ECMWF Newsletter

Number 60 - December 1992

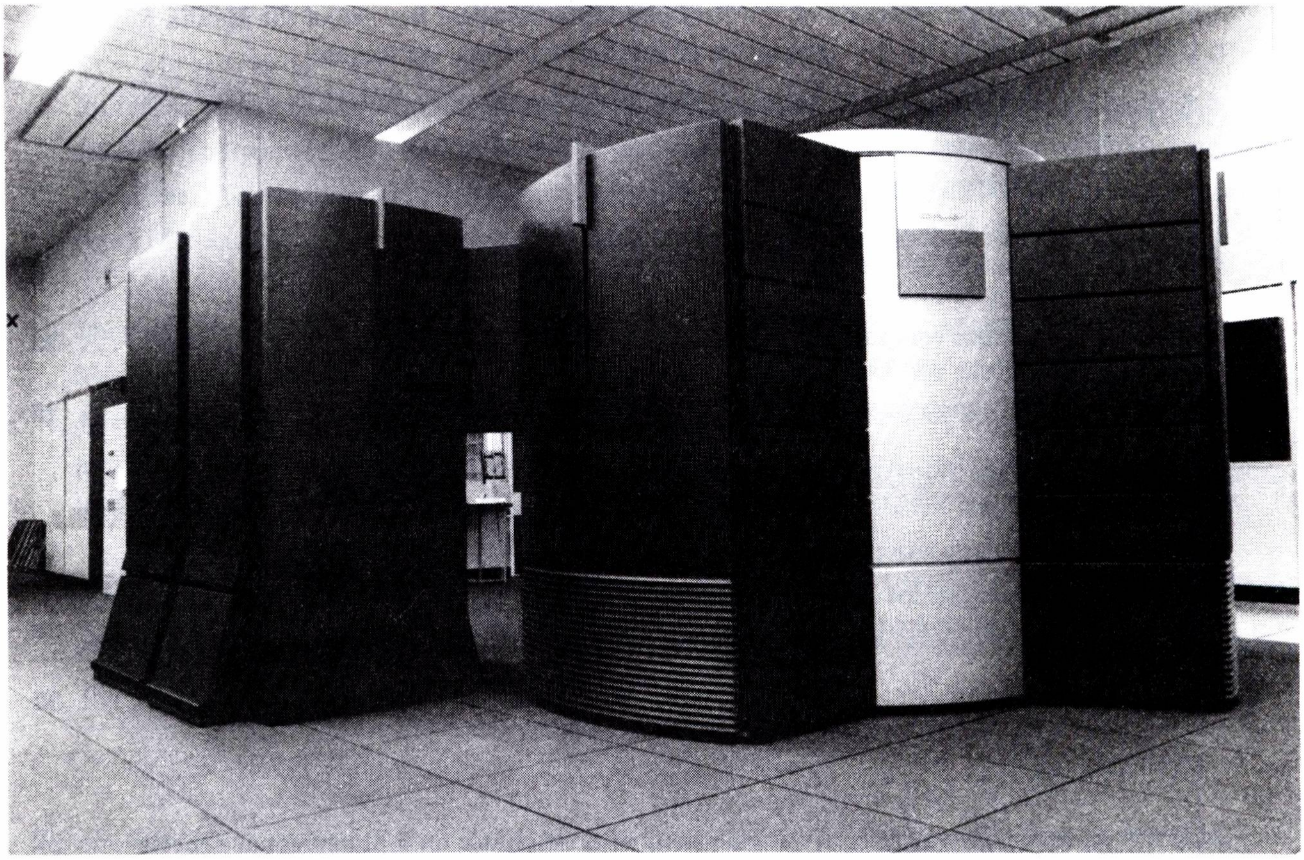

Shinfield Park, Reading, Berkshire RG2 9AX, England. Telephone: U.K. (0734) 499000,
International (+44 734) 499000, Telex: 847908 ECMWF G, Telefax (0734) 869450

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

IN THIS ISSUE

Editorial 1

METEOROLOGICAL

Changes to the operational forecasting system 2

The ECMWF multitasking model 3

Minimum temperature forecasts at the Regional Meteorological Service of the Emilia Romagna region (Northern Italy) by the application of the Kalman filter technique 9

COMPUTING

Installation of the CRAY YMP-C90 17

Fortran 90 and the CRAY compilers 22

Fortran 90 moves into compiling systems 23

Fortran 90 control constructs in CF77 5.0 25

COMPUTER USER INFORMATION

Computer resource allocation to Member States in 1993 30

Computer user training course 33

Still valid news sheets 35

GENERAL

ECMWF workshop on variational assimilation, with emphasis on 3-dimensional aspects . . . 37

The meteorological training course 38

ECMWF calendar 40

ECMWF publications 41

Index of still valid Newsletter articles 42

COVER: The CRAY YMP-C90, newly installed in the Computer Hall

---

This Newsletter is edited and produced by User Support.

The next issue will appear in March 1993.

An article on the ECMWF multitasking model in this edition describes the technical adjustments made to the current forecasting model to optimise its use of the Centre's available computer resources. The installation of the current supercomputer, the Cray YMP-C90, and the several years of planning and preparatory steps prior to that, are also described in a subsequent article.

This edition also comprises an article contributed by two visiting scientists from the Italian Meteorological Service, describing a study of the application of the Kalman filter technique. Another external contribution is the reprint of two articles on Fortran 90. These are the first of a series which have appeared in the NCAR SCD Computing News, and it is planned to continue with these reprints, as they are of considerable relevance for ECMWF computer users.

Standard items at this time of year are the announcements of the training courses (computer user and meteorological) for the next year, and the computer resource allocation to Member States in 1993.

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

CHANGES TO THE OPERATIONAL FORECASTING SYSTEM

Recent changes

A modification to the model post-processing was implemented on 7 December 1992, to improve the extrapolation of temperature and MSL pressure under model orography. The resulting extrapolated fields better continue the free atmosphere distribution of fields and show much less dependence on the overlying orography.

Planned changes

Improved representation of cloud/radiation interaction and of surface and planetary boundary layer processes are under development and will probably be implemented during the first half of 1993.

- Bernard Strauss

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

THE ECMWF MULTITASKING MODEL

Introduction

Today's production model has developed over a number of years and has run on a series of Cray machines ranging from the first Cray-1 to the current Cray YMP-C90. During this development, technical aspects such as multi-tasking and I/O have undergone considerable change. Some of these developments have been described in a series of Newsletter articles (March 1985, March 1986, March 1987). In September 1991, a substantial revision of the model code was introduced. The background to these developments has been described in the December 1991 Newsletter. This article outlines the technical changes needed to support this revision and indicates performance on today's computer (C90).

Two significant aspects, introduced in order to reduce the computational requirements and enable the forecast to be completed in an acceptable time, have a considerable impact on the technical framework of the model. These are:

(i) a semi-Lagrangian advection scheme;

(ii) a quasi-regular grid.

Semi-Lagrangian advection

Hitherto, data relating to a line of latitude was accessed only by the task dealing with that row during the grid point computations. Thus, there was no need for a task to have any knowledge of the location of data belonging to neighbouring rows. The semi-Lagrangian technique requires the calculation of trajectories, so that data must be available in a region surrounding each grid point. In particular, data to the north and south of the target point, generated by independent tasks, must be accessible to the task which 'owns' that grid point.

If the necessary data for each row is collected in a buffer, a set of such buffers must be available at any stage of the global scan, holding data for rows to the north and south of the current computational row. With a scanning order utilising north-south pairs (the original strategy), such a cluster of buffers would have to be maintained simultaneously for each hemisphere, resulting in a substantial memory requirement. By choosing a pole to pole scanning order and abandoning the hemispheric pairing, this problem is avoided. Independent tasks dealing with different row numbers can share a common pool of buffers.

The multitasking strategy is now simple and efficient. A lockstep mechanism (using EVENTS) keeps all active tasks working on neighbouring rows so that the predefined size of the buffer pool is always adequate. Otherwise, no synchronisation is necessary until the scan is completed at the South Pole. Scheduling is totally dynamic and takes advantage of any number of available processors. However, additional I/O has to be introduced.

In order to achieve this decoupling between hemispheres, the Legendre transform originally computed when each north-south pair was complete has to be delayed. This,in turn, requires that the Fourier components generated by the Fast Fourier Transform (FFT) on each row of data must be saved somewhere until all rows are completed. The primary memory requirement for this is large so a new workfile is introduced to store the components on secondary storage. An important memory economy is achieved because the space required to store the spectral data is not needed until the first scan has been completed and the direct Legendre transform commences.

Reduced grid

The reduced grid is defined such that east-west inter-grid distances remain approximately constant throughout the globe. Thus, polar rows contain only 12 points in comparison to equatorial rows which contain 640 points. There are now large imbalances in the quantity of work contained in many lines of latitude (see Fig. 1). This problem is minimised by the dynamic work sharing strategy and by the fact that the units of work are small in the run down phase as the South Pole is approached.

Note also that the imbalances due to the reduced grid are very much larger than those due to variations in the physical parametrization computations. The solid line in Fig. 1 shows variations in computation times (about 10%) when a full grid is employed.

Legendre transform

Since the Fourier coefficients have now been stored in a workfile, this gives an opportunity to reorganise the multitasking of the direct Legendre transform by splitting the work content according to the zonal wave number (M). Computation for each M is independent so that the reproducibility problem, present in the original strategy, is eliminated. An additional bonus comes in that the transform can be written as a product of matrices, allowing the use of a highly efficient matrix multiply library routine. It is also possible to increase vector lengths by careful data organisation allowing several meteorological fields to be transformed in a single matrix product. However, the I/O now required is more complex since, for each M, Fourier components for all rows must be available in memory. These exist as subsets of each record in the new workfile, calling for a 'transpose' access method (a compound read request).

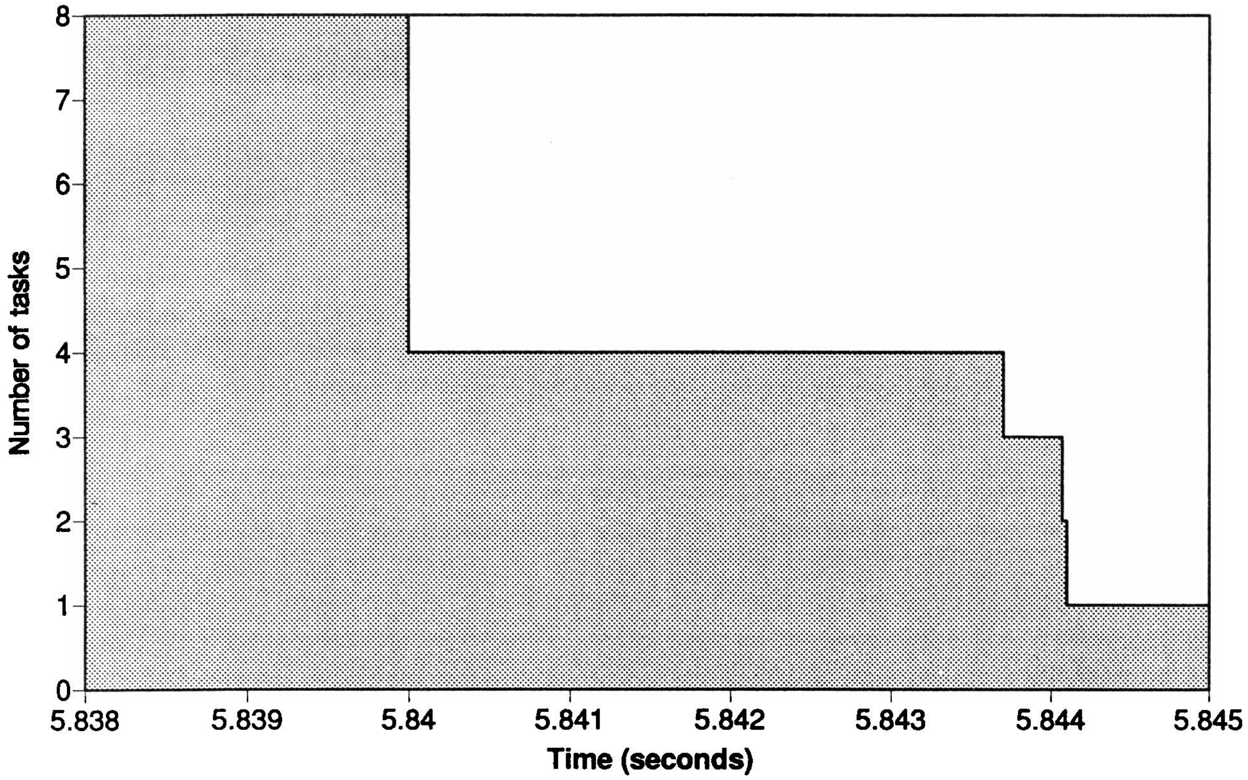

Fig. 1: Grid-point scan timings: time for each latitude row at T213

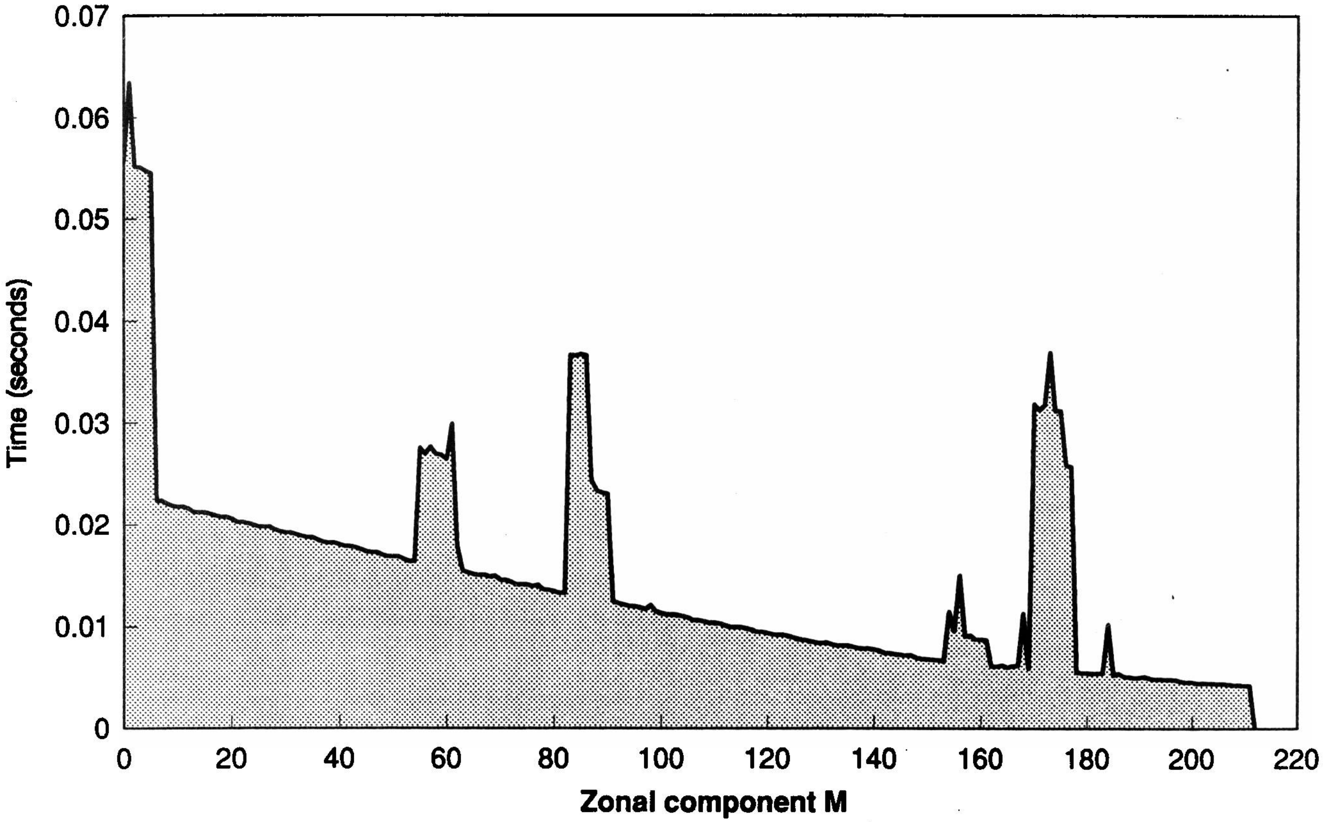

Fig. 2: Scan 2 timings: time for each M component at T213

As M increases, the work content reduces so that dynamic scheduling is again important. By commencing with M=0, the final work items (large M) contain the least computation and residual imbalance is minimised. The timings shown in Fig. 2 illustrate this strategy. The 'spikes' are caused by I/O activity reading Fourier and Legendre coefficients from work files.

I/O package

Given a desire to retain an out-of-memory solution for this model, the revised strategy for scanning and multitasking requires ever more complex access to the data in the work files. A rewritten model I/O scheme provides support for the new features and full flexibility for multiple buffer handling so as to allow (optional) asynchronous I/O. The design aim is to provide a clean and simple model-independent interface so that the package may be used, for example, by the IFS model. A small number of library calls are available to the model code so that data for a specific row can be requested and rewritten. All details of buffer handling are completely hidden from the application code. Varying record lengths such as those generated by a reduced model grid are catered for by means of an index. A further refinement allows all I/O to be inhibited so that a model can be made memory resident without code changes simply by providing sufficient buffer space to the I/O scheme.

The I/O requests made to the operating system can be changed to use any desired access method. The current implementation makes use of the Cray queued I/O library package because this offers an efficient access method coupled with an option for the compound I/O problem.

At the level of the operating system, there are several methods of handling the work files. Generally the most flexible is that of LDCACHE where active pages of a file are retained within SSD space. Much greater I/O speed can be achieved with secondary data segments (SDS) but it is somewhat less convenient to use from the scheduling point of view in a batch job system. However, current evidence points to a substantial system overhead in handling cached file systems. This overhead is almost eliminated with SDS.

For small model resolutions it is practical and efficient to consider running with memory resident work files. This can be achieved using the appropriate 'assign' statements and has the advantage of avoiding operating system calls and therefore minimising system overhead.

Further economy of SSD space and record lengths can be made by selectively packing work file data. In practice, this works well for grid point fields and can be accomplished very efficiently using library routines.

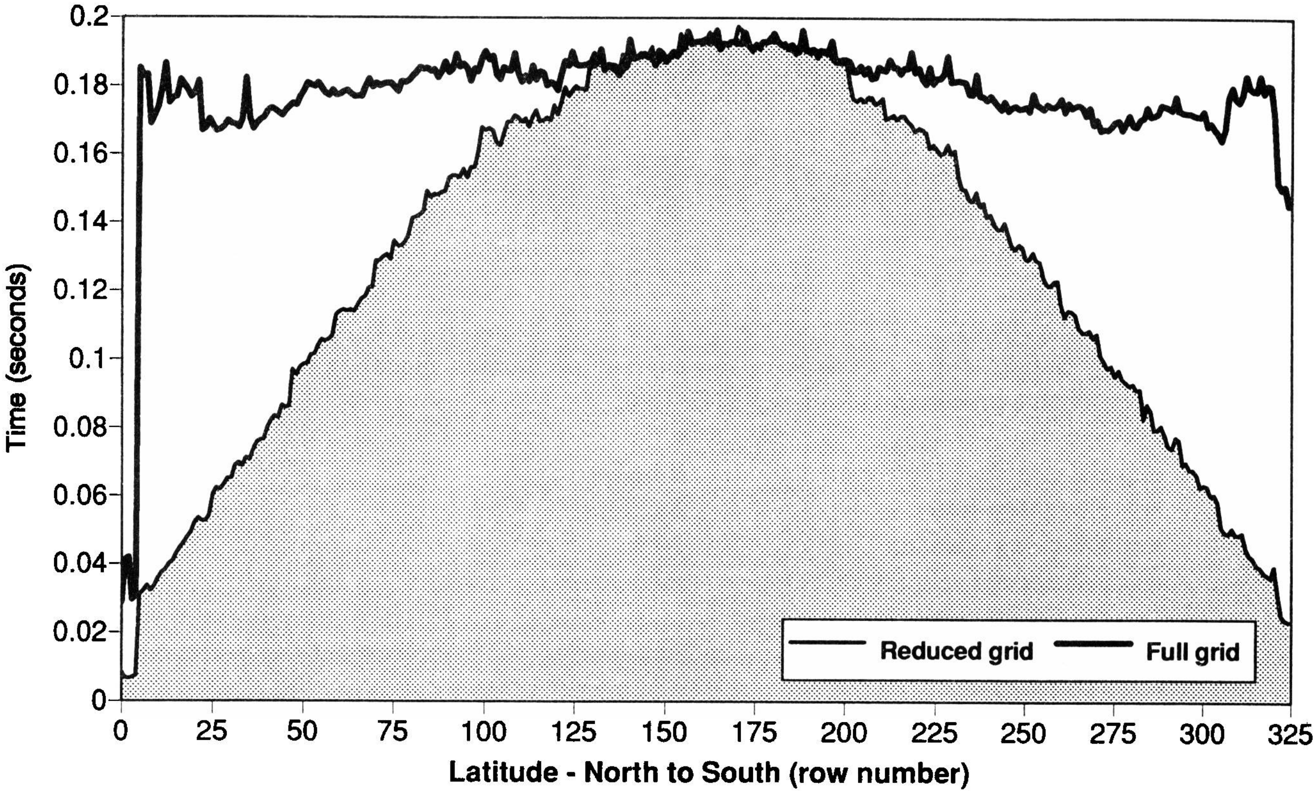

Fig. 3: Task imbalance: grid-point scan

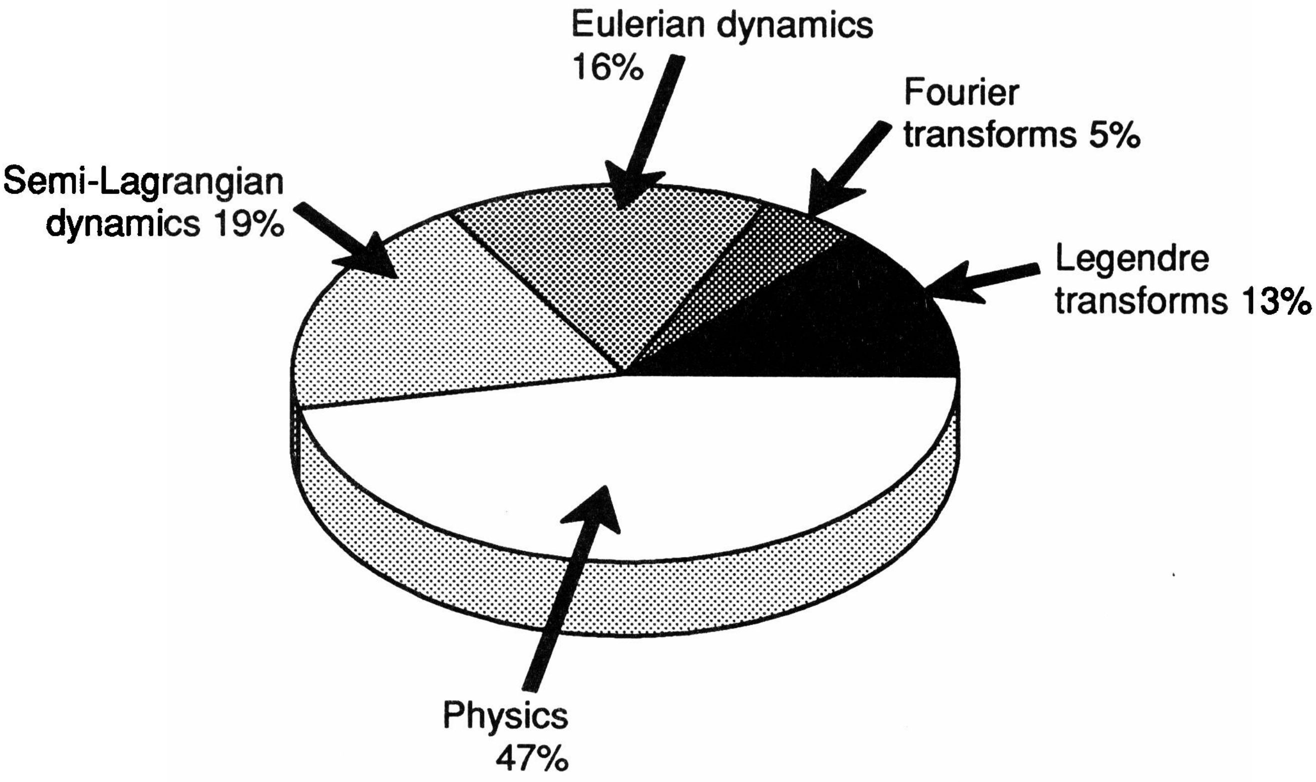

Fig. 4: Production model T213 L31: % CP costs

Performance figures

The following timing statistics relate to the current production model at resolution T213 L31 executing on twelve processors of a Cray YMP-C90:

elapsed run time (10 days)	2 hours
sustained flop rate	3 Gflops
multi-tasking speedup	10.5
cp performance relative to YMP	2.43.

Fig. 3 shows task activity in a small time interval during the completion of the grid-point scan. It illustrates the success of the dynamic load balancing strategy by showing the rundown process as the residual pieces of work complete. This takes about 4 milliseconds, or about 0.1% of the total time taken for the total scan.

IFS

The rewritten model known as the Integrated Forecasting System (IFS) is expected to become the operational model in the first half of 1993. From the technical point of view, most of the features have been carried over so that similar performances can be expected. There are several improvements:

(i) the work files may be individually disabled allowing 'in core' execution. Work file I/O is controlled by the I/O package described above;

(ii) the effect of the reduced grid has been minimised by blocking up short rows into a buffer to create chunks or 'super rows'. This improves the vector performance significantly but reduces the amount of parallelism. The blocking factor may be chosen at run time. The improved vector performance can be seen by comparing the sustained flop rate (5 Gflops) and the cp performance relative to YMP (2.66).

The relative costs of significant portions of the model are indicated in Fig. 4. The comparatively small cost of the semi-Lagrangian computation is a tribute to the success of significant optimisation effort.

- David Dent

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

MINIMUM TEMPERATURE FORECASTS AT THE REGIONAL METEOROLOGICAL SERVICE OF THE EMILIA ROMAGNA REGION (NORTH ITALY) BY THE APPLICATION OF THE KALMAN FILTER TECHNIQUE

Introduction

It is well known that the surface parameters, for example minimum and maximum two meter temperatures, are not well forecast by most "state of the art" global Numerical Weather Prediction Models (NWPM), such as the ECMWF operational model. This is essentially due to the unrealistic description of orography, especially in some regions where it is particularly complex and inhomogeneous. One such region is the Po Valley area of Northern Italy.

A post-elaboration of the Direct Model Outputs (DMO) of the NWP models is necessary in order to obtain realistic forecasts of surface parameters such as 2m temperatures, relative humidity, precipitation, cloudiness, etc.

Two general post-processing schemes have been proposed in the past and are used operationally in most weather services: Perfect Prognostic (PP), and Model Output Statistics (MOS) techniques (Glahn, 1980). Both these methods work in two steps: (1) evaluation of the statistical link between the upper-air meteorological parameters, the *predictors* (observed in case of PP or forecast by NWPM in MOS) and the surface parameters to be estimated, the *predictands*; and (2) utilization of these statistical relationships to the actual NWPM forecasts.

A great deal of literature exists on these two post-processing methods of NWP products. Both methods present two main common problems which can be summarized as

(i) the need for long data series of observed (or forecast by NWPM) weather data in order to obtain stable statistical relationships;

(ii) the impossibility of changing the statistical relationships obtained in step (i), except by repeating the statistical study from the beginning.

The second problem is particularly crucial and makes it difficult to adapt PP and MOS post-processing techniques to actual weather situations.

Kalman filtering (Gelb et al. 1989; Persson, 1989 and 1991) permits adaptation of the statistical link between predictors and predictands in a recursive manner very simply. The main idea is that the statistical relationships between predictors and predictands can be continuously and automatically updated taking into account the recent performances of the NWPM outputs.

In the case of a simple bias relation between predictors, *DMO* (Direct Model Output) and predictands, Y_t:

$$Y_t=X_0(t)+DMO \tag{1}$$

or in the case of a linear relation:

$$Y_t=X_0(t)+X_1(t)\bullet DMO \tag{2}$$

the time dependent coefficients $X_0(t)$ and $X_1(t)$ can be updated, according to an evolution equation (the state equation):

$$\vec{X}_{t+1}=F_t\,\vec{X}_t \tag{3}$$

where F_t is the linear evolution operator and X is either the coefficient X_0 in the simple model (1) or the vector (X_0,X_1) in the case of linear model (2). The scheme to iterate the estimates of X_0 and X_1 can be obtained by imposing that these estimates are unbiased and with minimum variance. More details on the procedure can be found in Gelb (1989).

In these "simple" Kalman filter models, (1) alone can reduce the DMO bias, and (2) could also further reduce other types of systematic errors (i.e. when they are dependent on the forecast DMO).

Results of the application of Kalman filter to the forecast of minimum temperature in Bologna

We have applied the two parameter Kalman filter (2) to obtain minimum temperature forecasts over Bologna (Northern Italy) using as predictor (DMO) the ECMWF 2m temperature forecasts (received at the Regional Met. Service of Emilia Romagna Region every day) on a grid point situated at 45°N, 10.5°E, i.e. more or less in the centre of Po Valley.

We started the application of the Kalman filter technique to the minimum temperature forecast because this parameter is of paramount importance to many users, and especially for farming (e.g. protection of crops from late frosts during the flowering season).

The period used to test the method covers the two years 1989 and 1990 when T106 was the operational model at ECMWF. We also tested the performances of the Kalman filter model in the period January-June 1992 when T213 was the operational version of the model at ECMWF. It should be remembered that the data was received in a 1.5° x 1.5° interpolation grid, and not in the original grid points.

In Fig. 1 scatter plots of +36 hours forecasts of minimum temperatures against observations are shown for the two spring periods March-May 1989 and 1990. DMO results are indicated by outline squares while DMO+Kalman filter are indicated by solid squares. It can be seen that the application of the Kalman filter drastically reduces the negative systematic error of DMO.

In Figs. 2a and 2b the mean monthly biases of minimum temperature forecasts are shown for +36 and +108 hours forecasts respectively for the two year period 1989-1990. It can be seen that the large (negative) bias of DMO is completely cancelled out by the application of the Kalman filter. In Figs. 3a and 3b the mean monthly RMS values for +36 and +108 hours forecasts of minimum temperatures are shown. DMO+Kalman filter gives very good results over the whole year (RMS around 2 - 2.5°C).

In Fig. 4a the RMS behaviour of +36 hours minimum temperature forecasts is shown for the six month period January-June 1992. Fig. 4b shows the results for the +108 hours forecasts. Comparing with Figs. 3a and 3b, we can see that T213 DMO gives better results than the T106 model. RMS errors of DMO are less than 4 degrees in the period January-April. The application of the Kalman filter technique improves the skill of the DMO forecast, especially during the months of May-June: at the short range (+36 hours, Fig. 4a) the reduction of the RMS error is about 1 - 1.5 degrees during the winter months (January-February) and more than 2 degrees during the months May-June. At the medium range (Fig. 4b) the application of the Kalman filter does not improve the results during winter and early spring very much; on the other hand, a considerable improvement of the skill is again visible during May and June. This might be due to less skill during winter in the synoptic scale medium range forecasts, compared to the less disturbed conditions during summer.

Conclusions

From these preliminary results it is evident that the application of a two-parameters Kalman filter to the ECMWF surface parameters forecast (DMO) permits reliable minimum temperature forecast over Bologna even at the medium range (+108 hours: 4-5 days).

T213 ECMWF-NWPM offers better DMO 2m temperature forecasts than the T106 version did; the application of Kalman filter improves the results again.

The technique is now operative at the Regional Meteorological Service of the Emilia Romagna Region and has been extended to 19 other weather stations.

Work is in progress to try different predictors, from the ECMWF model, with the Kalman filter (850 hPa temperature forecasts, low troposphere vertical stability and 850 hPa wind).

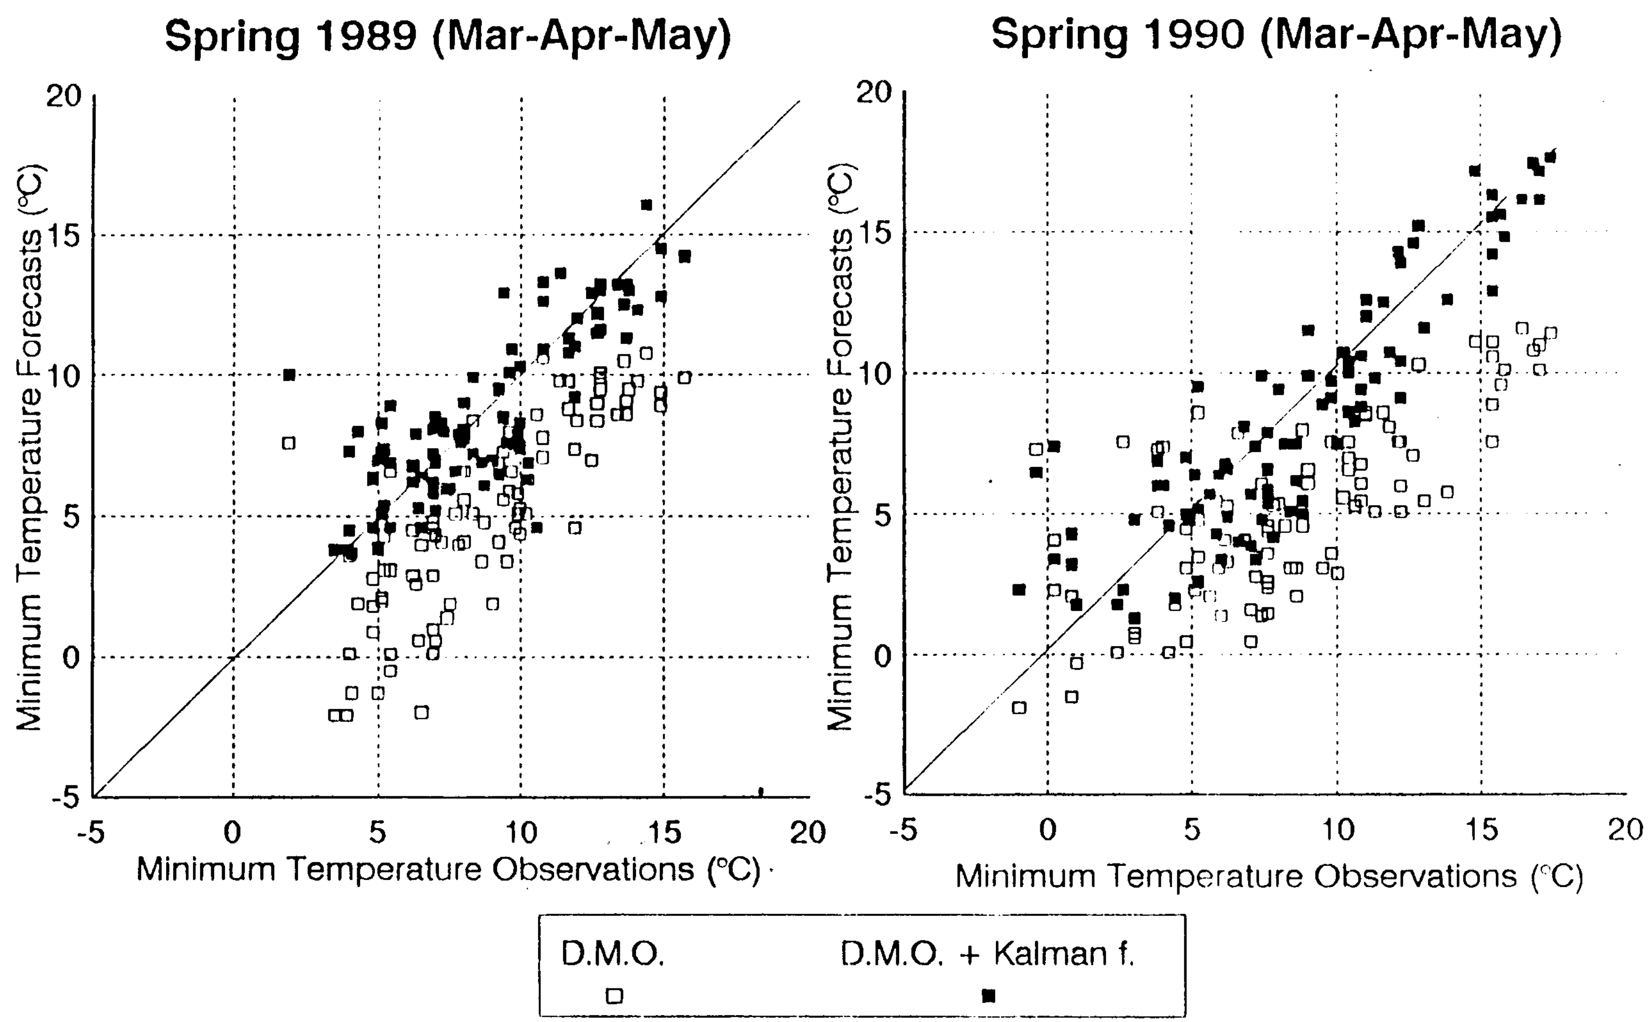

Fig. 1: **Scatter plot minimum temperature observations-forecasts**
Station: Bologna. Forecast time: + 36 hours

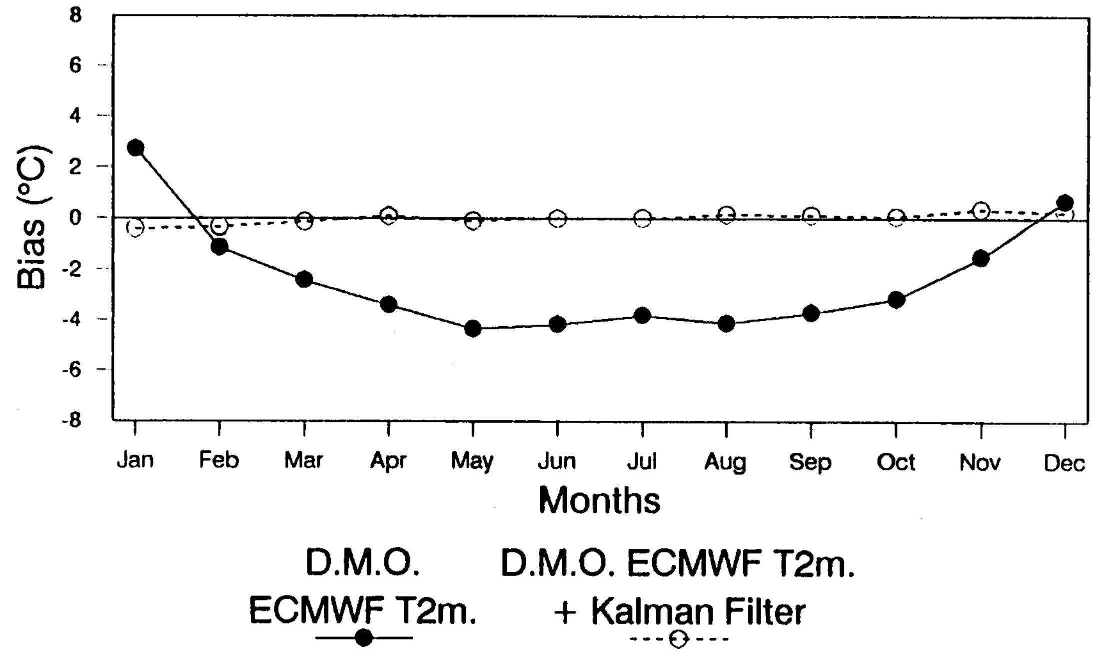

Fig. 2a: **Bias of minimum temperature forecast**
Station: Bologna. Forecast time: +36 hours
Period: Two years: 1989-1990

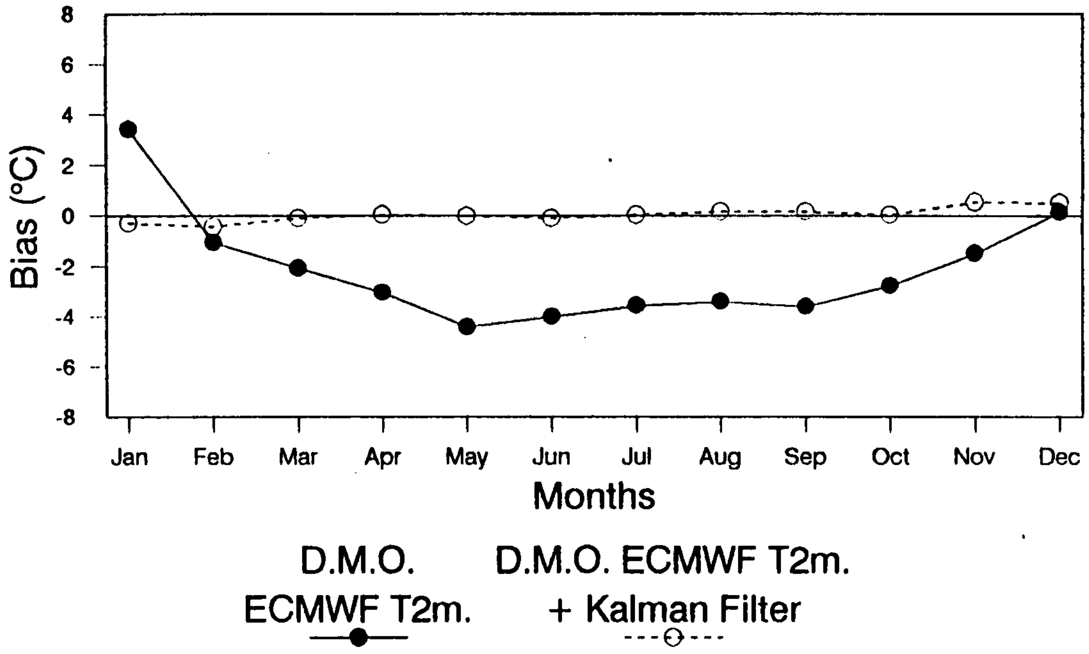

Fig. 2b: **Bias of minimum temperature forecast**
Station: Bologna. Forecast time: +108 hours
Period: Two years: 1989-1990

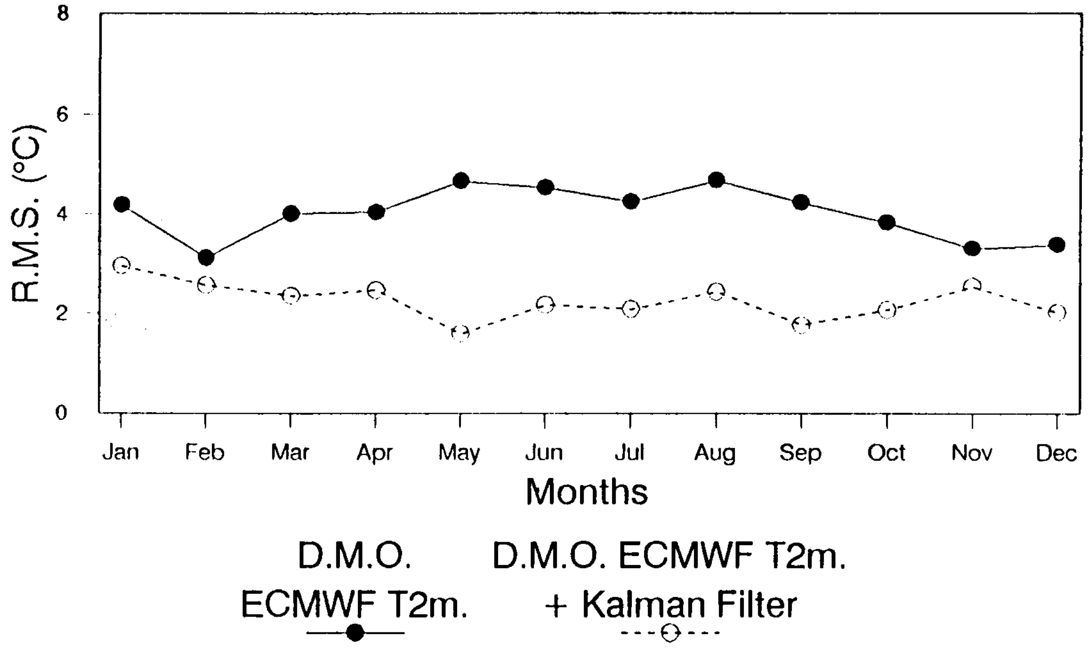

Fig. 3a: **RMS of minimum temperature forecast**
Station: Bologna. Forecast time: +36 hours
Period: Two years: 1989-1990

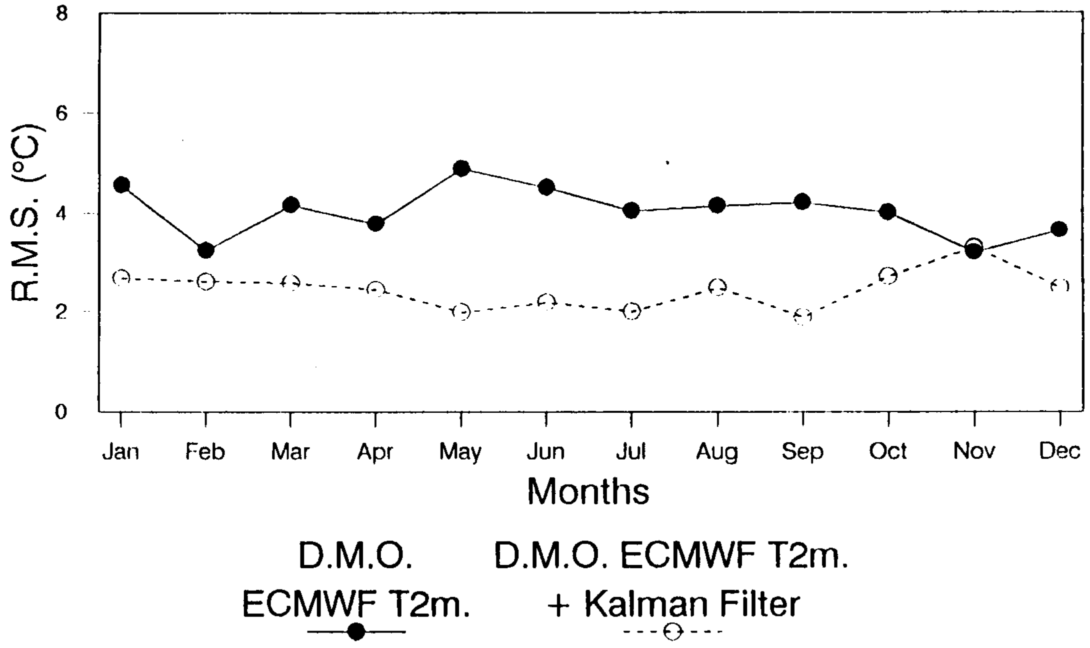

Fig. 3b: **RMS of minimum temperature forecast**
Station: Bologna. Forecast time: +108 hours
Period: Two years: 1989-1990

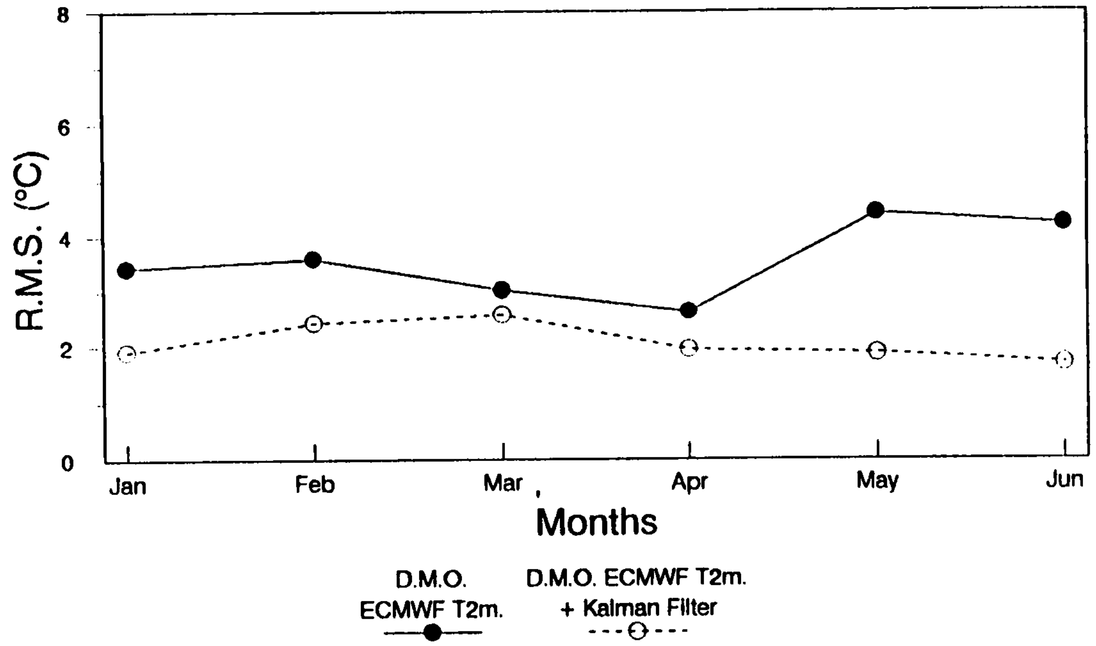

Fig. 4a: **RMS of minimum temperature forecast**
Station: Bologna. Forecast time: +36 hours
Period: Jan-Jun 1992

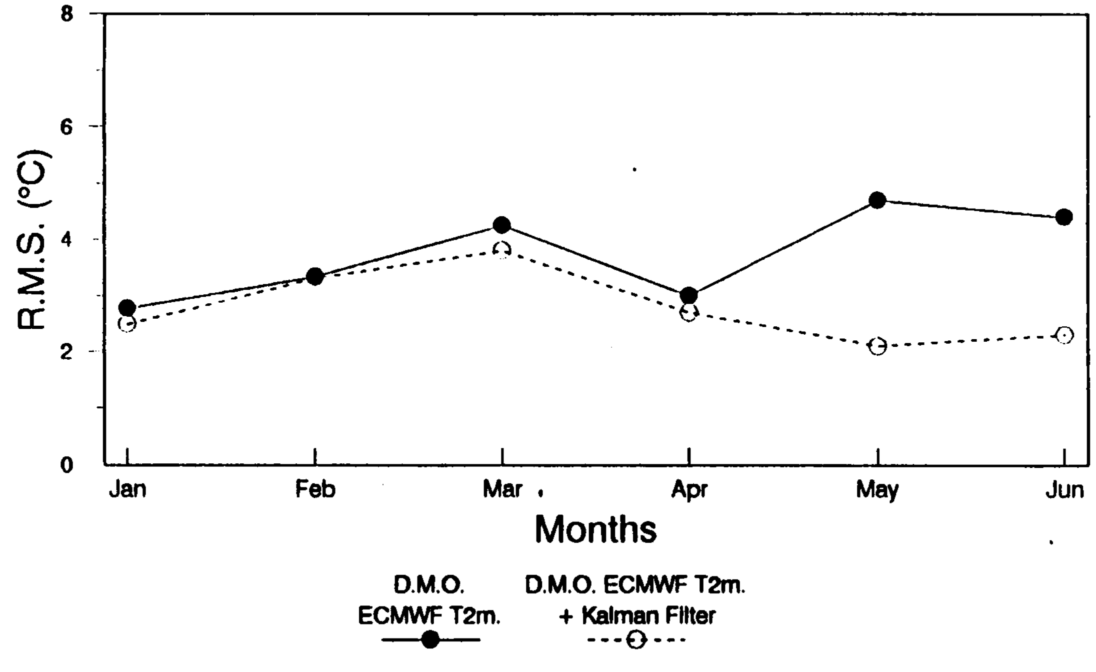

Fig. 4b: **RMS of minimum temperature forecast**
Station: Bologna. Forecast time: +108 hours
Period: Jan-Jun 1992

Acknowledgement

We thank Anders Persson of ECMWF for very useful suggestions about the application of the Kalman filter technique to the forecasting of surface meteorological parameters.

- Carlo Cacciamani, Servizio Meteorologico Regionale - ERSA, Bologna
- Costante de Simone, Servizio Meteorologico dell'Aeronautica Militare, Roma

References

Glahn, H.R., 1980: Methods and accuracy of statistical weather forecasting in the United States. WMO Symposium on Probabilistic and Statistical Methods in Weather Forecasting; Nice, 8-12 September 1980.

Gelb, A., J.F. Kasper Jr., R.A. Nash Jr., C.F. Price and A.A. Sutherland Jr., 1989: Applied Optimal Estimation. The M.I.T. Press; Massachusetts Inst. of Techn., Cambridge.

Persson, A., 1989: Kalman Filtering. A new approach to adaptive statistical interpretation of numerical meteorological forecasts. ECMWF Newsletter No. 46, June 1989.

Persson, A., 1991: Kalman Filtering. A new approach to adaptive statistical interpretation of numerical Meteorological Forecasts. In Glahn, H., et al.: Lectures presented at the WMO Training Workshop on the interpretation of NWP products in terms of Local Weather Phenomena and their verification. Wageningen, 29 July - 9 August 1991. WMO/TD No. 421.

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

INSTALLATION OF THE CRAY YMP-C90

Preparatory work for the Cray YMP-C90 installation actually began almost 4 years ago in late 1988. At that time the Computer Hall was full of equipment, the air conditioning system was almost life expired, the false floor was unsafe in some areas, and power and air conditioning provision at the site was inadequate for the expected equipment.

The choice of computers to succeed the Cray X-MP/48 was not made until late 1989 and it was therefore impossible to make detailed plans until then. However, we were able to begin planning work which has taken more than 4 years effort and which is still not fully completed.

Infrastructure work carried out over the 4 years includes:

- replacement of the Computer Hall air handlers, involving an airflow reversal to allow under-equipment cooling;
- replacement of the fresh air handling equipment;
- addition of a fifth chiller to increase cooling capacity;
- upgrading of chilled water pumps to increase cooling capacity;
- replacement of the false floor supports and the addition of horizontal stringers and spot floor tile replacement;
- installation of Euro Diesel electrical power conditioning equipment providing a much more reliable electrical supply;
- replacement of the primary electrical power distribution system to increase capacity.

Much of the above work would have had to be done whatever equipment had been selected, but all of it was necessary to enable the Centre to continue to provide a high quality computer service.

During 1990 the Cray Y-MP/8 was installed and we began to receive information from Cray about the needs of the C90. The first information received gave rise to concern in that initial estimates were that the power consumption, and thus also the cooling requirements for the system, would be in excess of 500 kW. This would clearly have overwhelmed the standby electrical power supplies and cooling capacity of the site and initial plans were made to resolve the problems.

Fortunately, as the C90 development progressed at Cray, estimates for power and cooling needs reduced and although we did have to install additional electrical equipment, it was possible to perform a relatively minor cooling system upgrade to accommodate the equipment.

Having ensured that the basic site infrastructure could support the C90 we were able to start detailed planning for installation late in 1991. A new computer has essentially three requirements:

- physical space;
- electrical supplies;
- cooling support.

Taking physical space first, a site was needed on the one hand on the north side of the Computer Hall, in order to keep refrigeration lines to the plant room short, but on the other hand on the south side of the Computer Hall where the floor is stronger! The solution was to install the system on the north side of the centre line, but with the heavy units set (just) on the stronger part of the floor. The central processing unit weighs almost 7 tonnes and is supported a few millimetres above the false floor on 12 very strong jacks which sit on the concrete floor.

We were aware that the IBM 3090 model 150E would be removed early in 1992 and since this was situated next to the Cray Y-MP/8 this was considered to be an ideal location. Various equipment had to be moved to clear the area and the false floor was replaced in readiness for the C90 installation.

The next problem was the provision of electrical supplies. The standard electrical supplies in the UK are 50 hz, 415 volt 3 phase, 240 volt single phase. These are slightly higher voltages than the rest of Europe but the same frequency. They are a different frequency as well as different voltages from those in the USA.

Traditionally, Cray computers have used 50 hz, 415 volt 3 phase supplies for their internal cooling systems, 400 hz 208 volt 3 phase supplies for the computer logic and 60 hz 208 volt 3 phase supplies for peripherals. Providing 400 hz and 60 hz supplies generally requires the use of rotating frequency converters.

Over the past few years efforts have been made with Cray to eliminate the need for 60 hz supplies, and now there has been partial success. No equipment needs 60 hz power though some needs American standard voltage (50 hz, 208 volt 3 phase). This is easier to provide as it needs only a transformer.

The 400 hz requirements for the C90 are almost twice those of the Cray Y-MP/8. It seemed that there would be a considerable problem here because the 150 kW 400 hz generator used to supply the Y-MP/8 takes a very large startup current which has caused problems on more than one occasion. Of particular note was the incident on Christmas Day 1991 when the 400 hz set was started while the site was being supplied by the old diesel generators. The net effect was that the entire site suffered a short power cut, and Christmas celebrations were blighted for a number of engineers!

Cray made two alternative proposals, either to run two 150 kw sets in parallel, or to use a new 300 kw set. There were advantages and disadvantages in both proposals. The parallel 150 kw approach was attractive because it was a proven machine and there would be the possibility of keeping an old unit as a spare. However, this did not eliminate the startup problems.

At the same time, Cray's suppliers, KATO, were working on a new 300 kw set which was physically only about 10% larger than the 150 kw set and, more importantly, required a much lower startup current.

At the time when the decision had to be made, in spring 1992, there were two other difficulties. The paralleling unit for the two 150 kw sets could not be made to work, and the newly-designed 300 kw set would not run cool enough to be reliable! Fortunately, both sets of problems could be resolved and the Centre elected to take a 300 kw set, with a second spare unit currently stored next to the production unit.

Finally came the problem of cooling. The capacity of the Centre's chilling system in late 1989 was clearly insufficient to cool two Cray systems running in parallel so it was necessary to uprate the system. This was done in early 1990, increasing the capacity from around 1140 kW cooling to around 1500 kW cooling. The system was uprated by adding a fifth chiller unit and uprating the pumps to increase the flow rate of the chilled water. This uprating provided just enough cooling capacity to allow the Cray Y-MP/8 and the Cray C90 to run, in parallel during the summer months, albeit at the cost of switching off the air conditioning in the Council Chamber and the Lecture Theatre!

The installation of the cooling system was done in a rather unusual way. The system requires an 80 ton capacity fridge (approximately 270 kW cooling capacity). Normally we would install such a unit in the plant room below the Computer Hall. On this occasion it was decided to install it in the Computer Hall next to the C90. This simplified the pipework and released valuable plant room space.

Having prepared the site, the actual installation was something of an anti-climax. Cray had chartered a Boeing 747 freighter to bring the Centre's C90, and a Cray Y-MP/8 which was being delivered to another UK customer, directly from Chicago to London Heathrow.

The actual delivery and installation schedule was as follows:

4 Aug 1992	Delivery began at about 1500. It took around 6 hours to unload the equipment and place it in position.
5 Aug 1992	The machine was powered up at around 1800. This was a few hours later than expected due to a refrigeration leak and a faulty thermistor. When the machine was powered up the engineers found one faulty CPU card and one faulty power supply which were both replaced. Otherwise the machine worked first time. Hardware checkout then began.
9 Aug 1992	UNICOS was loaded and the acceptance test jobs were used to exercise the system.
10 Aug 1992	Local Area Network connections and links to other mainframes were started. This work continued until 13 August.
14 Aug 1992	24 hours provisional acceptance test began.

At this stage it appeared that we had achieved an acceptance test 3 days earlier than planned and the system looked to be very reliable. However, this was not quite the case....

When we began to run "normal" work rather than the special test jobs which had been carefully debugged, a CPU design fault was revealed. The effect of the problem was that jobs with software errors could, under certain circumstances, overwrite the memory of other jobs in the machine. When the problem was isolated at around 2100 hours on 19 August, the situation looked very bleak. However, Cray were able to understand the problem quickly, redesign a CPU chip in a few hours, and manufacture sample replacements within a couple of days. This was truly a tremendous effort on their part.

At this stage it was agreed that the system would formally be "down" but that the Centre could make use of as much of the system as could be made available. All 12 CPU modules, and the spares, had to be returned to Chippewa Falls to have the redesigned chip fitted. The logistics of this exercise were huge because the modules are heavy and relatively fragile. Surprisingly it was possible to ship a module to Europe in about 24 hours, but returning modules to the USA usually took three to four days. Most of that time was spent at the airport in Chicago!

The whole process of returning the CPUs to Chippewa, three or four at a time, having them upgraded, and returning them to ECMWF took until 23 September when the machine was considered "up" again. 45 days later the machine successfully completed its final acceptance test and has continued to run reliably ever since.

- Peter Gray

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

FORTRAN 90 AND THE CRAY COMPILERS

Now that Fortran 90 has become an international standard, vendors are beginning to work on compilers to support it. Cray has a double approach to providing Fortran 90 support.

(i) A new compiler is being written to support the full Fortran 90 standard. Their current plan is to release it in 3Q1993.

(ii) Some elements of Fortran 90 have been incorporated into the existing cft77 compiler.

At ECMWF users may wish to begin to investigate these Fortran 90 extensions. To help in this process we are reproducing in the ECMWF Newsletter a series of articles by Jeanne Adams which first appeared in NCAR's SCD Computing News. These articles describe various Fortran 90 extensions to cft77 in a clear and concise way, giving a very useful introduction for those wishing to experiment. The author, Jeanne Adams, chaired the International Programming Languages Committee of ISO (International Standards Organisation), and chaired the ANSI committee that developed Fortran 90. She also recently co-authored the book "*The Fortran 90 Handbook: Complete ANSI/ISO Reference*".

Although Fortran 90 is now an international standard it was mainly designed for serial computer systems. A new initiative is underway to define a version of Fortran suitable for the highly parallel computer systems which are now coming on to the market. This initiative, known as High Performance Fortran, is an attempt by a mixed group of users and vendors to define a subset of Fortran 90 (in particular the array syntax) plus an agreed set of compiler directives as a "recommended" version of Fortran for use on highly parallel systems. Cray Research, amongst others, is part of this initiative. Hopefully, by the time highly parallel machines become more widely used, there will be an agreed "high performance Fortran" that will aid portability between different highly parallel systems.

Finally, one word of caution about using Fortran 90 at this time. Although Cray, and others, are producing compilers, there is at this stage not much in common between the various vendors. Hence, a Fortran 90 command implemented by one vendor may not yet be available in the Fortran of another vendor. Thus, although one of the major ideas of a standard is to aid portability, be aware that as of today programs containing Fortran 90 constructs may be less portable than those restricted to Fortran 77 only.

- Andrew Lea

The following articles are reprinted here by courtesy of SCD Computing News. The first article appeared in the November 1991 issue of that publication and the second in the December 1991 issue.

Fortran 90 moves into compiling systems

Fortran 90, the new international Fortran standard, is being incorporated by several vendors into their compiling systems. One company that has already released a Fortran 90 compiler is the Numerical Algorithms Group (NAG). The NAG library has been compiled on this software, and a number of releases have been sold. The compiler is UNIX-based.

Cray Research is at work developing a compiling system that complies fully with the new standard. Their goal is to provide a system that is efficient and fast, with both the upward-compatible features and performance of CF77 (their current compiling system) and the new features of Fortran 90.

In the meantime, Cray has incorporated a number of Fortran 90 features into CF77. In the Cray manual *CF77 Compiling System, Volume I: Fortran Reference Manual,* Version 5.0 (SR-3071), these features are described as extensions to FORTRAN 77. I am presently testing these features and comparing them with the function described in the Fortran 90 standard. Brief accounts of the results will be given.

The following is a list of Fortran 90 features supported in Cray's CF77 compiling system.

Control structures

- DO WHILE, END DO, nonlabeled DO. Some extensions to the FORTRAN 77 DO have been added, as well as the DO WHILE looping convention.

Subprograms

- RECURSIVE functions and subroutines. Both direct and indirect recursion are indicated by prefixing the function or subrouting with the word RECURSIVE.

- MIL-SPEC bit intrinsics. The military specification (MIL-SPEC-1753) intrinsics for bit data are: BTEST, IAND, IBCLR, IBITS, IBSET, IEOR, IOR, ISHFT, ISHFTC, and NOT.

Data

- Hexadecimal (Z) and octal (O) constants. Both hex and octal constants may be represented.

- IMPLICIT NONE. IMPLICIT NONE declares that all variables must be explicitly typed; this statement must appear first before other type specification statements.

- Character and noncharacter data may be mixed in the same COMMON block. The restriction has been removed that either character or noncharacter data, but not both, must appear in the same COMMON block.

Arrays

- Automatic arrays.

- Subset of array syntax (a subset of the array syntax, including whole and partial array assignment and subscript triplet notation). Cray Research currently does not support nonsequence arrays as actual or dummy arguments, passing vector valued subscripted arrays, or allocatable arrays (that is, when memory is allocated for the array only while the array is being used). **Note:** Cray does, however, support allocatable arrays in subroutines, in that when an array is declared solely in a subroutine, that array's needed memory is allocated only during the subroutine's execution. You can also pass into a subroutine an argument declaring the array's memory size at run time. In addition, you can allocate or deallocate memory by using the Cray-specific HPALLOC or HPDEALLOC calls.

Input/Output

- NAMELIST. A name may be given to a list, and subsequently used in input/output statements instead of the list.

- OPEN and INQUIRE specifiers. All OPEN and INQUIRE specifiers in Fortran 90 are included. (The INQUIRE statement provides various kinds of information about a file. It is documented in the *CF77 Compiling system* reference manual).

Miscellaneous and formatting

Note: More flexibility has been provided with new alternatives for names, delimiters, and repeat counts.

* INCLUDE line. The INCLUDE directive enables a file containing source code in Fortran to be included in a program where the INCLUDE line appears.

* Quotes, "", as string delimiters.

* 31-character names plus underscore, upper/lower case allowed. (Other normal Fortran rules for naming still apply.)

* Exclamation mark for inline comments.

* Character string edit descriptor can be delimited with quotes.

* Repeat count for slash editor (for example, 3/ means ///).

Cray Research Fortran has pointers, but they differ from the pointers defined in Fortran 90. Differences will be highlighted in future articles.

Jeanne Adams

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

Fortran 90 control constructs in CF77 5.0

Version 5.0 of CF77, the Cray Research compiling system, contains a number of extensions to FORTRAN 77 based on the new Fortran 90 standard. This article discusses the control constructs from Fortran 90 available on the CF77 5.0 compiling system. These constructs include the DO WHILE statement and a few of the DO loop extensions.

Two DO WHILE examples

The two examples of DO WHILE statements given below are valid on CF77 5.0. These code fragments have been tested on a CRAY Y-MP8/864 running UNICOS 6.1.4.

Example 1. DO WHILE loop

```
PROGRAM TESTLOOP

IMPLICIT NONE

INTEGER POWER_OF_2, I

POWER_OF_2 = 1  !THE ZERO POWER OF 2

  DO WHILE (POWER_OF_2 .LT. 1000)

                    PRINT *, POWER_OF_2

                    POWER_OF_2 = 2 * POWER_OF_2

  END DO

END
```

Example 2. Statement label on DO WHILE loop

```
      PROGRAM TESTLOOP

      IMPLICIT NONE

      INTEGER POWER_OF_2, I

      POWER_OF_2 = 1  !THE ZERO POWER OF 2

        DO 10 , WHILE (POWER_OF_2 .LT.1000)

                    PRINT *, POWER_OF_2

                    POWER_OF_2 = 2 * POWER_OF_2

10    END DO              !or 10 CONTINUE

      END
```

Example 3. END DO to end loop

```
      PROGRAM TESTLOOP

      IMPLICIT NONE

      INTEGER POWER_OF_2, I

      POWER_OF_2 = 1  !THE ZERO POWER OF 2

        DO I=1,100

                    PRINT *, POWER_OF_2

                    POWER_OF_2 = 2 * POWER_OF_2

                    IF (POWER_OF_2 .GE. 1000) GO TO 10

        END DO

10    CONTINUE

      END
```

Example 1 used END DO, the Fortran 90 form for the end of the loop. The loop will continue until the condition in the DO WHILE statement is met.

Note: 31-character variable names are allowed, as well as the use of underscore in names. These features enable code to be written with variable names that improve programm readability.

The label and the comma (,) are optional, but could be used as in Example 2.

IMPLICIT NONE, used in Examples 1 and 2, is a feature new to Fortran 90, and not included in FORTRAN 77. A program unit that contains an IMPLICIT NONE statement must have type declarations for all variables. It is safer programming practice to use this feature.

Extensions to the FORTRAN 77 DO loop

The FORTRAN 77 form of a DO loop is entirely valid in Fortran 90. However, there have been a number of extensions, two of which are already included in the CF77 5.0 compiling system. One is that the statement label required for FORTRAN 77 is now optional; the second is an END DO statement, which has been used in Example 3. (If a label were used in the DO statement, the END DO would have that label; a CONTINUE, along with that label, could replace the END DO.)

Note also that in Examples 1-3 the exclamation mark (!) is used for commentary. This is a Fortran 90 feature that has been in CF77 for a very long time. It is *not* a FORTRAN 77 feature, but a Cray extension. (Possibly the standards committee had this, among other conventions, in mind.)

Fortran 90 has added two other new statements that are not yet in the CF77 compiling system. These are the CYCLE and EXIT statements. If the full Fortran 90 were implemented, you could have said in Example 3

```
IF (POWER_OF_2 .GE. 1000) EXIT
```

and deleted the 10 CONTINUE statement. In Fortran 90 there is also a DO statement without any variable or conditional control.

Summary

To summarize, the syntax for the DO loop currently on the CF77 5.0 compiling system could be stated as in Fig. 1.

Figure 1. Syntax for CF77 5.0 DO loop

Syntax for the DO loop on CF77 5.0 can be stated as follows, using square brackets for optional items:

```
        DO  [label]  loop_control
```

where

```
        loop-control    is      [,]     do_variable = initial, limiting [,increment]

                        or      [,]     while  (scalar_logical_expression)
```

Jeanne Adams

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

COMPUTER RESOURCE ALLOCATION TO MEMBER STATES IN 1993

Table 1: Allocation of CRAY resources and data storage by Member State in 1993 (including a 10% reserved allocation for Special Projects)

Member State	Cray (kunits)	Data (Gbytes)
BELGIUM	111	20.2
DENMARK	94	17.2
GERMANY	480	87.3
SPAIN	169	30.8
FRANCE	394	71.6
GREECE	76	13.8
IRELAND	68	12.4
ITALY	343	62.3
YUGOSLAVIA*	84	15.3
NETHERLANDS	138	25.1
NORWAY	89	16.2
AUSTRIA	101	18.4
PORTUGAL	72	13.0
SWITZERLAND	125	22.7
FINLAND	92	16.7
SWEDEN	119	21.6
TURKEY	84	15.2
UNITED KINGDOM	331	60.2
SPECIAL PROJECTS	290	60.0
TOTAL	3260	600.0

* At its 37th session on 2-3 December 1992 the Council decided that the telecommunications link between ECMWF and Belgrade would be terminated with immediate effect. As a consequence access to the Centre's computer system is not available to Belgrade.

Table 2: Special Project Allocations for 1993

Member State	Institution	Project title	1993 Resources - proposed allocations		1993 Resources requested
			Cray Kunits	CFS Gbytes	CFS Gbytes
<u>Continuation Projects</u>					
Austria	Institut für Meteorologie und Geophysik, Vienna (Hantel)	Subsynoptic vertical heat fluxes: Comparison diagnosed vs modelled data	0.8	1.5	1.5
Finland	FMI, Helsinki (Lönnberg)	The HIRLAM 3 project	15	5	5
France	CNET/CRPE (Eymard)	Determination of ocean surface heat fluxes using satellite data and the ECMWF model	40	2	2
	Univ Science & Technology, Lille (Vesperini/Fouquart)	Use of earth radiation budget data for verification of ECMWF cloud and radiation outputs	20	1	1
	METEO France, Toulouse (Cassé)	AVISO. Study of surface winds and surface fluxes at the interface ocean/atmosphere.	5	1	1
	LMD, Palaiseau (Duvel)	Validation of spatial and temporal variabilities of the ECMWF model	3	2	2
Germany	Institute for Geophysics and Meteorology (Speth)	Interpretation and calculation of energy budgets	2	4	6
	GKSS, Geesthacht (Raschke/ Rockel)	Parametrization of radiation and clouds for use in general circulation models	0.5	0.2	0.2
	MPI Hamburg (Roeckner)	Modelling the earth's radiation budget and evaluation against ERBE data	40	5	5
	MPI, Hamburg (Bengtsson)	Numerical experimentation with a coupled ocean/atmosphere model	70	13	20
Italy	Istituto per lo Studio della Dinamica delle Grandi Masse, Venezia (Cavaleri)	Testing and applications of a third generation wave model in the Mediterranean Sea	3	1	1
Netherlands	KNMI, De Bilt (Janssen)	Testing and evaluation of a third generation ocean wave model at ECMWF	1	2	2

Table 2 (cont.):

Member State	Institution	Project title	1993 Resources - proposed allocations		1993 Resources requested
			Cray Kunits	CFS Gbytes	CFS Gbytes
	KNMI, De Bilt (Siegmund)	Analysis of a CO_2-experiment performed with a GCM	1	0.5	0.5
	KNMI, De Bilt (v Dorland)	CO_2 transient atmosphere model	2.5	1	1
	KNMI, De Bilt (Cuijpers/ Duijnkerke)	Large eddy simulation of (strato)cumulus clouds	1	1	1
	KNMI, De Bilt (Können)	Climatological scenarios	0.5	3.8	4
New projects					
Austria	Institut für Meteorologie und Geophysik, Vienna (Hantel)	Diagnostic tests of the ECMWF physics package	-	-	-
Finland	Fin.Inst.Marine Research (Kahma)	Implementation of WAM wave model to the Baltic Sea	1.3	2	3
France	L.A.M.P., Aubière (Chaumerliac)	Chemistry, cloud and radiation interactions in a meteorological model	10	4	5
Germany	MPI, Hamburg (Bengtsson)	Simulation and validation of the hydrological cycle	40	4	5
Total allocated/requested			256.6	54.0	66.2
Reserve (to be allocated by ECMWF)			32.6	6.0	6.0
Overall total allocated/requested			289.2	60.0	72.2

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

COMPUTER USER TRAINING COURSE

The Centre is offering a computer user training course for Member States' personnel and ECMWF staff from 1-18 March 1993. Full information and a request for nominations has been sent to all Member State meteorological services.

The course is divided into three one-week modules, and attendees may register for separate modules.

Week 1: An introduction to UNIX for those who have no knowledge of this operating system.

Week 2: UNICOS extensions to UNIX, ECMWF utilities and the ECFILE file system.

Week 3: MAGICS and MARS.

Each week will consist partly of lectures and partly of practicals. In more detail, the three modules are:

MODULE 1 (1 - 5 March 1993) UNIX

Introduction to UNIX history and basic structure
Introduction for the file system
Basic commands
File manipulation and attributes
I/O commands
Basic shell scripts

MODULE 2 (8 - 12 March 1993) ECMWF'S UNICOS SERVICE

System and hardware overview
UNICOS batch jobs
FORTRAN
ECFILE file storage system
Specialist file services, including **sendtm**

Those attending module 2 are expected to know basic UNIX commands and be able to use the vi editor.

MODULE 3 (15 - 18 March 1993) MARS & MAGICS

MARS

- Overview
- MARS data
    - Data format
    - Archive contents
- MARS utility
    - System description
    - User interface

MAGICS

- Introduction and overview
- Concepts
- Parameters
    - Subroutines
    - Action and pseudo-action routines
    - Data input
- Plotting features.

Those attending module 3 are expected to know basic UNIX commands, be able to use the vi editor and to be able to submit jobs to UNICOS.

- Andrew Lea

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

STILL VALID NEWS SHEETS

Below is a list of News Sheets that still contain some valid information which has not been incorporated into the Bulletin set (up to News Sheet 289). All other News Sheets are redundant and can be thrown away.

No.	Still Valid Article
204	VAX disk space control
205(8/7)	Mispositioned cursor under NOS/VE full screen editor
207	FORMAL changes under NOS/VE
224	Job information cards
235	VAX public directory - how to create
236	Alternative VAX graphics service for in house users
248	Changes to the Meteogram system
253	Copying/archiving NOS/VE catalogs to ECFILE Copying complete UNICOS directories to ECFILE
254	UNICOS carriage control
260	Changes to PUBLIC directories for VAX users
261	Meteogram system on UNICOS

No.	Still Valid Article
265	Lost UNICOS outputs submitted via RJE or VAX Microfiche changes
266	Reminders on how to import/export magnetic tapes
267	Checking on your UNICOS account usage
268	Changes to WMO FM 92 GRIB
270	Changes to the Meteogram system
271	New ECFILE features on UNICOS
276	Periodic deletion of all Cray /tmp files
280	UNICOS on-line documentation: docview
281	File transfer via FTP (possible problems)
283	New features for Member State batch users (RQS 1.1)
284	UNICOS 7 features & differences

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

ECMWF WORKSHOP ON VARIATIONAL ASSIMILATION, WITH EMPHASIS ON THREE-DIMENSIONAL ASPECTS

(9 - 12 NOVEMBER 1992)

ECMWF organises regular workshops to evaluate the current state of knowledge on relevant topics and to assist in its programme of research. The workshop on 9-12 November reviewed recent operational and research developments in the use of variational methods for data assimilation. This was a good time for the workshop to be held since it helped the Research Department both in the short term: the 3D-var configuration to be implemented this summer, and in the medium term: how to make optimal use of the possibilities of the variational approach. In addition to this the longer term was an entire area for discussion in the two topics of 4D-var and Kalman filtering.

The workshop consisted of 2½ days of lectures, followed by one day of working group sessions, and a final general session to discuss conclusions and recommendations. The three working groups addressed the issues of a) new observation types to be used in 3D-var, b) use of background information in 3D-var, and c) temporal dimension.

- Philippe Courtier

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

METEOROLOGICAL TRAINING COURSE

19 APRIL - 18 JUNE 1993

As in previous years, ECMWF will organise a meteorological training course in 1993. The objective of this course is to assist Member States in advanced training in the field of numerical weather forecasting.

The course is divided into four modules as follows:

Numerical Weather Prediction

MODULE 1 (19 April - 7 May 1993)

Numerical methods, adiabatic formulation of models, data assimilation and use of satellite data

MODULE 2 (10 - 21 May 1993)

Parametrization of diabatic processes

MODULE 3 (24 - 27 May 1993)

General circulation, systematic model errors and predictability

ECMWF Products

MODULE 4 (7 - 18 June 1993)

Use and interpretation of ECMWF products.

Students attending the course should have a good meteorological background, and are expected to be familiar with the contents of standard meteorological textbooks. Some practical experience in numerical weather prediction is also an advantage.

Students can attend any combination of the modules.

In each module there will be lectures, exercises and problem or laboratory sessions. Participants are encouraged to take an interest in the work of ECMWF and to discuss their own work and interests with the staff of the Centre. All the lectures will be given in English and a comprehensive set of Lecture Notes will be provided for modules 1, 2 and 3.

A booklet describing the course in more detail will be mailed to the Member States at the end of December. The booklet also includes information on the application procedures and application forms.

Further inquiries to

E. Kooij-Connally
ECMWF
Shinfield Park
Reading/Berks.
RG2 9AX
UNITED KINGDOM

Tel: UK 0734 499 000
Intern. +44 734 499 000
Telex: 847 908
Telefax: 0734 869 450
Email: els.kooij-connally@ecmwf.co.uk

- Els Kooij-Connally

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

ECMWF CALENDAR 1993

1 - 18 March	**Computer user training course**
10 - 11 March	**Finance Committee,** 50th session
8 (pm) - 12 April	**ECMWF HOLIDAY**
19 April - 18 June	**Meteorological training course:**
19 Apr - 7 May	**Met 1** Numerical methods, adiabatic formulation of models, data assimilation and use of satellite data
10 - 21 May	**Met 2** Parametrization of diabetic processes
24 - 27 May	**Met 3** General circulation, systematic model errors and predictability
7 - 18 Jun	**Met 4** Use and interpretation of ECMWF products
3 May	**ECMWF HOLIDAY**
28 - 31 May	**ECMWF HOLIDAY**
2 - 3 June	**Council,** 38th session
7 - 10 June	**Workshop** - Parametrization of the cloud-topped planetary boundary layer
30 August	**ECMWF HOLIDAY**
6 - 10 September	**Seminar** - Developments in the use of satellite data in Numerical Weather Prediction
27 - 29 September	**Scientific Advisory Committee,** 21st session
29 September - 1 October	**Technical Advisory Committee,** 18th session
5 - 7 October	**Finance Committee,** 51st session

15 - 17 November	**Workshop** - The role of the stratosphere in Numerical Weather Prediction
22 - 26 November	**Workshop** - Meteorological Operational Systems
1 - 2 December	**Council**, 39th session
24 - 28 December	**ECMWF HOLIDAY**

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

ECMWF PUBLICATIONS

Technical Memorandum No. 186:	A bias correction scheme for simulated TOVS brightness temperatures (October 1992)
Technical Memorandum No. 187:	Assimilation of TOVS radiance information through one-dimensional variational analysis (July 1992)
Technical Memorandum No. 188:	Ensemble prediction (July 1992)
Technical Memorandum No. 189:	Unstable perturbations computed using the adjoint technique (October 1992)
Technical Report No. 68:	Implementation of a third generation ocean wave model at the European Centre for Medium-Range Weather Forecasts (August 1992)

Forecast and Verifications Charts up to 31 July 1992

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

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in the ECMWF Newsletter series. As one goes back in time, some points in these articles may have been superseded. When in doubt, contact the author or User Support.

			Newsletter	
		No	Date	Page
GENERAL				
ECMWF publications - range of		26	June 84	16
Obsolete manuals, bulletins, and news sheets		53	Mar 91	30
	&	58	June 92	33
TAC and Computing Representatives, Meteorological Contact Points		57	Mar 92	30
COMPUTING				
CFT77		36	Dec 86	12
Computer Division's status and plans (1992)		56	Dec 91	14
Data handling system		57	Mar 92	16
ECMWF's Cyber 962		53	Mar 91	14
MAGICS - a new release		51	Sept 90	15
Massively parallel computing - a view		53	Mar 91	16
MicroMAGICS - an overview		49	Mar 90	11
Migration of ECMWF's meteorological operational system to UNICOS on the CRAY Y-MP		52	Dec 90	10
Multitasking ECMWF spectral model		29	Mar 85	21
	&	33	Mar 86	9
	&	37	Mar 87	5
NAG library name changes		58	June 92	26
Networks				
- ECMWF's internal network		54	June 91	20
- New LANs at ECMWF		59	Sept 92	20
Supervisor Monitor Scheduler (SMS)		59	Sept 92	13
Telecommunications				
- The ECNET system		54	June 91	16
- Digital links to Member States		57	Mar 92	19
UNICOS 7		58	June 92	23
UNICOS scripts for ECMWF's meteorological operational system		52	Dec 90	11
Workstations at ECMWF		55	Sept 91	7
Y-MP/C90: an introduction		57	Mar 92	10

METEOROLOGY

	No	Newsletter Date	Page
Bogus data - Use of to improve analysis and tropical forecast	45	Mar 89	3
Comparison between SSM/I and ECMWF total precipitable water	57	Mar 92	3
Data acquisition - ECMWF's new system	46	June 89	21
Development of the operational 31-level T213 version of the ECMWF forecast model	56	Dec 91	3
Envelope orography - discussion of its effects	33	June 86	2
ECMWF Analysis			
- New version of analysis system	35	Sept 86	16
- Divergent structure functions	42	June 88	2
- Revised use of satellite data	39	Sept 87	4
ECMWF Preprocessing - new scheme	43	Sept 88	3
Ensemble prediction	58	June 92	5
ERS-1 mission	54	June 91	8
Evaporation from tropical oceans	51	Sept 90	3
Forecast model			
- Revisions to physics	46	June 89	3
- Cloud cover scheme	29	Mar 85	14
- Convection - parametrisation of	43	Sept 88	6
- Diurnal radiation cycle - introduction of	26	June 84	1
- Increased resolution - studies of	38	June 87	10
- Initial conditions - the spin-up problem	39	Sept 87	7
- Parametrisation of gravity wave drag	35	Sept 86	10
- Surface and sub-surface scheme revised	38	June 87	3
Global forecast experiment at T213 resolution	41	Mar 88	3
Good prediction of a severe storm over southern Sweden	50	June 90	10
GORBUSH - a storm in the Mediterranean Sea	53	Mar 91	4
GTS: ECMWF grid code product distribution	27	Sept 84	6
MARS - the ECMWF meteorological archival	32	Dec 85	15
and retrieval system &	33	Mar 86	12
Monte Carlo forecast	49	Mar 90	2
Performance of the ECMWF model in tropical cyclone track forecasting over the western north Pacific during 1990-1991	58	June 92	16
Potential vorticity maps at ECMWF	50	June 90	3
Recent verification of 2m temperature and cloud cover over Europe	54	June 91	3
Skill forecasting - experimental system	40	Dec 87	7
Systematic errors - investigation of, by relaxation experiments	31	Sept 85	9
Use of reduced Gaussian grids in spectral models	52	Dec 90	3

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

USEFUL NAMES AND 'PHONE NUMBERS WITHIN ECMWF

		Room*	Ext.**
DIRECTOR	- David Burridge	OB 202	2001
HEAD OF OPERATIONS DEPARTMENT	- Michel Jarraud	OB 010A	2003
ADVISORY: Available 9-12, 14-17 Monday to Friday			2801
Other methods of quick contact:	- Telex (847908 ECMWF G) - Telefax (+44 734 869450) - VMS MAIL addressed to ADVISORY - Internet mail addressed to Advisory@ecmwf.co.uk		
REGISTRATION			
Project Identifiers	- Pam Prior	OB 225	2384
User Identifiers	- Tape Librarian	CB Hall	2315
COMPUTER OPERATIONS			
Console	- Shift Leaders	CB Hall	2803
Reception Counter	- Tape Librarian	CB Hall	2315
Tape Requests	- Tape Librarian	CB Hall	2315
Terminal Queries	- Norman Wiggins	CB 026	2308
Telecoms Fault Reporting	- Michael O'Brien	CB 028	2306
ECMWF LIBRARY & DOCUMENTATION - DISTRIBUTION	- Els Kooij-Connally	Library	2751
LIBRARIES (ECLIB, NAG, etc.)	- John Greenaway	OB 226	2385
METEOROLOGICAL DIVISION			
Division Head	- Horst Böttger	OB 007	2060
Applications Section Head	- Rex Gibson	OB 014	2400
Operations Section Head	- Bernard Strauss	OB 328	2420
Meteorological Analysts	- Andreas Lanzinger	OB 314	2425
	- Ray McGrath	OB 329	2424
	- Anders Persson	OB 315	2421
Meteorological Operations Room	-	CB Hall	2426

		Room*	Ext.**	Bleeper
COMPUTER DIVISION				
Division Head	- Geerd-R. Hoffmann	OB 009A	2050	150
Systems Software Sect.Head	- Claus Hilberg	OB 104A	2350	115
User Support Section Head	- Andrew Lea	OB 227	2380	138
User Support Staff	- Antoinette Alias	OB 224	2382	154
	- John Greenaway	OB 226	2385	155
	- Norbert Kreitz	OB 207	2381	156
	- Dominique Lucas	OB 206	2386	139
	- Pam Prior	OB 225	2384	158
Computer Operations Section Head	- Peter Gray	CB 023	2300	114
Security, Internal Networks and Workstation Section Head	- Walter Zwieflhofer	OB 140	2352	145
GRAPHICS GROUP				
Group Leader	- Jens Daabeck	CB 133	2375	159
RESEARCH DEPARTMENT				
Head of Research Department	- Anthony Hollingsworth	OB 119A	2005	
Computer Co-ordinator	- David Dent	OB 123	2702	

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (0734) 499000, international +44 734 499000, or direct dial to (0734) 499 + last three digits of individual extension number, e.g. the Director's direct number is (0734) 499001.

DEC MAIL: Contact scientific and technical staff via VMS MAIL, addressed to surname.

Internet: The ECMWF address on Internet is **ecmwf.co.uk**
Individual staff addresses are firstname.lastname, e.g. the Director's address is David.Burridge@ecmwf.co.uk