European Centre
for Medium Range Weather Forecasts

TECHNICAL NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England

Reading (0734) 85411 Telex 847908

Number 6 - December 1979

Was YOUR data on here ?

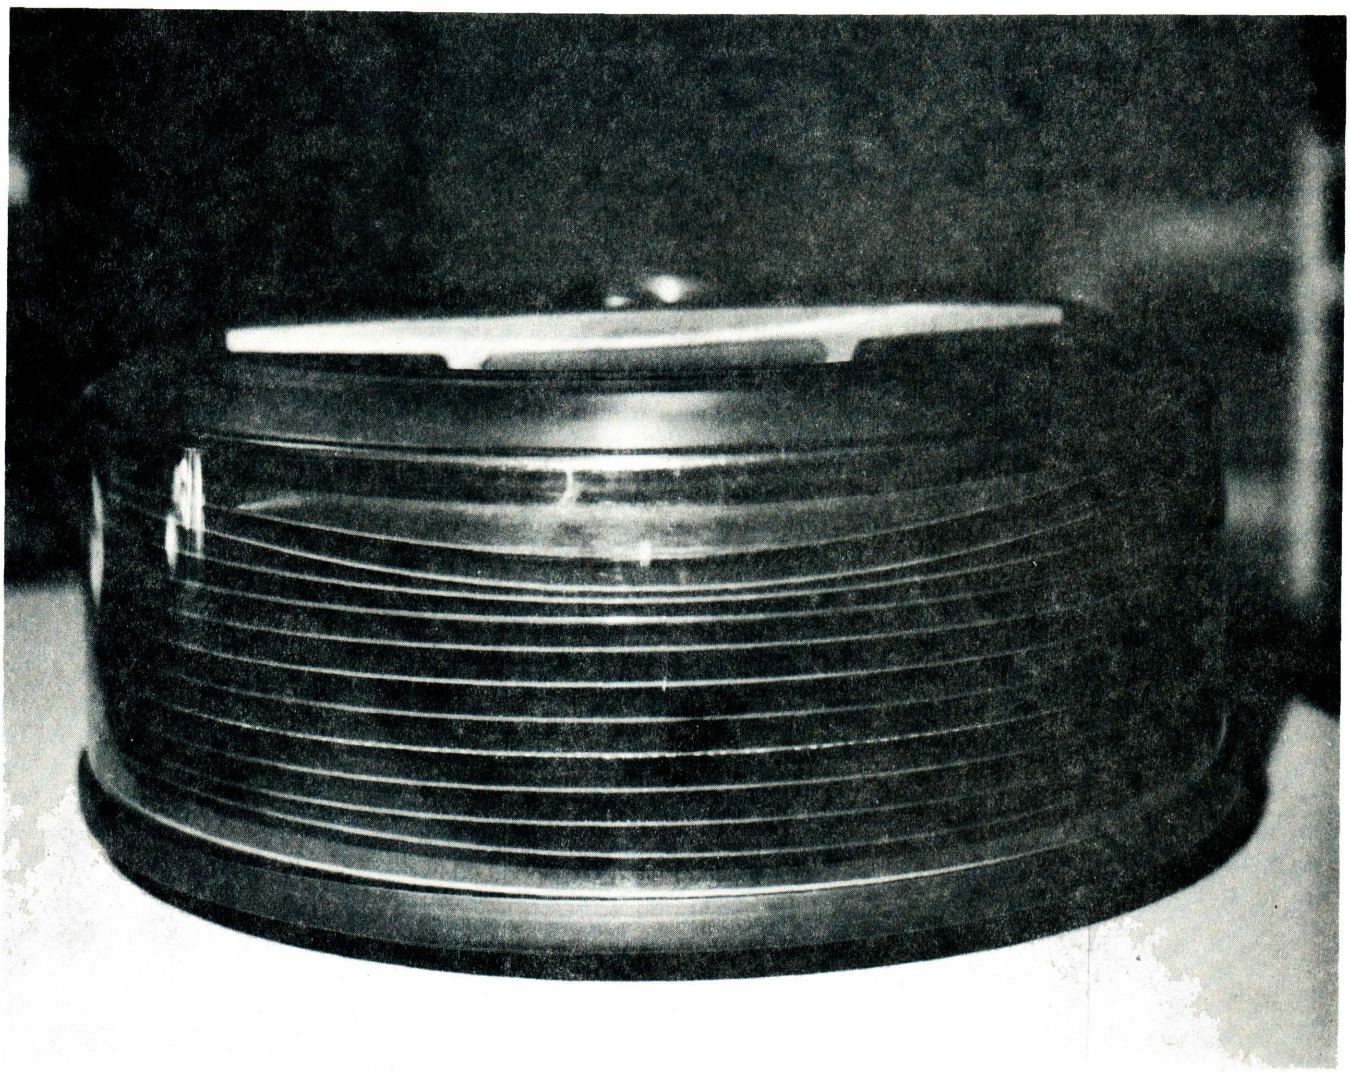

IN THIS ISSUE:

METEOROLOGICAL

COMPUTING

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

COVER: BACK UP YOUR PRIVATE DISCS!

The 'floppy' disc shown on the front cover is a Cyber private pack which was accidentally dropped. All files were lost as there was no tape copy. Backup of private packs is the owner's responsibility but a special service is offered (contact Graham Holt).

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

The next issue will appear in February 1980.

Status of ECMWF's Meteorological Operational System

ECMWF's operational forecasting system has now been running for three and a half months. It is encouraging that so far, not one planned forecast has been lost, although there have occasionally been delays of a few hours in forecast production, usually as a result of computer hardware problems. Therefore, we can say that the operational forecasting system is working satisfactorily and reliably with sufficient back-up and restart facilities included,to allow the operational system to survive the various computer failures.

At present ECMWF is receiving the input observational data in raw GTS form, copied to tape at the United Kingdom Meteorological Office at strategic times, and transported by road to the Centre. An important development in the next months will be the transfer of data over a telecommunication link between the United Kingdom Meteorological Office and ECMWF, leading to a continuous flow of input data in real-time.

Improvements to the existing system, optimisation and addition of various facilities are continuing. Some of the main work in hand:

(i) the development of more tools for monitoring analyses and forecasts such as observation plotting, cross-section display, and charts for additional meteorological parameters.

(ii) Generalised and comprehensive access to ECMWF's data bases - observational and processed data, on-line and archived - by use of directives specified on job control cards.

The overall plan for and structure of ECMWF's meteorological operational system was described in Technical Newsletter No.1. A series of articles describing the different sub-systems making up the complete operational suite in more detail will be presented in this and subsequent Newsletters. Included in this Newsletter, as part of this series, are descriptions of data acquisition and decoding subsystems, and initialisation.

J. Martellet

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

Acquisition of observational data and decoding in the ECMWF operational system

At present, as noted in the previous article, ECMWF acquires the input observational data in raw GTS form, copied to tape at the United Kingdom Meteorological Office at strategic times, and transported by road to the Centre.

In the longer term, the plan is that use will be made of the 2400 bps telecommunications link between ECMWF and the Bracknell Regional Telecommunication Hub (RTH) on the WMO Global Telecommunication System (GTS). In this way the basic data for ECMWF's operational forecasting will be acquired continuously in real-time. The format of the data will be that in which it is transmitted on the GTS but the line control protocols are not WMO protocols. The line operates under the Data Link Protocol of Level 2 of X25 (LAP-B), an End-to-End protocol based on the IFIP proposals and ECMWF's own File Transfer Protocol. This line will be handled at ECMWF by the Network Front End Processor (NFEP) which also handles the link to the Cyber Front End (FE).

For the real time data-acquisition, RTH Bracknell will send files which contain 1-250 meteorological messages each. The data will be accumulated at Bracknell until 250 messages are awaiting transmission or 20 minutes have elapsed since the last transmission, whichever is the sooner. Each of these transmission files will have a unique filename consisting of several fields. These include the transmission file sequence number, the sequence number of the first meteorological message in the file and the time of the start of transmission of the file. These fields will be used by the data acquisition program running in the FE and so the transmission file name is passed to the FE as well as the data.

When a file is transmitted to ECMWF it will initially be stored on the NFEP disc system. For input to the FE the file of data will be disguised as a job submitted from a remote site. This will be effected by inserting one logical record containing a constant set of job control statements at the beginning of each file. The transmission file name will be included on these as part of a COMMENT statement. The other logical records are of variable length and each will contain one meteorological message (from SOH to GS) in CCITT No.5 as received from Bracknell RTH. Normal data flow is shown in figure 1.

One of the job control cards will cause the data acquisition program to execute in the FE. This program already exists and works on the discrete batches on the magnetic tapes. For real-time data acquisition, it will use the sets of data passed to the FE from time to time. This program will perform only the basic functions of "capturing" the data and checking for missing data. Any statistics required are gathered at a later stage. Security of the data is vital and all files used at this point will be duplicated on completely independent mass storage devices.

The program copies an input data file of meteorological messages to the Bulletin File and the Back-up File and sets pointers to these data files in the Statistics file of the Decoding System. New files are created every 24 hours and the old ones purged when they are no longer required as back-up to the data base, by the data-base maintenance programs. For real-time acquisition, checks for missing data will consist of sequence number checks on transmitted files and messages to ensure that all the data sent by RTH Bracknell has actually been received by the FE. If data are missing the operators will be notified and they must arrange for retransmission.

The data acquisition program also controls the running of the Processing System (Decoding, Quality Control, Data-base storage). In the real-time phase, for most of each 24 hour period preprocessing of data will occur only when the acquisition phase has accumulated a sufficient number of bulletins, or at certain time intervals. The schedule governing the activation of the preprocessing will be on a disc file and can easily be modified without causing any disruption. As the cut-off time for the analysis phase approaches, the data acquisition program will start the preprocessing each time data is received. The preprocessing will be started by calls to the ECMWF Meteorological Operations Supervisor program.

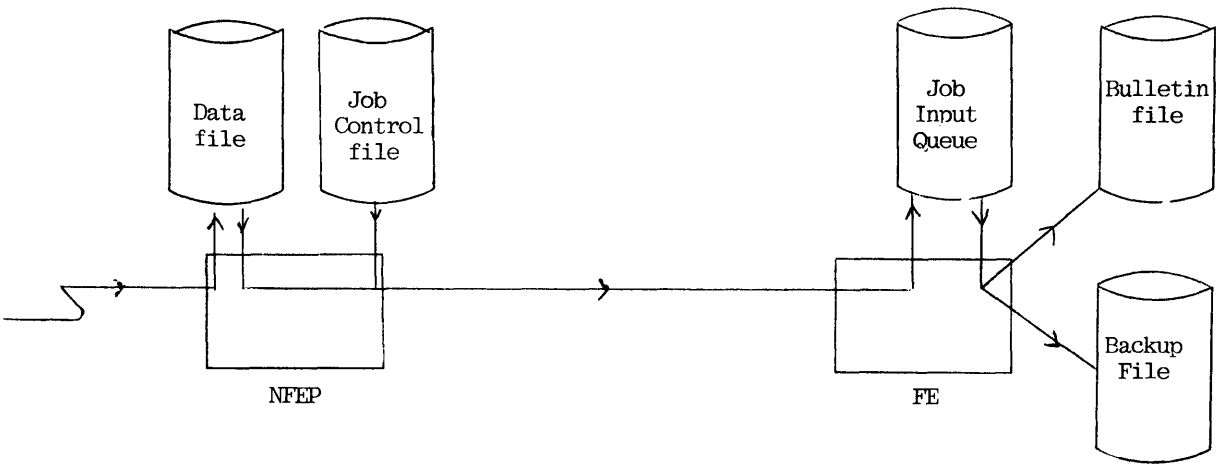

Figure 1: Normal data flow.

Decoding of GTS data

The decoding and data acquisition phases are inevitably closely linked. The Bulletin File written by the data acquisition program is the input file for the decoding and the pointers set by the acquisition phase are used to determine at which record (bulletin) the decoding program starts processing. The decoding program produces 7 files of data in floating point format for use by the Quality Control programs. These files contain types of data which correspond with the 7 input files for the analysis program. They are:

1. SYNOP, SHIP, SHRED etc.
2. AIREP, COLBA
3. SATOB
4. BATHY, TESAC, DRIBU
5. TEMP (all parts)
6. PILOT(" ")
7. SATEM

Separating the reports in this way facilitates further processing. These files are created on each run of the program and deleted when the Quality Control program has processed them.

The program also uses an Error File to which are written any bulletins or reports which cannot be handled by the program. An index file is maintained to optimise access to this file.

Statistics on numbers of bulletins and reports received are kept for each bulletin and code type. Rejected bulletins and reports are also counted along with the reasons for rejection.

The format of the input data is 5 characters (CCITT No. 5) per 60-bit CYBER word. Odd parity is used and each character has 4 leading zero bits. The first step in the decoding locates the Starting Line, Abbreviated Heading and end of each bulletin, and then identifies the bulletin type. The Abbreviated Heading is then checked for valid date and time and COR bulletins are specially flagged.

After this general bulletin format and header check each code type is handled by two special routines although these have roughly parallel structures for all code types. The first routine breaks the character strings (still in CCITT No.5) and converts them into an intermediate integer format. The second routine does the actual decoding and produces floating point values for the parameters of the report. Where floating point cannot be used (e.g. ship's name) conversion is to CDC Display Code.
e.g.:

		PPPTT		
Input format	- character	20516		
Intermediate format	- integer	205	16	2 words
Decoded format	- floating point	1020.5	16.0	2 words

These routines for the different code types use many common subroutines - character to integer conversion, knots to metres per second etc. When errors are found which cannot be handled, the original bulletin/report in CCITT No.5 is written to the Error File. See Figure 2.

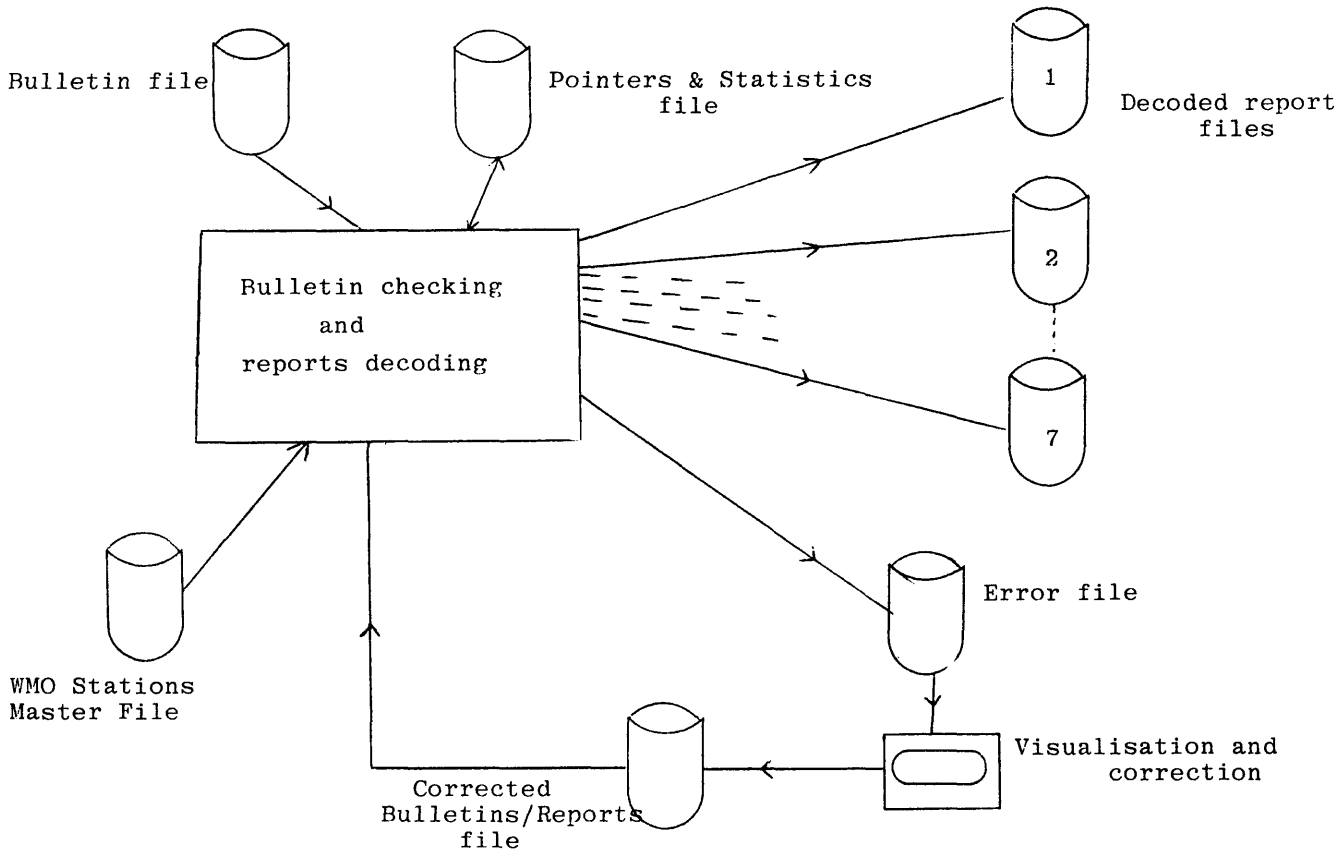

Figure 2 : Bulletin and Reports Decoding

Each decoded report has a Header which contains report type, IIiii (or ship's name, satellite number etc.), latitude, longitude and altitude (where applicable). For WMO land stations the necessary information is obtained from a direct access WMO Stations Master File. A Trailer is added to each report and gives date/time of bulletin, date/time of report, decoding flags, instrument indicators, WMO region number and country number. These fields are used by the Quality Control and Database Statistics programs.

The bulletins and reports which cannot be decoded automatically and are written to the Error File can be displayed on a VDU for manual correction. After corrections have been made the bulletins are re-submitted to the decoding process. A future article will describe the amount of data rejected at the various stages of the decoding etc.

When the decoding of each batch of data is complete the program calls the Supervisor to start the Quality Control and Database Storage programs.

- J. Hennessy

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

The ECMWF Initialization Scheme

The ECMWF analysis scheme (see Technical Newsletter No.3, June 1979) provides values of temperature, surface pressure, wind components and humidity at each gridpoint of the numerical weather prediction model (or corresponding spectral coefficients in the case of the spectral model). A forecast run directly from these initial conditions would be contaminated by spurious high-frequency gravity wave oscillations of much larger amplitude than are observed in the real atmosphere. (This problem is, of course, not peculiar to the Centre's analysis system and model; the oscillations result from imbalances between the initial mass and wind fields which are quite subtle and practically inevitable). From the point of view of the medium-range forecast itself, these oscillations may not be too serious, since they tend to die away slowly as a result of various dissipation mechanisms in the model without interacting very much with the low-frequency "meteorological" solution. However, they are likely to be detrimental to the analysis cycle itself, in which the six-hour forecast is used as a first-guess field for the next analysis. The synoptic changes over the first six hours of the forecast could well be swamped by spurious changes due to the oscillations, with the result that at the next analysis time, good data may be rejected as being too different from the first-guess field. For this reason, an initialization step is performed between the analysis and the forecast, with the object of eliminating the spurious oscillations.

The initialization scheme is based on the normal modes of the forecast model. To determine the normal modes, the model is first linearized, i.e. stripped down to a simple form which describes small-amplitude perturbations about a given basic state (in our case, a state of rest with the temperature varying only in the vertical, and zero winds everywhere). Given a set of initial conditions for this simplified version of the model, it is possible to write down the subsequent time evolution as a sum of normal modes. Each of these normal modes consists of a three-dimensional field of the model variables, with which is associated a characteristic frequency of oscillation in time. The amplitudes of the normal modes are determined from the initial conditions. The characteristic oscillation frequencies enable us to classify each normal mode as either a low-frequency "meteorological" (Rossby) mode or a high-frequency gravity mode. The slowly varying and approximately geostrophic nature of the real atmosphere implies that most of the energy should be found in the Rossby modes.

Linear normal mode initialization is carried out by determining the coefficients (amplitudes) of the normal modes from the analysed data, setting to zero those coefficients corresponding to high-frequency modes, and reconstructing gridpoint fields from the remaining low-frequency modes. If the simplified linearized model were then used for the forecast, the high-frequency modes would never reappear, and the problem of the spurious oscillations would have been solved.

However, the real forecast model is of course nonlinear, and although the high-frequency modes are absent from the initial data, they are soon regenerated by nonlinear interactions between the Rossby modes (and by the model "physics" which is omitted from the

linear model). The basis of nonlinear normal mode initialization is that it is possible to give each high frequency mode an initial amplitude such that the linear contribution to its rate of change in time just balances the contribution from the non-linear forcing; in other words, rather than removing the high frequency modes from the initial data we set their initial time tendencies to zero. As long as the nonlinear forcing does not change too rapidly, the high frequency oscillations are thereby removed from the forecast. Fig.1 shows a typical 24-hour surface pressure trace at a particular gridpoint during two forecasts, with and without a preliminary step of nonlinear normal mode initialization.

Despite its apparent complexity, nonlinear normal mode initialization is quite cheap; most of the work goes into calculating the nonlinear forcing, which in practice is done by integrating the forecast model through one timestep.

One possible cause for concern is that the initialization procedure inevitably results in changes to the analysis; as long as these are smaller than the expected analysis error then there is no harm done, but at present it is found that surface pressure changes occur on quite broad horizontal scales and with amplitudes which are a little too large for comfort. Possible ways of preventing or reducing these surface pressure changes are currently being explored.

- C. Temperton

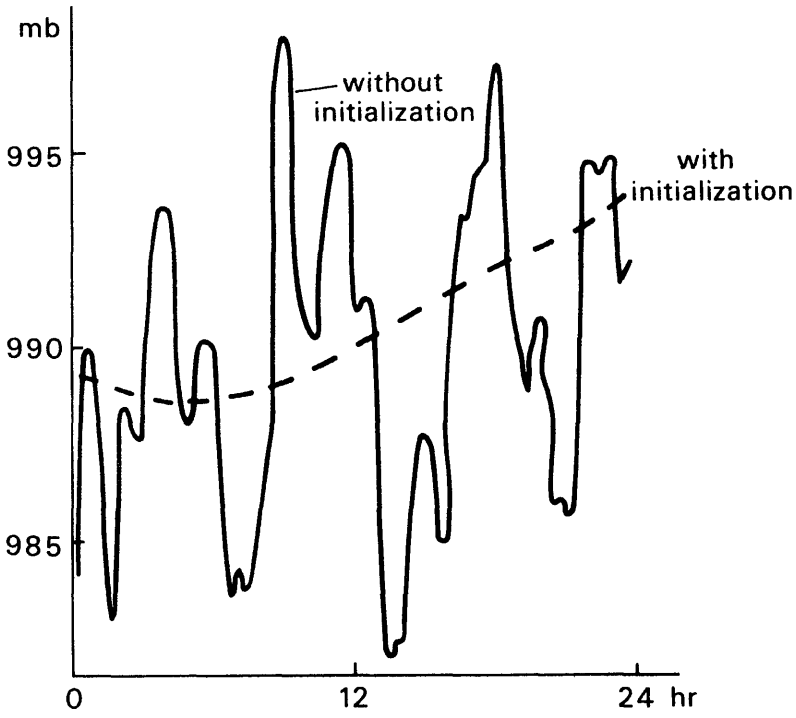

Fig. 1 Surface pressure at a particular gridpoint during two 24-hour forecasts, with and without a preliminary initialization step.

Figure 1

Computing Facility Status and Plans

The last status information about the computing facility was given in Newsletter number 4 in August. Since then, some changes have taken place and furthermore the Council, during its meeting in early November, has decided upon various recommendations concerning the computing facility and its use.

Cray-1

No major changes have taken place or will take place next year with regard to the Cray-1. At present, the system is processing 3000 to 4000 jobs/week, using around 80 hours CP time per week.

It has been decided that COS 1.0.6 and related products will not be implemented. COS 1.0.7 is planned to be implemented in Spring 1980. Another feature which will be implemented in early 1980 is a job class scheduler, which will greatly improve the scheduling facilities on the Cray-1.

Allocation of Cray resources as approved by Council is presented in another article in this Newsletter.

Cray Research has recently announced the Cray-1S, a system with the same processor as the Cray-1, memory up to 4Mw, dedicated I/O processor(s), a buffer memory, a high speed channel (850 Mb/s) into memory, and double density discs. Unfortunately, these new features cannot be applied to the existing Cray-1, although the technical possibility of upgrading the disc subsystem is being investigated (see later article).

Cyber 175

The initial Cyber configuration has been completed and the system is in full operation. At present, the system is processing about 16000 jobs and 60 CP hours per week.

In 1980, the memory of the Cyber will be incremented to the full complement of 256K words. This will not allow individual Cyber program sizes to be increased (the maximum program size remains 128K), but will provide an improved throughput, as it will allow a higher level of multiprogramming with fewer program swaps. In the autumn of 1980, half of the disc subsystem (8 of the 16 844-41 exchangeable disc drives) will be replaced by 4 855-12 double spindle non-exchangeable disc units. As the 855 drives have an increased capacity, this will approximately double the on-line data storage capacity on the system. At the same time, the number of access paths to the disc subsystem will be increased from 4 to 6. The installation of these disc upgrades will be dependent upon benchmark comparisons which are to take place from September to November 1980. The memory upgrade should take place around June 1980.

Allocation of Cyber resources as approved by Council is presented in another article in this Newsletter.

Telecommunication Subsystem

The Network Front End Processor was installed at ECMWF in June. During the past months, the contractor has completed the development work and performed exhaustive testing.

The system passed Provisional Acceptance (PA) in late October, and a nine months final acceptance period has started. (See later article).

Before PA, the planned low-speed connections to Member States were established, most of which are now used for operational purposes. Medium-speed lines have been established to U.K., Sweden and the Federal Republic of Germany. The U.K. Met. Office is well advanced in the development and check out of the software in their terminating equipment and expect to go live in December.

FRG, Sweden and Denmark are waiting for completion of the software developed for the RC3600 terminal by SIA.

The updated implementation schedule for the ECMWF network as approved by Council is presented below.

Table 1 - Updated implementation schedule for ECMWF telecommunication network

Member State	Date previously approved by Council	Revised date for medium speed lines	Low speed line (1 Aug.1979 or as soon thereafter as possible).
United Kingdom	January 1979	March 1979*	-
Sweden	April 1979	October 1979*	-
Denmark	November 1979	November 1979†	-
Germany (F.R.)	April 1979	November 1979*	-
France	April 1979	March 1980	Yes* (50 bauds)
Finland	October 1979	May 1980	-
Ireland	July 1980	June 1980	-
Italy	July 1980	July 1980	Yes* (50 bauds)
Austria	July 1981	August 1980	-
Greece	September 1980	September 1980	Yes* (100 bauds)
Netherlands	October 1981	October 1980	Yes (100 bauds)
Portugal	July 1981	December 1980	Yes (50 bauds)
Spain	December 1980	May 1981	Yes* (50 bauds)
Belgium	July 1981	July 1981	-
Yugoslavia	December 1980	July 1981	Yes* (50 bauds)
Turkey	November 1981	January 1982	Yes* (50 bauds)
Switzerland	January 1984	January 1984	-

* denotes circuit already implemented

† delayed to January 1980

Acquisition of a second RC 8000 computer (the kernel of the NFEP is a RC 8000) is planned to provide the basis for duplication of the telecommunication facilities. It will help in establishing the medium-speed lines and in the testing of terminal software of the Member States in 1980, without interrupting the general service. Further steps towards full duplication and integration of the local and remote network will be studied next year and are planned for 1981.

With regard to the local in-house terminal network it is planned to install an automatic switching system so that more terminals can be connected to the system than there are physical communication ports available. At LOGIN time a terminal is connected automatically to an available physical port. With these facilities the number of keyboard terminals will be increased to allow optimal access for the programming staff.

Graphics

Experiments and developments are planned to provide ECMWF with a versatile graphics facility. At present, all the graphics work is being processed on the Cyber, which constitutes a considerable load, and output is restricted to the plotters or Tektronix. The amount of data to be reviewed and the range of facilities required necessitates an investigation into, and experiments with, a device independent, distributed system driving a variety of devices.

- R. Brinkhuysen

Network Front-end Processor Provisionally Accepted

The end of October marks the successful completion of the Provisional Acceptance Trials for our Network Front-End Processor (NFEP). This is a Danish Regnecentralen 8000 computer which links to the Cyber 175 by a channel coupler and serves the network of telegraphic and telephone data lines to our Member States. The traffic on the telephone circuits (the "medium-speed" lines) will be governed by a layered set of communications "protocols" which conform to the latest developments in this area and are manufacturer independent. They permit a mixture of remote job entry, data dissemination, some interactive traffic, and other applications, to take place simultaneously on each line and provide scope for future applications, higher line speeds and replacement of the point-to-point connections by a public packet-switching service. A Technical Control Centre completes the system and includes line and modem switching and testing facilities.

After a tendering phase in 1977, the contract for delivery of a turnkey tele-communications system was awarded to the London based software engineering division "Ganymede" of Service in Informatics and Analysis Ltd. (SIA). Apart from their overall responsibility, Ganymede had to develop almost the entire software for this project. Regnecentralen developed the software for the data link protocol, the LAP B procedure, as it is called, of HLDC. This is one of the first working implementations of this procedure by a computer manufacturer, it was approved only two years ago as a standard by ISO (the International Standards Organisation) and CCITT (the UN sub-organisation for telegraphic and telephone traffic). A significant feature of the project is that the original specification remained more or less unchanged from design to final implementation which is unusual in projects of a similar nature. Ganymede faced a number of difficult problems, partly unforeseeable, during development and implementation. Delays were therefore incurred but the original objectives of performance and reliability were maintained and the programming standards are excellent. We are glad, with them, that the important milestone of Provisional Acceptance could be achieved and that the feasibility of a fully European solution to an advanced computer project could be proved.

Although ECMWF would now be fully prepared to provide remote usage of all planned facilities, there is no Member State yet fully ready with its medium-speed line termination. There exists a particularly close liaison between the Centre and the British Meteorological Office as this link will be used for the vital data acquisition by the Centre. The testing of this link is almost complete. A semi on-line operation can possibly start before the end of this year.

The joint project between Ganymede and the Centre and three Member States (Fed. Rep. of Germany, Sweden and Denmark) enters into the Provisional Acceptance Trials for the common software parts by the end of November. We can safely expect that Sweden will be the first full on-line user during December. With the special software for Deutscher Wetterdienst completed soon afterwards, Germany could come next by February 1980. The further sequence is at present difficult to predict.

On the low speed (telegraphic) lines, dissemination of our products had started when the Operational Suite went live, during August. Ganymede had made available the NFEP for this specific purpose before Provisional Acceptance. Routine 5-days per week dissemination has since been established to six Member States, i.e. Spain, Yugoslavia, Turkey, Italy, France and Greece. There are still Post Office problems to finalise the circuits to the Netherlands and Portugal. Feedback so far indicates that our products are used in these Member States, partly by computer, partly by manual methods. These first operational trials have required a lot of effort spent on expected and unexpected problems with the low speed line transmissions, which have been solved by the close co-operation of all parties involved.

- Fritz Königshofer

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

CPU Resource Allocation, Reporting and Control*

As the Centre's computing facilities become heavily used, especially as Member States begin to submit work, a more formal method of control is required. Beginning in early 1980, therefore, the following scheme is to be introduced:

- Council has allocated the major resources to Member States and the Centre itself for the calendar year 1980. These are detailed in a later article.
- The Computer Division will control at Member State level on behalf of Council. Within a Member State, control can be exercised down to the project level.
- Accumulated usage will be compared on a daily basis to the allocation, when the allocation is exceeded no further work will be accepted from that Member State (or project), except at base priority (see below).
- The allocations have been split 10%, 45%, 45% between the 3 user priority groups high, medium and low respectively (see ECMWF Technical Newsletter No. 4, p.13, for the definition of these priority levels). When usage exceeds allocation at a given level, no further work will be accepted at that level. Jobs attempting to run will be dropped with a message saying that no allocation is available. The job may be re-submitted at a level where an allocation remains.
- Base priority will be available where an allocation has been used up at all levels. However, work at base priority will only be run when all other work has been cleared. Although usage at base priority will not be taken from an allocation, it will be reported.
- At the beginning of each job the current state of allocations for that project will be reported.
- Reports on accumulated usage will be distributed every 4 weeks to Computing Representatives.

The various stages of this scheme will be introduced over the next few months.

- Andrew Lea

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

Cyber Permanent File Space Control *

The present scheme of allocating 100 RBs per individual user is no longer tenable, bearing in mind the differing needs of Member States and of projects within the Centre itself. Thus, the scheme outlined below will be introduced early in 1980:

- An allocation of PF space has been given to Member States by Council (see later article), this will be split into various projects. Within the Centre it will probably be best to consider allocations to Departments, and then to projects.
- The AC parameter on CATALOG will be automatically generated from the project identifier on the ACCOUNT card. Thus, every file will belong to a project and so to a Member State.
- Space occupied per Member State or Centre Department will regularly be compared with their allocation.
- On behalf of the Council, the Computer Division will control at Member State (or Centre Department) level when an allocation has been exceeded,by:
  - i) requesting the Member State or Centre Department to reduce its space occupied to below the allocation;
  - ii) giving them approximately 48 hours to do so, this period may be altered depending on circumstances;
  - iii) if no action has been taken within that period, the Computer Division will select those projects exceeding their allocation by the largest amount and dump all their files to tape until the balance has been restored;
  - iv) the Member State or Centre Department will be informed which files have been dumped to which tape. The tape will be held for a fixed period, e.g. one month, allowing users to copy files to elsewhere. The tape will then be recycled.
- The above procedure means that the tape dump taken is not an archive. Users with files which they wish to keep must move them to their own tapes. The above procedure has been adopted because it forces the user to really consider whether a file is to be kept or not. Any automatic archive scheme results in a large number of tapes holding a lot of unwanted material, usually causing the eventual breakdown of that scheme just through sheer volume of files held.

Further details of this scheme will be published in a bulletin. Dates for the introduction of various stages will be advertised via News Sheets.

- Andrew Lea

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

Cray and Cyber Allocations for 1980 *

At its 10th session, Council approved the allocation of the following resources to Member States for 1980:

- Cray CPU capacity units;
- Cyber CPU capacity units;
- Cyber on-line permanent file (PF) space.

Note that CPU capacity is allocated in units, rather than hours.

Because of the differing characteristics of jobs, it is not adequate to allocate on CPU time alone. For example, if charged on CPU time only, I/O bound jobs get an unfair advantage over CPU bound jobs. To try to charge equitably for all the resources a job might use is difficult, taking into account all possibilities. Hence, a compromise has been reached whereby only the 3 main functions are included, so leading to a resonably simple definition of this unit, as follows:

$$UNIT=K_1*CPUTIME+K_2*IOTIME+K_3*CMTIME$$

where CPUTIME is the time spent executing in the CPU.
IOTIME is the time spent performing input/output.
CMTIME is a memory/time integral, based on the memory used while executing in the CPU and waiting for input/output.
and K_1, K_2, K_3 are constants.

For both machines the constants have been set such that (excluding overheads), we use an average 10,000 units per day. For a rough comparison, taking an "average" job, i.e. not wholly CPU or I/O bound, then

- 1000 Cray units approximately equals 1 Cray CPU hour
- 1650 Cyber units approximately equals 1 Cyber CPU hour

Details of the definitions of the units are given in a bulletin to be published shortly.

The allocation of these units to Member States for 1980 is given in Table 1. Based on expected capacity, this leaves the following for use by the Centre:

- Cray CP 2,737.5 K units
- Cyber CP 3,334.85 K units
- Cyber PF space 52.9 M words

Action is underway to split this up between the various activities within the Centre.

- Andrew Lea

(see page 11 for table)

Table 1 - Allocation of Computer Resources for 1980

Member State	Allocation of units (1000's)		Allocation of PF space (Mwords)
	CRAY	CYBER	
Belgium	10	1.65	0.1
Denmark	56	26.4	3.2
FDR	255	0	10.0
Spain	25	0	0.5
France	204	41.25	10.0
Greece	0	0	0
Ireland	5	1.65	5.0
Italy	60	19.8	10.0
Yugoslavia	10	1.65	0.15
Netherlands	81	16.5	0.1
Austria	0	0	0
Portugal	0	0	0
Switzerland	3	1.65	0
Finland	5	1.65	7.5
Sweden	30	34.65	5.5
Turkey	0	0	0
U.K.	151	165	10.0
Special Projects*	18	3.3	0
TOTAL	913	315.15	62.1

* Special projects are those approved directly by Council. There is only one for 1980.

More News on User Memory Space on Cray

With effect from 1 November 1979, the maximum user program field length is increased from 3200_8 blocks to 3204_8 blocks (including Job Table Area). The JTA now occupies a minimum of 13_8 blocks ($13{,}000_8$ words) and expands according to several factors, the most important being the number and size of files being used.

The absolute maximum program field length is now 3171_8 blocks (848384_{10} words), assuming minimum JTA size, i.e. the maximum job field length remains 3204_8 blocks, so any JTA size greater than the minimum will reduce your user program field length.

It is now opportune to remind users that accurate estimations of field length use will allow the throughput of the Cray to be improved. In particular, requesting maximum possible field length is rarely necessary.

It is a feature of COS that it always automatically increases your field length when loading a program larger than the present field length in use, (that is the CM value you specify is the starting field length you will use, not the maximum). It is a deficiency in COS that it never automatically decreases your field length when loading a program smaller than the present field length. Also, there are bugs in CFT which prevent it from correctly increasing its field length when it exceeds its allocated field length.

Thus the recommendation is:

i) Do not make a CM request on the jobcard unless you know that CFT is unable to work with the default field length (200_8 blocks), In that case, request sufficient field length to make CFT work correctly (e.g. 500_8 blocks).

ii) In multi-job step jobs, use the RFL card after each user program execution to manually reduce your field length whilst processing standard JCL or when performing additional CFT compilations. An RFL card of the form:

```
RFL(M=200)
```

will normally be adequate.

- Peter Gray

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

Saving and Acquiring Multiple Files on Cray-1

It is ECMWF policy not to maintain a permanent file base of users' files on the Cray-1, due to the limited disc space available and the large demands made on that space by the operational suite of programs. User jobs requiring access to data, and jobs wishing to store data, must transfer such data to or from the Cyber 175 via the Cray-Cyber link. Many data sets of this type are large (e.g. 3.4 million words for a restart point for the ECMWF forecast program). In consequence, 6250 bpi magnetic tape is a useful storage medium for such data. In fact, one 6250 bpi tape is capable of storing data for 4 such restart points, provided all the data can be DISPOSED to tape in a single DISPOSE operation.

The restriction of DISPOSING one data set to, or ACQUIRING one dataset from, one single tape is a limitation of the current Cray-Cyber link software. Future editions of the software will contain facilities to process multi-file tapes. Until such facilities exist, there is a need for careful control in the way in which multiple files are saved and restored.

It has become common practice to copy several Cray datasets onto a single dataset using COPYD, then to DISPOSE the single large dataset to 6250 bpi tape. Should

one of the original datasets be required again, the single dataset is ACQUIRED, and broken down to individual files using COPYF. This system is excellent for archiving, as collections of files can be built up to occupy most of the available length of a 6250 bpi tape. The disadvantages of the system are associated with the recovery of small amounts of data. This can only be accomplished by an ACQUIRE of the whole tape, a COPYF sequence to extract the required sub-set of data, and a DELETE and RELEASE of the ACQUIRED dataset. The difficulties involved are:

a) ACQUIRING a nearly full 6250 bpi tape is a heavy load on the Cyber-Cray link. It results in considerable degradation of the I/O performance of both Cyber 175 and Cray-1. Several such transfers cause contention, and in practice it is found that more than two such transfers taking place simultaneously, degrade the Cray-1 to such an extent that operator intervention is necessary.

b) Tapes have often been constructed with the most likely required data as the first file. This has resulted in users ACQUIRING the whole tape and using the resulting data set as they would a data set containing the first file only. This practice causes the Cray discs to be burdened with millions of words of data which are not required, and has in the past brought about situations where other jobs have failed because of lack of disc space.

As a short term solution to these problems the Research Department at ECMWF have adopted the following guidelines:

i) Jobs which ACQUIRE more than 4 million words of data should be submitted to Cray with priority zero. Operators assume that all priority 0 jobs about which they have no notification are jobs requiring extensive link transfers, and schedule them at suitable times.

ii) Jobs which require only a sub-set of a large dataset should extract the sub-set required, and delete the remaining data.

iii) If such a sub-set of data will be required by several jobs, or repeatedly by the same job, it should be DISPOSED to either tape or Cyber disc (according to its size). This enables future use of the data with the minimum of transfer overheads.

iv) Never submit jobs simultaneously which ACQUIRE the same dataset. If several jobs ACQUIRE a dataset which is not available on Cray discs at the time the ACQUIRE command is processed, each job will originate a request for the transfer of that dataset. It has been known for up to 4 copies of a full 6250 bpi tape to be passed via the link to Cray disc because 4 jobs attempted to ACQUIRE the data before the first requested transfer was complete.

The long-term solution lies in improvements to software. Multi-file tape support in the link should overcome many problems, while the introduction of job classes could well assist in the scheduling of high link activity jobs. In the meantime, sensible use by users of the existing facilities can provide an environment in which all can benefit.

- Rex Gibson

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

The CRAY-1S (Sports model?)

In October 1979, Cray Research Incorporated announced the availability of a new range of computers, known as the CRAY-1S series, planned to be available late in 1980. Several different systems are possible, ranging from 250 K words to 4 M words of memory. Prices for these machines range from $4.8 million to beyond $15 million.

The CRAY-1S is essentially two sub-ranges of machine, each of which is entirely compatible with the CRAY-1 from the user programming point of view. The first sub-range spans memory sizes of 250K, 500K and 1 megaword and is entirely compatible in all respects with the present CRAY-1 system, although the new systems can be upgraded to the second subrange (unlike the present CRAY-1 system). The second sub-range, spanning 1 megaword, 2 megaword and 4 megaword memories has, in addition, an I/O sub-system

comprising between 2 and 4 separately programmable processors and a buffer memory ranging in size from 0.5 megawords to 2 megawords. The I/O subsystem is connected by a very high speed channel to the CPU. Parts of the operating system which currently run in the CPU, and which handle physical I/O processes, such as the station I/O package and the disc I/O package, will be run instead in the I/O subsystem. The remaining part of the operating system will continue to run in the CPU. Cray intend to offer IBM compatible channels on the I/O subsystem and will eventually provide support initially for 6250 bpi tape decks through that system to the CRAY.

CRAY also reported the availability of dual density disc drives, termed the DD29 disc and an upgraded controller, termed the DCU3 (used on systems without an I/O subsystem) and the DCU4 (used on systems with an I/O subsystem).

Machines incorporating an I/O subsystem will not in future incorporate a Data General Eclipse maintenance control unit. The Eclipse MCU has been changed to handle an 80 megabyte disc in the future.

The buffer memory, available with the new machines having an I/O subsystem, is used only as a buffer between the I/O subsystem processors and the CRAY-1. It is not used as a 2 level memory for CRAY-1 programs and cannot be addressed by CRAY-1 user programs.

Apart from the obvious advantages of more memory and more mass storage capacity, programs such as forecast models which perform large quantities of I/O should run more efficiently since most of the operating system overheads in handling I/O will be removed from the CP to the I/O subsystem.

The CRAY-1 S Series of Computer Systems

Model	S/250	S/500	S/1000	S/1200	S/1300	S/1400	S/2200	S/2300	S/2400	S/4200	S/4300	S/4400
CPU Central Memory size (64-bit words)	¼M	½M	1M	1M	1M	1M	2M	2M	2M	4M	4M	4M
Front-End Interfaces	1-3	1-3	1-3	1-3	1-3	1-3	1-3	1-3	1-3	1-3	1-3	1-3
I/O Subsystem												
I/O Processors				2	3	4	2	3	4	2	3	4
Buffer Memory Size				½ or 1M	½ or 1M	½ or 1M	½ or 1M	½ or 1M	½ or 1M	½ or 1M	½ or 1M	½ or 1M
DCU-4 Disk Control Units				1-4	1-8	1-12	1-4	1-8	1-12	1-4	1-8	1-12
DD-29 Disk Storage Units				2-16	2-32	2-48	2-16	2-32	2-48	2-16	2-32	2-48
Block Multiplexer Controllers					1-4	1-4		1-4	1-4		1-4	1-4
Block Multiplexer Channels					1-16	1-16		1-16	1-16		1-16	1-16
Mass Storage Subsystem												
DCU-3 Disk Control Units	2-8	2-8	2-8									
DD-29 Disk Storage Units	2-32	2-32	2-32									

- P. Gray
- D. Dent

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

How to avoid destroying valuable information stored on magnetic tape

On the Cyber there are two control cards for use with magnetic tape, namely LABEL and REQUEST.

In both these,if RING is specified a ring will be placed in the back of the tape, allowing it to be written to. The absence of this ring (NORING) prevents the tape being written on (NORING is default).

On the LABEL control card, there are additional parameters which specify whether or not a magnetic label is to read (R) or written (W). If W is specified on the LABEL card a label will be written to the tape regardless of any future actions concerning the tape. So if this parameter is left in, by mistake, after a previous creation, the information existing on the tape at that time will be overwritten.

In a case seen in Advisory recently, this was done from within a CRAY job doing an ACQUIRE for a file on a magnetic tape. The text field contained a LABEL card with the W parameter remaining from the previous DISPOSE. The Cyber rewrote the label and so destroyed the data the tape contained. It then sent to the CRAY an empty file!

Thus care should be exercised when using W on LABEL cards.

- J. Greenaway

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

Use of LOADPF and DUMPF

This article is reprinted from ECMWF Computer Newsletter No. 11 (Nov. 78)

As some users may be aware, the DUMPF and LOADPF utilities are described in the NOS/BE Reference Manual. However, due to a deficiency in these utilities, if two jobs attempt a dump on the same permanent file set simultaneously, one of the two automatically aborts. As the public permanent file set is dumped and archived by Operations, it is not normally necessary to do this yourself. If your file has recently been created or modified and is complicated or expensive of resources to recreate, you should copy the file onto tape yourself within the job which created the file (without using DUMPF). This should then avoid the above problem. The archiving of permanent file sets resident on private packs is wholly and exclusively the responsibility of the owner. Consequently, no archiving of these files will be done by operations. Owners of private packs should consult Tony Stanford or Jean-Luc Pepin concerning the use of DUMPF and LOADPF.

Reloading of files (LOADPF) is not as critical as DUMPF, described above. However, it is still not a straightforward procedure as it can mean the loading of multiple tapes (currently running at 5 or 6) or a search through an unsorted listing of files about 60 pages long, in order to reduce the number of tape mounts to one. Clearly, this process is not efficient, but without modification to the Dump utility (an important piece of software), no obvious remedy is available. Consequently, you are asked to take care when deleting your permanent files. If you make a mistake please do not make requests to User Support to reload files of a trivial size, unless they cannot be easily recreated. Please retype the file yourself. However, if you believe this is a lot of work and you do require a file reloaded on the public permanent file set from the archive, bring the following information (as a minimum) to User Support:

Permanent file name - PFN
Identifier - ID
Cycle - CY
Last date modified/
accessed.

We will attempt to recover it for you but cannot guarantee success, especially if it is an old file.

- J. Greenaway

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

NAG MK7 on the Cyber

As reported in the last newsletter, the NAG MK7 library documentation is available and the NAG MK7 Mini-Manual has been placed in the User Area. Since then we have received a tape containing the library for Cyber machines. This has now been made available on a TEST basis as a permanent file. Anyone who wishes to use it should include the following in their control card sequences:

```
:
:
ATTACH(NAGLIB,NAGMK7)
LDSET(LIB=NAGLIB.......)
```

Thus the only additional statement required is the ATTACH statement.

Due to the fact that MK6 is still the version of the library on the CRAY the Main manual has not been updated. Anyone wishing to consult the documentation relating to new routines should contact Advisory Office.

If anyone is interested in a presentation being given by the NAG staff on the contents of the library, would they please let me know so that we can assess the likely demand.

- J. Greenaway

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

File Cataloguing Under Intercom

Many users commonly believe that a "REQUEST,FILE,*PF." command is not required before a "CATALOG,FILE,---,ID=----." command, when working interactively. The absence of "REQUEST,FILE,*PF." qualifies FILE as SCRATCH to the system.

Everybody should be aware that scratch file space is available both on PF and Q devices. This an open door to two dangerous possibilities:

1) If a scratch file has been allocated space on a Q device, an attempt to catalogue that file would be followed by the message:

   ```
   FILE NOT ON A PF DEVICE
   NEW CYCLE CATALOG
   ```

   and the catalog would be honoured.

   An attempt to attach that same file would nevertheless fail, unless the "SN=SYSQUE" parameter had been specified in the ATTACH syntax. Also, that file would not be backed up, since DUMPF only runs against PF devices.

2) If a scratch file has been allocated space on a PF device, then a catalog would be accepted without problems, but there is a risk that the original scratch file overflows to a non PF device, so that only part of it gets catalogued.

This should discourage users from trying their luck. If the REQUEST,*PF has not been used, the STORE command would implicitly copy a file to a PF device, before making it permanent.

The only caution to be used with STORE: the file to be made permanent should not be a random UPDATE library.

- Luigi Bertuzzi

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

Question and Answer

Q. What happens if my program stops reading data from its input file before sensing a 7/8/9 card? Will my next job step find the input file positioned "after" the next 7/8/9 card?

- Eva Oriol

A. There are two answers: one for the Cyber and one for the Cray.

1) Cyber answer: It all depends on whether the last data card preceding the 7/8/9 card has been read into your input file circular buffer or not. If the last card your program reads from this buffer does not coexist with your last data card before a 7/8/9, then the answer is NO; your next job step will find the input file positioned "at" the first data card which was not YET read into your buffer, unless you skip forward to the next section with a "SKIPF,INPUT,1." control card. If the last card read by your program is held in the buffer with the last card you submitted in that same section of your input file, then the answer is YES; there is no need for you to skip the cards your program did not read, before initiating your next job step - Of course, the YES or NO to this answer depend on your buffer size and the amount of unread data your program would leave around. Isn't it clear? If it isn't, come and see me and we will talk it over.

- Luigi Bertuzzi

2) Cray answer: Answer is always NO. The 7/8/9 terminator is an end of file (EOF) in the Cray blocked dataset and is treated as a separate record. Hence if the program stops processing after reading the last data card, you are positioned ready to read the EOF in a subsequent program. To reposition use SKIPF. Alternatively, always read until you detect EOF, e.g.

```
1         READ(5,..)....
          IF(EOF(5).NE.0)GO TO 9
c         process data card
             =
             =
          GO TO 1
c         end of input data.
9         CONTINUE
```

- David Dent

Q. Why are the results I obtain from the execution of my program compiled with FTN,OPT=2, not correct, while when omitting OPT=2 everything is fine?

- U. Pilz

A. Because you are making life too difficult for the poor compiler (note: the stress you put in reading "the poor" is not my responsibility). Among the strategies FTN uses to optimise your program there is one called 'Dead Definition Elimination' (see CDC FTN4 User's Guide, page 3-4), which results in time saving by avoiding storing the result of a computation into a variable whenever the optimiser determines that the variable will no longer be needed, or that the same variable will soon be redefined. However, if the compiler finds your logic too difficult, and the decision as to whether one of your variables will be used again cannot be made,then the storing should always take place.

Well, sometimes it doesn't.

Take your program - you have:

```
        :
        :
a)      COMMON//A(1),B(1),C(1),D(1),E(1),BUF(700)
        :
b)      M1=100
        :
c)      DO 22 JL=1,32
        :
d)      BUF(M1+JL)=SQRT(....).......
e)      BUF(M1+JL)=AMAX1(BUF(M1+JL),....).....
f)  22  CONTINUE
g)      DO 23 JL=1,32
h)      PRINT*,JL,BUF(M1+JL)
        :
        :
```

This code is correctly compiled and executed at OPT=2: the 32 values you print as the contents of BUF(M1+JL) are those you would expect.

Now take a slightly different version of your program, where statement e) is replaced by

```
e1)     BUF(M1+JL)=AMAX1(C(3+M1+JL),....).....
```

this is an implicit case of redundant equivalencing. There is no EQUIVALENCE declaration statement, but you know that "location of BUF(M1+JL)"≡" location of C(3+M1+JL)", so the compiler does not even warn you. Fair enough: this code works without the optimiser. At OPT=2, this same code produces incorrect results and, after having read the simple description of the afore-mentioned Dead Definition Elimination, you understand why.

So you decide to straighten things up by replacing statements d) and e) with

```
d1)     C(3+M1+JL)=SQRT(....).....
e1)     BUF(M1+JL)=AMAX1(C(3+M1+JL),....).....
```

Oddly enough, this code produces just one correct result: the first one of your 32 elements of BUF.

It looks like a compiler bug. Especially after finding that

```
d1)     C(3+M1+JL)=SQRT(....).....
e2)     B(4+M1+JL)=AMAX1(C(3+M1+J),....).....
```

produces 32 good results!

So what? At OPT=2 a bug like this should practically be ignored. As I said: you are making the compiler's life too difficult.

This problem, submitted to CDC, would attract the following comment (more or less): the ability of the optimiser to eliminate stores (...) is limited by the availability of registers (errare humanum est). (...) Many dead definitions result from (....) redundant code. (....) The best advice is to keep program logic simple and avoid unnecessary use of (...) equivalence (...). (cf. FTN User Guide, page 3-5).

Should I say more?

- Luigi Bertuzzi

Q1 Can my program stop writing to a private pack when the device is full?

- Rauno Nieminen

A. No!

Q2. What happens to my data when they are written to a full private pack?

A. THEY GET LOST!

Q3. You are joking?

A. NO!

- Luigi Bertuzzi

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

Systems Software Section News

This article gives news about the work and the plans of the Systems Software Section. First of all, what do we do?

The section is currently 9 members strong, including the section head, and we have responsibility for the implementation, maintenance, support and enhancement of the systems software, such as operating systems, compilers etc. provided by CDC and Cray. We are also heavily involved in a few special projects which require particular systems programming experience.

Currently, the group is split into a number of sub-projects as follows:

- General Cyber Support
- General Cray Support
- Cyber Intercom and Cray Link Software Support
- Quality Assurance (i.e. ensuring software is reliable before implementation)
- Special Projects.

The majority of the section's time is spent handling day to day problems as they arise. Problems are passed to the section either from the Operations Section, or from the users via User Support. Most of the balance of the time is spent implementing corrective fixes to software generated either in-house or by the manufacturers. Both CDC and Cray provide onsite programming staff to assist in this work. (Currently somewhat less than 2 staff from each manufacturer).

Additional to the general day to day support activity, the section is currently engaged in the following items of work:

i) Upgrading the Cyber/Cray software to provide improved reliability and a few minor enhancements.

ii) Improving the Quality Assurance procedures which are undergone to ensure that software changes do not adversely affect system reliability.

iii) Planning for the Cyber hardware upgrades due in 1980. In addition to scheduling change-over procedures, benchmarking the disc configuration changes will also be necessary.

iv) Providing software to permit the production of alphanumeric microfiche output, using an external bureau.

v) Providing accounting (reporting and control) software to allow usage of the Cyber and Cray to be made according to Council's policy. This work is being jointly undertaken with User Support Section.

vi) Attempting to isolate various problems in the Cyber disc subsystem upgrades which were delivered in July 1979.

vii) Commencing work with Meteorological Operations Section on the second phase of the EMOS supervisor project.

During the forthcoming months, the following software upgrades are planned:

- An upgrade to COS 1.07, which will probably take place during February 1980.
- Implementation of a job class scheduler on the Cray, which will probably take place during March/April 1980.
- Implementation of COS 1.08, which will probably take place around June 1980.
- An upgrade from NOS/BE 1.3 to NOS/BE 1.4. This will probably take place around June 1980 and the primary purpose of the change is to provide software support for the Cyber 7155/885 disc hardware.
- An upgrade to the Cyber/Cray link software, to be implemented around January 1980 to provide improved reliability and a few minor enhancements.

More information about each of the above upgrades will be provided closer to the time.

In addition to the above work, effort will be put into the following major projects:

- Extending the accounting software to provide better facilities.
- Providing operator control facilities for Cray jobs, from the Cyber.
- Work on phase 2 of the EMOS supervisor project.
- Miscellaneous minor software changes and improvements.
- Further improvements to Quality Assurance procedures.

- Peter Gray

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

Saving Time and Reducing Memory Requirements for Graphics Jobs

As indicated in Newsletter No. 5, users are urged to use the new libraries of graphical software as soon as possible. Since the introduction of the Versatec Electrostatic plotters, some changes have been made to the new libraries to enable user jobs to run faster and to require less memory. These changes are detailed below, for more information users are referred to the Contouring Package User's Guide which has recently been republished (Bulletin B5.2/3).

Time Saving

The table below shows some times taken to plot examples from the User Guide. The columns correspond to:

- Version 101 with the original coastline data base (BACKBASE)
- Version 101 with the new coastline data base (COASTLINES)
- Version 106 with the new coastlines.

Version 101 was the first release after the introduction of the Versatec. Version 106 is the current production version. The times (in seconds) correspond to the usage of both the Contour Package and the Varian Basic Software and are for the complete INITT-CLOSS sequence.

	#101 BACKBASE	#101 COASTLINES	#106 COASTLINES
Example 1	3.0	1.9	1.6
2	3.0	1.7	1.5
3	2.7	1.9	1.7
4	6.0	4.7	4.4
9	6.0	4.6	4.4
10	3.5	2.1	1.9
11	3.4	2.0	1.8
Average	3.9	2.7	2.5

As can be seen from the above table, a major speed-up comes from using the new coastline data base

```
ATTACH,TAPE90,COASTLINES.
```

This is because the drawing of the coastlines on a map is quite time-consuming, and the new database has only 8,333 co-ordinate pairs compared to 22,383 for the old data-base. Apart from saving time, the new coastline data-base has other advantages such as no political boundaries, no ice-limits, fewer errors on the coastlines, etc.

Using the Segment Loader

For the examples shown above, the use of segmentation has virtually no effect on the CPU times, but does reduce the memory requirements. However, the actual memory saved will vary from user to user because it depends on how many routines of the Contouring Package he is using. A standard set of directives is available, but they assume that all routines of the Contouring Package will be used. However, the aim is to show by example how segmentation can be used, and it is suggested that users build their own segment loader directives based on the standard set:

```
ATTACH,SEGDIR.
SEGLOAD,I=SEGDIR.
LDSET,LIB=CONTLIB/VARLIB/ECLIB.
```

Note that the same directives may be used for the Tektronix.

For the examples in the User's Guide, the memory requirement (when using VARLIB) is reduced from 154000B to 123100B, saving 30700B.

Plotting Stream Diagrams

A new routine STRM1 has been introduced to enable STREAM diagrams to be plotted row by row (or column by column - the three matrix layouts definable via MATPOS can be used for stream diagrams). This could avoid having to store the entire matrix in memory prior to calling the stream routine, but is obviously only applicable in cases where the matrix can be read or computed row/column by row/column.

STRM1 is described on page 8-6/1* of the User's Guide. (* Not yet printed - will be distributed with next set of updates to Contour Package User's Guide.)

New Contouring Routine

In addition to the routine CONTOUR there is now another routine called FASCON. This is described on page 8-5/1 of the Contour Package User's Guide. The difference in appearance on a map can be seen by comparing examples C.6 and C.7 of the User's Guide; the major change is the treatment of contour labels.

The table below shows some timings from FASCON and CONTOUR. The numbers after FASCON indicate the increments for searching along each dimension of the matrix.

Time for CONTOUR with smoothed contours	3.2 secs
Time for CONTOUR unsmoothed	2.9 secs
Time for FASCON, 1x1	2.6 secs
Time for FASCON, 2x2	1.5 secs
Time for FASCON, 3x3	1.2 secs.

New Matrix Layout

A third matrix layout is now available for Contouring and Stream diagrams (see chapter 3 of the new User's Guide, and routine MATPOS, page 6-15). This third layout is the same as that used by the Forecast Suite; users therefore no longer have to exchange rows/columns or transpose the matrix before contouring or drawing stream diagrams. A second "work-space" array is thus no longer needed.

Using the Background File

Another method of saving time is to generate a background file and merge it in with other information. Such a background file might typically be coastlines, lat-long grid etc, and would be merged in with changeable information such as contours. Further information on the generation and usage of background files can be found in Chapter II of the User's Guide. The following timings exhibit the saving possible with background files:

	#106	
Times:	1.8 secs	generation of background for coastlines
	2.8 secs	using background with a contour field
	4.4 secs	same picture as above, but using coastlines directly.

Obviously, if only one map is to be produced, nothing is gained by using background files - but time savings can occur with as few as two maps. (1.8 + 2.8 + 2.8 = 7.4, compared to 4.4 + 4.4 = 8.8).

Vector-to-raster Buffer

When using the Varian Basic Software (i.e. when using library VARLIB, but not when using library TEKLIB), a large workspace is needed for the vector-to-raster conversion. This workspace is stored in common block /COMRAS/ and is 9,900 words long.

This common block contains no variables or constants needed by either the Contour Package or the Varian Basic Software. It is only used as workspace by routine CLS31 of the Varian Basic Software, which is called from routine CLOSS of the Contour Package. Thus any data put into /COMRAS/ by the user is safe from over-writing by all calls except CLOSS or CLS31.

This means that large arrays for contouring, stream diagrams, I/O etc. can be used by making reference to /COMRAS/. An easy way to do this with an existing program is via EQUIVALENCE statements.

```
COMMON /COMRAS/ DUMMY (9900)
DIMENSION CON(49,193), MAT(400)
EQUIVALENCE (CON(1,1),DUMMY(1)),(MAT(1),DUMMY(9458))
:
:
CALL CONTOUR (CON,49,193,O,IERR)
```

N.B. /COMRAS/ is not mentioned in the standard set of directives (SEGDIR) for the segment loader. This means that it is not present when using segmentation on the Tektronix, and is in the same branch as CLOSS when using the electrostatic plotters. If a user wishes to make use of /COMRAS/ he should make it a GLOBAL common block in order that it resides in the root segment. However, this will increase the memory requirements somewhat; in the case of the User's Guide examples, the memory requirement goes up from 123100B to 134000B (cf. 154000B without any segmentation).

- Howard Watkins

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

Debugging Graphics Jobs

MANTRAP symbol tables corresponding to the various graphics libraries are now available. Users who experience problems with their graphics jobs may find these symbol tables useful to find the cause of the error. The example below shows how to use the symbol tables; the libraries CONTLIB, VARLIB, etc. are the same as normally used, but note that segmentation should not be used with MANTRAP as MANTRAP is unable to decide which segment is loaded.

```
ATTACH,VSYM,VARZZZSYM.
ATTACH,CSYM,CONTZZZSYM.
ATTACH,CONTLIB.
ATTACH,VARLIB.
      :
      :
      :
FTN.....
COPY,VSYM,ZZZZSYM.
COPY,CSYM,ZZZZSYM.
LDSET,LIB=CONTLIB/VARLIB/ECLIB.
LGO.
```

Similarly, for the Tektronix a file TEKZZZZSYM is available, but it may not be possible to run the job without segmentation.

- Howard Watkins

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

Library Changes *

CERNLIB routines used by the graphics libraries have now been incorporated into ECLIB. Therefore, there is now no need for graphics users to specify CERNLIB.

You are reminded that the following library names will disappear after 31st December:

```
ECMWF, ID=EWP3
VARIANLIB, ID=EWPLOT
NEWCONTLIB, ID=EWPLOT
ECMWFLIB, ID=DUMP00
CONVLIB
CRAYLIB
```

- David Dent

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

STATISTICS

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

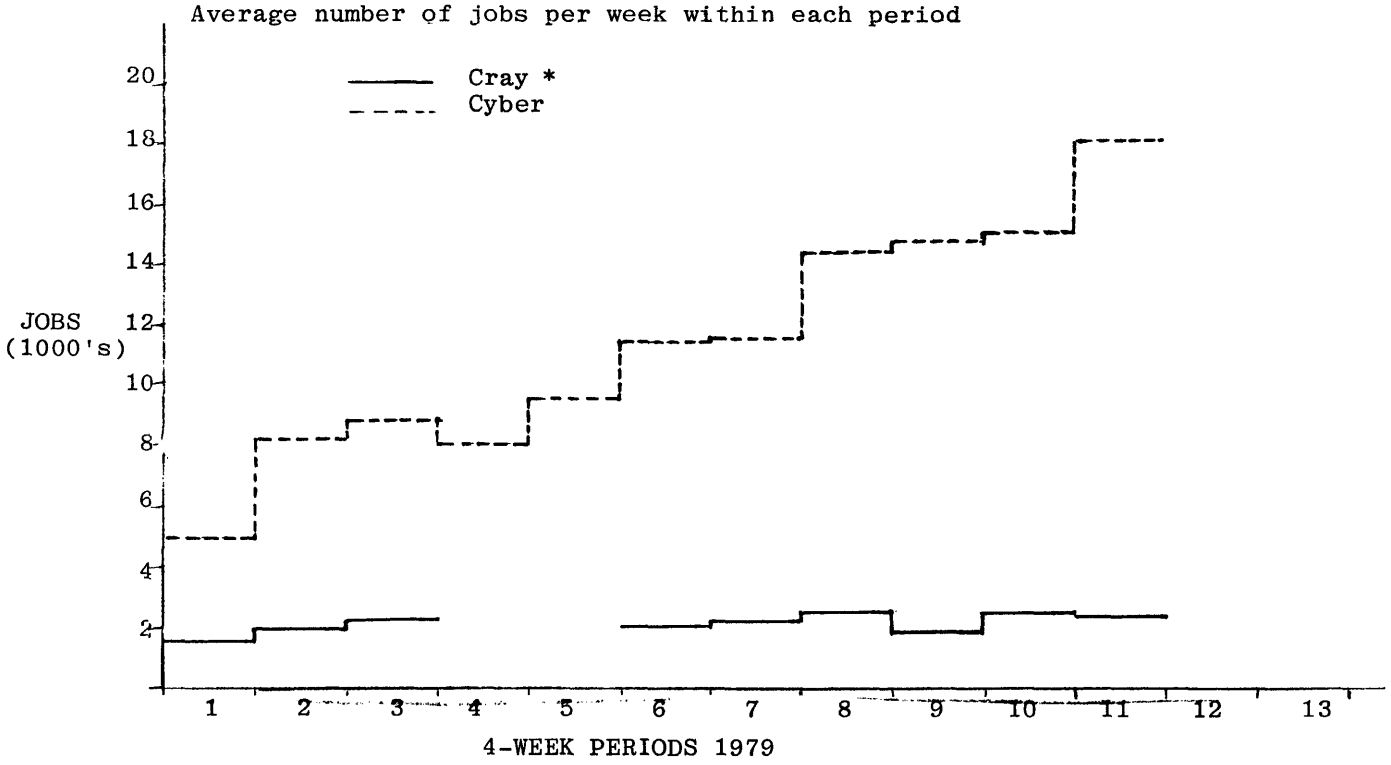

Average CPU hours used per week within each period

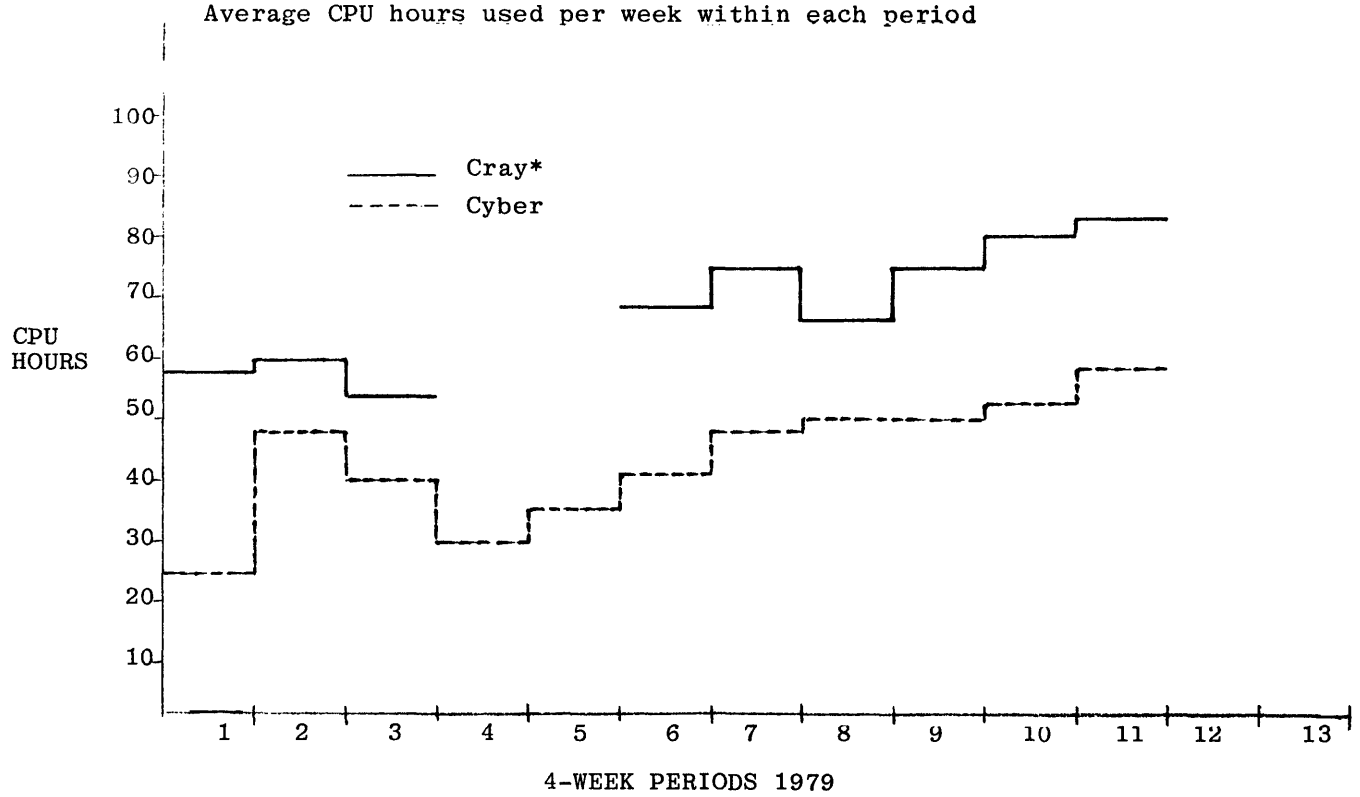

* CRAY statistics for periods 4 and 5 are not yet available.

Documentation *

Our paper output has been particularly high since the last newsletter! We issued the following new bulletins to internal Centre staff:

B0.1/1 (ECMWF Computer Division Management and Personnel List)

B0.2/2 (Levels of Software Support)

B7.1/2 (Cray Random I/O Package)

B1.0/1 (Introductory Guide to the ECMWF Computer Facility)
This is a completely new issue, though users who came on our computer courses here may have seen copies of our first draft issue.

Completely re-written versions of the following bulletins were also distributed:

B2.2/1 (Introduction to Control Statements on the Cray-1)

B5.2/1 (Varian Basic Software)

B5.2/2 (ECMWF Contouring Package)

B5.2/3 (Contouring Package User's Guide) - in re-issuing this document we have dispensed with the old pink cover and included it in the computer bulletin set, so that it re-appears in a blue binder.

Finally, we have sent out copyright pages to be inserted at the front of each of the binders, to conform to Centre documentation standards, a list of bulletins already produced, and title sheets for those large general bulletin binders which were issued some time ago.

Two sets of the above bulletins,plus one copy of the ECMWF Front-End Subsystem ERS to be included in the Reference set of manuals,were despatched to Member States in mid-November.

- Pam Prior

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

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 66). 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 & 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

News Sheets which can now be thrown away since this list was last published are numbers 23, 45 and 48.

- Andrew Lea

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

A Key to Scientific Research

What he said	What he meant
While it has not been possible to provide definite answers to these questions....	The experiment didn't work out, but I figured I could at least get a publication out of it.
The operant conditioning technique was chosen to study the problem.....	The fellow in the next lab already had the equipment set up.
Agreement with the predicted curve is	
excellent	fair
good	poor
satisfactory	imaginary
fair	imaginary
It is clear that much additional work will be required before a complete understanding.....	I don't understand
Unfortunately, a qualitative theory to account for these results has not been formulated.	I can't think of one and neither can anyone else.
Correct within an order of magnitude.	Wrong
Thanks are due to Joe Glotz for assistance with the experiments and to John Doe for valuable discussion.	Glotz did the work and Doe explained what it meant.

ACKNOWLEDGEMENT TO:

Milton Hodge, Univ. of Georgia
from the American Psychologist

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

INDEX

INDEX of Still Valid Newsletter Articles

This is an index of the major articles published in this ECMWF Technical Newsletter series. Articles in the original Computer Newsletter series which are still valid have all been reprinted in this Technical Newsletter, making the Computer Newsletter obsolete. Therefore, all Computer Newsletters can be thrown away.

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.

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	340/342
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 Officer	- Pam Prior	OB 016	355
Libraries (ECMWF,NAG,CERN,etc.)	- John Greenaway	OB 017	354
METEOROLOGICAL DIVISION			
Division Head	- Roger Newson	OB 008	343
Operations Section Head	- Austin Woods	OB 107	406
Applications Section Head	- Joel Martellet	OB 011	360
Meteorological Analysts	- Ove Akesson	OB 106	380
	- Veli Akyildiz	OB 104A	379
	- Horst Böttger	OB 104A	378
	- Rauno Nieminen	OB 104A	378
	- Herbert Pumpel	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
Tape Requests	- Pauline Litchfield - George Stone	CB 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