European Centre
for Medium Range Weather Forecasts

TECHNICAL NEWSLETTER

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

Number 2 - April 1979

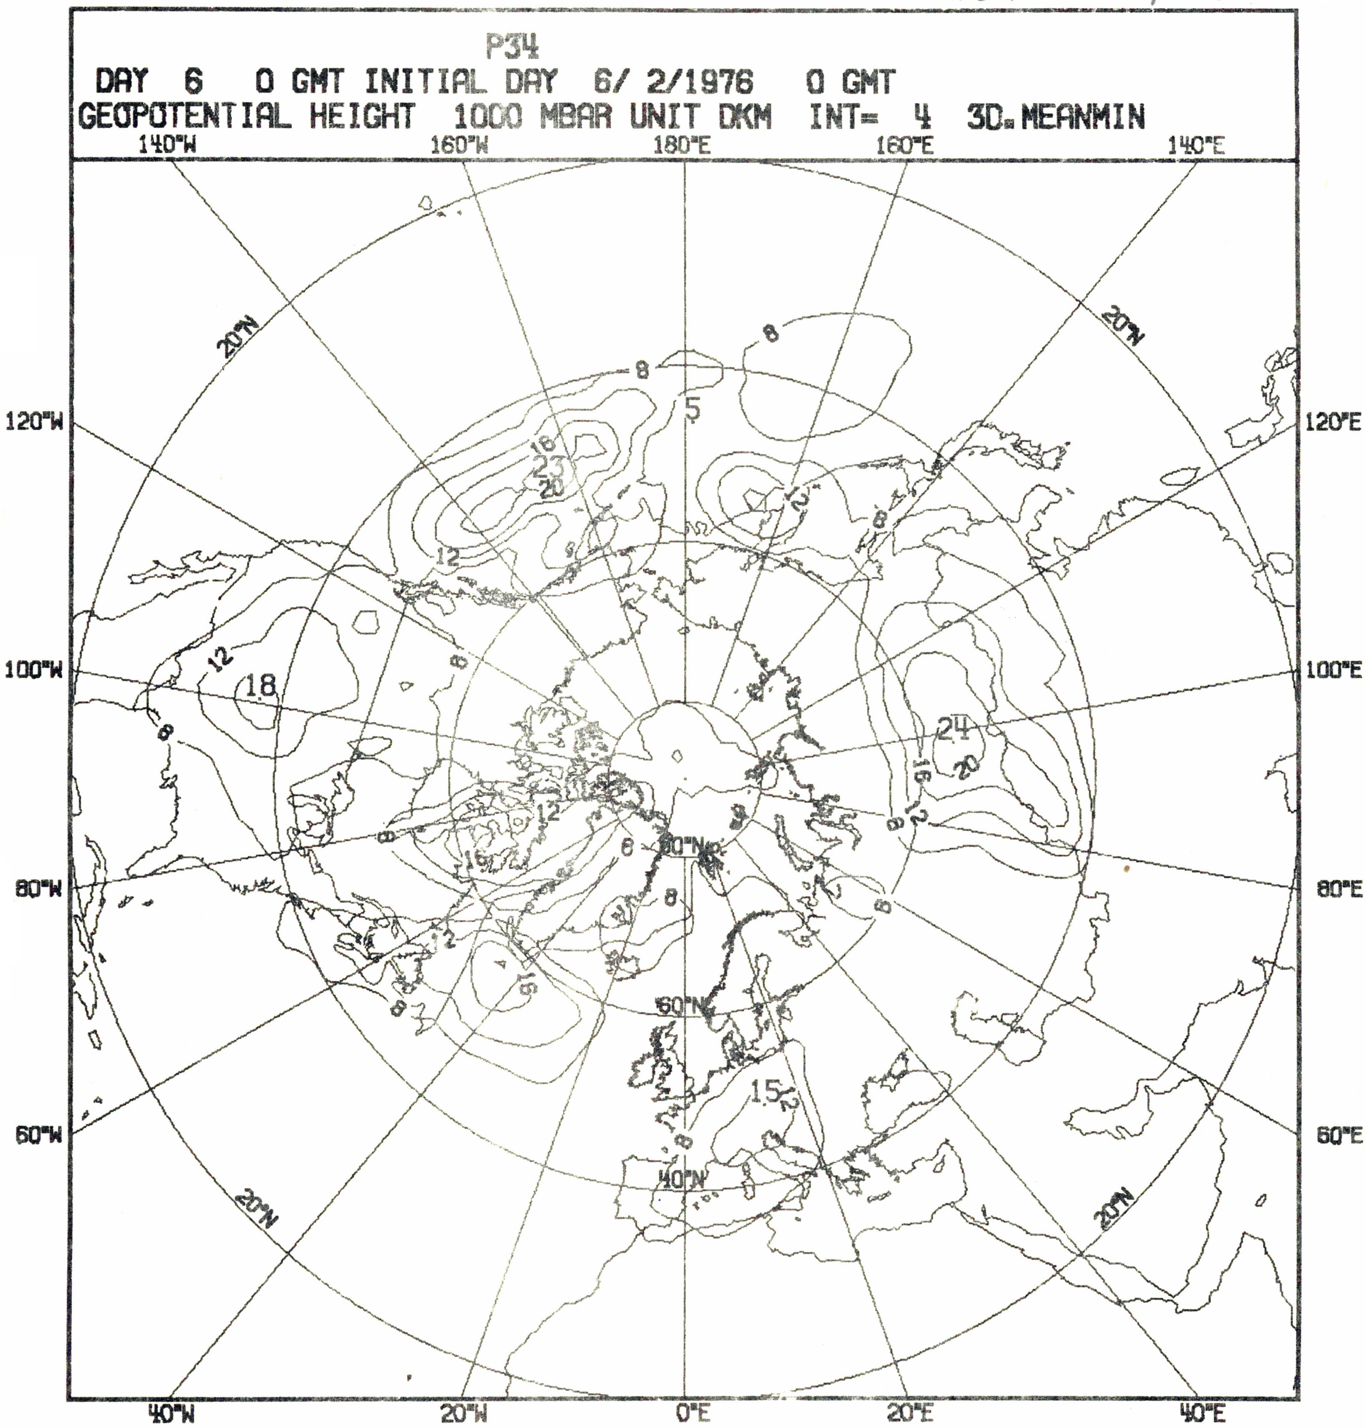

IN THIS ISSUE:

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

COVER : A three-day "mean-minimum" 1000 millibar geopotential height chart centrec on day 6 of a 10-day forecast. See article by Austin Woods and Herbert Pumpel beginning on page 4.

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

The next issue will appear in June.

INTRODUCTION

In the February issue (Number 1) of the ECMWF Technical Newsletter, the Centre's planned operational system (EMOS) was described, together with the work of the Meteorological Operations Section who are concerned with real-time and longer term monitoring of the analyses and forecast results. This month we try to give an idea of features of forecast results that might be available in the first phase of operations. This is based on an assessment of the results from a trial series of forecasts carried out by the Research Department during model development. This article has been written by Tony Hollingsworth of the Research Department. To complement this, the second article illustrates various ways of looking at the forecast results, especially applicable for later periods of the forecast, with a view to presenting the meteorologically valuable and interesting content in a concise way. The methods shown help to eliminate the small scale features from the presentations which we realise are not well predicted, but which can obscure predictability in the larger scale.

- Roger Newson

****

FEATURES OF FORECAST RESULTS IN WINTER CASES

Introduction and Summary

This note is a summary of a forthcoming technical report, where we consider the results of a series of forecasts on eight weather situations from February 1976 using two models which differed only in their physical parameterisation processes. One set of processes was developed at the Geophysical Fluid Dynamics Laboratory (GFDL) some years ago, the other at ECMWF more recently. The resolution of the models was fairly close to what may be used in the first phase of operations. The particular aim of the experiments was to study the importance of the differences in the parameterisation schemes for the models, but a general view of features of the forecast results that might be available in the first phase of operations was also obtained.

In summary, both models gave similar results in terms of forecast quality. When measured by the standard objective methods, the range of predictability seemed to be about 5 or 6 days; however a subjective synoptic evaluation could sometimes find useful indications as to the evolution beyond this time in this particular set of winter cases. A study of the systematic errors in the forecasts showed that these were mainly in the largest waves and the evolution of these systematic errors appeared to be roughly linear in time. This is not to say that the systematic errors arise due to linear mechanisms. Regarding transient phenomena, the downstream intensification of baroclinic waves appeared sometimes to be predictable as far as nine days. A study of the energetics showed that the major part of the loss of energy in the long waves was due to a failure to maintain the stationary part of the long wave energy.

The Models

It is not the intention in these notes to give a full technical description of the models used in the experiments. These are fully described in scientific literature and recent and forthcoming ECMWF Reports. Some aspects of the models may be the subject of future articles in these Technical Newsletters. The main features are that in the horizontal an enstrophy conserving finite difference scheme was used, and a semi-implicit time-stepping scheme. The basic grid was a latitude/longitude system with grid lengths $\Delta\lambda=\Delta\theta=1.875^{o}$. There were 15 vertical levels distributed as given below (using $\sigma=P/P_{surface}$ as the vertical co-ordinate).

Level No.	1	2	3	4	5	6	7	8	9	10
σ	0.025	0.077	0.132	0.193	0.260	0.334	0.414	0.500	0.588	0.678

Level No.	11	12	13	14	15
σ	0.765	0.845	0.914	0.967	0.996

The aim was to give sufficient resolution in the troposphere, near the lower boundary, and near the tropopause.

Initial data for the experiments

The initial data were taken from the global analyses for February 1976 made by NMC Washington during the "Data Systems Test" (DST) Experiment. This was supplemented by 10 mb analyses for the Northern Hemisphere available from the United Kingdom Meteorological Office and climatological 10 mb data for the Southern Hemisphere. Sea surface temperatures were taken from February normals; the topography used in these experiments was the (fairly smooth) field originally prepared by Berkofsky and Bertoni. The data were initialised using the technique of non-linear normal mode initialisation.

Initial dates for the experiments were 3rd, 6th, 9th, 12th, 15th, 22nd and 25th February.

Synoptic evaluation

By and large the synoptic evaluation of the experiments indicated that five to six days was about the time for which the forecasts were directly useful over Europe and the Atlantic. However useful information as to the general synoptic trends was sometimes available beyond this. In particular, the phenomenon of downstream intensification of baroclinic waves was forecast successfully to 8 or 9 days, on occasion. Overall it was noted that the quality varied from one forecast to another, and, for any given forecast, from one region to another.

Objective verification

The scores used for objective verification were root mean square (rms) errors of height and anomaly correlation coefficient of height for the troposphere north of 20°N. The anomaly correlation coefficient is the correlation between forecast and observation after the monthly normal has been subtracted. This measure has quite different properties from the tendency correlation, where the initial data is subtracted. The tendency correlation is used for studying short range forecasts while the anomaly correlation is more suitable for the study of medium range forecasts. Rms errors of height and wind, tendency correlation and S_1 skill score for an area covering Europe were also calculated.

There are no generally accepted rules for the interpretation of these scores. A criterion often used is that useful skill is much degraded when the hemispheric rms errors reach the level of the climatological variance or when the hemispheric anomaly correlation coefficients fall below 0.6. In the forecasts, the average anomaly correlation coefficient for the troposphere north of 20°N for the whole field and for wavenumbers 1 - 3 and 4 - 9 fell to 0.6 between five and six days. The rms errors for the total field and for the wave numbers 1 - 3 each reached the level of the climatalogical variance at about the same time. At 500 mb the scores reached these levels also between five and six days, but they reached them about one day sooner at the 1000 mb level.

The performance of both models, i.e. with GFDL and ECMWF physics, was very similar on average. One particular feature of interest in the results was that the model with ECMWF physics had skill (in the sense of rms error less than the climatological variance) in predicting 850 mb temperatures over Europe for six or seven days.

Systematic errors

By systematic errors is meant the difference between the average of the n^{th} day of the forecasts and the corresponding analyses. The evolution of these errors can be studied as n increases. For sufficiently large ensembles of forecasts one can speak of the averages of the forecasts as the "day n climatologies"; however the present ensemble of eight cases is too small to justify the use of such a term.

Even though the ensemble is small the averaged analyses for the eight cases showed a good deal of similarity to the monthly normal. Striking differences between forecast and analysis were the displacement of the Icelandic low into the Norwegian Sea and the over-intensification of the Aleutian low. The Siberian high was well maintained. Results for 500 mb had a similar character. Since differences between averages over several cases are being considered, the difference fields show energy only in the longest waves, and a study of these fields shows that the forecast errors were mainly in wave number two. From examination of the change of coefficients of the individual spherical harmonics, the growth of errors appeared to be approximately linear in time. The causes of the errors are not clear. The linear behaviour, with time, of the coefficients of the harmonics making up the error suggests a shortcoming in some forcing effect such as the interactive with topography. On the other hand the position of the maximum errors, on the eastern sides of the oceans, suggests that the over-development of surface lows, particularly over the oceans, which has also been noticed as a persistent feature in the forecasts, may make an important contribution to the systematic errors.

Energetics

For the purposes of discussion the average of the energy (kinetic or available) in the forecasts is referred to as the total energy, the energy in the ensemble as the "S energy" (for "should be stationary") and the difference between the two as "T energy" (for transient).

As the forecasts progressed there was a steady growth in the energy of the zonal flow, a decline in the total long wave energy (waves 1-3) and a growth in the energy of the intermediate waves (waves 4-9). The spectrum between wavenumbers 10-20 appeared to follow a k^{-3} law.

The drop in the total long wave energy seems to be due to a fall in the S energy. Indeed the S part of the flow is a major contributor to the available potential energy and to the energy conversion terms in this wave band. The level of energy in the baroclinic wave band became rather too high in the course of the forecasts; this was particularly marked at low levels, where the disturbances tended to over-develop.

The largest differences between the forecasts and analyses are found in the treatment of the zonal flow. In mid-latitudes the westerly zonal flow tended to accelerate by about 4m/s at all levels over the ten day forecast period and this effect was found in both versions of the model. At high levels in the subtropics the zonal flow was decelerated by up to 8 m/s in the runs with the ECMWF physics, but was slightly accelerated in the runs with the GFDL physics. The net result was that in the forecasts with ECMWF physics the main tropospheric jet tended to move northwards; in the forecasts with GFDL physics the jet remained in place but increased in speed.

Differences between the models

The differences between the models grew slowly in the first five days and more rapidly thereafter. The anomaly correlations between the forecasts from the two models dropped to 0.6 and the rms error reached the climatological norm in about ten days. This would indicate that if one model were used as a predictor for the other it would be quite good up to five or six days but it would be rather poor on average by the tenth forecast day.

- A. Hollingsworth

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

AN EXAMPLE OF PRESENTATION AND ASSESSMENT OF A MEDIUM RANGE WEATHER FORECAST

An article in the first ECMWF Technical Newsletter (February 1979) described the work planned for the Meteorological Operations Section at ECMWF. One particular aspect mentioned was the plan for developing chart formats suitable for use in the context of a medium-range forecast, which will help to exploit the very large amounts of information available from the Centre's analyses and forecasts. Meteorologists are familiar with the representation of many different parameters on weather charts, for example, geopotential, temperature, vorticity or stream-function charts in the case of short-range forecasts, or charts showing expected areas of above- or below-normal precipitation or temperature for monthly forecasts. Monthly mean charts are useful in the case of climatological studies. As with these specific types of charts for particular purposes it may be necessary to consider other methods of presentation of meteorological information appropriate for the results of medium-range weather forecasts.

First, it is worthwhile to consider for a moment the particular properties that we expect in a medium-range forecast (which for the Centre's purposes is defined as a four-to-ten day forecast). In the case of a short-range forecast - up to 72 hours - one is normally most interested in forecasting the development of existing synoptic-scale features on the chart. The monthly forecast concentrates on expected differences to climatology in the very large-scale features on the chart. In a medium-range forecast, however, we are extending our forecasts beyond the short-range limit, whilst still attempting to forecast in a deterministic way. For the range of prediction beyond a few days, the physical limit of the deterministic predictability for smaller scale disturbance (normally only 3-4 days) is exceeded. There is no <u>true</u> solution in a deterministic sense for the smaller scales of motion, and such systems are likely to be predicted only in a statistical way; thus in the later stages of a medium-range forecast only large-scale areas of activity (e.g. areas of cyclogenesis and preferred cyclone paths) would be indicated. Nonetheless the small-scale systems which are predicted will have the correct space and time-scales (and look like real meteorological systems) but will not necessarily be correct in timing or position. The implication is that many of the features of interest in a forecast for 4-10 days could perhaps be presented in, for example, the form of cyclone tracks (rather than the exact positions of individual fronts or synoptic-scale features), regions of expected cyclonic activity (deepening, filling or moving of lows) or changes in weather pattern, such as breakdown or establishment of a blocking high or of a mild (or cold) outbreak.
If then we are looking to the medium-range forecast for a somewhat different range of information from that in a short-range forecast, perhaps it may be desirable to try to develop charts which will complement traditional meteorological charts and which will help the forecaster to extract the meteorologically meaningful information from the forecast.

During the Centre's 1978 Seminar "The Interpretation and Use of Large-Scale Numerical Forecast Products", K. Arpe gave examples of a cyclone track chart, and of a chart showing areas of cyclonic activity based on closed contours surrounding individual lows. Examples of the interpretation and use of mean charts and minimum charts, and the charts which result from taking the difference between these two, are now also discussed as a way of presenting useful information in medium-range forecasts. Figure 1 a, b and c shows 1000 mb geopotential charts for days 5, 6 and 7 of a 10-day forecast. Figure 1 d shows the 3-day mean chart for this period. The appearance of the 3-day mean chart shows how the transient smaller scale features of the flow are smoothed out by the averaging process. Nonetheless it must be stressed that these transient features are sometimes of great synoptic interest, and this smoothing is a major disadvantage of mean charts as a forecasting tool. Figure 1 e shows a 3-day minimum chart for the same period (i.e. the minimum 1000 mb height reached at each point of the chart during the 3 days). It differs from the mean chart in areas where there has been a forecast of cyclonic activity, for example over the Pacific or over the south-western United States. Figure 1 f is a chart which shows the differences between the two - a "mean-minimum" chart - (i.e. the areas where the minimum 1000 mb height differed from the mean 1000 mb height during the 3 day period). Only mean-minimum differences of 8 geopotential dekameters or higher are shown. The centre of activity over the Pacific is associated with the forecast deepening of the low, which on day 5 (Fig. 1 a) was shown as a trough near 170°E, by day 6 (Fig. 1 b) had become a low centered near 170°W and on day 7 (Fig. 1 c) was a very deep depression forecast near 160°W. The closed centre over the north of America on the mean-minimum chart (Fig. 1 f) is due to the development of a trough in the area between days 5 and 7, while the centre over Europe on the mean-minimum chart is associated with the trough which can be seen most clearly on day 6 of the forecast.

We now examine the detailed synoptic evaluation rather more carefully, and compare it with the verifying analyses for the 3-day period that is being considered. During days 5, 6 and 7 in the forecast (Fig. 1 a, b and c) an outbreak of cold air into the Western Mediterranean led to the formation of a 1000-mb trough extending from a deep low between Iceland and the Norwegian Sea to North Africa. There was no further deepening of the trough over the Mediterranean after day 6, and the system propogated eastwards. In the verifying analyses (Figs. 2 a, b and c), it can be seen that the parent low was situated in the North Sea between Britain and Denmark on day 7, rather than near Iceland; also the gradients over southern Europe are weaker, however, the forecast of the development of this trough over Western Europe was reasonably successful. Over eastern North America on day 5 of the forecast (Fig. 1 a) a trough extended southwards to the Gulf of Mexico. In this trough, a low formed on day 6 (Fig. 1 b) and was forecast to move towards Greenland on day 7 (Fig. 1 c); over the central and midwestern states, a new trough was forecast to form and extend rapidly southwards on day 7. Here again, the principal features correspond reasonably well with the analyses (Fig. 2 a, b and c), which show that the trough over eastern North America moved north eastwards towards Greenland, although no closed low appeared. The second trough formed in a more easterly position and was much weaker than in the forecast. Indeed this is an example of a tendency which the model has shown on other occasions to overestimate surface developments. Another example can be seen in this case over the Pacific where the trough which deepened rapidly from day 5 to day 7 in the forecast, is seen in the verifying analyses to have been absorbed into the parent low between days 5 and 6 (Fig. 2 a and b) and this remained almost stationary throughout the 3 day period.

How is this synoptic information summarised in the mean, minimum and mean-minimum maps? Comparing mean forecast and mean analysis charts (Figs. 1 d and 2 d) we see both the trough over Europe and the northwesterly flow over Britain and Ireland. The European trough was noticeably stronger in the forecast, with the main low centred north of Scotland. The mean position of the Azores high was displaced slightly to the southwest, but the mean synoptic situation seems quite well forecast over most of the North Atlantic ocean. The forecast mean flow over the Pacific does differ from the analysis, with the strong westerly gradients between 40 and 55°N over the eastern Pacific clearly over-predicted, and the quasi-stationary low near 170°W in the analysis not appearing on the forecast mean chart. Over eastern Europe, the mean charts show the blocking Siberian high to have been displaced to the south in the forecast.

Turning our attention now to the 3-day minimum charts (Figs. 1 e and 2 e) these bear some resemblance to the mean charts (Figs. 1 d and 2 d) but are far less smooth and have stronger gradients which extend somewhat further south in places. However, there are noticeable differences from the mean charts in regions of cyclonic activity, for example over Europe and over the Pacific. Therefore, we study the mean-minimum charts in Figs. 1 f and 2 f, for the forecast and for the verifying analyses respectively. Good agreement is apparent in places where the forecast was successful, for example over Europe (where the difference in position of the centres is due to lower 1000mb heights over the Channel and to the trough extending further south in the forecast), the Atlantic and the eastern side of North America. On the other hand, the (incorrect) forecast of the Pacific low is highlighted in Figure 1 f. Further, the tendency which the model exhibits in over-estimating the intensity of 1000mb troughs and lows is shown by the greater number of closed contours on the mean-minimum chart for the forecasts. A common feature of both maps is the relative absence, with one or two exceptions, of cyclonic activity south of 40°N.

It can be seen therefore, that the mean-minimum charts are useful in that they identify clearly and summarise areas of cyclonic activity over a number of days. These charts isolate regions where the minimum height during a period is significantly different to the mean height. The regions where the flow patterns are stationary, and hence from the meteorological point of view usually less interesting, are evident by the lack of contours. Furthermore, as a verifying tool, comparison of the mean-minimum charts for the analyses and forecasts quickly highlights differences between the two, although judgement is still necessary here, analyses in data-sparse regions like the Pacific being perhaps less reliable than one would wish.

Further investigation and study of these and similar types of charts is continuing, with a view to the use of the charts as objective verification tools for medium-range forecasts, and interpretation of the forecasts. As we noted above, the properties and qualities of a medium range forecast differ considerably from those of a one, two or three-day forecast, and so the sort of chart such as "mean-minimum" described in this article could help to provide one useful method of getting at meteorologically valuable information contained in a medium-range forecast.

FIGURE CAPTIONS

Fig. 1 a,b,c:

Forecast field 1000mb geopotential height for days 5, 6 and 7 respectively, initial day 6.2.1976
Contour intervals 8 dkm

Fig. 1d:

3-day mean of 1000mb geopotential height, centered on day 6 00Z
Contour intervals 4 dkm

Fig. 1e:

3-day minimum of 1000mb geopotential height centered on day 6 00Z
Contour intervals 4 dkm

Fig. 1f:

3-day "mean-minimum" of 1000mb geopotential height centered on day 6 00Z
Contour intervals 4 dkm

Fig. 2 a,b,c:

NMC Washington analysis 1000mb geopotential height
Contour interval 8 dkm
corresponding to forecast days 5,6,7 respectively.

Fig. 2 d,e,f:

Same as Fig. 1, but for NMC Analysis

- A. Woods
- H. Pumpel

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

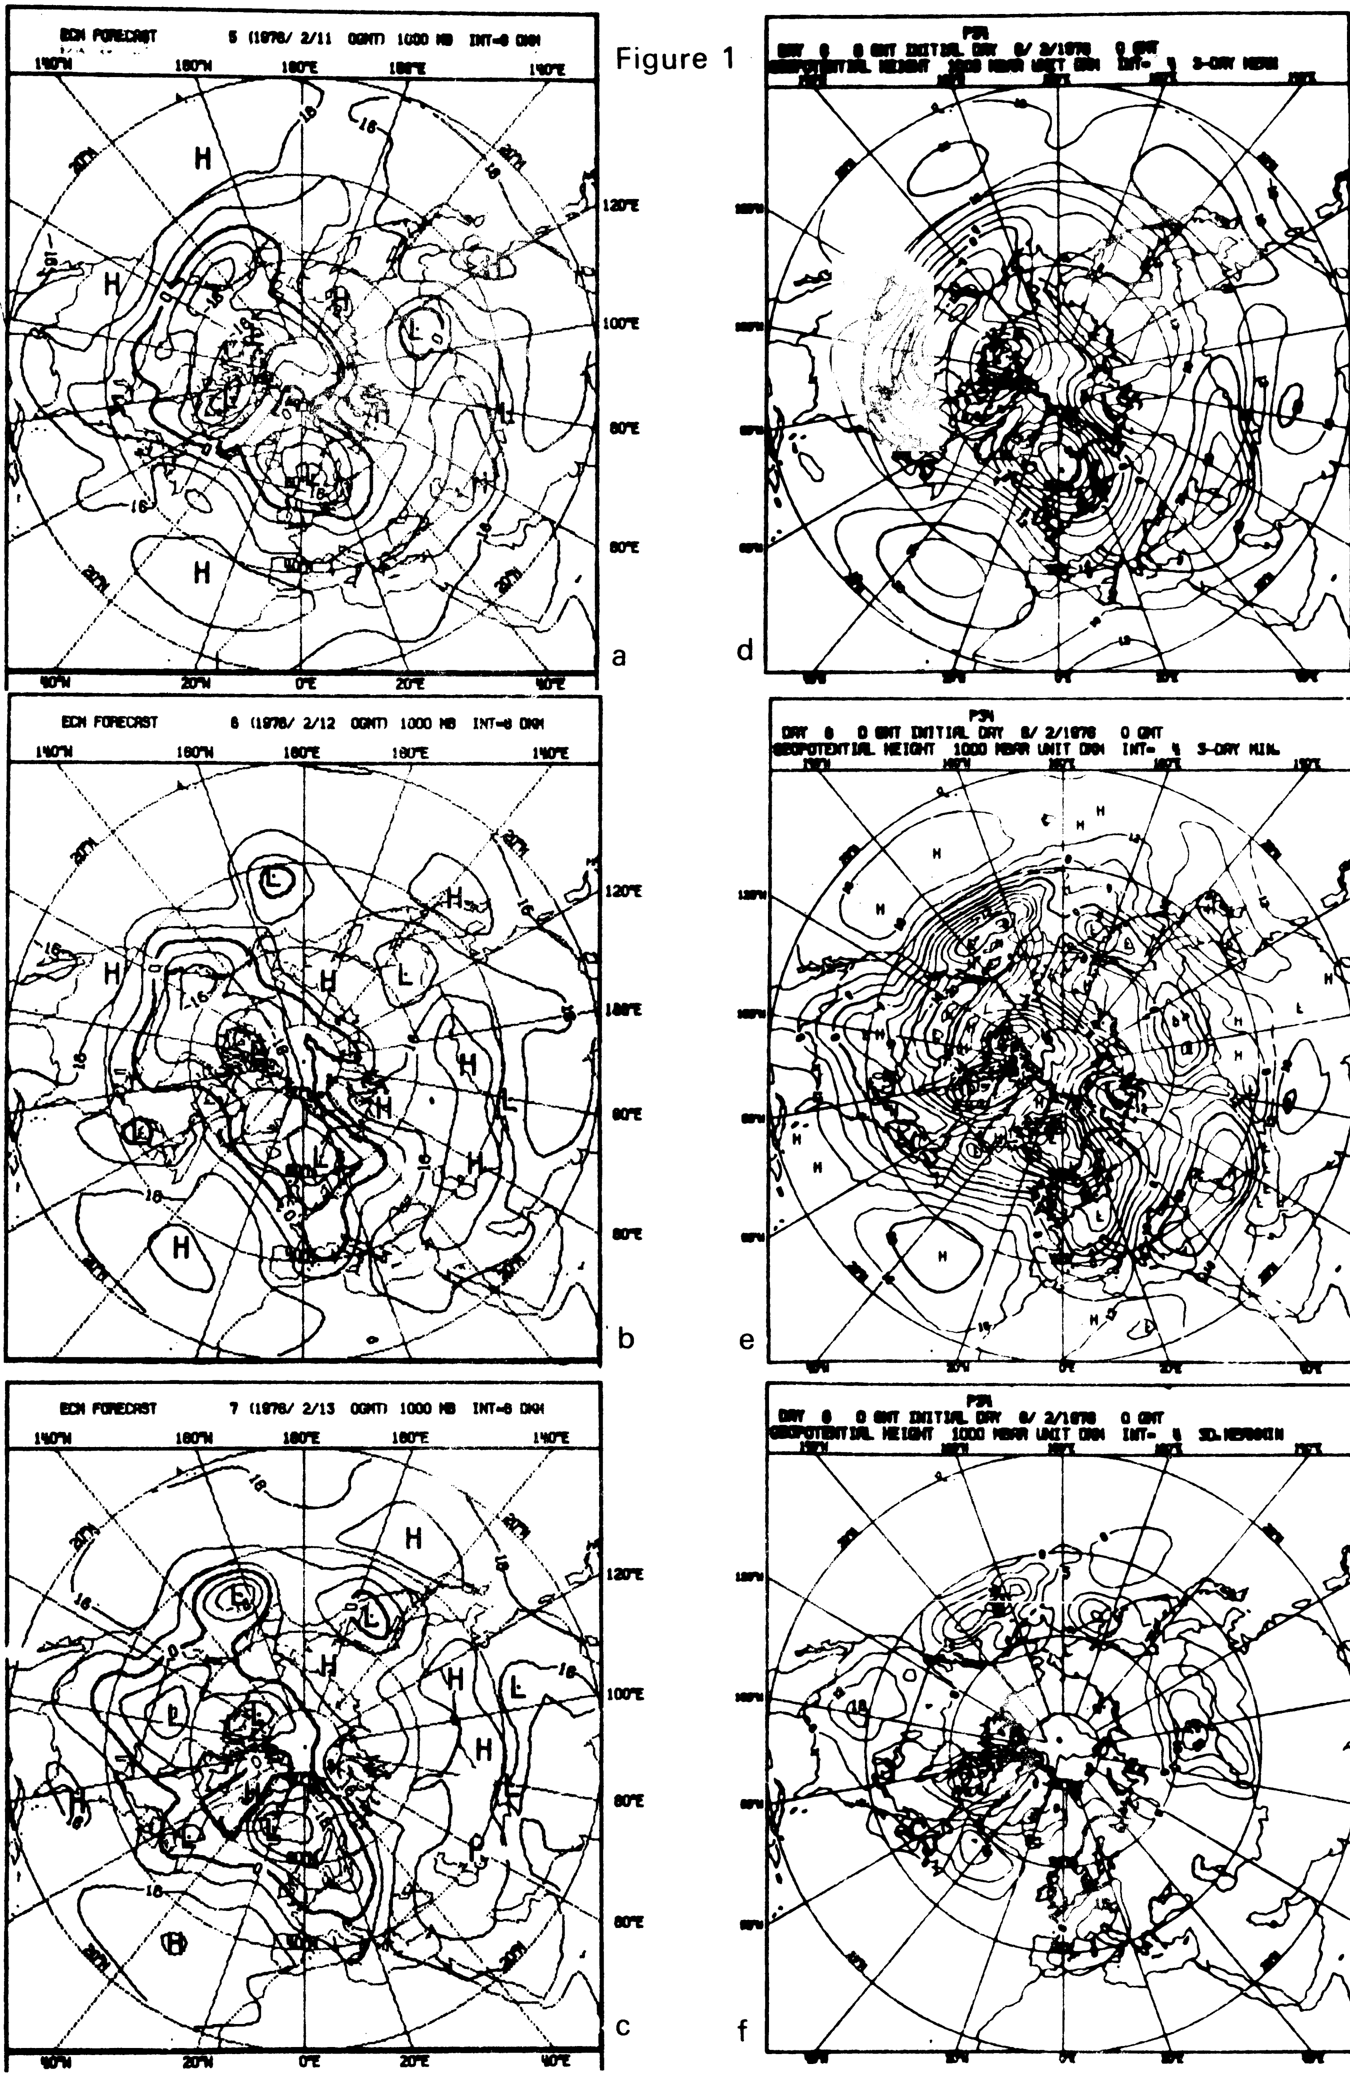
ECM FORECAST 5 (1978/ 2/11 0GMT) 1000 MB INT=8 DKM
ECM FORECAST 6 (1978/ 2/12 0GMT) 1000 MB INT=8 DKM
ECM FORECAST 7 (1978/ 2/13 0GMT) 1000 MB INT=8 DKM
a
b
c
d
e
f

Figure 1

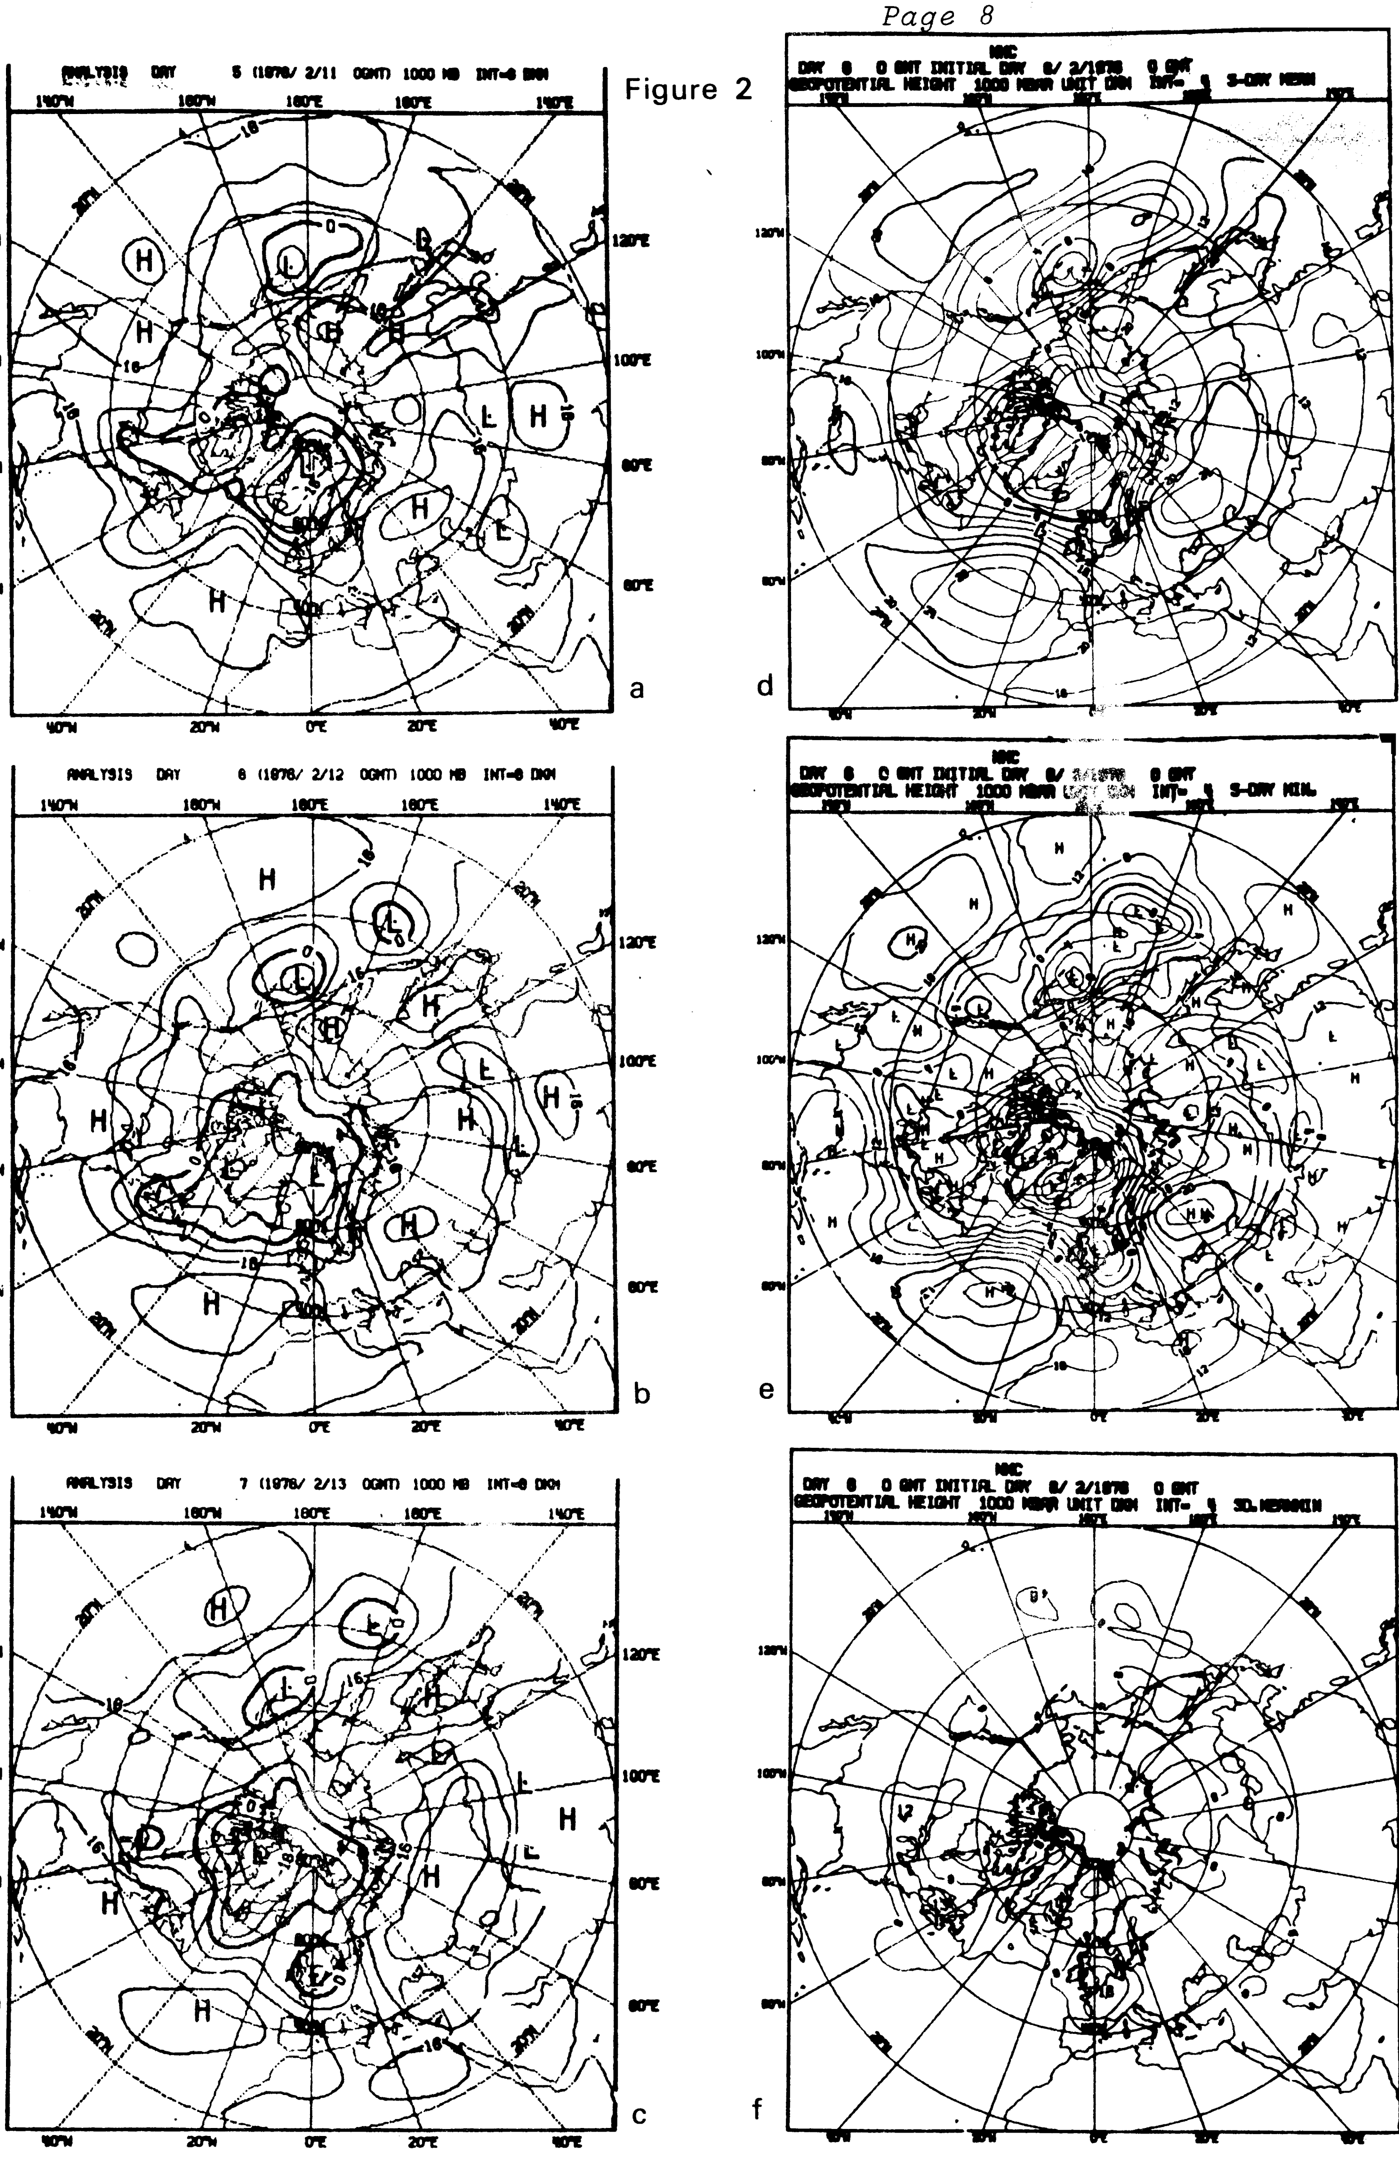

Figure 2

ACCEPTANCE TESTS FOR CYBER 175 AND CRAY-1

After a long period of preparation, during which various problems of installation and integration were solved, both systems have passed their final acceptance tests. Had any component proved unsatisfactory, formal acceptance could have become very complicated, since the Centre had a separate contract with each supplier yet acceptance of each system was dependent on its satisfactory performance linked into, and operating with, the total configuration.

Acceptance for both systems was phased; after installation, a provisional acceptance test checked features and performance during a 48 hour trial, this was then followed by a final acceptance trial of 3 months to check performance under normal operational conditions. The Cray-1 passed provisional acceptance on 8 November 1978, and the Cyber on 7 December 1978. On 1 January 1979, the link between the two (hardware from Cray, software from Control Data), capable of transferring jobs and files at high transfer rates, was put into full operation.

It was on 6 February 1979, that the Cray-1 successfully completed its final acceptance trial with an availability, over the 3 months, of about 99% and a MTBI of around 43 hours, and on 7 March 1979, the Cyber, too, completed its final acceptance period successfully, with an availability of about 98% over the three months and an MTBI of over 45 hours.

Now, the linked system is in full use and has already proved satisfactory to the user community, although some work still has to be done to tune the link software to achieve optimal performance.

- Rob Brinkhuysen

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

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.

Average number of jobs per week within each period

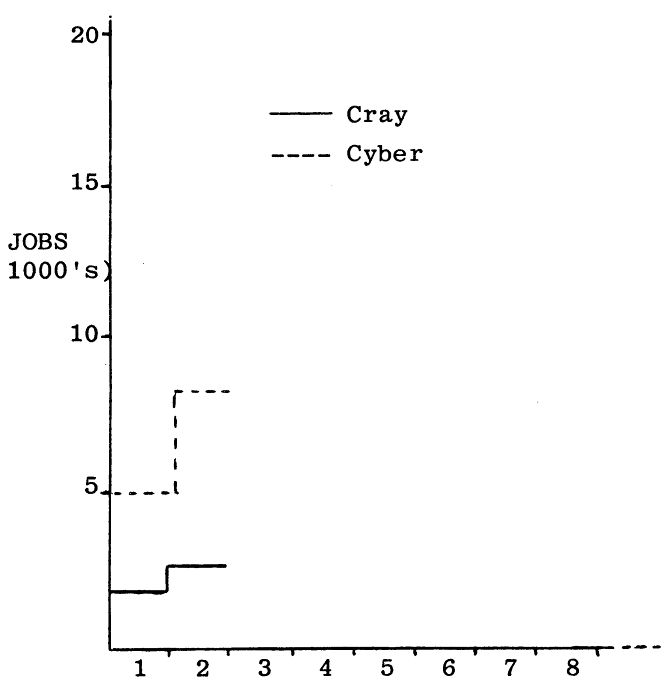

Average CPU hours used per week within each period

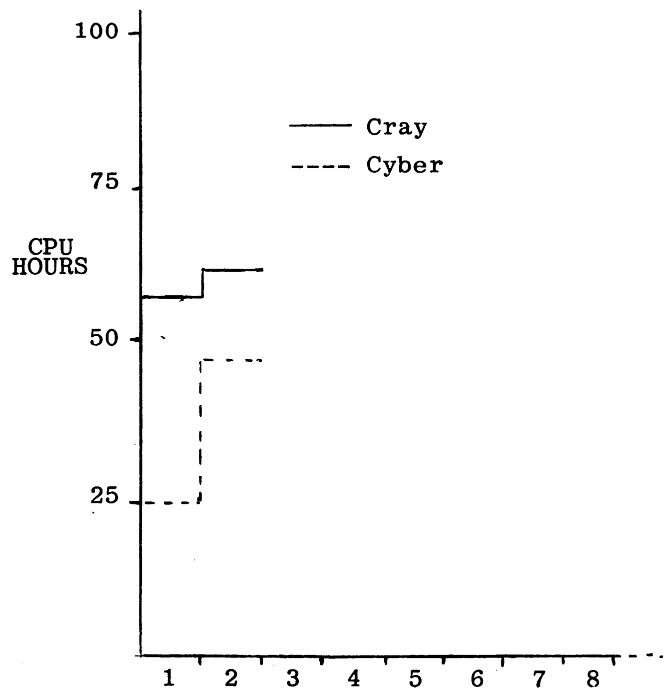

THE CRAY-1 ARCHITECTURE

This is the first of a series of articles describing the CRAY-1 in more detail. It has been adapted from the article "The CRAY-1 Computer System" by Richard M. Russell (Cray Research Inc.) published in the Communications of the ACM (January 1978)†

- David Dent

Introduction

Vector processors are not yet commonplace machines in the large-scale computer market. To date, customer installations exist for the Texas Instruments Advanced Scientific Computer, the CDC STAR 100 and the CRAY-1.

Newly announced machines include the Cyber 203 (an upgrade of the STAR 100), the Burroughs Scientific Processor (an array processor) and the ICL Distributed Array Processor. Within this market the CRAY-1 has proved to be a powerful processor providing a significant improvement in computing speed over the leading scalar processors (CDC 7600, IBM 360/195).

Independent benchmark studies have shown the CRAY-1 fully capable of supporting computational rates of 138 million floating-point operations per second (MFLOPS) for sustained periods and even higher rates of 250 MFLOPS in short bursts (1,2). Such comparatively high performance results from the CRAY-1 internal architecture, which is designed to accommodate the computational needs of carrying out many calculations in discrete steps, with each step producing interim results used in subsequent steps. Through a technique called "chaining", the CRAY-1 vector functional units, in combination with scalar and vector registers, generate interim results and use them again immediately without additional memory references, which slow down the computational process in other contemporary computer systems.

Other features enhancing the CRAY-1's computational capabilities are: its small size, which reduces distances electrical signals must travel within the computer's framework and allows a 12.5 nanosecond clock period (the CRAY-1 is the world's fastest scalar processor); a one million word semiconductor memory equipped with error detection and correction logic (SECDED); its 64-bit word size; and its optimizing Fortran compiler.

Architecture

The CRAY-1 has been called "the world's most expensive love-seat" (3). Certainly most people's first reaction to the CRAY-1 is that it is so small. But in computer design it is a truism that smaller means faster. The greater the separation of components, the longer the time taken for a signal to pass between them. A cylindrical shape was chosen for the CRAY-1 in order to keep wiring distances small.

Figure 1 shows the physical dimensions of the machine. The mainframe is composed of 12 wedge-like columns arranged in a 270° arc. This leaves room for a reasonably trim individual to gain access to the interior of the machine. Note that the love-seat disguises the power supplies and some plumbing for the Freon cooling system.

An Analysis of the Architecture

Table 1 details important characteristics of the CRAY-1 Computer System. The CRAY-1 is equipped with 12 i/o channels, 16 memory banks, 12 functional units, amd more than 4k bytes of register storage. Access to memory is shared by the i/o channels and high-speed registers. The most striking features of the CRAY-1 are: only four chip types, main memory speed, cooling system, and computation section.

Fig. 1. Physical organization of mainframe.

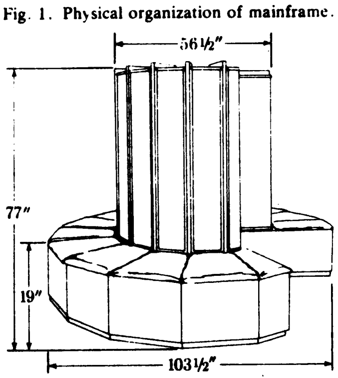

- Dimensions
  - Base-103½ inches diameter by 19 inches high
  - Columns-56½ inches diameter by 77 inches high including height of base
- 24 chassis
- 1662 modules; 113 module types
- Each module contains up to 288 IC packages per module
- Power consumption approximately 115 kw input for maximum memory size
- Freon cooled with Freon/water heat exchange
- Three memory options
- Weight 10,500 lbs (maximum memory size)
- Three basic chip types
  - 5/4 NAND gates
  - Memory chips
  - Register chips

Table I. CRAY-1 CPU characteristics summary

Computation Section
- Scalar and vector processing modes
- 12.5 nanosecond clock period operation
- 64-bit word size
- Integer and floating-point arithmetic
- Twelve fully segmented functional units
- Eight 24-bit address (*A*) registers
- Sixty-four 24-bit intermediate address (*B*) registers
- Eight 64-bit scalar (*S*) registers
- Sixty-four 64-bit intermediate scalar (*T*) registers
- Eight 64-element vector (*V*) registers (64-bits per element)
- Vector length and vector mask registers
- One 64-bit real time clock (*RT*) register
- Four instruction buffers of sixty-four 16-bit parcels each
- 128 basic instructions
- Prioritized interrupt control

Memory Section
- 1,048,576 64-bit words (plus 8 check bits per word)
- 16 independent banks of 65,536 words each
- 4 clock period bank cycle time
- 1 word per clock period transfer rate for *B*, *T*, and *V* registers
- 1 word per 2 clock periods transfer rate for *A* and *S* registers
- 4 words per clock period transfer rate to instruction buffers (up to 16 instructions per clock period)

i/o Section
- 24 i/o channels organized into four 6-channel groups
- Each channel group contains either 6 input or 6 output channels
- Each channel group served by memory every 4 clock periods
- Channel priority within each channel group
- 16 data bits, 3 control bits per channel, and 4 parity bits
- Maximum channel rate of one 64-bit word every 100 nanoseconds
- Maximum data streaming rate of 500,000 64-bit words/second
- Channel error detection

Four Chip Types

Only four chip types are used to build the CRAY-1. These are 16 x 4 bit bipolar register chips (6 nanosecond cycle time), 1024 x 1 bit bipolar memory chips (50 nanosecond cycle time), and bipolar logic chips with subnanosecond propagation times. The logic chips are all simple low or high-speed gates with both a 5 wide and a 4 wide gate (5/4 NAND). Emittercoupled logic circuit (ECL) technology is used throughout the CRAY-1.

The printed circuit board used in the CRAY-1 is a 5-layer board with the two outer surfaces used for signal runs and the three inner layers for -5.2V,-2.0V, and ground power supplies. The boards are six inches wide, 8 inches long.

All integrated circuit devices used in the CRAY-1 are packaged in 16-pin hermetically sealed flat packs supplied by both Fairchild and Motorola. This type of package was chosen for its reliability and compactness. Compactness is of special importance; as many as 288 packages may be added to a board to fabricate a module (there are 113 module types), and as many as 72 modules may be inserted into a 28-inch-high chassis. Such component densities inevitably lead to a mammoth cooling problem (to be described).

Main Memory Speed

CRAY-1 memory is organised in 16 banks, 72 modules per bank. Each module contributes 1 bit to a 64-bit word. The other 8 bits are used to store an 8-bit check byte required for single-bit error correction, double-bit error detection (SECDED). Data words are stored in 1-bank increments throughout memory. This organisation allows 16-way interleaving of memory accesses and prevents bank conflicts except in the case of memory accesses that step through memory with either an 8 or 16-word increment.

Cooling System

The CRAY-1 generates about four times as much heat per cubic inch as the CDC 7600. To cool the CRAY-1 a new cooling technology was developed, also based on Freon, but employing available metal conductors in a new way. Within each chassis vertical aluminium/stainless steel cooling bars line each column wall. The Freon refrigerant is passed through a stainless steel tube within the aluminium casing. When modules are in place, heat is dissipated through the inner copper heat transfer plate in the module to the column walls and thence into the cooling bars. The modules are mated with the cold bar by using stainless steel pins to pinch the copper plate against the aluminium outer casing of the bar.

To assure component reliability, the cooling system was designed to provide a maximum case temperature of 130°F (54°C). To meet this goal, the following temperature differentials are observed:

Temperature at centre of module	130°F (54°C)
Temperature at edge of module	118°F (48°C)
Cold plate temperature at wedge	78°F (25°C)
Cold bar temperature	70°F (21°C)
Refrigerant tube temperature	70°F (21°C)

References

1. Preliminary Report on Results of Matrix Benchmarks on Vector Processors: Calahan, Joy, Orbits, System Engineering Laboratory, Univerity of Michigan, Ann Arbor, Michigan 48109

2. Computer Architecture Issues in Large-scale Systems, 9th Asilomar Conference, Naval Postgraduate School, Monterey, California

3. Computer World, August 1976.

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

* CRAY PERMANENT FILES

It has already been agreed by the Computer Steering Group and by the Computer Users Liaison Committee, that permanent files can only be maintained securely on the Cyber, as the operational suite will require most of the Cray disk space daily.

Until now, small amounts of Cray permanent filespace (under the DUMPXX identifiers) have been backed up, and the reliability of the Cray disk system has made the loss of Cray permanent files a very rare occurrence. However, integration and testing of the operational suite is about to begin and so we must implement the agreed procedure to create the necessary Cray filespace.

Commencing Monday 2nd April 1979, backup of Cray files under the DUMPXX identifiers will cease. At the same time, a daily procedure will be put into effect which will involve deleting Cray permanent files in order to create disk space for the operational suite.

This means that no Cray Permanent Files will be secure and that the users must arrange to secure their own permanent files on the Cyber.

It is realised that this could cause some congestion with the available Cyber PF space, therefore, everybody is urged to review their permanent files, to copy or purge where possible and to keep large files on tape instead of disk.

The situation will improve over the next four months with the installation of another four disks in April, another four in July, and by allocation of PF space.

- Rob Brinkhuysen

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

New CFT and $FTLIB

The January release of CFT and $FTLIB (known as 1.0.4) is now the production version.

The main changes can be divided into four sections:

a) changes to the language

b) changes to the CFT call card

c) changes to subroutines

d) internal changes

a) Changes to the language

*i) Addition of END and ERR parameters to READ and WRITE statements.
e.g. READ(5,100, END=90, ERR=91) ...
will return control to statement 90 if an EOF is detected on input and will return control to statement 91 if an error is detected while reading.

ii) The variable on a DO statement and the iteration count are limited to 2^{23} -1.

iii) Two new compiler directives to selectively control instruction scheduling
CDIR$ SCHED and
CDIR$ NOSCH
(mainly for use when a compiler bug is suspected)

b) Changes to the CFT call card

i) New parameter TRUNC=nn
This will set the bottom nn bits of all floating points to zero. Values of constants and function results are not affected. (It may be useful for examining the instability of an algorithm).

ii) Three new options to the ON and OFF parameters

M to enable/disable the scheduler (default ON)
R round results of multiply operations (default ON)
W compiles all floating point operations as calls to user supplied routines (default OFF).

c) Changes to subroutines

*i) Two new intrinsic functions

INT24 converts 64 bit integers to 24 bits
LINT converts 24 bit integers to 64 bits

ii) Three new subroutines

REMARK2 like REMARK except that argument can be 9 words with 1st 5 characters being message identifier

ERREXIT terminates the job step. Job continues at the next EXIT statement. No traceback is printed.

TRBK writes traceback to file specified by 1st argument preceded by message specified by second argument. If 1st argument is minus or '$COS' output goes to LOGFILE.

d) Internal changes

i) Some changes have been made in the evaluation of expressions such that the results should be identical from either a vector or scalar loop

ii) The divide algorithm has been improved but this may give slightly different results.

Further details can be supplied and any problems answered by either myself or User Support. Updated manuals have been issued.

\- Gary Harding

* but see article on "CFT 1.0.4. PROBLEMS" below

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

* CFT 1.0.4 PROBLEMS

1. A bug has been found in the CFT compiler, which arises only under particular circumstances, when using the new features END= and ERR= , together, in a single READ statement. The exact circumstances in which this fails, are not clear, but in a case recently seen the bug disappeared or reappeared depending on the placement of a format statement. So , if possible, do not use END and ERR parameters together. When either was used singly, in the example seen, no problem occurred.

2. The compiler directive INT24 does not have the desired effect. It is apparently not possible to restrict integer computation to 24 bits by this means:

   ```
   e.g.   CDIR$   INT24    I, J, K, L
                  I=J*K+L
   ```

   results in 64 bit computation throughout !

3. Subscript expressions containing parenthesis inhibit vectorisation in an otherwise vectorisable DO loop.

\- D Dent
\- J Greenaway

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

* File transfer to a Named Device on the Cray

Due to inadequacies in the protocol, operating system (COS) and station, it is not possible to ACQUIRE a Cyber file and place it directly on to a particular disk unit. In the longer term we hope the manufacturers will provide this facility. In the interim, however, this can be accomplished as follows:

```
ASSIGN  (DN=XX, DV=yy)        where yy is a valid device name
ACQUIRE (DN=X, PDN= ...., UQ)
COPYD   (I=X, O=XX)
REWIND  (DN=XX)
SAVE    (DN=XX, PDN=....)
DELETE  (DN=X)
```

puts dataset XX on to device yy

To hide this from the user, the Cray subroutine ATTACH will be changed to perform this operation in the following circumstance:

```
CALL REQUEST (NF, LU,'DV=...')
CALL ATTACH  (NF, LU ....)
```

If LU exists but is empty, ATTACH will perform the operations mentioned above

```
i.e:        ACQUIRE
            COPYD
            REWIND
            SAVE
            DELETE
```

\- D. Dent

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

Cray-Cyber Data Conversion Utilities

These utilities (CONVERT package) were described in outline in the last ECMWF Technical Newsletter. They are now fully available, usable either as control cards or as subroutine calls. Full details are in the ECMWF Computer Bulletin B7.2/1, an update to which was issued recently.

- David Dent

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

Permanent File Control Cards as Subroutine Calls

Fortran callable subroutines to perform permanent file manipulation have been available for a long time on the Cyber. They have now been implemented on the Cray with the same subroutine interface. In addition, they can also be used to control file transfers across the link. The following routines are available:

ATTACH	REQUEST	CATALOG	PURGE	DISPOSE
RENAME	SUBMIT	EXTEND	RETURN	

Please contact User Support for full details.

- David Dent

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

* NAG Library on the CRAY

The Nag Library, Mark 6, has now been implemented on the CRAY, i.e. all the machine-dependent constants and routines have been provided and the stringent test programs run. A small proportion (approximately 5%) of the stringent test programs have failed outright, and they are still under investigation.

However, this version should be a considerable improvement on the previous version, which did not undergo stringent testing at all.

Because files with an ID of DUMPØØ will no longer be archived from the CRAY as from April 2nd, it has been decided to make Version 6 of the NAG library available with a public ID (i.e. no ID needed) and with the name NAGLIB. This means that it will no longer be necessary to use an ACCESS statement, prior to loading from it. As the old version is no longer protected, it could disappear at any time after an install. Therefore, you are urged to change over as quickly as possible.

i.e. use LDR(LIB=NAGLIB.....)

The routines that failed the stringent tests are listed below. If you have a problem using one of these (or any other) please contact User Support.

EØ4DEF	HØ1ADF	S13ACF	S17ADF
EØ4KAF	HØ1AEF	S13ADF	S18ACF
EØ4VBF	HØ2BAF	S14AAF	S18ADF
EØ4WAF	HØ3ABF	S17ACF	

- John Greenaway

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

* CYBER SCHEDULING - CM PARAMETER FOR LARGE JOBS

A job class scheduler was re-introduced on the Cyber on 12 March. This article gives information on how it will affect you. Initially, the class structure is based mainly on CPU time, with some account being taken of main memory requirements.

The main change directly affecting you, is the need to use the CM parameter, although initially there will be a large default setting (260K octal).

The Cyber scheduler is intended to:

i) improve the utilisation of Cyber resources
ii) permit faster turnaround of small jobs
iii) permit very fast initialisation of high priority jobs.

To achieve this, some experimentation will be done, thus the classes to be used will be decided by a phased program of experimental changes which will allow the structure to be optimised.

The rest of this article describes the various parameters affected by the above changes.

Time

This is measured in octal seconds, the default time limit will remain at 40 (octal) seconds, with the maximum set at 77777 (octal), i.e. 32767 decimal seconds. TØ is also now interpreted as the same as the above, about 9 hours maximum, i.e. 77777 octal.

The scheduling classes will tend to give a better turnaround to shorter jobs, because more short jobs will be allowed to run at one time than longer ones. Thus it will be in your interest to ask for as short a time as is practicable.

Memory

This is measured in octal words. The default (CM) value for jobs will be set at the high level of 260K octal. However, it would be in your interest to specify a CM parameter with a lower value as your turnaround will be improved. It is intended to reduce the default CM size to 60K octal in the near future.

STBIG users requiring more than 260K octal memory will now have to put a CM parameter on the job card. Unless you know how much you require, it is suggested that you put the maximum limit of 320K octal. Currently you must include both the CM and STBIG parameters. Later the STBIG parameter will become unnecessary.

Note that if a job tries to exceed the job card CM limit, then the job will be aborted.

If you wish to find out what CM parameter your job requires, you will need to determine which job step uses the maximum amount of core. Usually, this will be the loading/execution phase and you will then need to obtain a load map. From the load map take the LWA+1 figure (Last Word Address + 1) and add the space required for I/O routines (capsules).You will only have to consider I/O buffers and the communication area (FIT) if you call Record Manager Interface routines yourself, such as FILESQ, FILEXX, etc. Files from the Program card plus FIT are included in the load map. Provided the total figure exceeds "CM storage used" which, in most cases, it will, you have the maximum core used. If you have more than one loading/execution phase, clearly you should take the largest of them all.

It is not suggested that these values are easy to work out, but the description above is intended to give some ideas of the processes involved. Note that for all simple Fortran compilations, UPDATE runs, and uses of utilities, e.g. AUDIT, a CM size of 60K octal is adequate. It is hoped to produce more details on the sizes of I/O capsules at a later date.

- Peter Gray
- John Greenaway

****

* USE OF MAGNETIC TAPES

Here are some hints on how to make use of magnetic tapes on the Cyber effectively and safely.

1. Only request a magnetic tape for as short a period as possible

   When using magnetic tapes, request the tape immediately before accessing it, and RETURN or UNLOAD it as soon as possible after processing the data. Thus, you should not request a data tape before program compilation has taken place and you should release it immediately after the job step which processes it has completed. If possible, release the tape within the job step.

2. Use RETURN rather than UNLOAD wherever possible

   The HD, GE or PE parameter on the jobcard tells the scheduler how many tape drives of a particular type you job will use. The scheduler will not start your job until it is sure that your tape drive requirements can be satisfied without deadlocks occurring. The RETURN command used on a magnetic tape file tells the scheduler not only that processing of the file is complete, but also that the number of tapes required simultaneously by the job can be reduced by one. i.e. that one tape drive can now be released for another job to use. UNLOAD says that processing of the file is complete, but that the number of tapes required simultaneously does not change. Use of RETURN where possible will improve tape job throughput.

3. Request as few tapes as possible in the jobcard

   The jobcard HD, GE or PE parameter does not specify the number of magnetic tapes that will be used by the job. They specify the maximum number of tapes which can be used simultaneously by the job. Using the parameter correctly will improve tape job throughput, as will altering the job to process several tapes serially, one after the other, rather than simultaneously.

   Examples

   i.
```
SYA,CM40000,T100,GE1.
ACCOUNT,....
ATTACH,MYLIB,ID=id,MR=1.
LIBRARY,MYLIB.
LABEL,TAPE1,D=GE,VSN=vsn,R.
PROG1.                          - Load PROG1 from MYLIB
UNLOAD,TAPE1.
LABEL,TAPE2,D=GE,VSN=vsn,R.
PROG2.                          - Load PROG2 from MYLIB
RETURN,TAPE2.
```

The above example shows that 2 tapes can be processed one after the other when GE1 only has been specified on the jobcard. If the command

UNLOAD,TAPE1

was omitted, the second LABEL command would have aborted as it would have been requesting a second tape simultaneously with the first. If the command

UNLOAD,TAPE1

had been replaced by

RETURN,TAPE1

the scheduler would have been told to decrement the magnetic tape requirement by 1. (i.e. from 1 to zero). Thus, in this case, the second LABEL command would have aborted as it was requesting a tape when the scheduler was not expecting a tape to be requested.

```
ii. SYA,CM40000,T100,GE2.             - Use max of 2 tapes simultaneously
    ACCOUNT,......
    ATTACH,MYLIB,ID=id,MR=1.
    LIBRARY,MYLIB.
    LABEL,TAPE1,D=GE,VSN=1000X,R.
    LABEL,TAPE2,D=GE,VSN=1001Y,R.
    PROG3.                            - Process 2 tapes 1000X,1001Y
    UNLOAD,TAPE2.                     - Finished with 1001Y, but still require
                                        2 tapes simultaneously
    REWIND,TAPE1.                     - Will use 1000X again
    LABEL,TAPE3,D=GE,VSN=1002Z,R.
    PROG4.                            - Process 2 tapes 1000X,1002Z
    RETURN, TAPE1.                    - Finished with 1000X, only need 1 tape so use
                                        RETURN
    UNLOAD, TAPE3.                    - Finished with 1002Z, but will use 1 tape in
                                        future
    LABEL,TAPE4,D=GE,VSN=1003A,R.
    PROG5.                            - Process 1003A
    UNLOAD TAPE4.                     - Finished with 1003A, but will use 1 more tape.
    LABEL, TAPE5,D=GE,VSN=1004B,R.
    PROG6.                            - Process 1004B, will not do any more tape
    RETURN,TAPE5.                       processing.
    PROG7.
    PROG8.
      .
      .
      .
```

4. <u>Mount private packs before requesting tapes</u>

Any jobstep which requires both private packs <u>and</u> magnetic tapes should MOUNT the private pack <u>before</u> requesting the tape. This is because private pack mounts frequently take a <u>long</u> time to complete, whereas tape mounts are normally quickly satisfied, thus magnetic tape decks will not be tied up for long periods whilst a pack mount is awaited.

5. <u>Where possible 'stage' tapes to disc for processing</u>

The process of copying a tape file to disc for processing, then processing the disc copy and, if necessary, rewriting the disc copy to tape, is known as staging.

Where the quantity of staged data would be less than 100 million characters or 10 million words † the staging technique should be used. Use of this technique will make more effective use of tape decks and will mean that your job will only require 1 tape deck to execute. This will generally improve job turnaround, and will be most effective when your job would have tape decks for more than 15 minutes.

<u>Examples</u>

i. <u>The conventional approach</u>

```
SYA,CM40000,TI000,GE2.
ACCOUNT,.....
ATTACH,MYLIB,ID=id,MR=1.
LIBRARY,MYLIB.
LABEL,TAPE1,D=GE,NORING,VSN=1000X,R.
LABEL,TAPE2,D=GE,RING,VSN=1001Y,W.
PROG10.                                 - Updates from TAPE1 to TAPE2
RETURN,TAPE1,TAPE2.
```

ii. <u>The staging approach</u>

```
SYA,CM40000,T1000,GE1.                    Notice only 1 tape unit needed
ACCOUNT,......
ATTACH,MYLIB,ID=id,MR=1.
LIBRARY,MYLIB.
LABEL,TEMP,D=GE,NORING,VSN=1000X,R.
COPYBF,TEMP,TAPE1.                      - Use COPYCF if coded file
```

Contd.

† 50 million chars or 5 million words before July 1979

```
UNLOAD,TEMP.
REWIND,TAPE1.
PROG10.                                    - Updates from TAPE1 to TAPE2
RETURN,TAPE1.                              - Free up disc space
REWIND,TAPE2.
LABEL,TEMP,D=GE,RING,VSN=1001Y,W.
COPYBF,TAPE2,TEMP.                         - Use COPYCF if coded
RETURN,TAPE2,TEMP.
```

Notes:

i. Remember to use REWIND where necessary
ii. RETURN or UNLOAD disc copies as soon as possible
iii. The sizes of TAPE1 plus TAPE2 should not exceed the maximum for staged data.

6. Ensure the tape is written at the expected density

The LABEL and REQUEST statements provide parameters to specify tape density (e.g. D=GE,D=HD,D=PE). This might be expected to specify at what density the tape is to be written. - However it is not that simple! -

If there is an existing label on the tape, the tape is always written at the density of the label. If there is no existing label, or the label is to be rewritten, the density parameter is used to specify the density of the label and data.

The label and data are always written at the same density on 9 track tape.

All tapes blank labelled by operations are labelled at D=GE (6250 bpi) so there will not normally be a problem. However, it is recommended that you use the W parameter on the LABEL card when initially writing to tape. Using the R parameter on the LABEL card on successive writes will allow you to ensure that you are writing to the correctly labelled tape.

It is hoped that CDC will produce a modification to the system to output a dayfile message when the tape is written at a density other than that specified on the LABEL card.

7. Always use 6250 bpi tapes where possible

The new Group Coded Recording (GCR) decks which operate at 6250 bpi are better than the other 9 track magnetic tape recording methods because:

i. You get far more data on a tape.
ii. You get much higher transfer rates.
iii. You get much higher reliability.
iv. We have more GCR decks (4 now, 6 from July 1979) so you should get better turnaround.

So, use them for all internal tapes, and copy tapes from other sites to 6250 bpi before processing. Taking such copies will give back-up.

8. Never ever overwrite an input tape in a job

Jobs which read a tape, create an output file, then write that output file to the tape, which has just been read, run a high risk of losing valuable input data. If the system or the job, or the magnetic tape, fails whilst you are copying the output file to tape, you will lose both input and output data.

9. Only append data to a tape when absolutely necessary

A job which takes an existing tape skips to EOI then appends data will never be rerunnable automatically by the system if it fails whilst data is being appended. Also, some failures will cause the end-of-information marker to be lost, so that it will be impossible to skip to EOI and will therefore make recreation of the data difficult.

10. Use the L field in the LABEL card

Specifying the L field in LABEL will make doubly sure that you are accessing the correct tape.

11. Use the correct error processing option when writing 6250 bpi tapes

Group Coded Recording method (D=GE or 6250bpi) has the ability to correct two tracks in error when the tape is read. In comparison, Phase Encoded method (D=PE or 1600bpi) can correct one track in error when reading and NRZI method (D=HD, or 800 bpi) cannot correct any tracks in error.

Given the ability to correct two tracks in error on GCR decks, there may be circumstances where it is permissable to allow one track in error on writing. A parameter to LABEL controls this feature. The parameter is as follows:

EEC - Enable Error Correction on write to GCR decks. This allows single track errors on writing. Use of this parameter makes more efficient use of tape and will allow slightly faster writing to tape.

IEC - Inhibit Error Correction on write to GCR decks. Whenever a single track error is detected on writing, the area of tape is erased and the block is rewritten. This parameter makes less effecient use of tape and causes a holdup on writing when a single track error occurs.

It is not easy to make any recommendation on the use of these parameters. Since writing at 6250 bpi requires higher quality tapes, it is inevitable that there will be more single track errors than when writing at other densities. In deciding whether to allow single track write errors one would need data on the probability of a second track giving an error after writing in a block which had a single track error. Such statistics are not available but the probability seems likely to be very low.

However, to be extra safe it is probably advisable to use the IEC parameter when writing tapes which will contain data having a very long lifetime, and to use EEC at other times.

- Peter Gray

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

A TAPE ANALYSIS PROGRAM - LOOK9

We have now acquired and implemented the LOOK9 program from the University of London Computer Centre. The program is designed to help us identify the properties and contents of a 9 track magr tic tape.

As the output from the program requires interpretation by a skilled analyst, the program has only been made available to be called from the Cyber operator's console. The identification of the properties of a magnetic tape will be a joint exercise undertaken by the Magnetic Tape Librarian and the advisory staff.

To request an analysis to be done on a given tape, please contact the Magnetic Tape Librarian.

- Peter Gray

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

NOS/BE (473) PROBLEM HARVEST

Some further points to note, in addition to those in the last two Newsletters.

- I/O BOMBS OUT WHEN SEGMENT LOADING -

A number of segment loaded programs have failed for having omitted to declare the FORTRAN Common Library common blocks as global:

```
Q8.IO.
FCL.C.
STP.END
```

If unexpected input/output errors abort your program, please ensure that the above named blocks only appear in your root segment - for more information, the Fortran User's Guide (page 7-9) proves particularly enlightening.

- FIELD LENGTH OVERFLOWS MAXIMUM AVAILABLE AT EXECUTION TIME -

This problem may arise because of dynamic allocation of I/O buffers and dynamic loading of BAM/AAM capsules. It is diagnosed with the dayfile message "CMM ERROR NO. 2......", where CMM stands for Common Memory Manager.

Field length requirements can be reduced by:

a) squeezing code (obvious but painful);

b) decreasing the number of file names on your Program card;

c) setting the circular buffer size to zero for all files processed by BUFFER IN/OUT, or to the minimum allotted (PRU size) for the others, i.e. 64 for mass storage files or 512 for tape files (please refer to FTN Ref. Man. page 7.4 for greater detail);

d) limiting simultaneous processing of I/O streams to the bare minimum;

e) visiting the advisory office with some well documented evidence of your problem (as a sign of good will).

- BUFFER IN STEALS RECORD

BUFFER IN has been summoned for stealing the first record of any file which happens to be declared in a PROGRAM statement without its associated circular buffer size set to zero. It is expected it will plead guilty, in the face of the overwhelming evidence against it provided by User Support. The Systems group are currently pleading for the defendant.

- ERROR MODE1 WHICH APPARENTLY ORIGINATES IN COMMON MEMORY MANAGER IS CAUSED BY USER OVERINDEXING HIS BLANK COMMON AREA

So be careful - before assuming the system is the culprit, check all those indices of yours and make sure they do not exceed the limits imposed by your blank common size.

- ADDING MAIN PROGRAMS TO USER LIBRARIES

An attempt to add more than one relocatable main program to a user library can only be successful if compilations have been made with the SYSEDIT option specified in the FTN card, i.e.

```
FTN(...,SYSEDIT,...)
```

where commas separate an arbitrary number of optional parameters.

Failing to do so, results in the user library creation aborting, because of the duplicate entry points generated by the program file names - refer to the FTN Reference Manual page 10-8.

- Luigi Bertuzzi

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

THE CYBER FORTAN CALLABLE CONTROL CARDS - REQUEST, ATTACH, CATALOG PURGE, EXTEND, ALTER, RETURN

The Cyber Fortran callable control cards : REQUEST, ATTACH, CATALOG, PURGE, EXTEND, ALTER, RETURN have now been in existence for some time in an enhanced version, which has undergone extensive and (we hope) exhaustive testing.

The enhancements so far implemented allow for:

a) calling any of the routines without specifying the logical file name in the Program statement. This is achieved by replacing the LFN parameter in the subroutine, with the negative value of the user supplied FIT address. So, in addition to the still valid sample sequence

```
PROGRAM OLD(...,TAPExx,...)
 :
 :
CALL ATTACH(NFAIL,TAPExx,...)
 :
 :
```

the following is also available to BAM/AAM users:

```
PROGRAM NEW(....)
DIMENSION FIT(35)
 :
 :
CALL FILExx(FIT,....)
 :
 :
CALL ATTACH(NFAIL,-LOCF(FIT),...)
 :
 :
```

b) calling REQUEST with any valid choice of parameters among those needed for Magnetic Tape Requests, in addition to the previously existing *PF, *Q, VSN= and SN= parameters.

As from the date of issue of this Newsletter, the enhanced version of the above mentioned routines has been made available in the ECMWF,ID=EWP3 library.

For any additional information, please contact User Support.

- Luigi Bertuzzi

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

JOB PROCESSING HINTS ON THE CYBER 175

(Reprinted from ECMWF Computer Newsletter No. 7, April 1978)

1. Checkpointing and Program Termination

To facilitate a smooth changeover from STBIG sessions to normal sessions, it is important that in all cases, the operators should have information for the required method of job termination.

1.1 Some jobs, notably those put in for overnight runs, have a mechanism whereby the operators may terminate them cleanly by setting sense switch 1. The programs are iterative in nature, and as such, by inspecting sense switch 1 each time (or every 'n' times) round the main loop, they can determine whether or not the operator requires the job to terminate and so tidy up before stopping. A call to sub-routine SSWTCH(N,J) will set J to 1, if sense switch N is set, or 2 if it is not. Would those users making use of this facility and those who would like to make use of it, please ensure that the operators know that your program uses this feature by:

a) putting the comment
ONSWn TO TERMINATE where n is the sense switch number
after the '.' of the 'LGO.' card,

and b) putting the control statement:
COMMENT.*ONSWn TO TERMINATE*
immediately after the 'LGO' card.

Otherwise the operators may have to DROP your job, causing it to abort, and possibly losing output and leaving files in an unusable condition, due to buffers not being flushed.

1.2 A similar facility is available for jobs not using tape or private disk pack whereby these jobs can be locked out at the end of a STBIG session and re-activated for the next suitable session. The convention to use is:

a) put the comment:
LOCKOUT
after the '.' of the 'LGO' card,

and b) put the control statement:
COMMENT. *LOCKOUT*
immediately after the 'LGO.' card.

In the absence of either of the above, the operators will DROP the jobs concerned.

N.B. LOCKOUT will only be allowed if a job is not using tape or private disk pack.

2. CP TIME

A meaningful 'T' parameter value must be put on future job cards.

Jobs for STBIG session should have a 'T' parameter within the following limits

$0<T\leq1000_8$

Normal daytime jobs should bear a reasonable estimate of CP time. The following list may help in preparing 'T' parameters:

HOURS	MINS	SECS	OCTAL SECS.	HOURS	MINS	SECS	OCTAL SECS
	5	300	454	4.5	270	16200	37510
	10	600	1130	5	300	15000	43120
	20	1200	2260	5.5	330	19800	46530
	30	1800	3410	6	360	21600	52140
1	60	3600	7020	6.5	390	23400	55550
1.5	90	5400	12430	7	420	25200	61160
2	120	7200	16040	7.5	450	27000	64570
2.5	150	9000	21450	8	480	28800	70200
3	180	10800	25060	8.5	510	30600	73610
3.5	210	12600	30470	9	540	32400	77220
4	240	14400	34100	9.1	546.1	32766	77776

3. Use of Private Packs

3.1 Identification of jobs requiring a private pack

In order to allow private packs to be identified by our operators, jobs requiring a private pack should, if no other ST parameter is present, specify the parameter STPAK on the jobcard. This will allow the operations staff to schedule these jobs manually should it become necessary.

Any jobs which use private packs and which have not specified STPAK will be dropped by the operator.

Note that if your job fits in the STBIG category, the STBIG parameter must still be specified and STPAK will not be required.

During normal running, it is not intended to restrict private pack jobs to a particular time of day. Rather, this parameter allows private pack jobs to be held back in the input queue if a pack scheduling difficulty arises.

Example

a) A job requiring a private pack

```
USXYZ,T100,STPAK.
ACCOUNT,....
MOUNT(SN=ssss,VSN=vvvv)
    :
```

b) A large core job requiring a private pack.

```
USXYZ,T1000,STBIG.
ACCOUNT,....
MOUNT(SN=ssss,VSN=vvvv)
    :
```

Note that STBIG always takes precedence over STPAK.

3.2 Make private pack requests before magnetic tape requests

Jobs which use a private pack as well as magnetic tapes should always MOUNT the pack required before issuing a REQUEST or LABEL for the magnetic tapes.

Example

```
USXYZ,T100,PE1,STPAK.
ACCOUNT,....
MOUNT(SN=ssss,VSN=vvvv)
REQUEST,TAPE1,PE,NORING,VSN=nnnnnn.
```

3.3 Release a private pack as early as possible within a job

Jobs should DSMOUNT their private packs as soon as processing on that pack is completed.

Example

```
USXYZ,T100,STPAK.
ACCOUNT,....
MOUNT(SN=ssss,VSN=vvvv)
    :
      process files from private packs

DSMOUNT(SN=ssss,VSN=vvvv)
    :
      continue job processing
    :
```

3.4 Make local copies of small files on a private pack

If a long job uses a short file from a private pack, it will be more efficient to make a copy of the file onto the system packs, and then process the copy. If necessary, the updated copy can be rewritten to the private pack at the end of the job. This procedure will only be worthwhile for jobs exceeding around 400_8 seconds and for private packs totalling less than around 5 million characters.

Example

```
USXYZ,T600,STPAK.
ACCOUNT,....
ATTACH(XX,PERMFILE1,SN=ssss,ID=USX)
COPY(XX,YY)
DSMOUNT(SN=ssss,VSN=vvvv)
REWIND(YY)
   :
  process the copy of PERMFILE1 in file YY.
```

4. Multi-read access for Cyber files

Users are encouraged to use MR=1 when ATTACHing 'read only' files as its absence can generate scheduling problems. There is, however, an easier way to provide multi-read access. If a file is initially catalogued with an XR password multi-read permission is granted for subsequent use if the file is attached without this password. Hence:

```
CATALOG,lfn,MYFILE,ID=USX,XR=YZ. and
ATTACH,lfn,MYFILE,ID=USX.
```

allows multi-read access.

- Operations / Systems

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

MANTRAP

The FTN Compiler now has MANTRAP on by default. Those who are currently using the file

```
MAN472,ID=EWGH1
```

should remove this from their jobs as soon as possible. Users not wishing to use MANTRAP should specify

```
LTP=Ø
```

on the FTN statement.

For details about MANTRAP, see Bulletin B7.4/1.

- Gary Harding

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

ONLINE PLOTTER IMPLEMENTATION

(This article appeared in the last newsletter, but unfortunately two paragraphs were misplaced thus altering the intended meaning. To avoid confusion, the article is reprinted (with a little updating)).

As mentioned in ECMWF Computer Newsheet 39, the format of the plot files used on the online 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.
```

The changeover to the new plotter occurred on Monday 5th February. On Tuesday 13th February the new plotter suffered a hardware fault and plots were produced on the old Varian plotter by means of a conversion program. This conversion program is transparent to the user, and programs continued to use the modified graphics software as provided for the Versatec. Since Wednesday the 14th March, the Versatec has been in constant operation.
The rest of this article gives details of the changes made to the graphics software.

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)
```

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).

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.

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 (764decimal 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 online 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 number will be incremented every time the software is altered.

- Howard Watkins

INSTALLATION OF TELECOMMUNICATION LINES TO MEMBER STATES

At its meeting in November 1978, the Council approved an updated implementation schedule for the telecommunications network. Orders have now been placed for those lines due in 1979, the table below shows the expected delivery dates, together with Council approval dates for 1980 and beyond.

TABLE 1

Updated implementation schedule for ECMWF telecommunications network as approved by Council, 21/22.11.78

Member State	Council approved date for medium speed lines	Line delivery date
United Kingdom	January 1979	installed
Germany F.R.	April 1979	1 Oct 79
France	April 1979	1 July 79
Sweden	April 1979	1 Sept 79
Finland	October 1979	1 Dec 79
Denmark	November 1979	1 Nov 79
Ireland	July 1980	-
Italy	July 1980	1 July 79 (low speed)
Greece	September 1980	" "
Spain	December 1980	" "
Yougoslavia	December 1980	" "
Belgium	July 1981	-
Austria	July 1981	-
Portugal	July 1981	1 July 79 (low speed)
Netherlands	October 1981	" "
Turkey	November 1981	" "
Switzerland	January 1984	-

- Andrew Lea

QUESTION AND ANSWER

You are very welcome to submit questions for publication on any matter concerning the computer service. Please direct your questions to any member of User Support.

Q. - When processing a 1600 BPI and a 6520 BPI tape, one after the other (hence only 1 drive at a time), the job card parameters "GE1, PE1" force my job to be scheduled as if 2 drives were needed, thus delaying its execution - it is not fair.

Massimo Capaldo

A. - There are two things to say:

1) Make a copy of your PE tape at GE recording density, so that all your tapes have the same density.

2) We recognise there is both a user education problem and a scheduling problem in the efficient use of tape drives. The Newsletter article on the "Use of Magnetic Tapes" (page 17, this issue) is part of our attempt to educate users into better practices. The Centre is fortunate in having obtained the services of a consultant (Steve Jay, University of Arizona) to work on the scheduling problem.

- Luigi Bertuzzi

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

* DOCUMENTATION CHANGES

ECMWF Bulletins

Bulletin B7.2/1(1) (Cray-Cyber Data Conversion Utilities) was re-issued in an updated version and Bulletin B6.3/1 (The Cern Library) was issued, both at the beginning of March.

Manufacturers' Manuals

Cray Personal Set:

A new issue of CFT Fortran Reference Manual was issued at the beginning of March.

Cray User Office Set:

Copies of Update package F-01 to the COS Reference Manual were distributed.

Copies of Update package C-01 to the COS System Programmers Manual were also distributed to those who already have copies.

- Pam Prior

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

COMPUTER TRAINING COURSES

The response to the request for users (or potential users) to attend the training courses in March and April was varied (see ECMWF Technical Newsletter no. 1, page 27 for details). One course was oversubscribed, while another had very few applicants. Hence, it has been decided to modify our original plans; the following courses will now be run:

- Course A - "Introduction to the Centre"

  To be run April 17 as planned.

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

  To be run twice
  April 4-6 for ECMWF staff
  April 18-20 for Member State users

- Course C - "Detailed use of the Cray-1"

  To be run April 23-27 as planned.

All other proposed courses have been cancelled. Those who applied to attend any of these courses have been notified separately.

The three courses (B, C and D) will be offered again in September, the tentative dates being

Sept. 19-21	Course B	"Simple use of the Centre's computers"
Sept. 24-28	Course C	"Detailed use of the Cray-1"
Oct. 1-5	Course D	"Detailed use of the Cyber 175"

It has not yet been decided whether to offer in September Course A ("Introduction to the Centre").

Details of all these courses will be sent to Member States Computing Representatives and ECMWF Section Leaders in the early summer.

- Andrew Lea

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

VACANCY

PROGRAMMER / METEOROLOGICAL ASSISTANT

The Meteorological Division at the Centre has an opening for a programmer/meteorological assistant, grade B5. The post requires some familiarity with an operational meteorological environment, plus experience of the meteorological codes and formats, and the sort of errors which commonly arise in meteorological data. Some elementary computing experience would be useful. The closing date is April 13. For further information contact the Personnel Section, quoting reference number C607.

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

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 numbers 1 to 7 can be thrown away. Some back copies of Computer Newsletters number 8 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	340/342
Disk Space and Permanent File Problems		AS FOR ADVISORY	
DOCUMENTATION Officer	- Pam Prior	OB 016	355
Libraries (ECMWF,NAG,CERN, etc.)	- John Greenaway	OB 017	354
OPERATIONS	- Console/Shift Leader	CB Hall	334
	- Reception Counter	CB Hall	332
Operations Section Head	- Eric Walton	OB 002	349/351
Deputy Ops. Section Head	- Graham Holt	CB 024	306
METEOROLOGICAL DIVISION			
Division Head	- Roger Newson	OB 008	343
Operations Section Head	- Austin Woods	OB 107	406
Applications Section Head	- Joel Martellet	OB 011	344
Meteorological Analysts	- Ove Akesson	OB 106	380
	- Herbert Pumpel	OB 106	380
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	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