European Centre
for Medium Range Weather Forecasts

TECHNICAL NEWSLETTER

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

Number 1 - February 1979

CONFIGURATION OF THE ECMWF COMPUTER SYSTEM

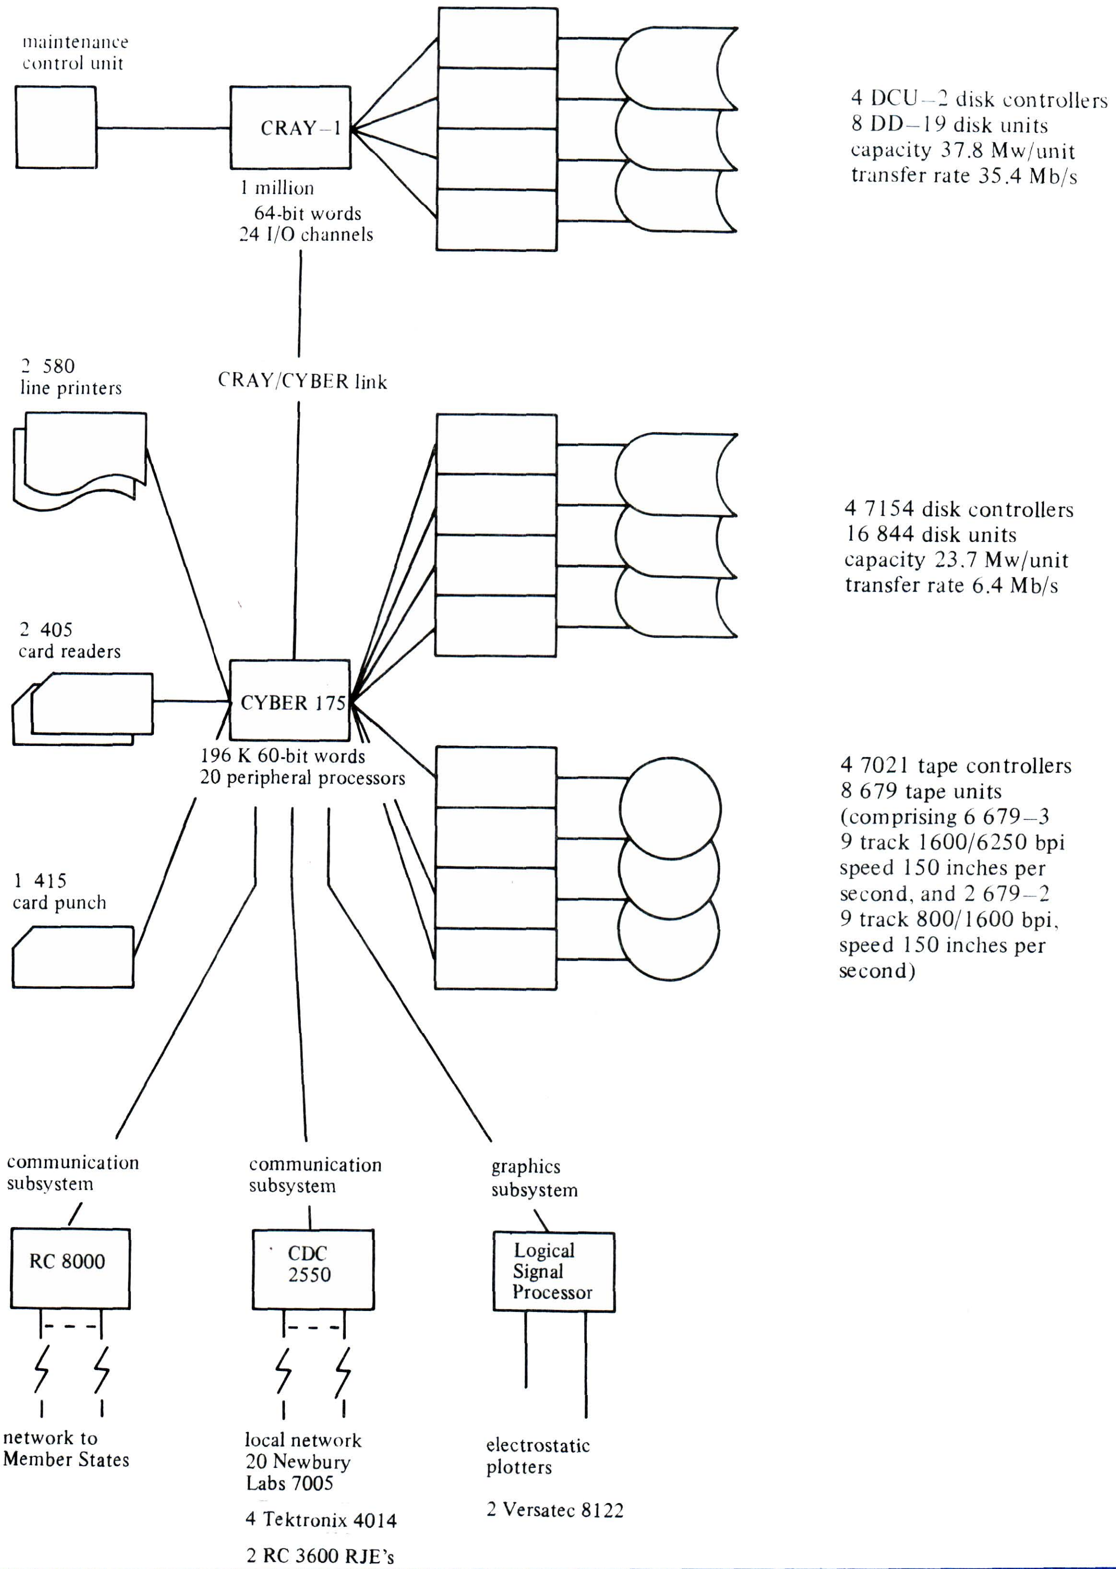

Number 1 - February 1979

IN THIS ISSUE:

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

COVER: A configuration diagram of the ECMWF Computer System

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

The next issue will appear in April.

ECMWF TECHNICAL NEWSLETTER

Since August 1977 the Centre has published a Computer Newsletter aimed at keeping the Centre's computer users aware of the current computer service situation. The readership was primarily the Centre staff.

During 1979 some significant changes will occur directly involving the Member States:

- the Centre will begin to produce an operational daily forecast.
- Member States personnel will begin to use the Centre's computing facilities, either by visiting the Centre or, later in the year, by RJE links.

To reflect these changes the Newsletter has been expanded to encompass information on both the meteorological work of and the computer service offered by the Centre. The Newsletter will be distributed automatically to all Member States, as well as ECMWF scientific staff and senior administration.

Because of its wider role the Newsletter has been renamed the "ECMWF Technical Newsletter". It will be published every 2 months, by the Operations Department of the Centre.

The readership of this Newsletter is expected to be those meteorologists and programmers actively using some part of the Centre's technical services. The aim is to keep them informed in such areas as:

Meteorology
- implementation of the operational forecast
- changes in the operational procedures
- description of forecast verification
- new presentations of forecast results

Computing
- past performance and usage
- problems encountered or outstanding
- changes just made or coming shortly
- future proposals and plans
- hints on using the Centre's facilities to maximum benefit
- special facilities offered periodically (e.g. training courses)

In addition the Newsletter will from time to time provide background articles on some of the major uses to which the Centre's services are put.

The first issue contains a number of articles describing the meteorological operational activities of the Centre and our computing system.

This Newsletter is just one part of the Centre's publications aimed at providing the overall community with a full understanding of the Centre's activities. The Centre hopes it fulfils the growing need for dissemination of direct technical information. Feedback to the Editors from readers on any aspect of this Newsletter is always very welcome.

- Aksel Wiin-Nielsen

Role of the Centre

To help new readers who may not be familiar with the role of ECMWF we print below an edited version of an article originally written in 1976 by Mr. E. Knighting, one of the members of the original Planning Staff.

\- Andrew Lea

Reliable weather forecasts for a week or so ahead would be of great economic value; almost every industry, including agriculture, transport, building, shipping, fuel and power, would benefit in the planning of operations and, of course, the general public would be better able to plan outdoor activities. And if, as has been estimated, the benefit to Europe would be more than £100 million a year at 1970 values, it is well worth setting up an organisation to conduct research into problems associated with medium range forecasting, and eventually to provide operational forecasts.

The resources of skilled manpower and equipment required for such a task, however, exceed those normally available at national level. This makes an international co-operative effort necessary, and in October 1973 sixteen European nations signed a Convention establishing the European Centre for Medium Range Weather Forecasts.*

The Centre's headquarters are at Shinfield Park, Reading, England; The Centre will eventually have a staff of about 140, including about 70 graduates, mostly in the research and operations departments. When the Centre is fully operational it is expected to have a yearly budget of about £6 million; the cost will be met by the member nations, contributions being scaled according to the gross national product of each nation.

The central premise of the development of these operational forecasts is that the atmosphere may be regarded as a compressible fluid, its behaviour being described by the Navier-Stokes equation and the thermodynamic equations concerned with sources, sinks and the transfer of energy. Considerable progress has been made in the last two decades in using these equations for short-range weather forecasts for a day or two ahead, and it is now common for national services, especially those in extra-tropical latitudes, to base their short-range forecasts on the numerical integration of the equations. The problem is the prediction, with considerable accuracy, of the motion and development of weather-bearing systems which have a lifetime of a week or so. It requires a fairly precise knowledge of the state of the atmosphere at the instant from which the forecast proceeds, while some slow acting physical processes may be neglected.

Forecasting for up to ten days ahead presents a difficult problem because both the initial details and the careful accounting for the physics have to be taken into consideration. The detail must be considered because the evolution of the atmospheric systems in existence initially has a considerable effect over the period; indeed, some of the systems may persist, although modified in form and position, for most or the whole of it. The changes caused by the slow acting physical processes can also be profound during the period; the effect of a change in sea surface temperature, for example, due perhaps to up-welling, can be remarkable. Over a few days the development of the weather bearing systems over Europe and the Atlantic can be affected by what is happening very far away, and forecasts for ten days ahead need to be global, taking into account the developments in the Southern Hemisphere, as well as in the Northern Hemisphere.

The equations describing the evolution of atmospheric motions are highly non-linear and whether any useful analytic solutions will ever be obtained seems doubtful. Certainly the only known methods of solution are numerical, and the arithmetic burden is enormous: operational numerical forecasts for a couple of days ahead require the use of large-scale computing facilities with speeds of operation of about 10 million instructions per second; computations associated with operational forecasting up to ten days ahead call for facilities about five times as fast.

To achieve daily operational forecasts there must be scientific research and development on three fronts. The first concerns the acquisition of the right daily data, its verification and reduction to a suitable form. The second concerns the important physical processes which often cannot be represented directly in the computations because the arithmetic has to be carried out for discrete grid points about 100km apart. The third concerns the selection of the mathematical methods of integration, where there will be a delicate balance between accuracy and economy.

An additional complication is the problem of predictability which is currently exercising meteorologists. It is the problem of determining for how long ahead useful forecasts can be made. Experiments carried out up to date show that there is value in forecasts up to ten days and this is confirmed by the few integrations carried out at the Centre. The Centre is co-operating as fully as it can in the First GARP Global Experiment, which is designed, among other objectives, to determine the limit of predictability.

Thus, Europe, through the member nations, has created the first scientific centre for the study and preparation of medium range weather forecasts.

---

* There are now 17 Member States:

Belgium
Denmark
Federal Republic of Germany
Spain
France
Greece
Ireland
Italy
Yugoslavia
Netherlands
Austria
Portugal
Switzerland
Finland
Sweden
Turkey
United Kingdom

The Meteorology Division

The Centre consists of 3 departments, namely:

- . Administration: to provide all the necessary administrative and personnel facilities.
- . Operations: to provide all the technical and operational facilities i.e operational forecast, and computer service.
- . Research: to pursue research into numerical weather prediction.

The Operations Department is further divided in 2 divisions, one concerned with the operational meteorological service, the other with the computer service. These are formally known as the Meteorology Division and Computer Division respectively.

The Meteorology Division is split into 2 sections - Meteorological Applications, and Meteorological Operations. These sections have the following basic tasks and responsibilities:

Meteorological Applications - programming and implementation of the full meteorological operational cycle required to sustain the daily production of medium range forecasts in real time.

Meteorological Operations - monitoring and evaluation of the operational cycle from a meteorological point of view, including checks of consistency of input data and forecast results.

The two articles following describe in more detail the work of these two sections, and are presented as a basis for future Newsletter articles on various particular aspects of the operational suite, progress towards implementation and a general indication of the operational results of the medium range forecast being achieved.

- R. Newson

The work of the Meteorological Operations Section during 1979

The objective of ECMWF to which highest priority has been given, is the development of numerical models of the atmosphere for the preparation of medium range weather forecasts on a daily basis, and the distribution of the resulting analyses and forecasts to the Member States. The efforts of the Research Department at the Centre have been devoted to developing the appropriate numerical models, the efforts of the Operations Department have been devoted to establishing an appropriate computing system and preparing programs making up the complete operational suite. As a result of this concentrated effort, ECMWF expects to begin its operational forecasting activity in the second half of this year. So where does the work of the Meteorological Operations Section and its staff of Meteorological analysts come in?

In fact the Section's work will really get underway this year, comprising for the first part of 1979 monitoring the trials of the Centre's operational suite and later the full daily operational suite. This operational work will be centred in the "Meteorological Operations Room", the Centre's nearest approach to a conventional forecasting office. It is a room of about 50m^2, with visual display terminals (both alphanumeric and graphical), light tables and work surfaces and ample hanging areas for displaying charts. Fundamentally the work of the Meteorological Operations Section will be to monitor the daily production and control of the Centre's analyses and forecasts. The visual display terminals and charts produced from the Centre's electrostatic plotters are the basic tools in this operational monitoring.

It is planned that the Centre's daily operational cycle will begin at 20Z with a re-analysis of the 18Z data for the previous day and then analyses will be carried out for each 6 hours to 12Z (or 18Z) for the current day. Before then a meteorological analyst will monitor, making use of the display terminals, the incoming data during the day, correcting errors where possible in bulletins and reports which have not been automatically corrected. Also, he will check the data coverage in real time to assess the number of new observations going to be used in this analysis, and particularly to be aware of areas of data inadequacies so that the analysis in these areas can be especially carefully examined. Then as the analyses themselves are produced, they will be checked for any such inconsistencies. The forecast will probably begin at roughly 22Z (using either 12 or 18Z of the current day as initial data) and is expected to proceed at the rate of one forecast day in about 45 minutes. As this forecast proceeds, it will be monitored for meteorological consistency and its realism assessed. In examination of the forecast results, a range of charts and diagnostics will be produced. A relevant sample of these will later be displayed in the Operations Room, and a daily briefing is planned, at which the previous night's analyses and forecasts will be discussed, noting particularly seemingly significant or unusual evolutions. Data inadequacies will be identified and previous forecasts which verify each day will be considered.

As well as this real-time operational function, the Meteorological Operations Section will also have a number of longer term responsibilities. The section will evaluate the Centre's analyses and forecasts synoptically in order to identify shortcomings in the Centre's products from the users point of view. In particular, it is hoped that unusual or unrealistic synoptic developments and consistent or recurring errors in the forecasts can be isolated, perhaps leading to improvements in the model. The ECMWF forecasts will also be compared, as far as possible, to the products of other major forecasting centres.

The Technical Advisory Committee, established by Council at its eighth session, is likely to take a considerable interest in and make proposals for liaison and co-operation with the Member States in the question of verification of the Centre's results. Also in this connection, it is hoped that the Meteorological Operations Section will have contacts in each of the Member States in order to channel information to the users of the Centre's products and to standardise quality control procedures in the Member States. It is planned, in the longer term, to develop chart formats including mean charts, summary charts, cyclone track charts and other charts suitable for use in the context of a medium-range forecast, and which will help to exploit the very large amounts of information available from the Centre's analyses and forecasts. It could well be useful to arrange visits to the Centre by forecasters from the Member States, in order to ensure that they are aware of the full range of products available from the Centre and of the uses to which they can be put.

In order to ensure the most efficient working of the operational suite, there will obviously be continuing liaison with other staff at the Centre, including the staff of the Meteorological Applications Section on the development of the ECMWF Meteorological Operational System (EMOS). (See also the article on 'EMOS' in this present Newsletter.) The meteorological analysts will need to become familiar with EMOS and to develop improved methods of data control and diagnostic routines to intercept errors and inconsistencies in data, analyses and forecasts. The Meteorological analysts will assemble long-term statistics on the amount of data received and try to assess if data coverage can be improved.

- J.A. Woods

ECMWF'S Meteorological Operational System

ECMWF plan to start medium-range operational forecasting in the second half of 1979. This bald statement conceals the tremendous amount of work that is necessary to conceive and develop, not only the analysis and forecasting systems at the heart of the operation, but all the supporting programmes involving file handling, data processing and facilities for monitoring and scheduling the whole suite. All in all, the complete operational suite is a highly complex system of interacting and simultaneously running programmes with asynchronous data flows. This present article gives a basic description of the system being implemented at ECMWF and some of the design concepts involved, mainly from a computing point of view.

The basic flow of meteorological information in ECMWF will be similar to the flow of meteorological information which any meteorological centre (World, Regional or National) has to manage as seen from a computer system point of view, though there are obviously design constraints imposed on the system by the scale of the forecasts that ECMWF plans to carry out (i.e. for 10days over a global domain). ECMWF's Meteorological Operational System (EMOS) can be broken down into a number of logical sub-processes, viz. data acquisition, pre-processing, analysis, forecast, post-processing and dissemination of results. Additional to this there are such aspects as the real-time operational supervision, control and scheduling of all the programmes in the system, the graphical display of results, and systematic archiving of the data. The analysis and the forecast needing to take advantage of the computing power will execute on the CRAY-1 The other sub-processes, more involved with data manipulation, file handling and maintaining data integrity, are being implemented on the CDC CYBER 175 system to take advantage of the much more proven reliability and sophistication of the software on this system. The CYBER-CRAY link will allow the appropriate data transfer for input to the analysis on the CRAY, and also the transfer of the products generated during the analysis and forecast period from the analysis and the forecast model on the CRAY-1 to the CYBER for post-processing and dissemination. The various sub-processes and data flows involved in EMOS are illustrated in the diagram.

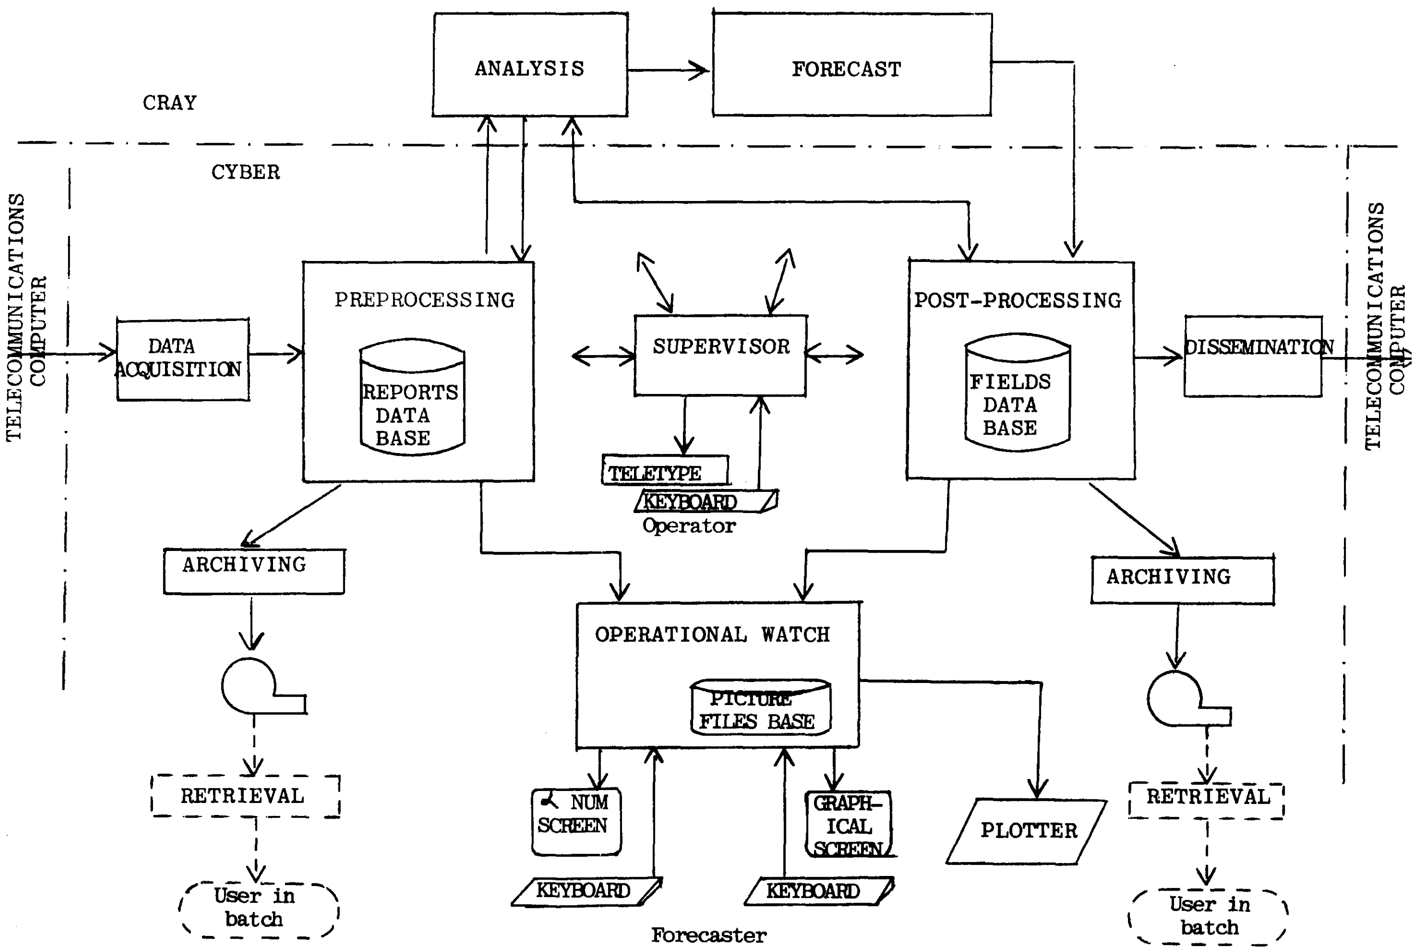

This program will be invoked every time a group of G.T.S. bulletins gathered in a file sent from Bracknell Regional Telecommunication Hub arrive on the CYBER via the telecommunications subsystem. To ensure maximum security of observational data for the operational process a back-up copy of the bulletins will be made and updated as each new file of data is received. Checks will be made for missing files and requests for repeat transmissions if necessary.

The Preprocessing

The main tasks of preprocessing are:

- checking the input GTS data, decoding of information, control of the quality of data;
- creating the file of input data for the analysis;
- providing statistical information on coverage and quality of data received.

It will be observed that these three flows of information are quite asynchronous and they will be handled by three independent systems centred around a data base in which will be stored meteorological reports and operational information on those data. When the data acquisition phase has accumulated a sufficient number of bulletins the decoding program starts. The variable quality of GTS data implies a rather complicated program (scanner, syntactic analyser) to recognise and decode the information from a bulletin. Even with the most sophisticated software, one cannot recognise automatically and decode all the contents of all GTS bulletins. Hence a few per cent of the bulletins are rejected by the program and written in a file where they wait for manual examination and correction if possible. This is performed by a meteorological analyst using an alpha-numeric VDU terminal. (see the article on the description of the work of the Meteorological Operations Section also in this Technical Newsletter). The decoding program is then called to reprocess any corrected bulletins.

The decoding program also keeps a record of input data in terms of statistics on the time of arrival of each bulletin, or each report, presence of bulletins and reports, quality of bulletins and reports. Those statistics will be saved providing later useful information for monitoring data receipt and scheduling of ECMWF's data acquisition/preprocessing system.

The last step before entering the report into the data base is the control of quality. The internal consistency of each report is checked, and each meteorological parameter verified according to climatological limits and in some cases physical laws. The quality control program will interrogate the data base to check the presence of a similar observation. The program will flag each suspicious meteorological parameter and, if necessary, will propose a substituted value However, both the original and the substitution will be kept in the data base. The program will, at the same time it adds or updates a report into the data base, register the time of arrival and the quality of that report in separate operational files, providing statistics and valuable information for later use by the meteorological analyst.

The reports data base is at the heart of the preprocessing and its efficiency is crucial since the execution of all other subsystems will depend on the ease and speed of access to the information in the data base. Data should also be read-accessible from any user's program of ECMWF's computer system. Careful study was undertaken to determine the optimum organisation and physical structure of the data base, which needed to provide:

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

The type of information likely to be required by the different users accessing the data base was identified and the structure and format of the data needed by the "user" was analysed. Following this, the ideal organisation and structure of the information for a perfect data base was defined and from this the most feasible and realistic organisation was constructed with the inevitable compromises between use of sophisticated facilities and efficient use of computer resources (mass-storage space, I/O and CPU overhead). It was thought that indexed sequential organisation would be the best choice to satisfy the various requirements, but with a specific direct access structure to allow fast access and construction of the operational files for the analysis.

Regarding the access methods, the design constraints are that the data base is neither for general purposes, nor a general transaction oriented system. The nature of users as well as the nature of the data manipulated are very well defined and limited to meteorological applications oriented towards weather forecast studies. Therefore, considerations of sophisticated data definition languages, data manipulation languages, for a sophisticated real-time data base monitor which would add useless overhead can be omitted. What is necessary is a batch data base where programs will access the meteorological data through a call to a subroutine with the arguments identifying the set of data which is needed for the computation. Tools providing retrieval of reports over areal elements from the data base are available as well as tools for access to single reports. The access software is designed so that any change to the data base organisation, Cyber Record Manager or the Operating System can be incorporated with the minimum of effort, with such changes being totally transparent to the user. A modular structure has been adopted which is defined by a number of different "levels", each level calling a number of subroutines from a level below. There will also be a control card system by which the user will be able to access data. A packed format has been adopted within the reports data base itself because tests have shown that I/O time is considerably reduced with less volume of data to handle. Extra CPU time is used in conversions between the packed data base formats and the unpacked formats normally wanted by users, but this increase can be accommodated on the CYBER 175.

Analysis-Forecast Systems

The analysis scheme needs all the reports organised in intervals of 6 hours for up to 6 such intervals at one time. The following steps are necessary to provide input data for the analysis for each 6 hour time interval:

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

This input file will be generated on the CYBER and passed to the CRAY when needed by the analysis program. The forecast to 10 days will follow the analysis, starting around 22Z and running for four hours. It is not the purpose of this article to describe the sophisticated analysis and forecast Systems, which deserve on their own a contribution. It will be simply said that the important operational constraint for those systems should be their ability to be restarted after interruption for any reason or loss of data files. They have to be designed in such a way that they can be considered as a set of "black boxes" able to be restarted, able to know the status of their permanent files, able to reposition themselves if necessary in the files and able to regenerate local files if necessary. Macro and/or control cards will be implemented in the programs to receive or send permanent files from or to the CYBER, as well as to send messages to the central supervision program. Some permanent files will be disposed to the CYBER for back-up reasons and put directly on 6250 bpi tape on the CYBER, for restart facility in case of data loss on the CRAY mass storage. The checkpoints and copies of back-up files on tape must obviously be carefully organised, providing safety but not with too heavy and long procedures for back-up or restart in terms of overhead and operator manipulations.

The Post-Processing

The post-processing handles, organises and distributes the analysis and forecast results. Fields of data, covering the globe, arrive on the CYBER from the CRAY computer at the end of each 6-hour analysis or forecast step. A very large amount of data is involved - each 6-hour step over $3x10^6$ numbers are passed from the CRAY and there are 44 such steps in a complete analysis/forecast cycle. The system has to transform and re-structure the data into forms suitable for systematic visualisation on a screen, automatic plotting of charts and transmission of files to the telecommunications computer according to the pre-defined schedule of Member States' requirements (several thousand fields of analysis and forecast are to be disseminated every day to the Member States.) The asynchronism of all these tasks necessitates the staging of data at different levels within the system. Also, various transformations have to be applied to the data. These considerations lead to the construction of different data bases corresponding to different stages in the process of successive transformations which are needed to satisfy the various output requirements. The "global fields data base" (GFDB) will contain the output history files from the analysis and forecast. The "Polar stereographic data base" (PSDB) will contain polar stereographic projections of Northern or Southern Hemisphere. (1400 such fields will be produced every day). The dissemination data base (DDB) will form the base for staging fields ready for transfer to the NFEP and transmission to Member States, in the form required for transmission to Member States.

The access to the data bases will be provided by a general subroutine called 'FIDABA' (FIELDS DATA BASES) of which the logical structure will be similar to the reports data base access software. Given the large quantity of data involved, an efficient and fluent back-up system has to be designed to prevent, in case of loss of data, the need to re-run too long or too many programs, or to remanipulate too large a volume of data. The system adopted is that the "GFDB", "PSDB" and "DDB" will be stored on different mass-storage drives; "PSDB" is a back-up for the "DDB", "GFDB", a back-up for "PSDB". The history files from the analysis and forecast are retained as back-up of the GFDB up to 8 steps, but for every 8 steps, the last 8 steps will be copied to tape. The corresponding 8 history files will then be purged.

The Dissemination

One program activated either at the end of each post-processing sub-step and/or according to a pre-defined time schedule, will collect all the products from the DDB (accessed through FIDABA) necessary to build up a file to be sent at that time to a Member State. The file built up is then forwarded to the telecommunications computer, with its destination indicated in a header for transmission to a Member State.

The Archiving

The archiving system will run once every day to archive:

- the 3 days old reports data base files;
- the operational analysis files;
- the GFDB files containing the forecast and analysis results generated during the previous operational phase (these results will mainly be archived in spectral form).

8 to 10 6250 bpi tapes will be mounted every day to archive these data (in duplicate for data security). Weekly or monthly for extra safety a third version will be generated which will be kept in a safe place outside ECMWF's Computer Centre. The identification and the location of the files on the tapes is kept in a directory written on mass-storage with two identical versions on two different drives. To retrieve data from the analysis a control card will be supplied (RETECAR-RETRIEVE ECMWF's ARCHIVES) and use of this will provide a direct copy of an archived file or a set of archived files on disk; there will also be facilities for using MEDABA or FIDABA sub-routines or control cards to access data from a regenerated data-base environment. Data bases will be regenerated within the definite limit of the cache mass-storage memory system. Flags and size of the file will be bit fields added to the file pointer in the directory to enable the retrieval program to manage the cache mass-storage memory.

Supervisor

Within EMOS, many basic tasks are asynchronous, e.g. continued arrival of observational data, analysis and forecast fields produced at regular steps, dissemination of the products following a predefined schedule. In addition to this, one must be able to follow in detail the progress of the operational suite. An internal system of synchronisation, or better, a task acting as a scheduler-monitor-supervisor (EMOS/SMS) is needed. This system will:

- schedule and initialise tasks;
- synchronise tasks;
- allow man-machine interaction;
  (by centralising messages from programs by taking into account manual requests)
- help restart procedures.

A file containing directives and job control cards scanned by the supervisor program, will control the progress of the operational suite. The basic principle used in developing a supervision method for EMOS is that jobs within the operational suite can be grouped in "families". A "family" contains a set of jobs for which the executions are connected. If job 1 of family F is executed, then necessarily job n of the same family F must be executed sooner or later. However, jobs belonging to different families are not necessarily linked at all. This concept of "family" assists in the management of the few hundred jobs which make up the operational suite. Families could be run independently, separately or in parallel. They can be restarted independently as well. The supervisor program will also collect operator's request or program messages to centralise the management of the suite, and facilities specifically to display supervisor messages will be implemented. Obviously, what is needed from the system point of view is a means of communication between tasks and the possibility for the central supervision program to activate tasks. The appropriate facilities are available within the CYBER NOS/BE computer operating system - specifically the system control point facility for the former requirement, and a 'ROUTE' macro for the latter.

The Operational Watch

This will be independent from the operational suite in terms of scheduling. This subsystem will provide information to the forecaster on duty while the operational suite runs. Information is produced in reply to requests typed on a keyboard from an alphanumeric or graphical visual display terminal. Basically the role of the operational watch is to allow the status of, and data in, the reports and the post-processing data bases to be monitored. A general"intercom-type" utility analyses the requests and calls the appropriate overlay to visualise information on screen or invokes a separate job to plot information. The subsystem is divided into a certain number of independent programs. The functions of each program can be designed independently and even developed and implemeted separately. At a later stage it is hoped that a more integrated system can be implemented where graphical images may be retained in a picture file base to be recalled later for display.

Practical Aspects for the Daily Operation of EMOS

A library will hold centrally all the programs of EMOS. Any modifications in these operational programs will be implemented after quality assurance by one single person responsible for the library management. A system for the naming of files and programs has been carefully set up enabling daily running of EMOS with possible retrieval of old archived files and interrogation of on-line data bases, as well as simultaneous test runs of future EMOS versions.

- J. Martellet

*The Computing Facility: Status and Plans

The start of the external distribution of this newsletter is a good occasion to present the latest plan for the computing system and its present status.

The period of scattered interim facilities is over now; staff and computers have moved into new headquarters at Shinfield Park. The two main computing systems, the Cray and the Cyber, have been installed successfully, passed their respective provisional acceptance trials successfully and are both well into their three months final acceptance period.

The configuration diagram of the computer system is shown on the front page of this newsletter. It reflects the three layered approach with specialized subsystems for communications and graphics, the Cyber as general support system and the Cray as number cruncher.

CRAY-1

Until the move to Shinfield Park we had the prototype computer at our disposal. This system is used elsewhere now, and ECMWF has its final system, serial no.9, with a million words of memory, four disk controllers and eight disks with a total storage capacity of 300 million words. The intention is to use the storage on the Cray for temporary files only, permanent storage will be provided via the mass storage (tapes and disks)on the Cyber.

The system is performing very well, hardware and software are proving reliable. The overall scheduled availability since November has been around 99%. There are currently no plans for enhancement to the configuration, the only planned changes in the software are updates to new releases of operating software (COS) and the Fortran compiler (CFT).

Control Data CYBER 175-300

During the interim period this new model from the Cyber range was installed. With the move to Shinfield Park the configuration was enhanced with additional memory, peripheral processing units (PPUs), disks and tapes. The present configuration has 196K words memory, 20 PPUs, 24 channels, 2 disk controllers and 8 disk drives, 2 tape controllers and six tape drives, 2 line printers, 2 card readers and a card punch. In April a further four disks will be installed to ease the development of the operational suite of programs. In July, prior to the start of the operational trials, the peripheral configuration will be enhanced to four disk controllers with sixteen disk drives having a total storage capacity of about 400 million words (4 billion characters) and four tape controllers with eight tape drives.

With the move to Shinfield Park a major upgrade of the software was undertaken to NOS/BE 1.3 with Intercom 5.0 as the subsystem for terminal facilities.

The system is performing very well, the overall scheduled availability of hardware and software since November has been well over 99%. Apart from the configuration enhancements mentioned above, there are developments going on to link the tele-communication subsystem and the graphic subsystem to the Cyber.

Cray-Cyber Link

The Cray and the Cyber are now linked to allow the Cyber to act as general support system, whereby all access to the whole computing system is handled via the Cyber and all permanent storage of information is managed on the Cyber.

The link facilities have been introduced in phases. The first phase was to allow transfer of jobs and input/output to and from the Cray; the second, to allow transfer of files between the Cray and the general mass storage on the Cyber, and lastly, the transfer of files between the Cray and private disks and tapes on the Cyber.

All the facilities are used extensively and perform satisfactorily. The last step is to show that transfer rates over the link are as expected and present experience indicates that no problems should be expected in this area.

The conclusion is that the two suppliers (Cray provided the link hardware and Control Data the link software) have brought this difficult project to a successful end and the linked system is available to the user community.

Communications subsystems

There are two communications subsystems, one for the internal terminal facilities and one to control the remote network to the Member States. This year a study of the desirability, and technical effort required, to integrate both networks will be undertaken. Internally 10 alphanumeric visual display units (VDUs), one graphics VDU and two remote job entry (RJE) terminals are installed. The network will be extended in March to 20 alphanumeric VDUs, 4 graphics VDUs and 2 RJEs. These terminals are used for program development by the Centre's staff and some will be used to control the execution of the operational suite.

The external network will be a star network of dedicated telephone lines from the Centre to the Member States. Initially this network will be a mixture of medium speed and low speed lines, gradually to be updated to all medium speed lines. This network will be used by the Centre for the acquisition of observational data and the dissemination of forecast results. The network can also be used by the Member States for remote job entry.

The present schedule assumes in 1979 medium speed lines to the United Kingdom, the Federal Republic of Germany, France, Sweden, Finland and Denmark, low speed lines to Italy, Greece, Spain, Yugoslavia, Portugal, Netherlands and Turkey. The network will be controlled by a subsystem, based on a Regnecentralen 8000 computer. The software for this system, to drive the lines using a protocol based on the standard CCITT X25 protocol, and to transfer files to and from the Cyber, is being developed by SIA. The system is scheduled to be operational in early June.

In a joint venture with FRG, Sweden and Denmark, the Centre concluded a contract with SIA for the development of software on a Regnecentralen 3600 minicomputer to be used by these Member States for the reception of forecast results and for remote job entry.

Graphics Subsystem

Graphical output can be produced via plotters, VDU or film. At the moment the Centre has an off-line Varian electrostatic plotter. In February an on-line Versatec electrostatic plotter (22 inches wide, 100 dots/inch, paper speed 2 inches/sec.) will become operational and in June a similar on-line plotter will be added. The graphics software which is in use, including a contouring package, is presently being converted to allow output to a VDU, so that during the operational trials charts can be previewed by the user before being plotted for longer term use.

Later in the year development will start on a truly device-independent graphics system, in which the application programs generate device independent picture files before the actual output device has to be chosen.

Also later in the year, a feasibility study will be undertaken for the use of computer output on microfilm (COM) where output, graphical or alphanumeric, is output on microfiche or microfilm.

Operations Schedules

From the start of service at Shinfield Park, a full service of 24 hours, 7 days per week has been in operation. Throughout the year an average of 20-25 hours per week on the Cyber and 10-20 hours on the Cray, will have to be reserved for dedicated sessions. These will be used to perform hardware maintenance, to implement and check out upgrades, changes and corrections to the systems software, to prepare and check out integration into the system of the graphics subsystem, telecommunication subsystem and additional peripherals, to implement the operational forecast suite of programs.

Allocation of Capacity

From mid-1979 computing capacity can be made available to the Member States. It has been decided that 50% of the Cray capacity will be used for operational purposes, 25% for research by the Centre and 25% for research by the Member States. Time has been allocated to the Member States according to the results of a questionnaire produced in 1978. It is assumed that the Member States only require capacity on the Cyber to support the work they are doing on the Cray, so no formal allocation of Cyber capacity has been undertaken.

Practical allocation of capacity will be done on a project basis, every project will be assigned an account code and an allocation. Regular 4-weekly summary reports will be produced. The Member States will be approached shortly to put the allocations into practice. Around June another questionnaire will be issued, asking for requirements for the next year.

- Rob Brinkhuysen

Forecast model I/O times on the CRAY-1

Initial runs of the grid point model on the Cray-1 indicated unacceptably long elapsed times. These times were dominated by the time necessary to perform the I/O. An explicit model executing on serial 1 using COS 1.0.0 and 2 disk controllers was not able to achieve a transfer rate higher than 625 Kwords/sec (c.f. theoretical 1000 Kwords/sec).

Since then, improvements have been introduced into the operating system. An explicit model run on the current system (COS 1.0.3) can achieve a peak transfer rate of 1768 Kwords/sec (c.f. theoretical 2000 Kwords/sec) i.e. 88% efficient using 4 disk controllers. The semi-implicit model requires that data be scanned backwards during part of the algorithm. Hence the I/O situation is much more complex. Backward scans may be accomplished by BACKSPACE statements or by random access methods. Both methods interrupt the normal streaming mechanism by which the high sustained transfer rates are achieved. Trials suggest that random access techniques are slightly inferior in terms of transfer rates.

In order to minimise the disruptions to I/O streaming caused by backspacing, a technique of merging the semi-implicit records has been tried. Thus, there are fewer, larger records and therefore fewer interruptions to streaming. Large improvements in elapsed times are possible, at the expense of extra memory to hold these 'super records'. The extreme case is to hold all semi-implicit records in memory for that part of the algorithm which requires backward scans.

A further technique which can significantly reduce model I/O times involves packing the field before writing to the work datasets. Meteorologically, this is satisfactory as sufficient precision can be retained within 32 bits provided that all computation within a time-step takes place at full precision (i.e. data is unpacked after being read in).

I/O timings for a 10-day forecast with an N48 grid are about 3 hours without any of these improvements. Combining all time saving techniques together suggests that I/O times of less than 1 hour are possible. Considerable reductions in the computational requirements of the model are of course essential in order to approach this figure for a real forecast.

- David Dent

*Cray-Cyber Data Conversion Utilities

The representation of characters and numbers on the Cyber differ significantly from those on the Cray-1. These differences may be summarised as follows:

	Cyber	Cray
word length (bits)	60	64
character code	6 bit display code	8 bit ASCII
negative integers	ones complement	twos complement
floating point exponent (bits)	12	16
floating point exponent bias (octal)	2000	040000

In addition there are special forms on the Cyber for -0, ± infinity, indefinite.

As a result of these differences, it is necessary for data generated on one system to be converted before it may successfully be processed on the other system. Note, however, that data may be stored on the alternative system without conversion provided it is never used on that system - e.g. Cray-1 datasets backed up on the Cyber.

The Link (or station) software running on the Cyber provides automatic character conversion between display code and ASCII for job input and output and for complete files. The station also handles conversion between Cyber record blocking and the Cray blocked format. However, it does not provide any integer or floating point conversion capability, nor does it handle the difference in word size, in this case.

A set of data conversion utilities to provide that additional conversion capability is now being tested. These utilities execute on the Cray-1. They can be used to convert integers, real numbers, characters, etc., and may be applied to words, records, files or complete datasets. For most operations the implementation is performed by vector instructions and is therefore very efficient.

These utilities convert a given Cray dataset,into a second Cray dataset, they do not transfer the dataset to or from the Cyber. Transferring datasets is done via the ACQUIRE and DISPOSE control cards.

Data types which can be handled include the following:

REAL
INTEGER
BINARY (bit string)
COMPLEX
DOUBLE
LOGICAL
CHARACTER

Conversion may be performed on a complete dataset, on a file, on a record or on individual words. The following examples illustrate the general use of the utility:

Ex: 1. CONVERT,I=X,O=Y,FROM=CYBER, TYPE=REAL.

Converts a data set of real numbers in Cyber floating point format held in dataset X to the corresponding Cray-1 floating point format, in dataset Y.

B.B. X must have been pre-staged to the Cray-1 e.g. ACQUIRE (DN=X,.......)

Ex: 2. CONVERT,I=X,O=Y,TO=CYBER,TYPE=INT:REAL, NF=1:EOI.

Converts the contents of the Cray-1 formatted dataset X to the corresponding Cyber format, and stores in the dataset Y. The first file of X contains only integers. Subsequent files (unspecified number) contain only real numbers.

Normally Y would then be DISPOSE'd to the Cyber e.g. DISPOSE,DN=Y,.....

Ex: 3. CONVERT,I=X,O=Y, TO=CYBER,TYPE=CHAR:INT:REAL,NR=1:2:5.

Converts one file of a Cray-1 format dataset X. The first record is formatted (8 characters per word), the next 2 contain integers and the next 5 contain real numbers. After conversion the dataset X is positioned after the 8th record. A further CONVERT statement could now be used to perform different conversions on subsequent records and files of the dataset X.

Ex: 4. CONVERT,I=X,O=Y,FROM=CYBER,TYPE=CHAR:INT:LOG, NW=100:22:EOR,NR=1

CONVERT,I=X,O=Y,FROM=CYBER,TYPE=REAL,NW=EOR,NR=EOF.

(a) converts one record consisting of 100 character words, 22 integers and the remainder logical. Leaves dataset positioned after first record.

(b) converts the rest of the dataset consisting of real data.

Equivalent conversions by subroutine call will also be available. A Computer Bulletin B7.2/1 describing the utilities has been distributed.

The utilities are currently being tested. They are available in the direction to=Cyber for general use. They are also available in the reverse direction on a trial basis. Contact User Support for details.

- David Dent

*Cray to Cyber Coded File Transmission

Until the Link became available the normal method of sending a coded file to the Cyber was by use of the DSOUT utility. There is no longer any need for this facility and it will be withdrawn shortly. Users should change to

DISPOSE (DN=x,SDN=pfn,MF=CY,ID=id,DC=ST)

(See Computer Bulletin B2.7/1)

- David Dent

*Audit of Permanent Files on the Cray-1

The version of AUDIT released with COS 1.0.3 contains a new feature which has led to some confusion.

When a Cray-1 dataset is catalogued (by the SAVE control statement), the account parameter of the job is recorded in the catalogue (i.e. the dataset is assigned to the same account as the job which catalogued it.) For all Cray-1 jobs submitted via the Cyber, this account field is taken from the Cyber ACCOUNT statement. However, for Cray-1 datasets created before the ACCOUNT statement came into use on the Cyber, a 'zero' account field has been recorded.

AUDIT compares this dataset account field with the account field of the AUDIT job itself. Only if the fields agree will AUDIT report on the dataset.

Ex: 1.

```
ABC,STCRA.
ACCOUNT,ECXYZ.
AUDIT,ID=ABC.
```

reports only on datasets catalogued under the account field ECXYZ.

Ex: 2.

```
ABC,STCRA.
ACCOUNT.
AUDIT,ID=ABC.
```

reports only on datasets catalogued before ACCOUNT cards came into use (zero account field)

* Note that, eventually, this account statement will be illegal as it contains no parameter.

- David Dent

*CFT Bugs

Users of COMPLEX or DOUBLE PRECISION beware. The BUFFER IN/OUT I/O statement does not handle the data length correctly.

e.g.
```
COMPLEX C(100)
BUFFER OUT (1,0)(C(1),C(100))
```
will transmit 199 words only.

- David Dent

Disappearing Trick

Shortly before Christmas it became evident that the ECMWF model was failing, but not consistently. Usually, when it was rerun, the problem disappeared. After some detective work by David Dent it was found that part of a word was being over-written. In some cases this was an instruction and only caused a failure if it was executed. If it was a data value it sometimes gave incorrect results, or caused an arithmetic error, or even went unnoticed, depending upon which variable was affected. As only a partial word had been overwritten and the computer model does not usually manipulate parts of words, the most likely culprit seemed to be the operating system. After some investigations by Neil Storer three possibilities were suggested. Some work by the Cray analysts then confirmed the error to be in the disk queue manager section of the operating system and that it only occurred when a file spanned more than one disk and hence was dependent on I/O activity. This bug has now been fixed. It has been present since COS 1.0.3 and could explain other 'mysterious' errors that came and went since we moved to Shinfield.

- J. Greenaway

Advice for Users of CRAY UPDATE

(Reprinted from ECMWF Computer Newsletter No. 5, January 1978)

CRAY UPDATE looks like CDC UPDATE. This is misleading. There are so many important differences that each should be approached as a different utility. The CRAY-1 COMPUTER SYSTEM UPDATE REFERENCE MANUAL 2240013 should be used in conjunction with CRAY UPDATE, and is, in general, correct. In this article, we introduce CDC UPDATE users to CRAY UPDATE (- it is not intended as an authoratative document, nor does it seek to amend or contradict information contained in the CRAY MANUAL, which should at all times be consulted when using CRAY UPDATE).

Invoking CRAY UPDATE

UPDATE is called in a similar way on both systems. There are normal, Q, and F forms of UPDATE, and most of the parameter key letters fulfil similar functions. CRAY UPDATE produces by default no output listing. Types of printed listing must be requested using the LIST parameter, as opposed to the L parameter on CDC. There is an L parameter on CRAY UPDATE, but this controls the data set to which printed output is directed.

Q mode CRAY UPDATE requires ALL decks including COMMON DECKS to be requested with *C cards.

Types of CRAY UPDATE

There are two types of CRAY UPDATE:

a) INITIAL RUN

This generates a library from an input data set. CRAY UPDATE recognises any run as an initial run if the first input card begins with either *DK or *CDK.

b) MODIFICATION RUN

This modifies an existing library. CRAY UPDATE recognises any run as a modification run if the first card on the input data set begins with *ID.

The INPUT data set

a) NULL FILE not permitted

Since the type of CRAY UPDATE run is determined by the first directive on the input data set, this data set must not be empty. If an empty file is encountered on the input stream, CRAY UPDATE attempts to find directives in the NEXT FILE!

b) Short form directives only permitted

Only *DK, *CDK, *ID are permitted in CRAY UPDATE. The long forms *DECK etc. allowed in CDC UPDATE are not supported.

c) ADDFILE not supported

*DK on a modification run is the accepted method of inserting a new deck in CRAY UPDATE. The CDC facility of being able to direct new decks to be added from a source other than the input data set can be effected by using the *READ directive.

d) Positioning deletes and insertions

The long form for specifying ranges of cards on a *B, *I, or *D directive are supported by both CDC and CRAY UPDATE. There are many short forms available on CRAY UPDATE; they should be avoided as they do not obey the same rules as short forms in CDC UPDATE.

e) PURGE, PURGE DECK, YANK, SELYANK not supported

There is provision in the structure of cards stored in CRAY UPDATE LIBRARIES for some of these facilities to be added later, but they are not supported as yet.

f) EDIT

The *EDIT directive in CRAY UPDATE enables a single deck or range of decks to be written to the new program library with all inactive cards purged. Re-sequencing does not take place.

g) CDK

A COMMON DECK is identified to CRAY UPDATE by *CDK, not *CD.

Moving CDC libraries to the CRAY-1

This can be accomplished as follows:

a) use CDC UPDATE to obtain a SOURCE file;

b) pass the SOURCE file through a program to change *DECK to *DK, *COMDECK to *CDK, etc.

c) use the modified file as input to CRAY UPDATE(N).

Conclusion

It is easy to make errors using CRAY UPDATE. These errors usually occur when trying to use CRAY UPDATE as though it were CDC UPDATE. On most occasions, working through the CRAY UPDATE MANUAL has led to detection of such errors. Though not as sophisticated as CDC UPDATE, CRAY UPDATE has provided a useful means of library maintenance.

- Rex Gibson

*COS 1.04

We have very recently received the latest version of the Cray System (COS 1.04, CFT 1.04 etc.). At present we have not yet decided when, or even if, the system will be implemented at ECMWF. This system does not include very many new features; the significant features are listed below for your interest.

1. The instruction scheduler in CFT can be enabled or disabled.
2. Unrounded multiply in CFT generated code may be selected.
3. Truncation of floating point operations in CFT generated code may be selected.
4. A DO index when used in non-subscript expressions is treated as a 64-bit value by CFT.
5. The divide algorithm used by CFT has been changed to use an unrounded multiply for the intermediate step; this may produce different results.
6. Expressions like B(I)*A are evaluated the same way by CFT whether in a vector or scalar loop.
7. An option may be selected to have CFT process all floating point arithmetic as calls to user supplied subroutines.
8. ERR=sn in READ/WRITE statements are now implemented.
9. Implied DO in READ/WRITE statements are done as vector transfers.
10. The capability to read a null dataset is added (see caution below).
11. Text mode in COMPARE is added.
12. New routines added to $SCILIB: MAXA, MINV, CFFT2, CRFFT2, RCFFT2, PACK, UNPACK.

The following caution should be noted:

Implementation of the capability to read null datasets (this completes the changes started in $FTLIB 1.03) requires that programs run under COS 1.04 use logical I/O routines no older than 1.03; older versions will cause output to disappear.

- Peter Gray

*NOS/BE(473) Problem Harvest (continued)

Some further problems or points requiring explanation have been discovered, in addition to those described in the last ECMWF Computer Newsletter.

- RMS ERRORS WHICH ARE NOT CORRECTABLE - This message indicates the existence of disc hardware problems, leading to a loss of RMS information. The job is then aborted. Operators and/or engineers intervention is require before any such aborted job can be successfully resubmitted.

- JOB DROPPED BY DEADSTART RECOVERY - The operating system may "crash" for diagnosed reasons, or it may "hang-up" for some inexplicable reason (to be later investigated). A common operational procedure called D/S Recovery allows the system to be reinitialised without losing the I/O queues and the PF base. Any jobs caught executing at a control point by a D/S Recovery will either be RERUN and reinitiated from the input queue automatically, or DROPPED. In this latter case a job was flagged as not capable of being rerun because it performed some PF functions (catalog, purge, alter, rename, extend, modify) whilst executing.

- BACKSPACE WITH CARE - Writing unformatted records, backspacing and rewriting them can be dangerous. Files generated with this procedure should be read immediately and any errors should be reported to USS.

- FTN MEMORY OVERFLOW - FTN is a 2 pass compiler. Only the second pass has a field length management capability. In pass 1 there are tables with a fixed length of 4095 words. Should compilations abort for lack of space an RFL control card will provide more of it (values from 60000B upwards will do). E.g.

```
RFL (70000)
FTN.
REDUCE.
```

Note: REDUCE is required to restore dynamic memory management.

- PERMANENT FILE "DUM, ID=PUBLIC" - Private pack owners who try to get rid of this Permanent File without success are advised that its existence is required by a utility called DUMPF, which is provided as a tool for backing up a PF base. See the last ECMWF Computer Newsletter.

- L. Bertuzzi

*"RT=W,BT=I file structure and the dustbin analogy"

- or - "what one should know before using the Cyber COPY utilities"

A Fortran unformatted write statement specifies a list of variables to be written. The contents of these variables are treated as one record of binary data, of record type W(RT=W).

Each record is disposed of by the operating system, analogous to the disposal of one's daily rubbish into a number of fixed size dustbins.

NOS/BE calls a "dustbin" a "physical block", (or just a "block") of block type I (BT=I). Each block can hold 512 central memory words of data, i.e. 5120 characters.

On the lids of its dustbins NOS/BE writes a sequential number (the block number), starting with one, plus the number of daily amounts of rubbish already thrown out (i.e. records), since the first dustbin was started. The dustbin lid is called BCW (Block Control Word) and it reduces the block capacity for the user now to 511 words.

Suppose now that the NOS/BE dustbins can be filled upside down, so that the lid becomes the bottom of the whole thing. For each record NOS/BE introduces into the dustbin another lid, thus reducing the capacity of the container again by one word. This record lid is called RCW (Record Control Word) and its main function is to keep track of how much block space has been taken by the previous record, and how much is going to be used by the new one.

Should a record exceed the capacity of a block, the RCW will flag the space to be used in that block as "initial"; the first RCW of the next block will tell how much more space is being taken by that same record. Once the last record has gone into the last block NOS/BE appends a special RCW to separate any empty space left in the block from the last record.

At the end of a weekly collection of rubbish a set of dustbins would then be lined up waiting for the dustman (or refuse collector, as they are now known). On the Cyber the same could be achieved by compiling:

```
      PROGRAM RUBBISH (BIN,  TAPE1=BIN)
      DIMENSION TRASH (...)
      INTEGER DAY (7)
      DATA DAY/..,..,..,..,..,..,../
C
C     PRODUCE SEVEN RUBBISH RECORDS WITH AN
C     RT=W, BT=I STRUCTURE
C
      DO 1 J=1,7
      JJ=DAY(J)
      WRITE(1) (TRASH(JJJ), JJJ=1,JJ)
    1 CONTINUE

      STOP
      END
```

and executing the above program.

When the dustman comes around, he does not collect more than a week's load of rubbish, nor does he want to bother checking what is written on each lid, so there is a need for an eye-catching mark, delimiting one week's set of dustbins from the next or from a neighbour's set.

In a NOS/BE environment this is achieved by adding to the last block a string of 48 bits: The mysteriously famous "Level Zero End of Record or "EOR" (everybody knows it exists, but nobody tells you where it is).

Now we can cast some light on the double meaning of "record".

1. for a fortran unformatted Read/Write statement a record is delimited by RCWs;

2. for NOS/BE and its file handling utilities (SKIPF,COPYBF,COPYBR,COPY) a record is delimited by a level zero EOR.

If Program Rubbish executes twice in the same job run, without rewinding file BIN between the first and second LGO, the final contents of BIN will be: 14 RT=W records and 2 NOS/BE records (by the way they are RT=S records).

In no way would COPYBR read from file BIN a specific RT=W record, but it would retrieve all those written by the first LGO and/or all those written by the second LGO. There is no way out of this first limitation, other than a made to measure fortran program. Conversely, a "Read" and even an "If (EOF(-))" statement would be unable to tell which records were written by the first LGO and which by the second.

There are two ways out of this second limitation:

1. A W type delimiter can be generated by setting a flag in the last RCW of a group of RT=W records; the flag is set with an Endfile statement and the resulting RCW is called End of Partition (EOP). Note: EOPs are not sensed by the SKIP/COPY utilities - they are RCWs.

2. a "level 17 EOR" (EOF) can be forced after a level zero EOR with a dummy COPYBF, i.e.

```
REWIND, BIN.
COPYBR,BIN,NEWBIN.   Copies all RT=W records written by first LGO.
COPYBF,DUM,NEWBIN.   DUM is empty; EOF is written onto NEWBIN.
COPYBR,BIN,NEWBIN.   Copies all RT=W records written by second LGO.
REWIND, NEWBIN.      Writes second EOF at the end of NEWBIN.
```

Note: "IF (EOF(-))" statements sense both EOPs and EOFs.

Conclusion

The aim of this prose exercise has been to provide some food for thought on a very simple, but sometimes tricky, subject. Those who might have cared to read it are now expected to flood the advisory office with clever questions. Failing this, the author will carefully restrain from conceiving more analogies of any nature whatsoever.

- L. Bertuzzi

FORTRAN 77

Reprinted below is an article from the ECMWF Computer Newsletter number 2 (Sept. 1977) outlining the new Fortran standard, known as Fortran 77.

Both manufacturers (CDC and CRAY) are working on providing compilers which will include all these enhancements. We will keep you in touch with developments through this Newsletter.

- Andrew Lea

Most current FORTRAN compilers adhere to, or extend the scope of the last American National Standards Institute (ANSI) FORTRAN language definition of 1966, known as X3.9-1966. Now, after eleven years, the ANSI X3J3 committee responsible for the FORTRAN standard has arrived at another definition informally known as FORTRAN 77. The X3J3 committee met 62 times and consumed 4127 meeting man-days during this time. The public were not forgotten and during the public review period that began March 1, 1976 and ended September 28, 1976, 289 people and organisations sent 1225 pages of comments!

As an introduction to FORTRAN 77, there follows a copy of text that highlights the major differences between the old and new standards. The text was taken from "FOR-WORD" (Fortran Development Newsletter - August 1977) which in turn, obtained it from the "ANS X3.9 FORTRAN Revision - Final Report" by the then X3J3 Chairman, Frank Engel, Jr.

Major Differences between FORTRAN 77 and the Previous Standard, ANS X3.9-1966

NOTE: An extremely important consideration in the development of FORTRAN 77 was the minimization of conflicts with the previous standard. The differences listed here represent (with only two exceptions) extensions to, rather than conflicts with, ANS X3.9-1966. It should also be noted that FORTRAN 77 consists of full language and a subset; differences noted in this list refer to the full language.

1. "Structured" branching statements. The following statements have been added to the language:

```
IF (e) THEN
ELSE IF (e) THEN
ELSE
END IF
```

For each IF-THEN statement, there must be a corresponding END IF statement. Between the IF-THEN and the corresponding END IF there may appear any number of ELSE IF-THEN statements, and at most one ELSE (which must not precede any of the the ELSE IF-THEN statements). Groups of statements delimited by IF-THEN and END IF must be properly nested, both with respect to other such groups and with respect to DO loops. Transfer of control into such groups is prohibited.

2. Character data type. A new data type, consisting of character strings of fixed declared length, has been added to the language. Included are character constants, character variables, and arrays of character data. Operations on character data include concatenation and designation of substrings. Intrinsic functions for conversion between single characters and small integers, for pattern matching, and for determining the length of a string are included.

The Hollerith data type of ANS X3.9-1966 has been deleted. Because this introduces a conflict with the previous standard, it is anticipated that some processors will wish to retain Hollerith data as an extension to FORTRAN 77; accordingly an appendix has been included with recommendations for the form such an extension should take.

3. DO loop changes. A DO statement specifying a terminal parameter whose value is less than that of the initial parameter is no longer prohibited. If the incrementation value is positive, such a statement specifies a loop to be executed "zero times". Negative increments are also permitted. The DO variable remains defined at completion. Transfer of control into a DO loop is prohibited (in conflict with the previous standard).

4. List-directed input and output. A form of input and output is provided which does not require an explicit format specification. The form of the external representation is determined by the input or output list item.

5. Expressions. An arithmetic expression may include subexpressions of more than one type. (If an operator has two operands of different types, the operand whose type differs from that of the result is converted before the operator is applied). A subscript expression may be any integer expression. A DO parameter may be any expression of integer, real, or double precision type.

6. Compile-time constants. A PARAMETER statement has been provided, which declares the value corresponding to the symbolic name of a constant. Such a name may be used in an expression, in a DATA statement, or in following PARAMETER statements.

7. Implicit type declaration. An IMPLICIT statement may be used to declare implicit types for variables and array names beginning with certain letters.

8. Generic intrinsic functions Many intrinsic (predefined) functions produce a value whose type depends upon the type of the function arguments.

9. Subprogram reference. Subroutines and functions may contain ENTRY statements, and subroutines may have alternate returns.

10. Array bounds. An array declaration may include both upper and lower dimension bounds; if the lower bound for a dimension is not specified the default is one. Arrays may have up to seven dimensions. The upper bound for the last dimension of a dummy argument array may be an asterisk, designating that the size of the array is to be determined from the actual argument.

11. Computed GO TO default. If the control expression of a computed GO TO is out of range, execution continues with the statement following the computed GO TO.

12. Input and output statements. The following features have been included:
- An output list may contain constants and expressions.
- An input or output statement may contain a character string to be used as the format specification.
- End and error condition control for input and output are provided.
- Tab format edit descriptors have been added.
- Direct access input and output are provided.
- A character array may be used as an internal file.
- OPEN, CLOSE, and INQUIRE statements are provided.

13. SAVE statement. Values of entities in a subprogram may be preserved during the time when the subprogram is no longer being referenced, if their names are specified in a SAVE statement.

14. Fortran character set. The apostrophe and the colon are added to the Fortran character set. The collating sequence is only partly specified.

15. Comment lines. An asterisk or a C in column 1 designates a comment line.

- Mostyn Lewis, now with Cray Research(UK)Ltd.

Ill-conditioning in FORTRAN Programs

(Reprinted from ECMWF Computer Newsletter No.4, Nov. 1977)

The following is adapted from an article that appeared in one of CDC's technical publications. In some cases, as outlined below, over-sensitivity to compiler or math library changes may be indicative of bad programming practice, or the use of an improper computation algorithm.

"The execution of a FORTRAN program is affected by a number of variables: the word-size of the machine on which the program runs, the code generated by the compiler, the presence or absence of rounding in the evaluation of arithmetic operations, the precision of the input data, the precision of storage of temporary results, the object-time routines employed, and any system-supplied data. If the output from execution of a FORTRAN program undergoes large changes when small changes are made in any single one of these variables, the program is called ill-conditioned with respect to that variable changed. In the case of interest, a program will be proved ill-conditioned with respect to the mathematical routines if traps in both old and new versions of some math routines indicate an acceptable relative error (say about 1.E-13 for single precision, 1.E28 for double precision) when run against the binary produced by the compiler, but the output shows an appreciably bigger variation.

We list some methods for detecting and eliminating ill-conditioning in programs:

1. Ill-conditioning due to inaccurate representation of input and intermediate result. This will be noticed when the program runs against different sets of data which are nearly equal, producing output differing appreciably more than input differences. Such ill-conditioning may be eliminated by increasing the precision of the data being input, or the precision of intermediate results.

2. Ill-conditioning due to inaccurate computation of arithmetical results. This is most conveniently detected by compiling the program, once with the FTN option ROUND=+*/ specified, and once without, then looking for differences in the program's output from execution. This type of ill-conditioning can arise in a variety of ways. One common way is the evaluation of a long sum to the same precision as each of its components, e.g., in the computation of inner products, variances. (The current FTN compiler inserts code inline for exponentiation with a constant fixed exponent between -16 and 16, and the multiplications are rounded or not according to the ROUND=* specification's appearance on the control card). Another common way is in the subtraction of nearly equal quantities, whose difference scales some output data. This may be important in certain iterations where the iteration count is a function of the data. Ill-conditioning arising from inaccurate differences may be checked by casting relevant variables and operations into higher precision.

3. Ill-conditioning due to the library of mathematical routines in use during execution. If the random number generator is called, check that the same initial seed and multiplier are being used. If the random number generator is not called, then changes in output across executions with two different libraries are due, in almost all cases, to one of the other forms of ill-conditioning affecting computation using a mathematical routine. To determine which routine is affected, substitute new math routines for old routines, one by one, when executing repeated runs with the same data. When the affected routine is detected, determine the occurrence affected by the ill-conditioning by replacing occurrences of that function, one by one, with calls for the function in higher precision, during repeated executions with the same data. An example of ill-conditioning associated with the library of math routines is as follows: If a program calls for computation of cosine near pi/2 (or sine near pi), and the result is used to scale output quantities, the program may be ill-conditioned. The reason is that the function has a zero at the point, but the result of the function evaluated near the point may lose significance because of the great magnification of machine error there.

4. Classes of programs known to be ill-conditioned. Included here are the programs to invert matrices, solve systems of linear equations, extract roots of polynomials, etc. An adequate literature exists to enable the program's user to draw conclusions on the degree of significance of his output, for most cases here. Certain techniques may be safer than others, even though the cost in execution time is higher. For example, the Crout reduction is preferred (for error control) to the Gauss-Seidal reduction in the inversion of matrices, although the Gauss-Seidel is quicker on all programs. Other methods of detecting ill-conditioning in such programs include replacement by equivalent expressions, and forward error analyses.

We give some examples of FORTRAN sequences where ill-conditioning can occur.

<u>Example 1</u> (Insufficient precision in a stored constant)

```
      PRINT 10, (I-1, TAN((I-1)*3.1416/180.),I=1,90)
   10 FORMAT(1H1,10X,*TANGENT FOR ANGLES 0 TO 90 DEGREES*//
     + 11X,*ANGLE*,SX,*TANGENT*//(11X,13,7X,F10.4)
```

In this example, pi is given (as 3.1416) to insufficient accuracy, since TAN has a singularity at pi. The approximation of 3.141562654 lowers the relative error of the last output item by several orders of magnitude. Instead write

```
      PRINT 10, (I-1,TAN((I-1)*3.141592654/180.),I=1,90)
```

Example 2 (insufficient precision in temporaries)

```
         REAL, SX,SXX,FCN
         SX=SXX=O
         DO 10 I=1,100000
         X=FCN(I)
         SX=SX+X
      10 SXX=SXX+X*X
         AV=SX/100000
         SD=SQRT(SXX/100000-AV**2)
         PRINT 20,SD
      20 FORMAT(11X,E17.10)
```

This example of a sequence to compute standard deviations may be ill-conditioned from two causes. First, the accumulation of running totals in SX and SXX may allow machine error to build up to significant levels. Second, if the standard deviation SD is small compared to the mean AV, significant accuracy loss may occur in the evaluation of the expression SXX/100000-AV**2 which is then small in comparison to subtrahend and minuend. A better sequence is:

```
         DOUBLE SX,SXX
         REAL FCN
         SX=SXX=O
         DO 10 I=1,100000
         X=FCN(I)
         SX=SX+X
      10 SXX=SXX+X*X
         AV=SX/100000
         SD=SQRT(SNGL(SXX/100000-DBLE(AV)**2))
         PRINT 20, SD
      20 FORMAT(11X,E17,10)
```

Example 3 (Precision loss in arithmetic operations)

```
         REAL MX,MY,VX,VY,MASS
         READ*,SPEED,APEX,MASS,NO
         DO 10 I=1, NO
         ANG=APEX*(RANF(O)*2.-1.)
         VX=.5*MASS*(SPEED*(1.-COS(ANG)))**2
         VY=.5*MASS*(SPEED*SIN(ANG))**2
         MX=MX+VX
      10 MY=MY+VY
         PRINT *,MX,MY
```

For values of APEX less than 1.E-7 in absolute value, any precision of MX is completely lost. This occurs in the evaluation of the expression 1.-COS(ANG), since for small x, $\cos(x)=1.-x^2$. However, some equivalent expressions preserve precision. For instance:

```
         REAL MX,MY,VX,VY,MASS
         READ *,SPEED,APEX,MASS,NO
         DO 10 I=1,NO
         ANG=APEX*(RANF(O)*2.-1.)
         VX=.5*MASS*(SPEED*2.*SIN(ANG*.5)**2)**2
         VY=.5MASS*(SPEED*SIN(ANG))**2
         MX=MX+VX
      10 MY=MY+VY
         PRINT *, MX,MY            ..."
```

- Richard Friedman
now at Lawrence Berkeley Laboratory

ECMWF Plotter Plans

Reprinted below is an article describing the ECMWF plotter plans which appeared in the ECMWF Computer Newsletter number 6 (March 1978). The current status of these plans is that one Versatec and the SRL hardware have been installed. Acceptance tests are planned for late January, operational use will hopefully begin in late February.

> ECMWF is buying an on-line plotting system comprising two Versatec 8122 electrostatic plotters connected to the CYBER 175 through a Logical Signal Processor (LSP).
>
> The LSP is a specialised programmable mini for signal controlling applications, such as interfacing two or more incompatible pieces of hardware (e.g. computers, peripherals, telephone system, manufacturing machines, etc.). In ECMWF's case the problem is interfacing two plotters to a computer. The instruction set is principally bit-oriented in order to set the input and output signals, 0 or 1. Level changers, purpose built hardware units, will ensure that the electronic signal level is right. Furthermore, word-oriented (16 bits) instructions are available to make a limited degree of normal data processing possible. The instruction time is 125ns, i.e. the LSP is an 8 MIPS computer! This makes it quite obvious that two fast electrostatic plotters (2 inches per sec.) can be driven at full speed. The instruction set works on two sets of one-bit registers, signal registers associated with the signal lines and the data registers used as internal storage. An additional 16-bit X-register is used for loop-counting, addressing and data-store. A hardware stack mechanism provides means for subroutine and interrupt handling. Furthermore, 16-bit word random access scratch pad memory is available to hold buffered data.
> The project of implementing this on-line system is tripartite. Sintrom Electronics, Reading, will supply the Versatec plotters. Systems Reliability Ltd. (SRL), Luton, will supply the LSP and the LSP software for receiving raster data and outputting to the right plotter. ECMWF is committed to supply the CYBER PP-driver software necessary to communicate with the LSP. Only one Versatec plotter has been installed initially, the second is expected to be available in mid-1979.
>
> \- K. Petersen

*Online Plotter Implementation

As mentioned in ECMWF Computer Newsheet 39, the format of the plot files used on the on-line Versatec plotter differs slightly from those used on the offline Varian plotter. Provided one uses the Varian Basic Software from the library VARIANLIB, ID=EWPLOT, no changes are required for either the program or the control cards. This applies also to those using the Varian Basic Software indirectly (e.g. with the contouring package or Neil Storer's GPINIT routines).

As an example, the following control cards are used for the contouring package:

```
ATTACH,LIB1,NEWCONTLIB,ID=EWPLOT.
ATTACH,LIB2,VARIANLIB,ID=EWPLOT.
ATTACH,LIB3,ECMWF,ID=EWP3.
LDSET,LIB=LIB1/LIB2/LIB3.
```

N.B. If you use absolute binary programs, or your graphics program does not use the above software, your job will not work on the Versatec, nor will it be possible for your job to continue to use the Varian offline plotter. If this causes any undue problems please contact User Support immediately.

The changeover to the new plotter is expected sometime in February. A Newsheet will give the exact date of introduction of the new service.

The rest of this article gives details of the changes that will be made.

Implications of the changed file format

As already mentioned, no changes are necessary to any programs in order to use the online Versatec plotter, (provided they use the library VARIANLIB,ID=EWPLOT). However, it is worthwhile describing the implications of the changed file format.

One important change is that the Raster file no longer contains Header information for each vertical stripe on the paper. This means that Raster files will be about 18% smaller than with the offline Varian plotter.

Another important change concerns the way Raster files will be sent to the Versatec. To understand this, consider two methods of producing graphical output with the Contouring package. In these examples a 'picture' may consist of one or several maps.

Method 1

Three'pictures' are produced, the initialisation and termination routines being called for each 'picture'.

```
CALL INITT(LVEC)
        :
 picture 1
        :
CALL CLOSS(LRAS)
CALL INITT(LVEC)
        :
 picture 2
        :
CALL CLOSS (LRAS)
CALL INITT(LVEC)
        :
 picture 3
        :
CALL CLOSS(LRAS)
```

Method 2

The same three pictures are produced but with only one call to the initialisation and termination routines.

```
CALL INITT(LVEC)
        :
 picture 1
 picture 2
 picture 3
        :
CALL CLOSS(LRAS)
```

In this method the'pictures' will always be plotted next to each other in the order in which they were generated. No other users 'pictures' will be interleaved. The disadvantage of this method is that large raster files will be produced (although they will be 18% smaller than those created with the offline Varian plotter). Large raster files are a disadvantage for two reasons:-

i) Although a large raster file takes up no more disc space than a number of smaller raster files containing the same information, a large raster file implies a large vector file. With Method 1 the same, smaller, vector file can be reused for each raster file. In other words, Method 2 needs more disc space during the creation of the graphical output.

With this method each 'picture' will be sent to the plotter as it is created (specifically CLOSS sends the raster file to the plotter queue). Consequently, 'picture 1' may well be plotted before 'picture 2' or 'picture 3' have been created. This gives rise to a disadvantage; namely the possibility that the 'pictures' may be interleaved with output from another job (but all the maps that make up what we have called 'picture 1' will be together and plotted in the same sequence in which they were generated).

It is possible that the 'pictures' generated as in the above example will be plotted out of sequence. Furthermore they may be on several pieces of paper due to the necessity of cutting out the 'pictures' of other users. Although each 'picture' will be identified with the job name as at present, this may be insufficient to uniquely identify each map. Consequently one should either provide adequate labelling via titles (e.g. routines TITLE, SUPTEXT and SUBTEXT of the contouring package) or plot a series of related maps together (e.g. 'picture' might be all the 1000mb charts corresponding to a complete 10-day forecast; the time-series of the charts would then be obvious).

ii) If plotting has to be repeated (e.g. if there is a Cyber systems crash or a tear in the paper), plotting will have to be restarted from the beginning of the raster file. This will obviously waste paper, and may affect turnround time. However, the speed of the plotter is 2 inches per second (10ft/minute or about 300cm/minute) and so the effort on turnround time should not be great.

Changes to the Varian Basic Software

(This section is only of relevance to those who use the Varian Basic Software directly).

The routine which has changed most is CLS31. This is smaller than before (764 decimal instead of 787) and will also run faster. CLS31 no longer produces an output tape but it does, of course, still produce an output file which is sent to the plotter. The documentation (Bulletin B5.2/1) should be read with this in mind. No changes are required in the users program:-

```
PROGRAM TEST(INPUT,OUTPUT,TAPE11,TAPE12)
LVEC=11
LRAS=12
CALL OPN31(LVEC)
      :
      :
CALL CLS31(XLIM,LRAS,LVEC,......)
```

The routine SLEW was provided in the original version of the Varian Basic Software to output about 6 inches of blank paper following a plot. This was needed in the original specification because the plotter did not automatically space paper between the plots. The on-line Versatec plotter will automatically space paper between plots, and the function of SLEW is therefore considered redundant.

Also, three routines - STOPS,RWDTP and ENDPLOT - were provided for controlling the tape motion, and are also considered redundant.

Routines no longer required

SLEW
STOPS
RWDTP
ENDPLOT

Dummy routines for all of these have been provided for backwards compatibility.

Changes to the dayfile

The format of the messages showing that plotting has started, and that the file has been sent to the plotter queue have been changed slightly, A new message has been added to reflect the version number of the software used. This version will be incremented every time the software is altered.

- Howard Watkins

Computer Training Courses for Member States Users

In order to help Member States use the Centre's computing facilities, training courses are being provided for Member States personnel. The first courses will be given in April at the ECMWF Headquarters in Shinfield Park. Information concerning these courses has been sent to the Member States, and nominations to attend are now invited.

This first round is devoted to four courses. In summary they are:

Course A "Introduction to the Centre'

This aims at giving Member State managers and project leaders:

- an insight into the Centre and its objectives;
- details of the computer facilities available;
- information on how to estimate computing requirements;
- an idea of the Centre's meteorological forecast plans, and products to be disseminated;
- a resume of the Centre's research plans.

Course B "Simple use of the Centre's computers.

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.

Course C "Detailed use of the Cray-1"

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 "Detailed use of the Cyber 175"

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 disks), and program preparation facilities (UPDATE, debugging etc.). Intending participants will be expected to know the basics of a Cyber NOS/BE system.

In September courses B, C and D will be repeated. It is also planned to offer a similar set of courses in 1980. For Centre staff initial versions of courses B, C and D will be run during March. Nominations will shortly be invited via the Section Leaders.

- Andrew Lea

*Documentation Changes

ECMWF Bulletins

Bulletins B1.1/1 (Naming Conventions), B2.7/1 (Cyber-CRAY Job Transmission & file staging) and B6.1/2 (Program Library Documentation Standard) were issued in December. Bulletins B7.2/1 (Cray-1 job submission and file transfer via the Cyber 175) and B8.1/1 (CRAY-1 Audit of Permanent Data sets) were removed from the sets. Bulletin B7.2/1 was then reissued in late January. (Cray-Cyber Data Conversion Utilities).

Manufacturers' Manuals

Cyber User Office Set:

We have at last received updates for the NOS/BE Reference Manual and NOS/BE Diagnostic Handbook. Unfortunately, due to CDC production difficulties, both have been sent at update level F, that is at PSR level 481, which is not yet implemented on our system (we are currently at level 473). However, differences in the manuals are minor and should not cause you problems.

- P. Prior

INDEX of still valid Newsletter Articles

This is an index of the major articles published in this ECMWF Technical Newsletter series, plus those in the original ECMWF Computer Newsletter series. Articles in the original series which are still valid will eventually all be reprinted in this Technical Newsletter, making the Computer Newsletter then obsolete. Currently Computer Newsletter number 1 to 6 can be thrown away. Some back copies of Computer Newsletters number 7 to 11 are still available, please apply to Mrs. P. Prior (ext. 355).

As one goes back in time some points in these articles may no longer be accurate. When in doubt, contact the author, or User Support (ext. 347).

*C indicates article appeared in the previous ECMWF Computer Newsletter series

USEFUL NAMES AND PHONE NUMBERS WITHIN ECMWF

		Room*	Ext.**
ADVISORY OFFICE Open 9-12, 14-17 daily		CB 037	308/309
Computer Division Head	- Rob Brinkhuysen	OB 009A	340/341
Disk Space and Permanent File Problems		AS FOR ADVISORY	
DOCUMENTATION Officer	- Pam Prior	OB 016	355
	- David Dent	OB 004	347
INTERCOM - registering new users			
	- Jean-Luc Pepin	CB 132	326
Libraries (ECMWF, NAG, CERN, etc.)			
	- John Greenaway	OB 017	354
OPERATIONS	- Console/Shift Leader	Computer Hall	334
	- Graham Holt	CB 024	306
	- Eric Walton	OB 002	349/351
METEOROLOGICAL DIVISION			
Division Head	- Roger Newson	OB 008	343
Operations Section Head	- Austin Woods	OB 021	303
Applications Section Head	- Joel Martellet	OB 011	344
Meteorological Analysts	- Ove Akesson	OB 020	302
	- Herbert Pumpel	OB 020	302
Registration (User and Project Identifiers)			
	- Pam Prior	OB 016	355
Research Department Computer Co-ordinator			
	- Rex Gibson	OB 126	384
Tape Requests	- Pauline Litchfield	Computer Hall	335/334
User Support Section Head	- Andrew Lea	OB 003	348

---

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (0734) 85411