ECMWF Operational Schedule, Data and Dissemination | 12 | Dec. 81 | 1 |
| ECMWF Production Schedule | 6 | Dec. 80 | 5 |
| Facilities to verify and diagnose forecasts provided by the Data & Diagnostics Section | 8 | Apr. 81 | 3 |
| Forecast products of various centres decoded and plotted at ECMWF | 9 | June 81 | 3 |
| Forecasting: development of the new system Meteorology Division | 15 | June 82 | 1 |
| Operational Forecast Suite (EMOS) | T1 | Feb. 79 | 4 |
| - general description | T1 | Feb. 79 | 6 |
| - data acquisition and decoding | T6 | Dec. 79 | 1 |
| - initialisation | T6 | Dec. 79 | 4 |
| - quality control | 1 | Feb. 80 | 3 |
| - bulletin corrections (CORBUL) | 2 | Apr. 80 | 1 |
| - archiving | 3 | June 80 | 4 |
| - post processing | 4 | Aug. 80 | 3 |
| - significant change made | 12 | Dec. 81 | 3 |
| Pseudo "satellite picture" presentation of model results | 1 | Feb. 80 | 2 |
| Research Department activities | 13 | Feb. 82 | 3 |
| Retrieval of data from the Centre's data bases | 5 | Oct. 80 | 3 |
| Spectral model | 7 | Feb. 81 | 4 |
| Weather-routing of ships based on ECMWF forecasts | 10 | Aug. 81 | 3 |

* T indicates the original Technical Newsletter series

USEFUL NAMES AND PHONE NUMBERS WITHIN ECMWF

| | <u>Room*</u> | <u>Ext**</u> |
|---|--------------|--------------|
| Head of Operations Department - Daniel Söderman | OB 010A | 373 |
| ADVISORY OFFICE - Open 9-12, 14-17 daily | CB 037 | 308/309 |
| Other methods of quick contact: | | |
| - telex (No. 847908) | | |
| - COMFILE (see Bulletin B1.5/1) | | |
| Computer Division Head - Geerd Hoffmann | OB 009A | 340/342 |
| COMPUTER OPERATIONS | | |
| Console - Shift Leaders | CB Hall | 334 |
| Reception Counter) - Judy Pearce | CB Hall | 332 |
| Terminal Queries) | | |
| Tape Requests) | | |
| 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 - Claus Hilberg | 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) 876000
international +44 734 876000