European Centre
for Medium Range Weather Forecasts

ECMWF NEWSLETTER

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

Number 1 - February 1980

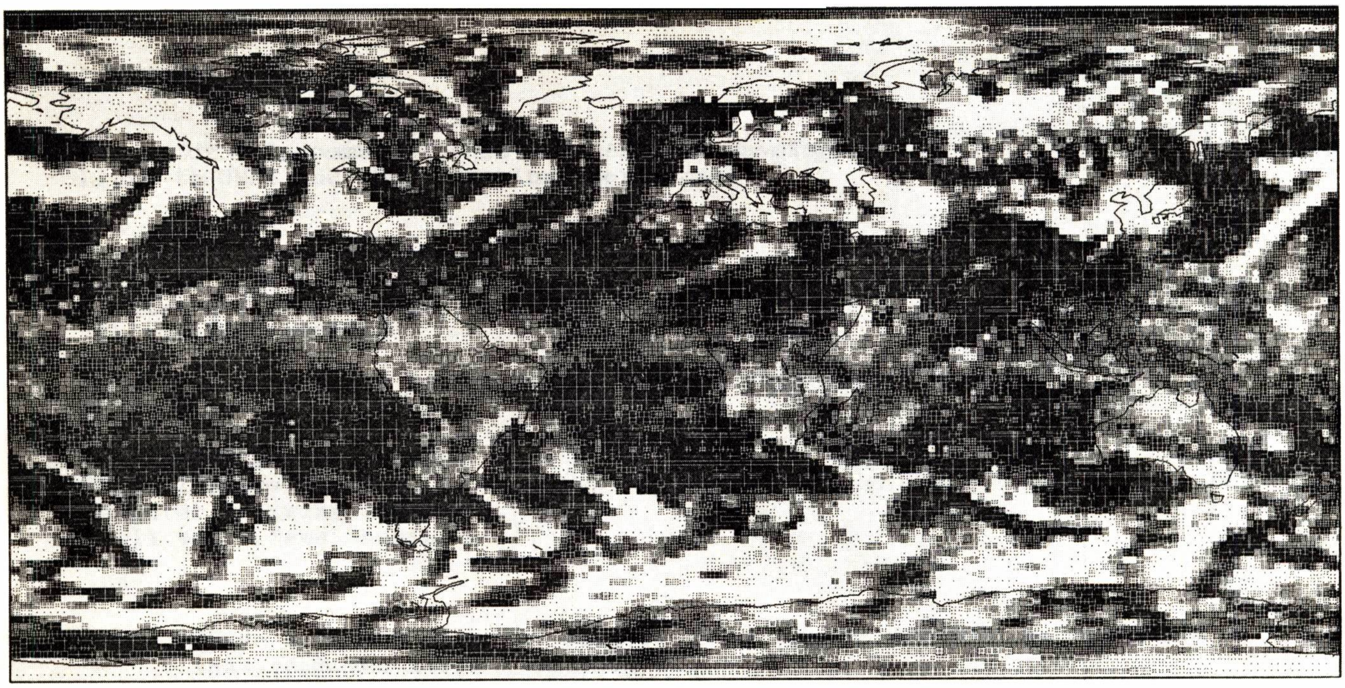

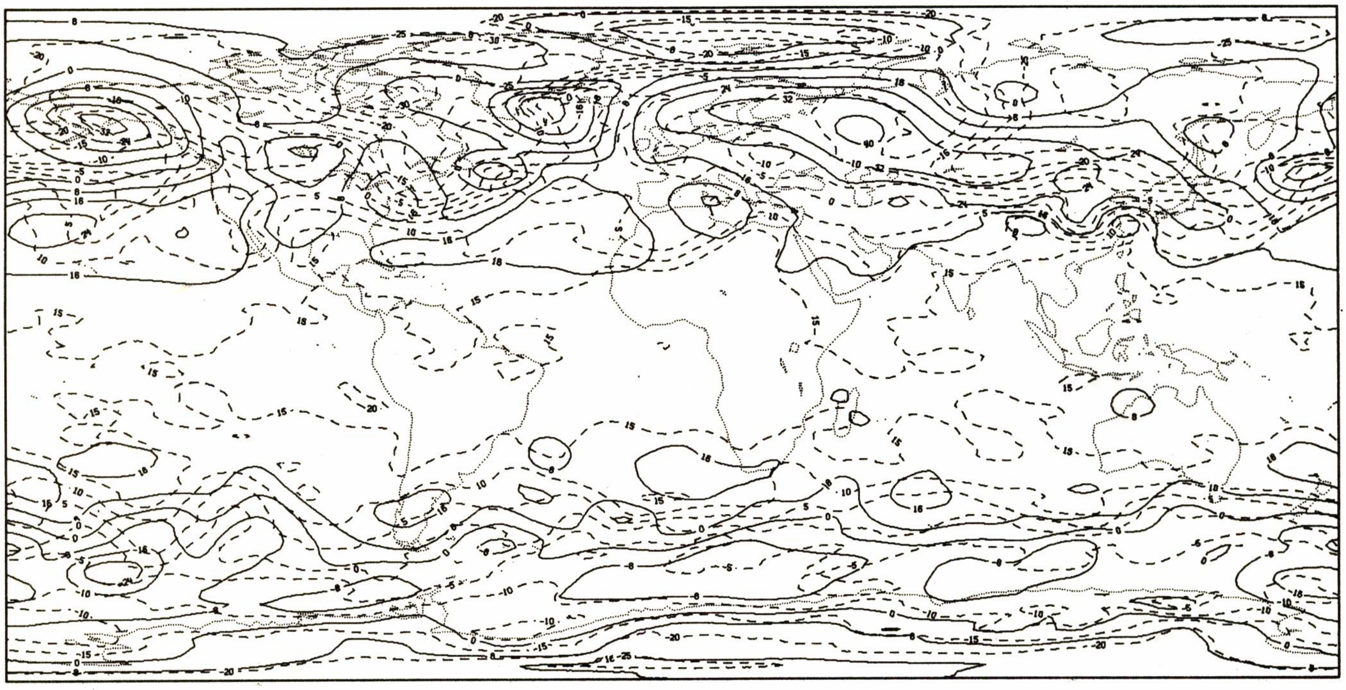

IN THIS ISSUE:

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

COVER: Pseudo "satellite picture" presentation of the model's cloud cover (as used in the radiation computation) after a trial 10-day forecast from 15.2.76 and corresponding 10-day forecast map of 1000 mb heights (8dm intervals) and 850 mb temperatures (5°C intervals).

See article on page 2.

This ECMWF Newsletter replaces the previous Technical Newsletter series. It is edited and produced by User Support.

The next issue will appear in April.

A LOOK AT ECMWF OPERATIONAL AND COMPUTING ACTIVITIES

In a surprisingly short time, ECMWF has built up an efficient and well-balanced computer complex with a CRAY-1 "number cruncher", a CDC CYBER 175 "front-end" and a Regnecentralen RC8000 tele-communications system as main components, all being linked via high-speed channels. Telecommunication links between ECMWF and the meteorological offices of the Member States are currently being set up. The first of the medium-speed links, to the Swedish Meteorological and Hydrological Institute in Norrköping, was fully implemented in December 1979 and operates, according to reports received from Sweden, to the complete satisfaction of the users of both the forecast products and the remote batch services of the Centre. Low-speed links have also been established to Yugoslavia, Turkey, Italy, Spain, France, Greece, and the Netherlands.

In parallel with the design, acquisition and implementation of the Centre's computer system, a complex operational forecasting suite, including a sophisticated global forecasting model, has been designed and set up. Operational forecasting tests started in August 1979 and 10-day forecasts have, since then, been run five days per week. A selection of these forecasts has been disseminated to Member States, where they have generally been received with appreciation. The forecasts have also been evaluated internally by Centre staff in order to identify and eliminate deficiencies and bugs in the operational suite.

It must be stressed at this time that the ECMWF forecasts of today should not be seen as "the final product". At present the data cut-off time is very early (around 17.30Z). Studies are currently being undertaken to determine the optimal cut-off time, with a view to including 18Z and possibly even some 00Z data. In addition, many improvements to the various modules of the operational suite are being implemented. Important target dates are 1 August 1980, when operational forecasting seven days per week is to be initiated, and 1 January 1981, when dissemination of a selection of ECMWF products to non-Member States via the GTS could start, in accordance with the decision of Council at its tenth session. As well as working to improve the operational suite itself, staff at the Centre are also themselves introducing refinements to the analysis and forecast, in an attempt to improve the quality of the forecasts and overcome some of the known shortcomings. However, in view of the fact that many meteorological services are already now using the Centre's products on a routine basis, the assistance of Member States in identifying problems will be very valuable. With this in mind, a number of additional products, including humidity and vertical velocity, will be released for dissemination in the very near future.

The expected reliability of the forecasting system is another important element, particularly for Member States planning how to use ECMWF products. During the operational testing period the whole system has been encouragingly reliable. In fact, at the time of writing, only one forecast (17 January 1980 12Z) has been lost since August 1979. In addition some forecasts have been delayed by a few hours. This experience may, perhaps, give a slightly too optimistic view for the future. It should be assumed that at least a few forecasts per month will be lost or seriously delayed because of hardware or software problems at the Centre or problems with the telecommunications links.

Even if operational forecasting is the most important activity of the Centre, and therefore must have higher priority than any other application, only some 50% of the computer resources are normally required for these activities. The rest is to be shared primarily between the Centre's research projects and the Member States, mainly for research in the area of numerical weather prediction. Some problems can be foreseen here in the autumn of this year when daily operational forecasting will have begun and with many Member States having started to use actively the resources allocated to them. At present both the CRAY-1 and the CYBER 175 are already used almost to their full capacity. As it is likely that the research requirements for computer resources will increase rapidly, the problems of allocating and scheduling the available resources will become more serious, and turn-round times will increase.

Finally, I would like to stress two main aspects of ECMWF - <u>operational</u> and <u>international</u>. As already stressed above, the operational routine will have the highest priority, and all efforts will be made to produce forecasts whenever possible; on some days this may mean that the computer system is not available to remote batch users in Member States or users within the Centre. As regards the international aspect, it should be recognised that the majority of the users of our products and services will be found in the Member States.

- Daniel Söderman

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

PSEUDO "SATELLITE PICTURE" PRESENTATION OF MODEL RESULTS

The cover picture of this issue of the ECMWF Newsletter, shows a 10-day forecast 1000mb height / 850mb temperature for the whole earth and a graphical representation of the associated cloudiness field. This field is the one used in the radiation computation, or more precisely, 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 temperatures predicted by the model at that particular time; the vertical combination is done by assuming maximum overlapping of adjacent cloudy points and random positioning of distinct cloud layers against each other.

The graphical representation is intended to simulate a satellite picture: "Zero" cloudiness should appear black and complete cloud cover as "white". Each model's grid point is associated with a small square area of the picture, a dot density is computed from the cloud cover value and dots are plotted (or not) in a regular pattern according to this density inside the small square.

One can judge from the cover picture and the associated map that, at least for our latitudes the synoptic agreement is fairly good with fronts and occlusions easily recognisable in the picture. A 10-day forecast cannot be realistically compared with a satellite picture verifying at the same time. However, such a comparison was made between a 1-day forecast from the 19th September 1979, and a satellite picture for 20th September. The forecast and satellite picture can be seen in figure 1 (beware of the different projection and scales!)

The agreement between the two pictures is rather encouraging; one can, for instance, clearly see the occluded frontal system and vortex near Scotland. Also, over the Western Mediterranean the cloud system is well represented in the pseudo satellite picture, as are cloud masses over the Caspian Sea, Azores and Labrador Straight.

Trials will soon be undertaken to produce these pictures for different projections and directly from the model's history files to allow quasi-operational production.

- Jean-Francois Geleyn

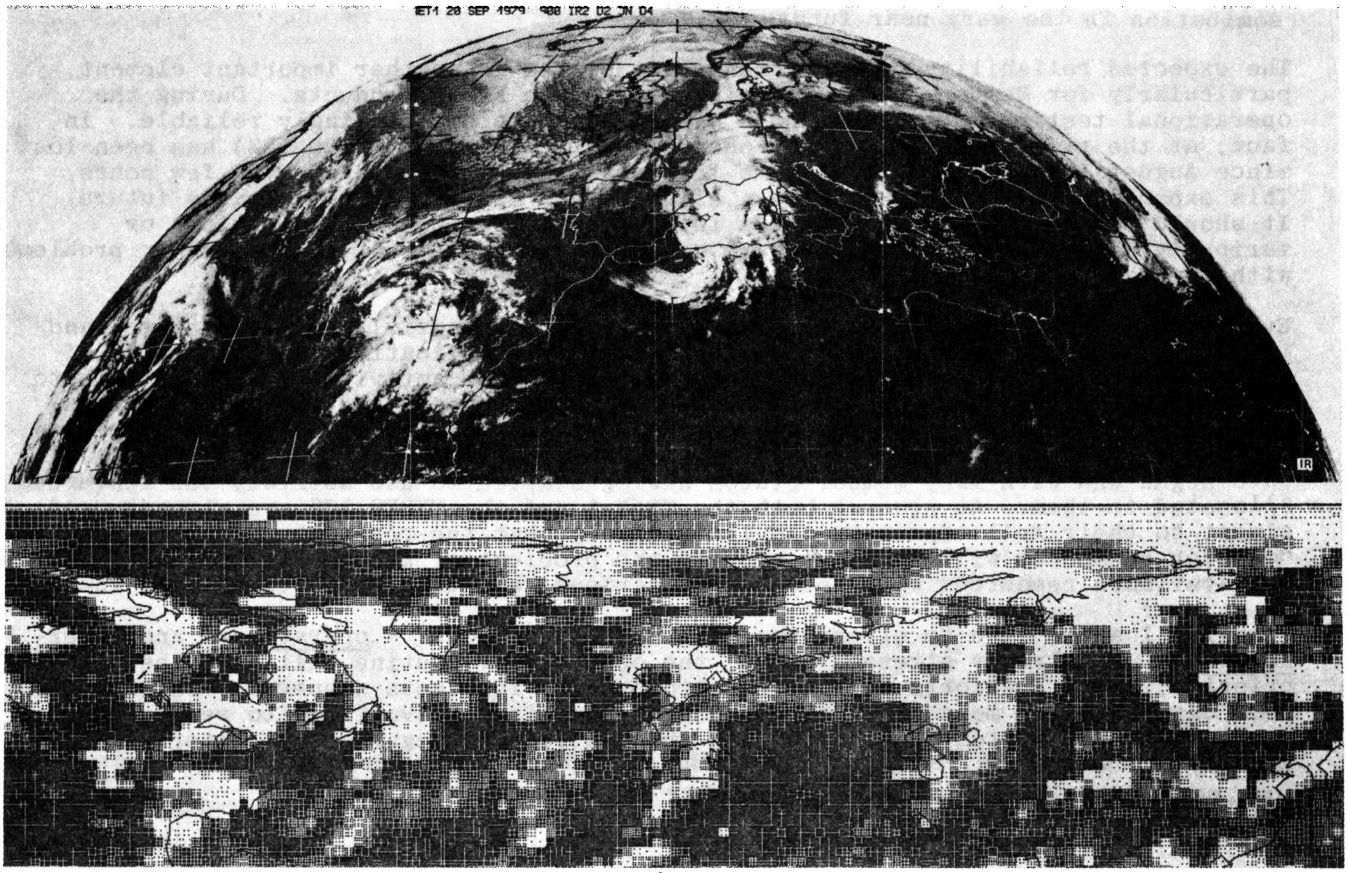

Figure 1

A LOOK AT ECMWF OPERATIONAL AND COMPUTING ACTIVITIES

In a surprisingly short time, ECMWF has built up an efficient and well-balanced computer complex with a CRAY-1 "number cruncher", a CDC CYBER 175 "front-end" and a Regnecentralen RC8000 tele-communications system as main components, all being linked via high-speed channels. Telecommunication links between ECMWF and the meteorological offices of the Member States are currently being set up. The first of the medium-speed links, to the Swedish Meteorological and Hydrological Institute in Norrköping, was fully implemented in December 1979 and operates, according to reports received from Sweden, to the complete satisfaction of the users of both the forecast products and the remote batch services of the Centre. Low-speed links have also been established to Yugoslavia, Turkey, Italy, Spain, France, Greece, and the Netherlands.

In parallel with the design, acquisition and implementation of the Centre's computer system, a complex operational forecasting suite, including a sophisticated global forecasting model, has been designed and set up. Operational forecasting tests started in August 1979 and 10-day forecasts have, since then, been run five days per week. A selection of these forecasts has been disseminated to Member States, where they have generally been received with appreciation. The forecasts have also been evaluated internally by Centre staff in order to identify and eliminate deficiencies and bugs in the operational suite.

It must be stressed at this time that the ECMWF forecasts of today should not be seen as "the final product". At present the data cut-off time is very early (around 17.30Z). Studies are currently being undertaken to determine the optimal cut-off time, with a view to including 18Z and possibly even some 00Z data. In addition, many improvements to the various modules of the operational suite are being implemented. Important target dates are 1 August 1980, when operational forecasting seven days per week is to be initiated, and 1 January 1981, when dissemination of a selection of ECMWF products to non-Member States via the GTS could start, in accordance with the decision of Council at its tenth session. As well as working to improve the operational suite itself, staff at the Centre are also themselves introducing refinements to the analysis and forecast, in an attempt to improve the quality of the forecasts and overcome some of the known shortcomings. However, in view of the fact that many meteorological services are already now using the Centre's products on a routine basis, the assistance of Member States in identifying problems will be very valuable. With this in mind, a number of additional products, including humidity and vertical velocity, will be released for dissemination in the very near future.

The expected reliability of the forecasting system is another important element, particularly for Member States planning how to use ECMWF products. During the operational testing period the whole system has been encouragingly reliable. In fact, at the time of writing, only one forecast (17 January 1980 12Z) has been lost since August 1979. In addition some forecasts have been delayed by a few hours. This experience may, perhaps, give a slightly too optimistic view for the future. It should be assumed that at least a few forecasts per month will be lost or seriously delayed because of hardware or software problems at the Centre or problems with the telecommunications links.

Even if operational forecasting is the most important activity of the Centre, and therefore must have higher priority than any other application, only some 50% of the computer resources are normally required for these activities. The rest is to be shared primarily between the Centre's research projects and the Member States, mainly for research in the area of numerical weather prediction. Some problems can be foreseen here in the autumn of this year when daily operational forecasting will have begun and with many Member States having started to use actively the resources allocated to them. At present both the CRAY-1 and the CYBER 175 are already used almost to their full capacity. As it is likely that the research requirements for computer resources will increase rapidly, the problems of allocating and scheduling the available resources will become more serious, and turn-round times will increase.

Finally, I would like to stress two main aspects of ECMWF - operational and international. As already stressed above, the operational routine will have the highest priority, and all efforts will be made to produce forecasts whenever possible; on some days this may mean that the computer system is not available to remote batch users in Member States or users within the Centre. As regards the international aspect, it should be recognised that the majority of the users of our products and services will be found in the Member States.

- Daniel Söderman

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

PSEUDO "SATELLITE PICTURE" PRESENTATION OF MODEL RESULTS

The cover picture of this issue of the ECMWF Newsletter, shows a 10-day forecast 1000mb height / 850mb temperature for the whole earth and a graphical representation of the associated cloudiness field. This field is the one used in the radiation computation, or more precisely, 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 temperatures predicted by the model at that particular time; the vertical combination is done by assuming maximum overlapping of adjacent cloudy points and random positioning of distinct cloud layers against each other.

The graphical representation is intended to simulate a satellite picture: "Zero" cloudiness should appear black and complete cloud cover as "white". Each model's grid point is associated with a small square area of the picture, a dot density is computed from the cloud cover value and dots are plotted (or not) in a regular pattern according to this density inside the small square.

One can judge from the cover picture and the associated map that, at least for our latitudes the synoptic agreement is fairly good with fronts and occlusions easily recognisable in the picture. A 10-day forecast cannot be realistically compared with a satellite picture verifying at the same time. However, such a comparison was made between a 1-day forecast from the 19th September 1979, and a satellite picture for 20th September. The forecast and satellite picture can be seen in figure 1 (beware of the different projection and scales!)

The agreement between the two pictures is rather encouraging; one can, for instance, clearly see the occluded frontal system and vortex near Scotland. Also, over the Western Mediterranean the cloud system is well represented in the pseudo satellite picture, as are cloud masses over the Caspian Sea, Azores and Labrador Straight.

Trials will soon be undertaken to produce these pictures for different projections and directly from the model's history files to allow quasi-operational production.

- Jean-Francois Geleyn

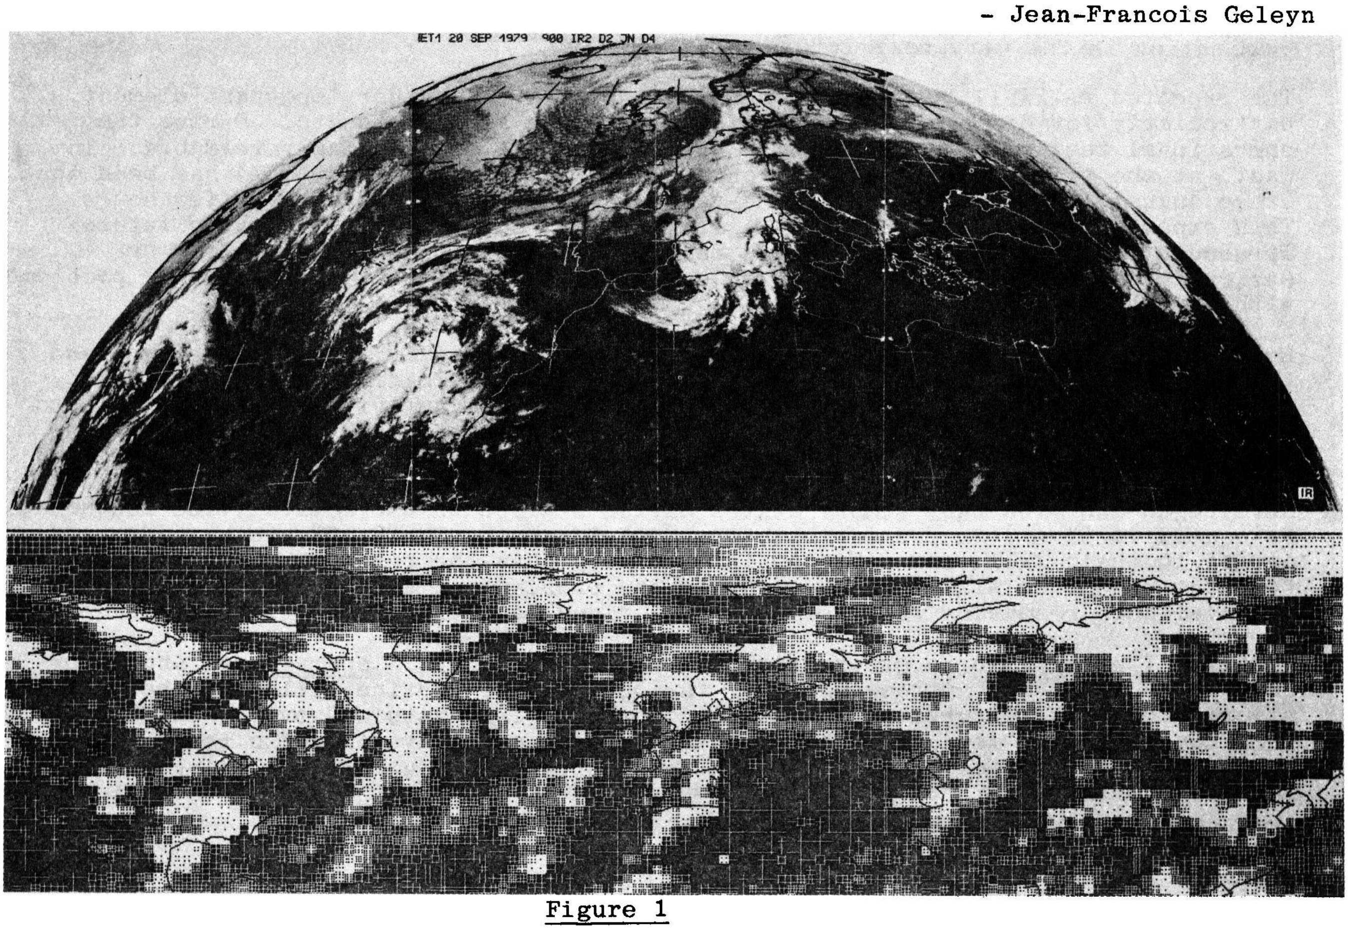

Figure 1

QUALITY CONTROL OF OBSERVATIONAL DATA AT ECMWF

This article, another in the series describing the subsystems making up the complete operational suite, outlines the procedure of checking observational data received in *Global Telecommunications System form (via RTH Bracknell). The first step is decoding;* (the system for decoding at ECMWF was described in the Technical Newsletter for December 1979). Next comes quality control of the decoded data, the subject of this article. Following this, the observation is inserted in the Reports Data Base, from where it can be later extracted for use by other subsystems, e.g. objective analysis, data coverage, etc.

For each parameter value checked, a flag is assigned indicating a degree of confidence:

0 : value correct
1 : value probably correct
2 : value probably wrong
3 : value wrong

As a report often contains some information allowing cross-checking and at least an indication of the value of some other values in the report, the quality control can sometimes suggest its own value (substitution) for a parameter which is wrong or missing. In this case, both the old value and new value are fed to the Data Base, though the new value only is fed to the analysis.

At present, each report is checked on its own. No horizontal checking is performed at this stage, though it is implicit in the analysis, and feedback flags from the analysis will shortly be returned to the Reports Data Base. However, it is planned to introduce a check of a report with the report from the same station for a few hours before, for position, pressure tendency and variability of parameters.

'Report' here means information from a particular location for a particular time. The quality control merges parts of multi-part reports (e.g. TEMP, PILOT) and checks the whole.

The checks vary depending on the type of observation. For all observation types, values are checked against climatological limits so that if a value falls outside the bounds, it is considered as wrong. For upper air observations, the limits depend on pressure level or height.

> e.g. for a pressure between 500 and 600mb, temperature should be within the range -90°C to +13°C, wind speed <103 m/s. At present, the limits do not depend on latitude/longitude position.

In addition, for surface reports, where information is present enabling cross-checking of various parameters, a large number of internal consistency checks are carried out.

> e.g. if temperature = 11°C and present weather = snow, there is an inconsistency and we know one of the two parameters must be wrong. By subjecting all the parameters to similar tests, we can assign flag values depending on the number of tests failed.

TEMP (radio-sonde) messages also contain information allowing cross-checking and which is used to construct an efficient scheme for vertical consistency control. The TEMP checking programs are based on those in use at the Swedish Meteorological and Hydrological Institute.

The first check is the lapse rate of vertical temperature profiles. Given temperature values at two consecutive pressure levels, an estimate of the temperature at the upper level is computed from the temperature at the lower level and the dry adiabatic lapse rate. If the reported value is less than the estimate, then the observation implies a super-adiabatic layer and flags are set, though a slightly superadiabatic layer is permitted near the surface.

Secondly, standard level data are recomputed from significant level data, assuming a linear variation in ($\log_n p$) between significant levels and compared with reported standard level data. If the difference exceeds certain limits, e.g. 1.5°C for temperature below tropopause, flags are set and substitutions provided for missing or erroneous values.

./.

Thirdly, the hydrostatic control provides a check between temperature and geopotential data at standard levels. Again, substitutions are provided, by interpolation from adjacent standard levels.

Finally, wind data from adjacent standard levels undergo shear control which includes a check on speed shear, and a check of a combination of directional shear and sum of speeds.

The routines for checking PILOT data are a subset of those used for radio-sonde. As well as the climatological limit check, PILOT's undergo the wind shear check.

For each report from a WMO land station which is added to the Data Base, statistics are incremented. A record for each WMO block, for each code type and for each hour contains numbers of reports received, numbers of reports with important parameters (pressure, temperature, wind) flagged 0, 1, 2 or 3, numbers of reports with other parameters flagged 0, 1, 2 or 3, etc., the worst flag being taken from each report.

As an example, the following figures are quoted for 3rd November 1979, 00Z, for data from WMO Blocks 1 to 20:

		Important parameters flagged				Other parameters flagged			
	Total received	0	1	2	3	0	1	2	3
SYNOP	815	810	5	0	0	748	62	1	4
TEMP	89	67	9	1	12	65	4	0	20

- Brian Norris

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

* PLANS FOR TELECOMMUNICATIONS AND GRAPHICS DEVELOPMENTS

The Council has approved for 1980 the purchase of a second Regnecentralen 8000, to provide the basis for duplication of the telecommunication facilities. This computer configuration - the exact details of which have still to be worked out - will mainly comprise a CPU, one peripheral processor, a multiplexor and a console. We plan to take other peripheral equipment from the present 8000 system, preferably making it switchable. Initially, this system will be used to support the test periods with new medium-speed connected Member States and to facilitate software maintenance and development. We are also thinking of integrating the second 8000 into our developments towards decentralisation of graphical software which, if successful, would add a most important facet to this machine.

Software will have to be adapted to enable the linking of the two 8000's but also software "instrumentation" will be developed in order to support the testing with Member States even more and enable line utilisation and performance measurements to be carried out. During 1980 we plan to undertake several smaller scale studies such as the possible integration of our local and remote network (at present the two are handled by different network front-end processors), the definition of a "Virtual Terminal Protocol" (to facilitate later the attachment of arbitrary interactive terminals and other facilities) and on new communications media (with a view to public packet switching services). Verification of our transmission protocols would also be a very useful task to undertake.

In Graphics, the provision of interim software will be completed by supporting the contouring package on the Cray. From 1980 onwards, there will be two major developments, viz. the graphical experiments and the adaption/implementation of device independent graphical software. The Council has approved some funds to acquire a colour raster display with significant local intelligence, this will be used for graphical experiments. Apart from the colour aspect,this would also permit trials involving animation. The Regnecentralen 8000 could then act as a first level, with the Cyber as a second level support computer providing everything from local editing to substantial calculations. Computer interfacing problems will have to be tackled and it is very likely that in this area our efforts towards local networking and distributed graphics will merge.

The device independent graphical software will be defined in a layered way. It is again probable that our protoools for data communications can be augmented here by the higher (application) level graphical layers, and that their topological distribution can be supported without major problems.

- Fritz Königshofer

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

LINKS TO SMHI AND BMO BECOMING OPERATIONAL

As previewed in the last Technical Newsletter, the Joint Terminal Project's Provisional Acceptance trials were held at ECMWF in November 1979. Based on Regnecentralen (RC) model 3600 minicomputers, the software had been developed by SIA Ganymede, London, the company which was also responsible for the software developments for our Network Front End Processor.

In early December, installation of this software at the Swedish Meteorological and Hydrological Institute (SMHI) was achieved without problems. Since then, forecast data have been regularly disseminated via the 4800 bits per second link from ECMWF to SMHI and the Institute has also started remote job access to the Centre's computer resources. This now constitutes our first operational medium-speed connection to a Member State and there have been very few problems with it so far. The quality of the line proved to be very good and the software at both ends works reliably and efficiently.

Another partner in the Joint Project, the Danish Meteorological Institute, has to adapt this software to their RC8000 hardware, which requires considerable effort. Therefore, it will take a few months until they become a full medium-speed line user. In the meantime, they plan to set up an interim configuration which will permit them to receive our products on a restricted daily schedule.

./.

The Joint Project also includes special software for the Deutscher Wetterdienst (DWD) in Offenbach to permit the passing of ECMWF's products into their Telefunken System via an RC3600 in an on-line fashion. This combines with the stand-alone function of the RC3600 as an RJE terminal. Unfortunately, the special software for DWD was most affected by some delays incurred in the final stages of the Joint Project. In the meantime, however, coding for this part has also been completed and on-site testing in Offenbach was scheduled to start on the 22nd January 1980, with the Provisional Acceptance trials of the special software scheduled in early February.

Developments on our link to the British Meteorological Office (BMO) have also entered a crucial phase. From the 10th of January, we have started daytime transmission of the GTS observational data as recorded by incremental tape dumps on their COSMOS system. The tapes are then read on their Ferranti 700E computer and the data transmitted over the link to ECMWF. On the 22nd of January, the switch to 24 hours per day, 7 days a week data acquisition over this link is scheduled, with a four week trial period to follow. The daytime transmission trials revealed a number of operational problems which required attention but the stability and performance of the telecommunications software on both ends proved to be acceptable. By its nature, data acquisition is a more difficult task than the other functions fulfilled by our network. The very close co-operation between ECMWF and BMO in this area has therefore been essential in achieving the present advanced state.

- Fritz Königshofer

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

THE CYBER 175

Previous Newsletter articles have described the Cray-1 architecture in some detail. This article covers the Cyber processor unit and is intended for those not familiar with the CDC style of architecture.

The ECMWF front-endcomputer is a Control Data Cyber 175-300. This is one of a large family of machines with similar instruction sets, which includes the Control Data 6000 series, the 7600 and the Cyber 70 and Cyber 170 series. The various machines differ in speed and memory size, technology and, of course, cost.

A notable feature of the design is the peripheral processor (PP). In addition to a central processing unit (CPU), all the machines have several PPs, which are independent computers with separate memories. They can handle many system tasks including input/output.

The CPU of the 175-300 is one of the most sophisticated in the line. Many advanced concepts that have reached full fruition on the Cray-1 appear on the Cyber. In particular, on this model, computations are done in nine separate functional units which can operate simultaneously. Nearly all the units operate in a pipeline fashion, that is, a new calculation can be started while a previous one is in progress. The operands for a calculation come from 8 'X' registers. The X registers also receive the results. These registers are 60 bits long, giving the same floating point precision as the Cray-1 (about 14 decimal digits), with a somewhat smaller range of allowed values. Associated with the X registers are 8, 18 bit, 'A' registers, which hold memory addresses. Altering an A register value results in a memory load or store from the corresponding X register (A1 → A5 fetch values to X1 → X5, A6, A7 store from X6, X7, A0 and X0 are exceptions). For indexing, counting, etc., there are 8 'B' registers, also 18 bits long. The 175-300 CPU has a 25×10^{-9} sec. clock. Simple operations such as integer adds, shifts, etc. take two clock periods. Floating point multiplies take 5 clocks, additions 4. Division is relatively slow, requiring 20 clocks, and the divide unit does not pipeline.

A most important feature is called the instruction stack. This is a set of high speed registers which acts like a moving window following the code in execution. For loops which fit in-stack (up to 10 words long) no memory fetches are necessary to obtain instructions during the loop execution, thus producing a significant increase in speed. The FTN compiler informs the user that a loop does or does not fit in stack. The CPU organisation is summarised in the following diagram.

./.

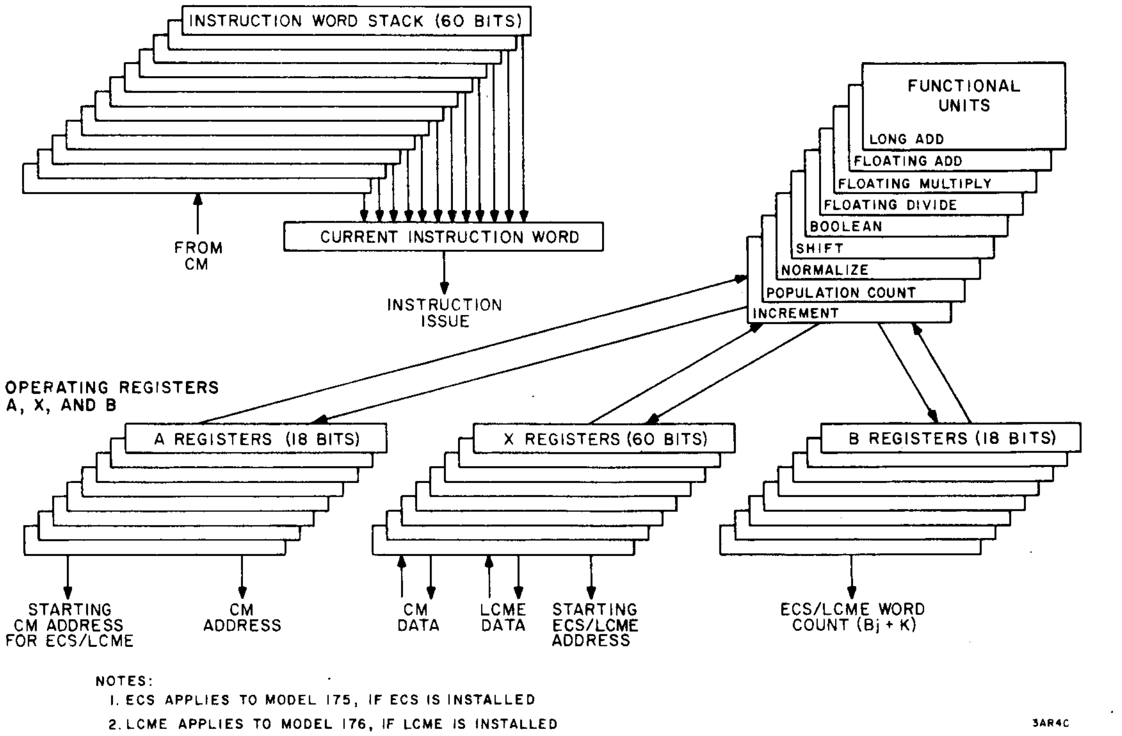

Models 175 and 176 CPU Information Flow

The central processor memory is built from metal-oxide semiconductor (MOS) chips. The present machine has 192 K words which is to be expanded to 256 K words (the maximum) this year. The memory is arranged into 16 independent banks, which can operate simultaneously. The memory cycle time is 300 x 10^{-9} secs and the maximum transfer rate is one word every 50 x 10^{-9} secs, or every 2 clocks. In memory, a 60 bit word is accompanied by 8 check bits which enable the correction of single bit errors, and the detection of double bit errors, (SECDED). (There is an optional second level of bulk memory for a 175 called ECS, which is not present on the ECMWF machine).

The CPU and memory together produce quite a powerful machine, which should be capable of peak rates of a few million floating point operations a second.

The ECMWF machine has 20 PP's arranged in two systems of 10. Logically each is an independent computer with 4K of memory (12 bit words, parity checked). In fact, the relative speeds of various devices are such that each PP in a system time-shares part of some common hardware. The PP's can read or write their memory or any word in the central memory. They can communicate with each other or peripheral devices through the I/O channels. A PP has the power to cause the CPU to change the program being executed and to specify the new program to be commenced. These powers of the PP's allow them to execute many of the functions of an operating system, such as multi-programming control, scheduling, and input-output handling.

A sizeable body of software is available for the Cyber series, including operating systems, language processors, and application packages. A full range of I/O devices may be connected to the mainframe. These factors, together with the good performance of the CPU on scientifically oriented tasks are some of the reasons it was chosen as the front-end of the ECMWF system.

- Dale Robertson

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

THE ECMWF PROCEDURE LIBRARY

In common with most other computer sites, a large number of ECMWF computer users have partially or completely abandoned use of stored punched cards in favour of files kept on mass storage devices. This change in habit requires adequate interactive facilities for file editing, batch job submission and retrieval.

Under the CDC NOS/BE operating system, these facilities are provided by INTERCOM. Although most of the basic operations are available, quite frequently the commands are cumbersome, lengthy and, therefore, subject to typing errors. An example is the command sequence necessary to examine a job output on the terminal-

```
BATCH,TRXAB19,LOCAL
PAGE,TRXAB19.
```

where TRXAB19 is the jobname and, hence, the name of the output file (routed to the remote output queue). Note that although the user defines up to 5 characters of the jobname (in this example TRXAB), the last 2 are always supplied by NOS/BE and are essentially random. Therefore, the user must first discover the complete jobname (usually by the FILES command) and then type the 7 characters as a parameter to 2 commands.

Fortunately, NOS/BE supports a command language (CCL) which makes it simple to execute a sequence of control statements by invoking a 'procedure'. The example could then become:

```
TYPE,TRXAB19
```

where the procedure called 'TYPE' executes the BATCH and PAGE commands with the appropriate insertion of the given parameter.

It has become quite common for individual users to develop their own collection of procedures for this and more complicated tasks. Obviously, it would be more sensible to have a centrally supported library of procedures for the common operations. Wasteful development effort is then minimised and new facilities can be immediately made available to all.

E.g. TYP

behaves similarly to TYPE but discovers the file-name for itself (based on your logged-in identifier).

Such a procedure library was established last year and has been in limited use; it was formally released a few weeks ago. The aims of the library are:

a) to minimise typing
b) to provide additional facilities.

These aims are satisfied under three general headings.

1. Permanent File Procedures

ATT, CAT, PUR - extensions to ATTACH, CATALOG, PURGE allowing multiple operations in one statement. PUR includes deleting Cray datasets.
SELPUR - selective PURGE of files under one ID.
REPLACE - CATALOG a new version of a file, and PURGE a previous one.
MAKFILE,GETFILE,ADDFILE,REPFILE,DELFILE,LOOK - a set of procedures which allow files to become members of one larger file or "library", thus saving much disc space wasted when small files are catalogued individually.
AUD - selective permanent file AUDIT.

2. Batch job submission and retrieval

SUB - submit a job to either machine. Optional generation of JOB and ACCOUNT statements.
TYP - retrieve job output and initiate PAGE.
DAYF - display the dayfile of a remote output job
PRINT - send a local file to be printed.

3. UPDATE procedures

MAKEPL,GETDECK,ADDDECK,REPDECK,DELDECK,DECKS - a set of procedures to handle UPDATE decks.
GETPL,UPD,COMPDK - various UPDATE operations.

./.

The library has the important ability to 'remember' the user's ID and project identifier as supplied on the Intercom LOGIN command. Hence, it is NEVER NECESSARY to repeat these parameters to any of the procedures unless you wish to override the initial values. Another useful option is the possibility of merging private procedures with the library, thus 'getting the best of both worlds'.

If you have not already used the procedures, why not give them a try? You might be agreeably surprised! If you are not, let me know please! To start up, type

```
                ATTACH,PROCIN.
                PROCIN.
and use         HELP.
    or          HELP,PROCNAME.   to find out about particular procedures.
```

Complete documentation in the form of a Bulletin is being prepared.

I shall be happy to discuss the library with anyone who is interested.

- David Dent

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

* MAGNETIC TAPE BACK-UP COPIES

Magnetic tapes can be duplicated on the Cyber by a method which only needs minor changes to handle each of the following tape types:

1) standard unlabelled
2) standard labelled
3) standard labelled multiple set
4) stranger (short and long)

Here, "standard" means "tapes generated on a Cyber computer for Cyber use, including any back-end processors linked to the Cyber."

By requesting that tapes be read as if they were stranger tapes, the system will treat each block of information as a record and the 'End of Information' as the one and only 'End of File'. Coded tapes and binary tapes can be copied alike-

Example 1 - Standard Unlabelled Tapes

```
:
REQUEST,TAPE1,GE,NORING,S,VSN=REEL1A.
REQUEST,TAPE2,GE,RING,S,VSN=REEL1B.
COPYBF,TAPE1,TAPE2.
UNLOAD,TAPE1,TAPE2.
REQUEST,TAPE1,GE,NORING,S,VSN=REEL2A.
REQUEST,TAPE2,GE,RING,S,VSN=REEL2B.
COPYBF,TAPE1,TAPE2.
:
```

and so on for as many tape copies as required.

If the contents of REEL1A and REEL2A are to be copied on to just one reel (REEL1B), then COPYBR is used:

```
REQUEST,TAPE1,GE,NORING,S,VSN=REEL1A.
REQUEST,TAPE2,GE,RING,S,VSN=REEL1B.
COPYBR,TAPE1,TAPE2,NB.      (NB = No.of blocks before End of Information - see
UNLOAD,TAPE1.                                     note 1 overleaf)
REQUEST,TAPE1,GE,NORING,S,VSN=REEL2A.
COPYBR,TAPE1,TAPE2,NB.
UNLOAD,TAPE1,TAPE2.
:
```

./.

Note 1:To determine the number of blocks before EOI, ask the operators to run LOOK9 against the original tape.

Note 2:Back-up tape contents can be checked within the same job but they must be unloaded and requested again, without the 'S' parameter. Use the 'IU' (Inhibit Unload) parameter to prevent the tape from being physically unloaded, so that operator intervention is unnecessary.

Example 2 - Standard Labelled Tapes

The same as example 1, but the LABEL control card should be used, instead of REQUEST:

```
LABEL,TAPE1,R,NORING,L=LABELNAME,
D=GE,F=S,VSN=REEL1A.

LABEL,TAPE2,W,RING,L=LABELNAME,D=GE,
F=S,VSN=REEL1B.

COPYBF,TAPE1,TAPE2.
:
```

and so on. Similarly, COPYBR must be used when merging more than one reel on to a single one.

Example 3 - Standard Labelled Multifile Sets

First, find out how many files are in the set, by running LISTMF against your tape. Since each file is preceded and followed by a tape label, the number of files listed by LISTMF should be multiplied by three, to allow for a full copy in stranger tape mode. Note that all labels are going to be treated as files, thus avoiding the need for as many label cards as there are files. Only the first label needs to be rewritten, since it must contain the correct tape VSN.

The following control cards duplicate a multifile tape containing 49 files.

```
:
REQUEST,MFILE,NS,S,GE,NORING,VSN=REEL1A.
REQUEST,MFCOPY,N,RING,S,GE,IU,VSN=REEL1B.
COMMENT.
COMMENT. DECLARE FIRST LABEL NAME AS FROM
COMMENT. FIRST LABEL IN MFILE.
COMMENT.
LABEL,MFCOPY,L=ANAME,W.
COMMENT.
COMMENT. NOW SKIP FIRST LABEL IN MFILE.
COMMENT.
COPYBF,MFILE,WASTE,1.
COPYBF,MFILE,MFCOPY,146.                    (49*3)-1=146
RETURN,MFILE.
COMMENT.
COMMENT. NOW CHECK TAPE COPY.
COMMENT.
UNLOAD,MFCOPY.
REQUEST,MFCOPY,MF,GE,RING,E,VSN=REELIB.
LISTMF,M=MFCOPY.
```

Example 4 - Stranger Tapes

The examples given apply to stranger tapes as well, provided the data block length never exceeds 512 Cyber words, but one file only will be copied.

For blocks greater than 512 words the examples can still be used if the paremeter 'S' in all REQUEST/LABEL cards is replaced with the parameter 'L' (for Long).

- Luigi Bertuzzi

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

COMPARING I/O EFFICIENCY ON CYBER AND CRAY

Recently, I examined the performance of a small job often used to locate the occurrence of specified variables in libraries. As the job took 23.27 CP seconds on the Cyber, it was thought that a considerable saving might be achieved by running it on the Cray-1. Not so - it took 19.36 CP seconds!

The reasons for the poor performance on the Cray would appear to be the following:

a) the job reads 38681 records with a FORMAT of 90A1. This is not the type of work one would normally expect a Cray-1 to perform, and it is probable that the system routines involved in the formatting and conversion of the input data are not very efficient;

b) the Cray-1, unlike the Cyber, does not have peripheral processors available for this type of work. Consequently, there is a CPU charge on Cray-1 which would be a PPU charge on the corresponding Cyber run.

To verify the above, runs were made on both machines by-passing the computational part of the program. The results were 7.12 CP secs on the Cyber and 12.42 CP secs on the Cray-1, giving times for the computational part of the program of 16.15 CP secs on the Cyber compared with 6.94 CP secs on the Cray-1.

The Cray-1 is essentially a back-end machine capable of producing answers to numerical computations at high speed. It is not, and should not be expected to be, a super high speed general purpose machine. In particular, some types of I/O are comparatively inefficient when compared to its performance in other fields. Care must always be taken when transferring tasks from Cyber to Cray to ensure that real savings are achieved.

- Rex Gibson

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

THE COMPILER OUTLAWS

The Question and Answer section of ECMWF Technical Newsletter No. 6 pointed out a problem which occurs when one uses FORTRAN statements containing equivalences, e.g. array elements referenced as overindexed elements of a preceding array. This practice has been reported as leading to incorrect results when the code is compiled by FTN at OPT=2.

Unfortunately, the test case used to reproduce the problem was incorrect and when the programmer bug was removed, good compiler code was produced. A better test case, producing incorrect code at OPT=2, is now given:

```
      :
      COMMON//A(1),B(1),C(1),D(1),E(1),BUF(700)
      :
      M1=100
      :
      DO 22 JL=1,32
      :
      C(3+M1+JL)=SQRT(---)----
      B(4+M1+JL)=AMAX1(C(3+M1+JL),----)----
      BUF(M5+JL)=BUF(M5+JL)+(BUF(M6+JL)-B(4+M1+JL))
22    CONTINUE
      DO 23 JL=1,32
      PRINT*,JL,BUF(M1+JL)
23    CONTINUE
      :
```

./.

The results produced by this code on the Cyber are wrong only when OPT=2 is used. Correct results can be obtained at OPT=1 and on the Cray, with or without vectorisation.

Even if the Cyber Fortran optimiser can cope with such equivalencing to a certain extent, one cannot consider its failure to handle any number of equivalences of this kind as a bug. In fact, the practice of coding equivalences in this way should be discontinued on the grounds that it is a form of illegal Fortran use, and one can, therefore, never guarantee a future compiler will support it. Indeed, the FTN Manual, in the section on EQUIVALENCE and COMMON (p. 3-12) gives a simple example of a program in which such a trick causes different answers to be produced under different optimisations.

- Luigi Bertuzzi
Dale Robertson

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

1979 CRAY-1 USAGE BY MEMBER STATES SINCE 1 JULY 1979

Member State	Allocation (Cray-1 CP hours)	Usage (Cray-1 CP hours)
Belgium	-	-
Denmark	12	0
Germany	100	4
Spain	-	-
France	100	23
Greece	-	-
Ireland	3	0
Italy	70	0
Yugoslavia	-	-
Netherlands	12	5
Austria	10	0
Portugal	-	-
Switzerland	-	-
Finland	5	0
Sweden	6	7
Turkey	-	-
U.K.	175	76
TOTAL	493	115

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

STATISTICS FOR 1979

The tables below show the weekly average for the number of jobs and CP time used for both systems. They are presented as averages over 4-week periods, to smooth out random week by week variations. Cray statistics do not include the figures for background diagnostic jobs. Cyber statistics include an approximate average of 5000 diagnostic jobs, and just over 2 CPU hours, per period.

Average number of jobs per week within each period

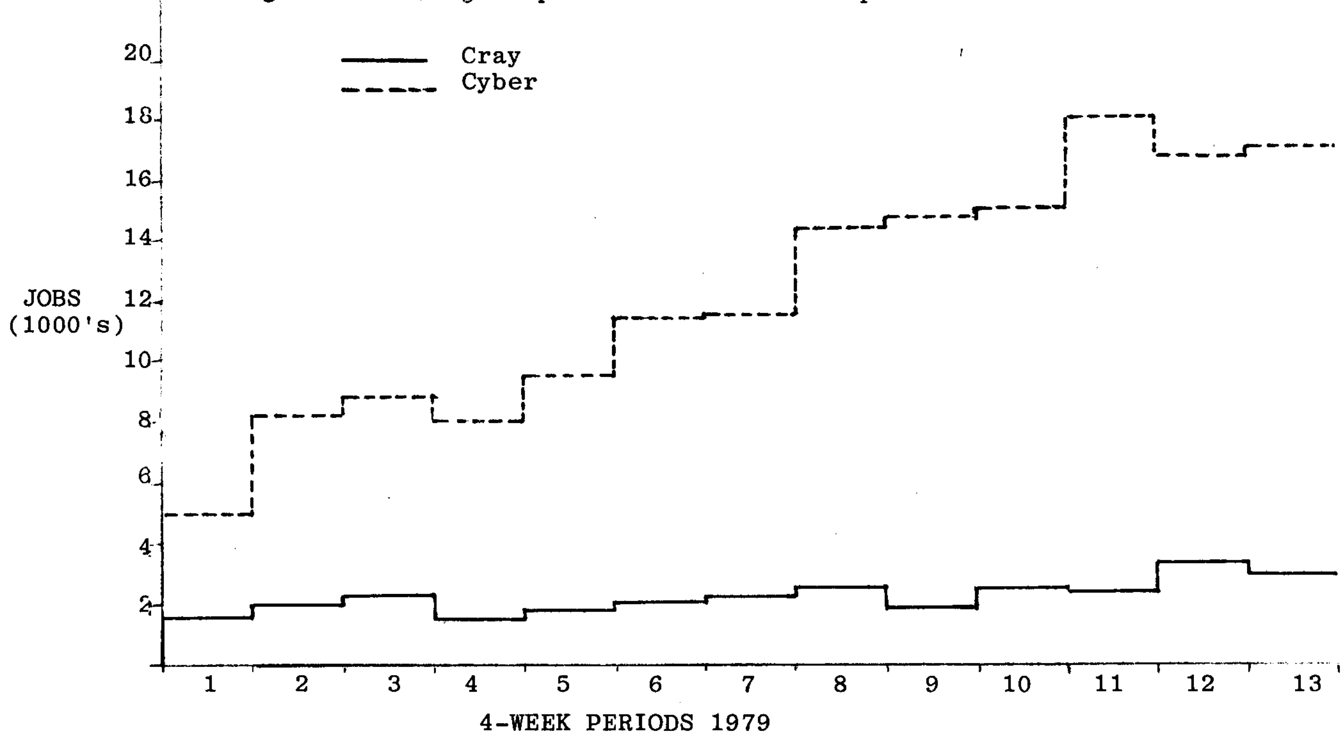

Average CPU hours used per week within each period

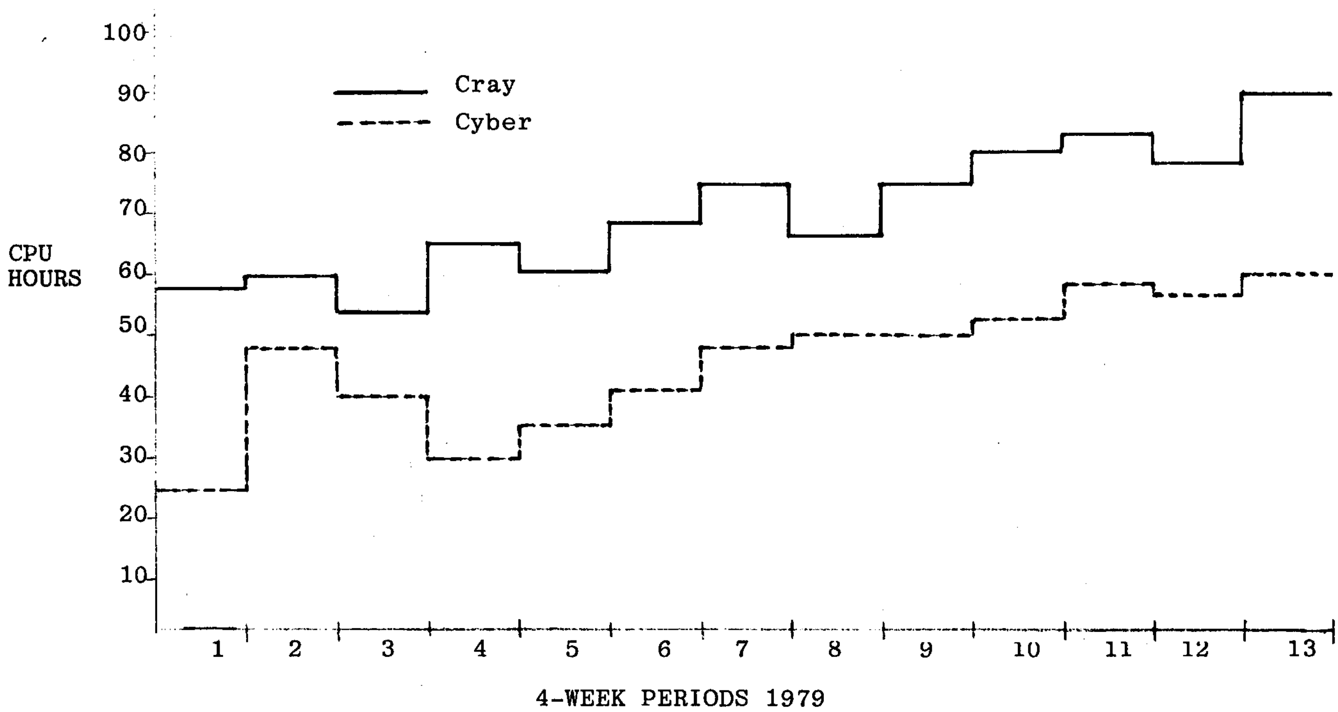

THE CRAY BLOCKING PROBLEM

Currently, there is a problem when files of certain lengths are ACQUIREd by the Cray. When the station generates the Cray block control words (i.e. the data format is NOT TRansparent) there is a chance that the end of data (EOD) word will become the last word of a 512 word sector. Such a dataset when read on the Cray, will cause a failure with the well-known message, "BLOCK NUMBER ERROR" even though the blocking is correct according to the Cray definitions.

E.g:
```
ACQUIRE,DN=A,PDN=TEST,ID=TRX,DF=BB.
            :
            2048 64 bit words DF=BB
            :
CONVERT,I=A,O=B,FROM=CYBER,TYPE=BINARY.

            xxxxx BLOCK NUMBER ERROR
```

The problem is caused by a 'look ahead' mechanism, in this case to a non-existent block control word. The problem can be avoided by terminating the read operation before EOD is detected, e.g. at EOF. In the above example:

```
CONVERT,I=A,O=B,FROM=CYBER,TYPE=BINARY,NF=1.
```

Note that this is not a station bug, but rather a discrepancy between the Cray operating system and the documentation describing it.

- David Dent

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

ECLIB LIBRARY CHANGES

The current version of ECLIB on the CRAY-1 contains a very early and slow version of the FFT package, with principal entry points FFT8 and FFT3. These routines have been superceded by a much faster package which will shortly be made available in ECLIB. At the same time, the old routines will be removed. Please contact me if this change will cause you any problems. Note that the new package is the multiple FFT version. A fast single FFT version will be added later.

For compatability reasons, the Cyber ECLIB now contains entry points NUMARG and NOARG. These are identical to NUMARG on Cray-1 and NOARG in CERNLIB and enable a subroutine to discover the number of parameters it has received.

- David Dent

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

COMPUTER 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.

In summary, the courses being offered are:

<u>Course B</u> <u>BASIC USAGE</u> <u>17 - 20 March</u>

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 (21 March) is devoted to explaining how to use ECMWF's meteorological database and archive system.

./.

Course C CRAY USER COURSE 24 - 28 March

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 21 - 25 April

Again an in-depth course, this time for those who will make detailed use of the Cyber 175 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.

Courses B, C and D will be repeated in September / October. It is also planned to offer a similar set of courses in 1981.

- Andrew Lea

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

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

No.	Still valid article
11	FTN Rounding Option
15	Private Packs on the Cyber (MOUNT/DISMOUNT)
16	Checkpointing and program termination
17	Private packs and interactive jobs.
19	CRAY UPDATE (temporary datasets used).
31	Fortran Callable Tape REQUEST.
37	IN trays for Cray and Cyber Jobs.
42	Cyber scheduler (see News Sheet 59 also)
43	Cray AUDIT.
	Transfer of coded files.
47	Libraries on the Cray-1.
50	8 Disc Cray System.
	Terminal Procedure.
51	Cyber Disc Reconfiguration.
53	Cyber Job Card Priority Usage.
	Writing 6250 bpi Tapes (EEC Parameter).
	Punching Conventions (Coding Forms).
54	Things not to do to the Station.
55	New Cyber Peripherals.
56	DISP
59	New Cyber System (Scheduler Changes)
63	Daily Schedule for Operational Suite.
64	New Version of Graphics Software.
65	Data Security on Cyber and Cray.
66	New Cray Audit.
	Cyber Accounting
67	New COS 1.05 System.
	Attention Cyber BUFFER IN Users.
68	Protected Files on the Cray.
70	Cyber/Cray Station.

The News Sheet which can now be thrown away since this list was last published is number 69 only.

- Andrew Lea

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

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in this issue 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.

GENERAL

METEOROLOGY

* T indicates the 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	- Rob Brinkhuysen	OB 009A	340/342
COMPUTER OPERATIONS			
Console	- Shift Leaders	CB Hall	334
Reception Counter) Terminal Queries)	- Judy Herring	CB Hall	332
Operations Section Head	- Eric Walton	OB 002	349/351
Deputy Ops. Section Head	- Graham Holt	CB 023	307
DOCUMENTATION	- Pam Prior	OB 016	355
Libraries (ECMWF, NAG, CERN, etc.)	- John Greenaway	OB 017	354
METEOROLOGICAL DIVISION			
Division Head	- Roger Newson	OB 008	343
Applications Section Head	- Joel Martellet	OB 011	360
Operations Section Head	- Austin Woods	OB 107	406
Meteorological Analysts	- Ove Åkesson	OB 106	380
	- Veli Akyildiz	OB 104A	379
	- Horst Böttger	OB 104A	378
	- 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
Tape Requests	- Pauline Litchfield - George Stone	CB Hall	335/334
TELECOMMUNICATIONS			
Fault Reporting	- Pierre-Pascal Regnault	CB 028	397/375
Section Head	- Fritz Königshofer	CB 130	310
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