ECMWF Newsletter

Number 61 – March 1993

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

COVER: The Silicon Graphics IRIS Indigo XS Workstation - see article page 21

This Newsletter is edited and produced by User Support.
The next issue will appear in June 1993.

The illustration on the front cover of this edition shows one of the new Silicon Graphics Indigo Workstations acquired by ECMWF; these, and the accompanying CRIMSON servers, are described in detail in the article "New Workstations at ECMWF".

On a somewhat different scale, the ongoing investigations into massively parallel processing are the subject of another article in this issue.

We continue with reprints from the series of articles on Fortran 90 by Jeanne Adams which first appeared in the NCAR SCD Computing News. Also in this edition are articles on new features in CFS and ECFILE.

In the meteorological section, steps to improve the performance of the model by refining the treatment of clear sky and cloud radiative properties are described.

As is customary in the March edition of the Newsletter, an updated table of Member States' representatives and contact points is given.

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

CHANGES TO THE OPERATIONAL FORECASTING SYSTEM

Recent changes

On 7 December 1992, a modification to the extrapolation of temperature below the lowest model level was implemented in the model post-processing. Before that change, the 850 hPa temperatures extrapolated under the model orography were anomalously low under winter anti-cyclonic conditions, reflecting the cooling of the model surface. Now the extrapolated low-level temperature and MSL pressure fields continue the free atmospheric distribution of fields and show less dependence on the overlying orography.

On 1 February 1993, changes were introduced to the horizontal diffusion and to the cloud/radiation scheme (inclusion of shortwave optical properties for ice and mixed phase clouds, and revision to the clear-sky absorption coefficients). Experimentation with these changes had shown a small but consistent positive impact on the quality of the medium-range forecast, particularly in summer situations.

Planned changes

An improved representation of surface and planetary boundary layer processes will be implemented.

- **Bernard Strauss**

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

REVISION OF THE CLEAR-SKY AND CLOUD RADIATIVE PROPERTIES IN THE ECMWF MODEL

Introduction

Recent diagnostics have indicated that the relatively poor performance of the model after day 5 of the forecast since the introduction of the high-resolution model in September 1991 was linked to an excessive radiative cooling below the tropopause at high latitudes. Sensitivity studies have shown that clouds play a sizeable role in this cooling. Another recurring problem is the excessive heating of continents in summer. Comparisons with ground station measurements of the solar radiation at the surface, particularly during the FIFE (First ISLSCP Field Experiment) have shown that the overestimation of the downward solar radiation can reach 8 percent (Betts et al., 1993).

A solution was proposed to improve these two aspects of the cloud and radiation formulation. This note documents the modifications to the cloud model, and to the clear-sky absorption coefficients. Results obtained during the experimental testing of the modifications in the ECMWF forecast system are then presented. These modifications were introduced in the operational model on 1 February 1993.

Ice and mixed phase clouds

Before 1 February 1993, the optical properties of the clouds were typical of water clouds, with a liquid water content diagnosed from the saturation mixing ratio for stratiform clouds and given a fixed value for convective clouds. The shortwave parametrisation followed the analysis of in-situ measurements of radiative properties of stratocumulus decks of Bonnel et al. (1983) and Fouquart (1987), whereas the longwave parametrisation followed the approach of Smith and Shi (1992). The recent revision considers ice cloud optical properties as derived by Ebert and Curry (1992).

Following a formulation first proposed by Slingo and Schrecker (1982) for water clouds, the scattering properties are parametrised in terms of the Ice Water Path (IWP) and effective radius of the particle size distribution (r_e):

shortwave optical thickness

$$\tau_i = IWP\ (a_i + b_i / r_e)$$

single scattering albedo

$$\omega_i = c_i - d_i\, r_e$$

assymetry factor

$$g_i = e_i + f_i\, r_e$$

a,b,c,d,e,f are all positive coefficients

These relationships simply reflect that, for a given water loading, a cloud with smaller particles is more radiatively active (larger τ), scatters more (larger ω, and smaller g, i.e., less energy available in the forward scattering peak - directly transmitted -, and more energy available for scattering and potential absorption).

At temperatures where both ice and water clouds may be present, the "mixed phase" cloud has the following optical properties:

shortwave optical thickness

$$\tau_{mixed} = \tau_{water} + \tau_{ice}$$

single scattering albedo

$$\omega_{mixed} = \frac{\omega_{water}\tau_{water} + \omega_{ice}\tau_{ice}}{\tau_{water} + \tau_{ice}}$$

assymetry factor

$$g_{mixed} = \frac{g_{water}\omega_{water}\tau_{water} + g_{ice}\omega_{ice}\tau_{ice}}{\omega_{water}\tau_{water} + \omega_{ice}\tau_{ice}}$$

The main differences between the old and new formulations are found for: ω_2 the single scattering albedo in the 0.68 - 4 µm interval of the radiation scheme, where the new value will give higher absorption (ω_2 decreases from 0.998 to around 0.97), and for the assymetry factors in both spectral intervals where the new values will decrease the fraction of the solar energy available for scattering in the forward peak of the phase function. Therefore the main impact of these optical properties is an increase of the heating by absorption of solar radiation in the higher level clouds (where the ice fraction is predominant) and a decrease of downward solar radiation under these ice clouds. The local heating effect of the ice clouds clearly appears in Fig. 1, which compares the temperature error for 2 sets of T106 L 31 simulations for the 15 January and 15 July 1992 using the control cycle 44 (OSA or ORN) and modified shortwave properties for ice clouds (Q1G or Q1I). The reduction of the net radiative cooling in the high clouds, particularly in the summer hemisphere, tends to reduce the level of kinetic energy with an overall positive impact on the model scores as illustrated in Fig. 2 for a set of 12 T106 L31 forecasts run for the 15th day of each month between November '91 and October '92.

Revised clear-sky absorption coefficients

At the core of any GCM-type radiation code are some parametric expressions for the optical thicknesses of the various radiatively active atmospheric components (gases: H_2O, CO_2, O_3, CH_4, N_2O, CFCs; ice and water clouds, aerosols). The absorption coefficients for transmission through the gases are generally computed from spectroscopic parameters compiled by various groups (the HITRAN compilation from the U.S. Air Force, and the GEISA spectroscopic database at LMD, Palaiseau). The data banks are regularly updated following new measurements and theoretical calculations by spectroscopists. Coefficients in the previous version of the ECMWF radiation scheme were derived around 1983 using what can be considered as relatively old versions of the HITRAN compilation: the 1980 version for the longwave coefficients and the 1982 version for the shortwave coefficients. The availability at ECMWF of a more recent (April 1991) version of HITRAN has made possible the verification and revision of these absorption coefficients.

Different points have been addressed regarding the clear-sky parametrisations:

* the role of missing water vapour bands in the compilation of line parameters used in the computation of the shortwave absorption coefficients of the present radiation scheme;

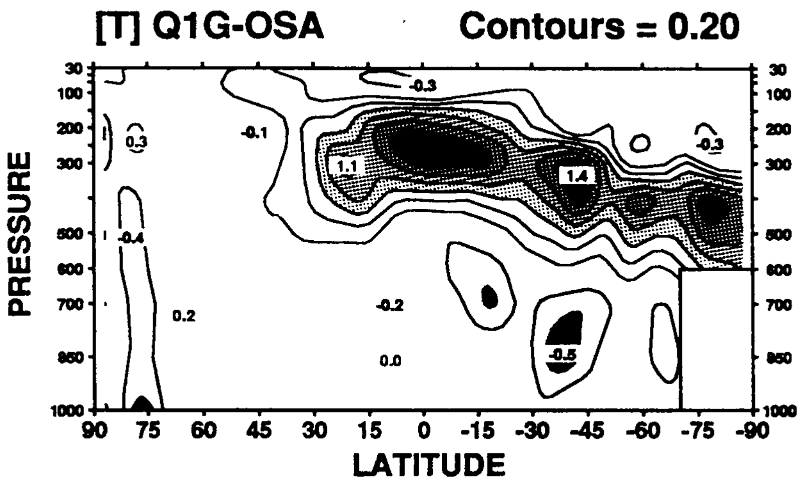

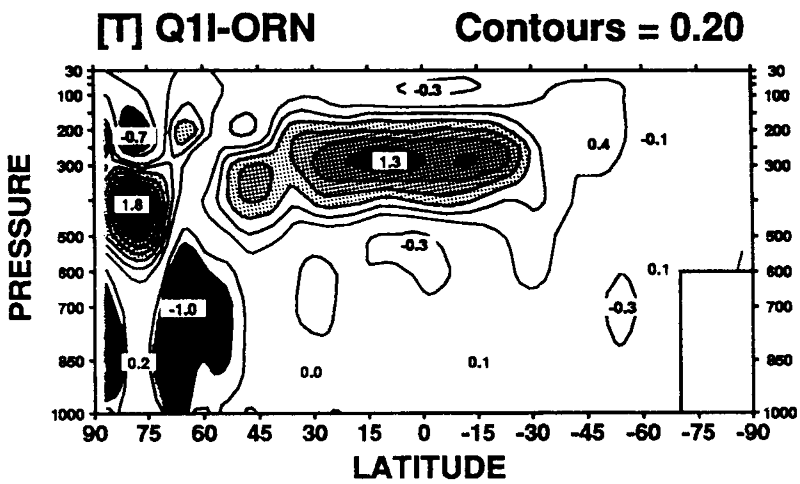

Fig. 1: The height/latitude cross-section of the zonal mean temperature difference between 10-day forecasts including or not the parametrisation for mixed phase clouds.
a) 15 January: control OSA, modified Q1G
b) 15 July: control ORN, modified Q1I

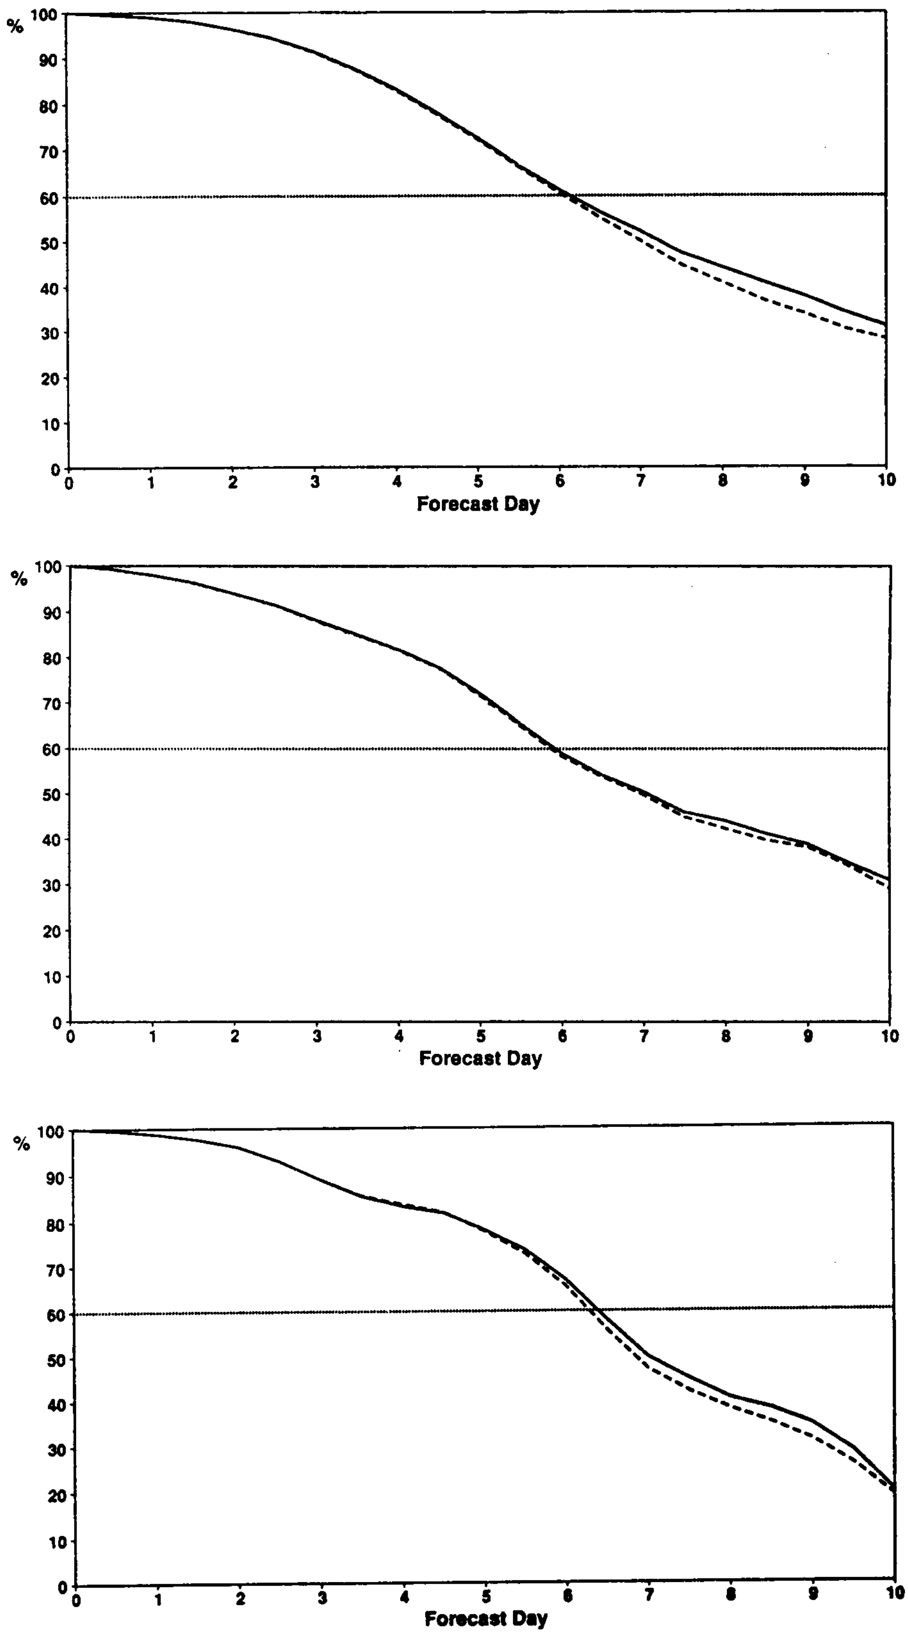

Fig. 2: Anomaly correlation of geopotential at 500 hPa over the Northern Hemisphere (top), Southern Hemisphere (centre) and European area (bottom) for two sets of 12 forecasts including (solid line) or not (dashed line) the parametrisation for mixed phase clouds.

* the introduction of the shortwave absorption effect of uniformly mixed gases other than CO_2 and O_2 (namely adding absorption by CH_4, N_2O and CO);

* the temperature and pressure dependence of the shortwave absorption (with reference values of p and T corresponding to higher levels in the atmosphere);

* the longwave and shortwave radiative effects of a climatological distribution of aerosols;

Three other points concern the parametrisations in presence of clouds:

* the effective zenith angle for O_3 shortwave absorption in cloudy atmospheres (assuming a diffuse rather than a directly transmitted beam);

* the water vapour absorption within the clouds (taking the saturated instead of the grid-averaged humidity);

* the anvil clouds over deep convective towers are given a liquid water content similar to that for stratiform clouds (5% of the saturation mass mixing ratio instead of the specified 1×10^{-4} kg/kg of the convective clouds).

Apart from the effect of the aerosols, each of these points has a tiny influence (0.1 to 0.5 %) on the SW absorption when taken separately. However, all of these changes are cumulative as they all lead to an overall increase in atmospheric SW absorption (by 0.5 to 1.5% depending on the atmospheric profile), i.e., a decrease in downward SW radiation at the surface.

The aerosols may play a bigger role. Much concern has been expressed recently at the overestimation of the downward shortwave radiation at the surface ($SW^-(0)$) by the ECMWF model, following, in addition to more indirect proofs, the comparisons of model surface shortwave radiation with ground station measurements carried out during FIFE. A recent paper by Bruegge et al. (1992) throws some new light on the problem, as it reports measurements of the temporal variability of the aerosol optical thickness over the FIFE site during the five Intensive Field Campaign periods. A sizeable quantity of aerosols is present most of the time, which is likely to have some impact on the surface shortwave radiation. Fig. 3 presents a typical diurnal cycle of $SW^-(0)$ for different aerosol loading between a no aerosol case and a τ=0.25 case, for July conditions over the FIFE site. This comparison shows that neglect of the radiative effect of aerosols can explain a large part (at least 3 of the 8%) of the overestimation of $SW^-(0)$ by the model, noted by Betts et al. (1993).

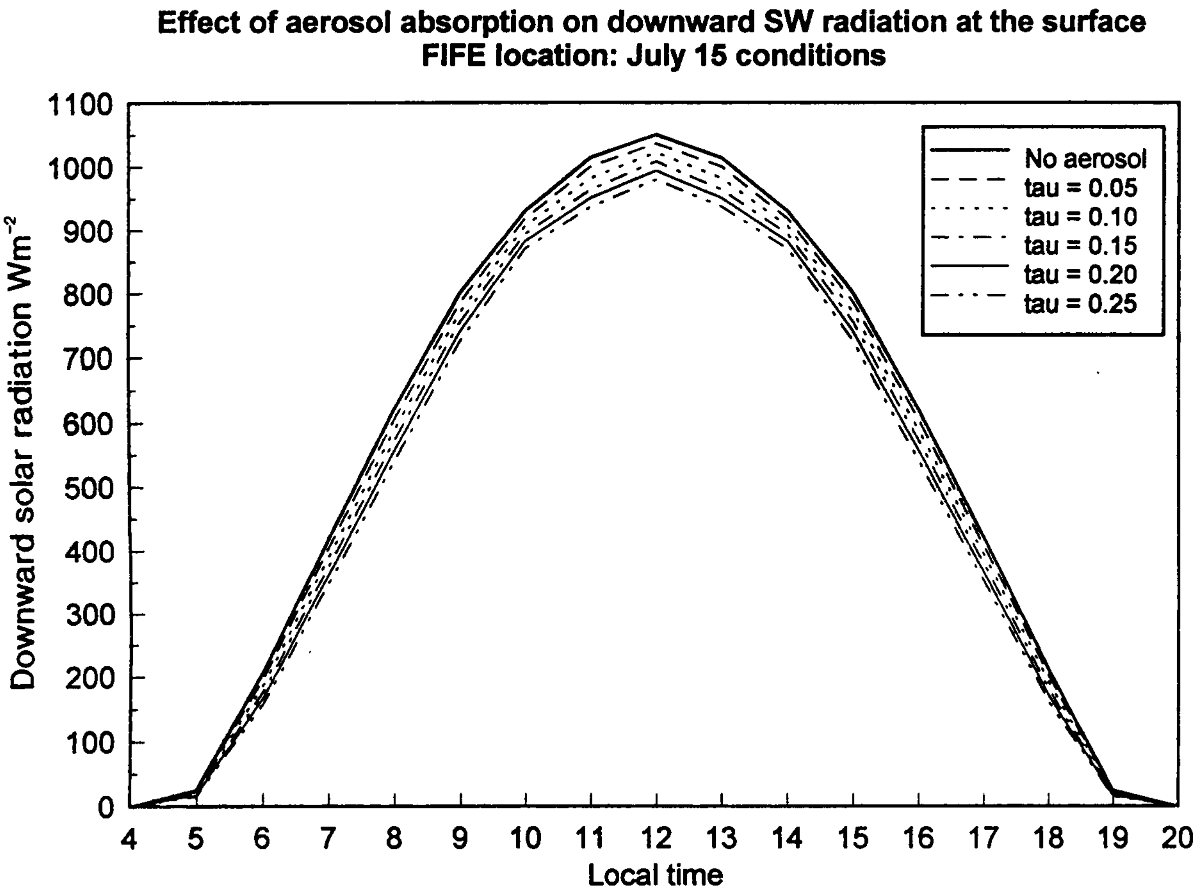

Fig. 3: The impact of the radiative effects of aerosols on the diurnal cycle of the downward shortwave radiation at the surface for 15 July conditions over the FIFE site (39.1°N). Different lines correspond to a clear-sky atmosphere with various aerosol loadings.

Climatological distributions of aerosols as implemented by Tanré et al. (1983) have been reactivated in the ECMWF model. Four types of aerosols are considered for their longwave and shortwave radiative effects: continental-desert, maritime, urban, and stratospheric background. Although present as a possible 5th type in the climatology, the active volcanic aerosol is not accounted for as proper information is required for specifying its 3-D distribution.

Impact on analyses and subsequent forecasts

A total of 32 days of assimilations including the proposed modifications was run, taking in a 10-day period starting 07 September 1992 12 GMT expected to be representative of mainly summer conditions, and another 16-day period starting 21 November 1992 for more wintery conditions. Results are presented mainly for the November-December assimilation and subsequent forecasts.

The impact on the analysed fields looks very small when seen as zonal means of the differences between the new and operational analyses. The only noticeable effects are a small temperature increase below the tropopause, a slight decrease in the intensity of the Hadley circulation, and a small (3 to 6 %) decrease in the poleward transport of heat and moisture. However, the agreement between observed and either first-guess or analyzed geopotential fields for the Northern and Southern Hemispheres, and the tropical belt is clearly improved by the new parametrisation, as seen in Fig. 4, averaged over 9 days of analysis (days 2 to 10) for the 00 GMT analyses of the September assimilation. Generally, the modified assimilation system accepts a few more data and displays noticeably smaller biases above 500 hPa.

For the November-December period, results are presented in Fig. 5 for a set of 9 forecasts between 921122 and 921130. Overall, the impact of the proposed modifications is positive, more so in terms of r.m.s. errors than in terms of anomaly correlations.

In Fig. 5, the comparison between dashed and solid lines shows the impact (generally positive over the Northern and Southern Hemisphere, and over Europe) that the modified analyses have on the subsequent forecasts.

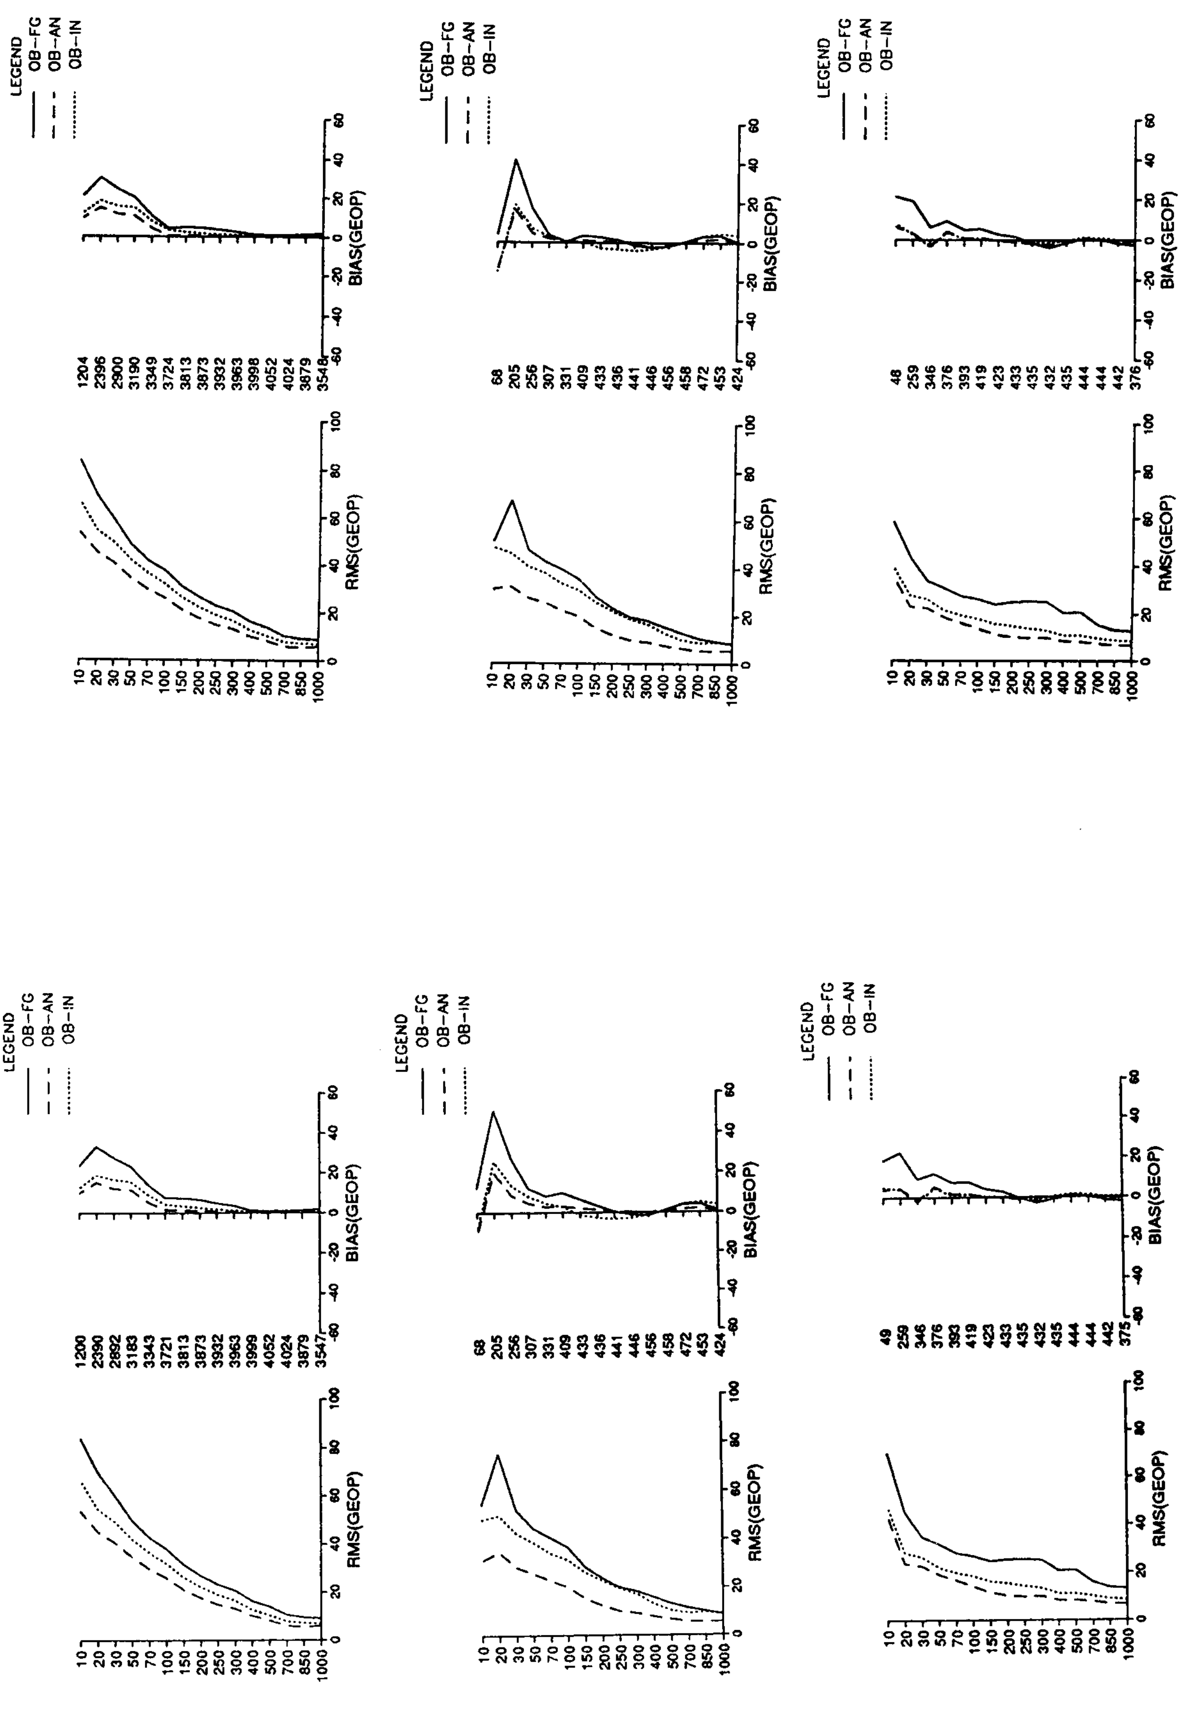

Fig. 4: The RMS errors and bias in geopotential at 500 hPa averaged over the period 920908 - 920916 in the Northern Hemisphere, tropical belt and Southern Hemisphere. Panel (a) is the operational analysis at the time (cycle 44); panel (b) the results for the modified analysis including the cloud radiation changes.

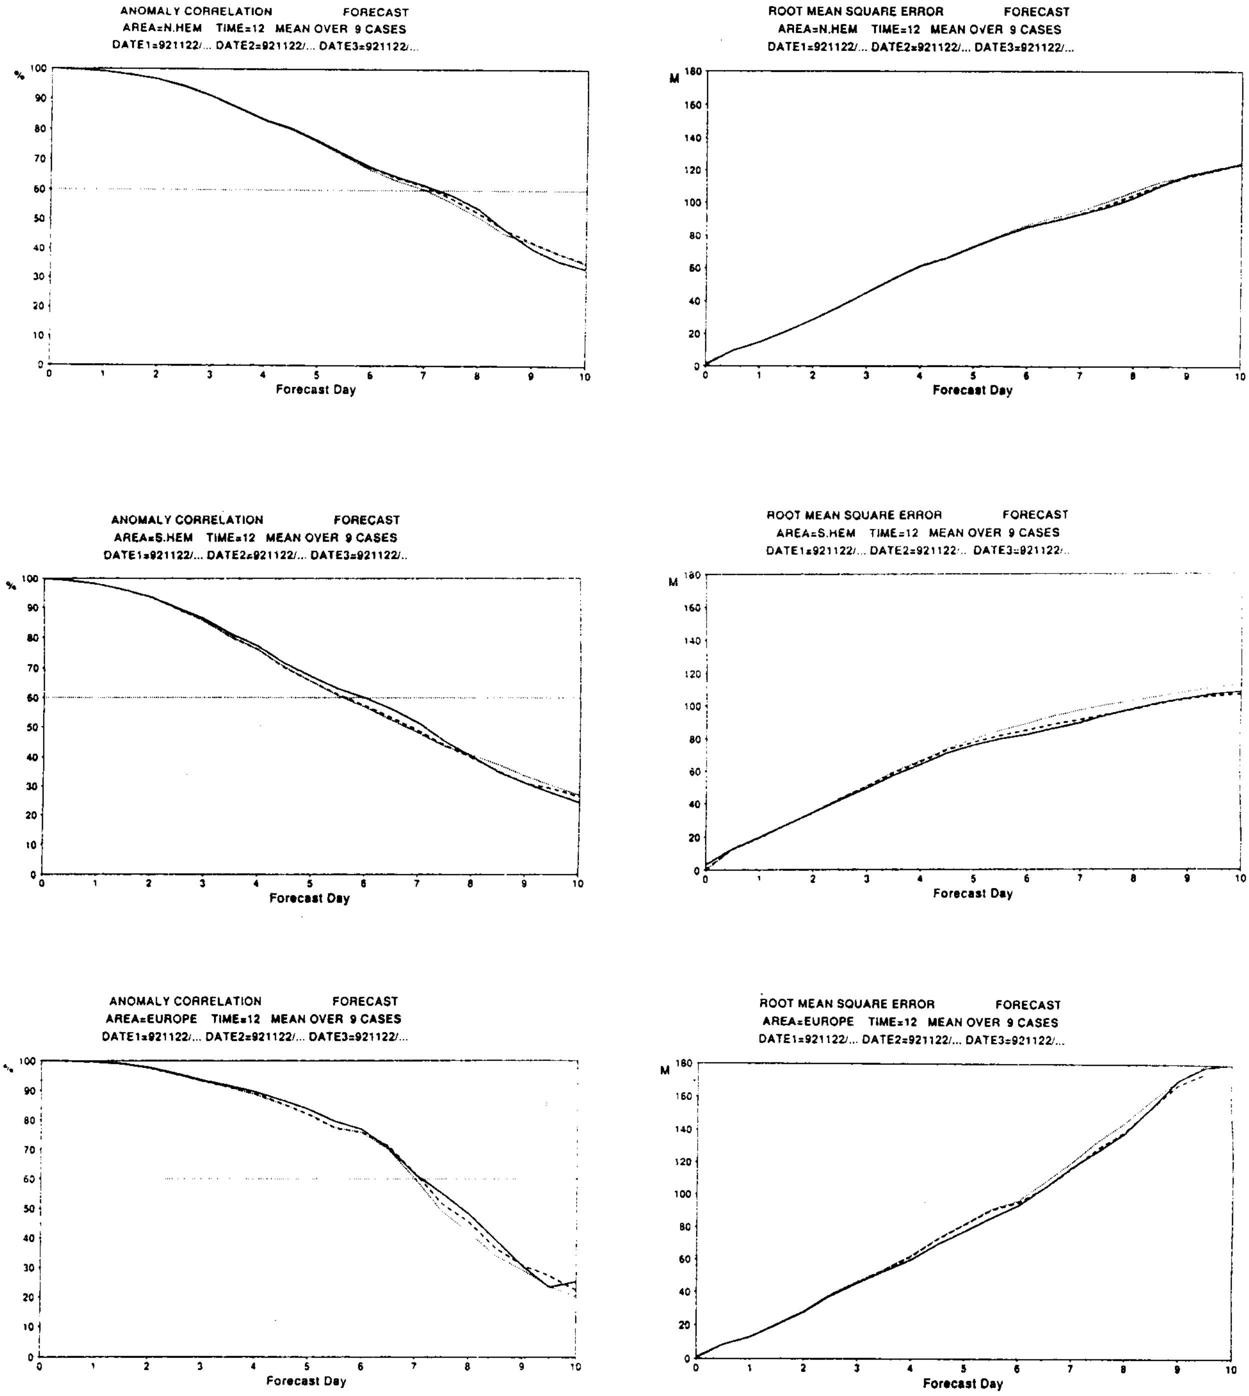

Fig. 5: The anomaly correlation and corresponding RMS errors in geopotential at 500 hPa over the Northern Hemisphere, Southern Hemisphere and European area for two sets of 9 forecasts between 22 and 30/11/1992. The dotted line corresponds to results with the operational system at the time (cycle 44). The dashed line shows results when the forecasts are run with the model including the cloud and radiation changes. The full line corresponds to the results when these changes are used in both the analysis and forecast systems.

Conclusion and future developments

As stated in the introduction, this recent change to the cloud-radiation parametrisation addresses two of the problems that the ECMWF model has been displaying: the long-known excessive surface shortwave heating (whose effect is now known to be emphasized by deficiencies in the land-surface parametrisation), and the more recent (since Sept.'91) excessive radiative cooling under the tropopause. The solution proposed here provides some physically-based improvement to deal with these two problems. However, these modifications must be seen as the first stage in an on-going revision of the clear-sky / cloud radiative properties as part of the experimentation with the new prognostic cloud scheme.

Now that the aerosols are back on the scene, an effort will be made to improve definition of their optical properties, particularly as regards the differences in the two spectral intervals of the shortwave scheme. Further study of the dependence of the aerosol optical properties on relative humidity would also be of interest.

- Jean-Jacques Morcrette

References

Betts, A.K., J.H. Ball, and A.C.M. Beljaars, 1993: Comparison between the land surface response of the ECMWF model and the FIFE-1987 data. Quart.J.Meteor.Roy.Soc., submitted.

Bonnel, B., Y. Fouquart, J.-C. Vanhoutte, C. Fravalo, and R. Rosset, 1983: Radiative properties of some African and mid-latitude stratocumulus clouds. Beitr. Phys. Atmosph., 456, 409-428.

Bruegge, C.J., R.N. Halthore, B. Markham, M. Spanner, and R. Wrigley, 1992: Aerosol optical depth retrievals over the Konza Prairie. J.Geophys.Res., 97D, 18743-18758.

Ebert, E.E., and J.A. Curry, 1992: A parametrisation of ice cloud optical properties for climate models. J.Geophys.Res., 497D, 3831-3836.

Fouquart, Y., 1987: Radiative transfer in climate modelling. NATO ASI on Physically-Based Modeling and Simulation of Climate and Climatic Changes. Erice, Sicily, 11-23 May 1986, M.E. Schlesinger, Ed.

Slingo, A., and H.M. Schrecker, 1982: On the shortwave radiative properties of stratiform water clouds. Quart.J.Roy.Meteor.Soc., 108, 407-426.

Smith, E.A., and Lei Shi, 1992: Surface forcing of the infrared cooling profile over the Tibetan plateau. Part I: Influence of relative longwave radiative at high altitude. J.Atmos.Sci., 49, 805-822.

Tanré, D., J.-F. Geleyn, and J.M. Slingo, 1984: First results of the introduction of an advanced aerosol-radiation interaction in the ECMWF low resolution global model. *Aerosols and Their Climatic Effects*.H.E. Gerber and A. Deepak, Eds., A. Deepak Publ., Hampton, Va., 133-177.

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

ECMWF'S CURRENT INVESTIGATIONS INTO MPP SYSTEMS

Introduction

An MPP (Massively Parallel Processor) system is one where many cheaper (but slower) CPUs are connected together, to give the same or higher processing power as the traditional supercomputer, which uses a few high performance CPUs. As an example, the Intel Paragon XP/S has at each processor node a 75 MFLOP/sec (peak performance) processor with up to 128 Mbytes of memory per node. Up to 4,000 nodes are then connected in a 2D grid using one iMRC (Intel Message Routeing Chip) per processor. The iMRC Routeing Chip has a raw throughput of 200 Mbytes per second (see Fig. 1). The distributed nature of the memory, each portion associated with one processor, is a main distinction from the shared memory of the Cray C90.

Why is ECMWF interested in MPP systems? The traditional shared-memory supercomputers are believed to be close to their ultimate performance. Their clock speed is very high, circuit sizes are very small, removing the heat is a problem and wiring is becoming difficult. All this makes the traditional supercomputer very expensive to design and build. Using a processor which is (say) 100 times slower than ECMWF's Cray C90, reduces the cost by a factor of at least 1000. Therefore, by combining 1000 such processors, for the same price as the traditional supercomputer one now has a machine with a performance 10 times that of the supercomputer, at least provided the cost of an efficient communications network can be neglected. The main problem then is how to use all these processors simultaneously to solve a problem.

ECMWF has a high operational performance requirement, that is a few highly optimised CPU demanding programs that must run each day in a time critical window. The more power that can be made available in that time window the better. This type of situation suits the MPP system well. For this reason ECMWF, other atmospheric and indeed a number of modelling groups around the world, are now studying such systems seriously for their potential for operational use.

ECMWF already has considerable experience with Cray macro-tasking on systems with up to 16 processors and, in fact, gets very good parallelisation of its model code. Now we must gain experience on systems of up to several thousands of processors.

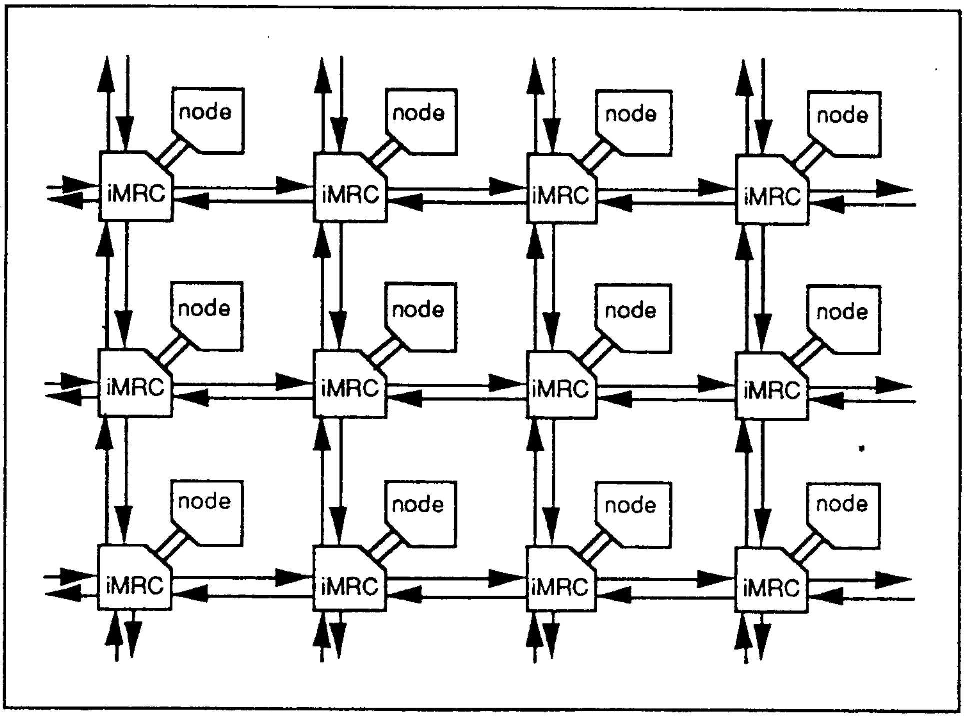

Fig. 1: The Intel Paragon XP/S 2D mesh network. Each node (processor) has up to 128 Mbytes local memory and 75 MFLOP/sec peak processing power. Maximum configuration is with 4000 processors.

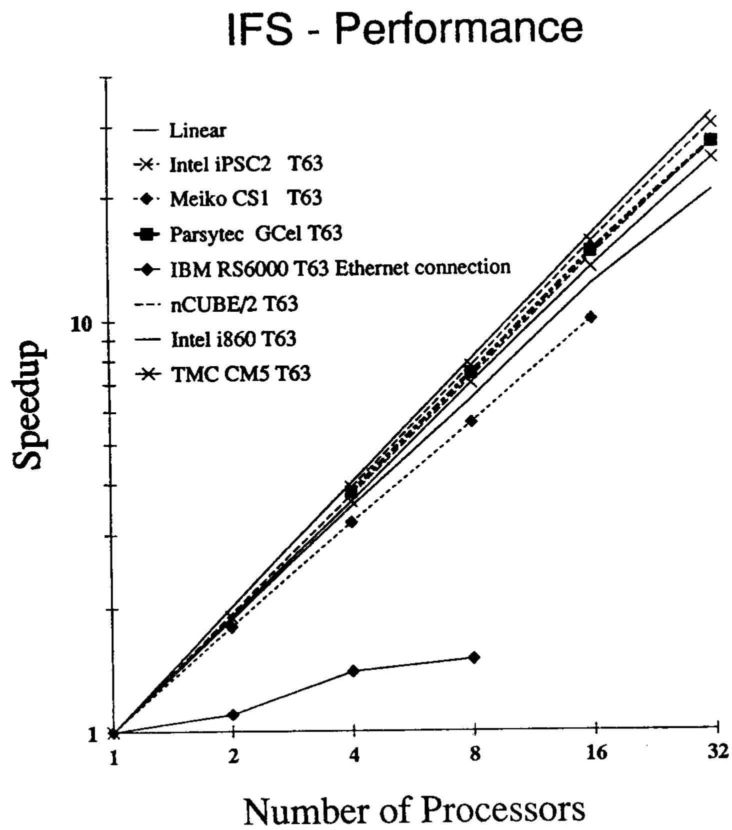

Fig. 2: The performance of the distributed memory version (Eulerian shallow water IFS) on 4 parallel computers. Triangular truncations with 21 and 63 waves respectively are used.

Today the hardware of parallel systems is becoming reliable, but software difficulties still remain. ECMWF's plan is therefore to study and, it is hoped, solve some of those software difficulties.

The Centre's current goal is to define a good strategy (both algorithmically and programming-wise) for introducing MPP systems into operational weather forecasting. To reach this goal the Centre is taking the following steps:

(i) analyse the computational requirements and data structures/flow of global weather models;

(ii) calculate the communication costs of different parallelisation strategies, for 1D, 2D, 3D and hypercube connections of processors;

(iii) implement simplified models, such as spherical Helmholtz solvers and shallow water models, on existing parallel computer systems;

(iv) extend the above techniques using more realistic codes such as the Optimum Interpolation (OI) analysis code, and the Integrated Forecast System (IFS);

(v) evaluate the effort required to program MPP systems and the usefulness of the tools provided.

Work to date

There are in the market place quite a few hardware vendors. Those marked * indicate systems on which we have run some trials already:

Intel iPSC2*	Meiko CS-1*, CS-2
nCUBE-2*	Parsytec*
Thinking Machines CM-2, CM-200, CM-5*	Cray MPP
Intel iPSC/860*	Intel Paragon XP/S
IBM RS 6000 network*, SPI	Fujitsu VPP 500
SUN network*	KSR.

The Centre has been involved in several MPP-related projects, both internal and external, to gain some initial experience.

ECMWF's first external collaboration was to participate in the EC funded ESPRIT 2 project known as GENESIS (1989-1991). The primary goal of the work conducted at ECMWF with this GENESIS funding was to define a good strategy for algorithms and programming with a view to introducing future massively parallel supercomputers into operational global weather forecasting. Results have been encouraging: it seems that the next generation of massively parallel supercomputers will be able to run ECMWF's spectral model at T213 resolution with a sustained performance far exceeding the Cray C90 performance. The evaluations do, however, employ raw performance characteristics for future computers provided by the manufacturers and, therefore, contain elements of uncertainty.

Initially we have mainly been investigating parallelisation strategies with simplified versions of ECMWF's operational model, to check if it is possible to utilise the present spectral transform method on an MPP system. Initial results look promising. It seems possible to implement an efficient solution where the communication between the processors is isolated to a few interface routines, and then retain most of the meteorological model parts unmodified.

From all the work done so far at the Centre we can draw up the following list of difficulties and points to watch:

* It is more difficult to write programs for MPP machines that have distributed memory, than for traditional supercomputers. Data structures and algorithms must be understood and organised to a greater extent.

* Use standards: Fortran 77 whenever possible, together with portable parallel extensions where necessary.

* Load balancing becomes an important issue. If you have a few grid columns assigned per processor, then much time can be wasted because all processors have to wait for the calculation at the most complex parts of the globe to be completed.

* The costs of rewriting code should be kept as low as possible. The IFS code is currently 1000 subroutines and 240,000 lines of code, having taken approximately 50 man-years to write. Clearly it is desirable to change as little of this as possible.

* The data communication should be structured so that it can be isolated as much as possible from the rest of the code. This should ease problems when porting it either to a shared memory or to a distributed memory system.

* The Cray macro-tasked code can often be used as a guideline when converting the code for efficient operation on an MPP system. However, the parallelism provided by the macro-tasked implementation may not be sufficient for an MPP system with hundreds or thousands of nodes.

* Single processor performance is an important issue. Optimal utilisation of RISC processors must be aimed for, especially because of the local data caches, which must be efficiently used in order to minimise the access time to local memory.

Continuing work

The main task is to make a portable parallel version of the IFS suitable for distributed memory MPP systems. (See D. Dent's article in the previous Newsletter for details about IFS.)

The work began in July 1992. At the moment, we have a portable distributed memory version of the Eulerian shallow water code, suitable for up to around 100 processors. Some results are presented in Fig. 2. (The workstation version performs poorly because of the slow Ethernet connection.) This version is now being extended to the full 3-D Eulerian model using static scheduling of work. When this version is ready, probably in May or June 1993, we will address the problems associated with

* semi-Lagrangian advection
* more parallelism up to around 10000
* load balance in the 'physics'
* overlapping I/O and computations.

Three of the current projects we are involved in are PPPE, RAPS and GP-MIMD2.

PPPE (Portable Parallel Programming Environment) is an EC-funded ESPRIT 3 based project to produce software tools to help people use MPP systems. For example, one will need a good debugger, while debugging on a vector system can be difficult, it can be an order of magnitude worse on an MPP system! Also, one will need performance analysers and data flow monitors to help optimise codes. Vendors (e.g. Meiko, IBM, Intel, etc.) produce the tools, ECMWF's role is to recommend which tools are required and then to test them.

RAPS (Real Applications on Parallel Systems) is an ECMWF-originated project whereby various organisations interested in running real applications get together to port large production codes on to various MPP systems. The participants in RAPS include DWD, Météo-France, UKMO, CERN, AVL, ESI, GMD, DKRZ, University of Southampton, and others. RAPS provides a forum for joint consultation with, and influencing of the various vendors, as well as offering tutorials and experiences amongst the participants. ECMWF's contribution to RAPS will be the big IFS code. The resulting benchmark studies will then be assessed and summarised.

GP-MIMD2 (General Purpose Multiple Instruction Multiple Data 2) is another EC-funded ESPRIT 3 based project which involves providing two 64 processor Meiko CS-2 systems (at CERN, Geneva, and CERFACS, Toulouse) for users to gain experience in actual use of an MPP system on real production codes. Our involvement is to use the CERFACS system to develop the parallel processor version of the Centre's IFS code and to coordinate efforts by UKMO, DWD, Météo-France and DKRZ to port other atmospheric and ocean models to this platform. In addition, we are required to study and advise on the operational aspects of the machine, covering such areas as ease of use, availability, reliability of service, etc.

As an aside it should be noted that we use 'compiler directive like' calls when we rewrite the sequential code. They are called PARMACS (PARallel MACros) which consist of a set of simple macro extensions to Fortran 77. These macros are essentially a set of message passing commands (send, receive, barrier, etc.) to move data between processors. PARMACS is now available on many MPP machines, and a simulator is available on the Cray C90. So the same program can be run on a whole range of MPP systems. Another programming model is HPF (High Performance Fortran). This is a set of more advanced features to be added to Fortran 90, and should be available sometime in the 1993/94 time-frame. These features should be easier to use than PARMACS, as much of the detailed work is done behind the scenes. HPF is a joint effort between most parallel machine vendors and many of the large users. It would be, however, naive to believe that HPF solves all problems. It is still necessary to understand the data structures/flow and algorithms.

ECMWF will continue to benchmark the PARMACS based shallow water code on as many MPP systems as possible, to study further parallelisation and data communication of 2 and 3D semi-Lagrangian advection codes, produce a parallel version of the IFS benchmark code, assess their performance, and to migrate to new standards (e.g. HPF), if appropriate.

The work described here has been carried out in co-operation with Tuomo Kauranne, Saulo Barros, Fritz Wollenweber, Adrian Simmons, Geerd-R. Hoffmann, David Dent, Mats Hamrud, Clive Temperton, among others at ECMWF, and Ute Gärtel, Tony Schüller and Wolfgang Joppich from GMD in Germany.

- Lars Isaksen

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

NEW WORKSTATIONS AT ECMWF

Background

The Centre has been in the process of gradually moving from an environment based on IBM-compatible PCs together with the CYBER NOS/VE interactive service to a UNIX-based workstation environment. Over the last 3 years a number of SUN SPARC stations were installed, which enabled development of the services required.

In order to provide all scientific and technical staff and consultants with workstations, at least 50 more workstations were required. An Invitation to Tender for the supply of Scientific Workstations was issued in November 1992, specifying the requirements for the workstations and servers to be acquired.

In February 1993, the Director approved the recommendation of the Tender Evaluation Board to award the contract to Silicon Graphics Ltd, for the installation of 52 Indigo XS workstations and 2 CRIMSON servers.

Indigo XS Workstations

The Silicon Graphics Indigo workstations to be installed have the following configuration (note that KB means KiloBytes, MB means MegaBytes and GB means GigaBytes):

* MIPS R4000 cpus
* 32 MB of memory
* 16" colour monitors with a resolution of 1280 X 1024 pixells
* 540 MB (unformatted) local disk drive
* 3.5" floppy disk drive.

The R4000 cpu is almost the latest offering from MIPS, being announced in July 1992 (the latest offering is the R4400, which was announced in December 1992). It is a high performance cpu, including two on-chip caches of 8KB each (one for instructions, one for data), plus an external 1MB cache. It implements two instruction sets, MIPS I, which is the instruction set for the MIPS R3000 cpu, and MIPS II which is the native R4000 instruction set, and allows use of the 64-bit capability of the R4000 cpu.

The systems have 516 MB (formatted) disks, which will initially be partitioned as follows:

* root partition (/) - 16 MB
* swap partition - 80 MB
* usr partition - 420 MB.

The usr partition includes the rest of the available disk space, which includes /usr/tmp. After the files comprising /usr have been installed, there will be at least 270 MB of disk space available for use on the local workstation, via /usr/tmp. The normal /tmp will be linked to point to /usr/tmp, rather than being space within the root partition. This maximises the amount of space available for temporary files on the local workstation disk. The "owner" of the workstation will in principle be responsible for managing the use of this space on his system.

The Indigo's system units are quite large for desktop systems, particularly in comparison with the Sun IPC SPARCstations. This is partly because there is space in the system to install extra graphics boards and up to 3 SCSI devices (the systems installed at ECMWF will have 2 SCSI devices installed, leaving one slot free).

The systems to be installed at ECMWF will have standard "PC"-style keyboards, unlike the keyboards supplied with the Sun SPARCstations, which include a number of Sun-specific keys for use with the Window system (Open, Cut, Paste etc.). They will also have a mechanical mouse, in contrast to the optical mouse that comes with a Sun SPARCstation (which also needs a special metal mouse pad).

The Indigos will have XS graphics boards with a screen resolution of 1280 X 1024 pixells. The 16" colour monitors are identical to those used on the Sun SPARCstations (they are in both cases manufactured by Sony). The graphics performance on the XS Indigo is generally better than on the entry level system. In particular, the graphics performance for 3D has been improved, whereas 2D graphics packages developed at ECMWF (MAGICS, METVIEW/ws) will gain less.

CRIMSON servers

In addition to the desktop workstations, two Silicon Graphics CRIMSON servers will be installed. These are "deskside" systems which will be installed in the Telecomms and servers area in the Computer Hall. They are basically the same as the desktop workstations, but have a larger cabinet, and have much more scope for expansion. For example, up to 74 GB of disk can be supported by a CRIMSON server, whereas an Indigo can currently only support up to 8.4 GB. The servers will each have:

* MIPS R4000 cpu, to be upgraded to the MIPS R4400 cpu (3Q93)
* 128 MB memory
* 1.2 GB system disk for the operating system
* 1.2 GB disk for software, man pages etc. not stored on the Indigos
* 2 X 2.4 GB disk for SCRATCH space
* FDDI connection.

The servers will initially have the same MIPS R4000 cpu as the Indigos, but will be upgraded later to the recently announced R4400 cpu, which provides 50% more power (by using a 50% higher internal clock rate).

The servers are primarily intended to be used as compute servers for applications which can take advantage of using local disk space, rather than running all applications on the office workstations. As mentioned earlier, only about 270 MB of local disk space will be available on the local workstation, so running applications which use a lot of non-local data on the office workstations could cause a large increase in the amount of network traffic, which can cause serious degradation of the performance of the whole network.

Each server will also have an FDDI connection, to allow higher bandwidth connections in particular to the IBM CFS system and also to the Cray UNICOS system. This will allow fast file transfers between the servers and both CFS and UNICOS.

The servers will also store software, applications, utilities, man pages etc. which are too large to be usefully stored locally on the office workstations. This aspect of the servers will be fully duplicated, so that if one of the servers is unavailable, the other will be able to provide the service.

UNIX

Both the Indigos and the Crimsons will run IRIX, Silicon Graphics' version of Unix, which is based on AT&T System V Release 3, with the usual "BSD extensions". The version to be installed initially is SGI's latest release, version 4.0.5F.

For most purposes, the IRIX system is very similar to the versions of Unix used on the Suns and the Cray, although there are differences once you delve below the surface. As well as the C shell and the Bourne shell, the KORN shell is available, which may be used if wished. (This shell is also available under UNICOS on the Cray, but not on the Suns at the moment.)

The new R4400 cpu that will be installed in the Crimsons requires the next major release of IRIX, 5.1, which is based on SYSTEM V Release 4. That is expected to be available in May 1993, so users can expect to see this new version of IRIX being introduced at ECMWF during the May-July time-frame. Although it is a major new version, it is fully binary-compatible with the current 4.0.5F version of IRIX, so the upgrade should be accomplished fairly easily.

Windowing environment

The Indigos will only have a MOTIF-based window manager. This is implemented by Silicon Graphics by a Window Manager called 4Dwm, which is actually a superset of MOTIF. It gives a somewhat improved appearance compared to Mwm, the standard MOTIF window manager (which is also available for use, if desired), and also easy access to a useful set of tools, called the Toolchest.

Although the Indigos only run the MOTIF window manager, it is generally possible to run OPEN WINDOWS based GUI (Graphical User Interface) applications on a Sun system, using the Indigo to display the application (i.e. using the Indigo as an X-Terminal). For example, the GUI version of WordPerfect 5.1, available for SUNs but not for Indigos, can be used in this fashion.

PC-emulation

A PC-emulation called SoftPC will be made available. SoftPC on the Indigo will emulate a 80268 PC with the performance of a slow 80486 PC. However, it must also be noted that SoftPC does not provide access to the NOVELL PC file server. Files have to be copied from the NOVELL server to a UNIX file system before they can be accessed from SoftPC on the Indigo. This can be done by using FTP to the PC file server. However, a new version of the NOVELL system is being introduced which will allow NFS access to the PC file server. This will make it easier to move files between the UNIX world and the PC network.

SoftPC allows UNIX files to be accessed as though they are MS-DOS files. However, since MS-DOS filenames are much more limited than UNIX filenames, SoftPC translates UNIX filenames to a form which is acceptable to MS-DOS. This can result in strange filenames.

One purpose in making SoftPC available is to provide access to WordPerfect 5.1. This is identical to running WordPerfect 5.1 on a PC.

Since SoftPC includes Windows 3.1, it is also possible to use the Windows version of WordPerfect 5.1. However, this has proved to be too slow for practical use.

New Software

There are a number of new software items being provided with the Indigos:

* WorkSpace - icon based file manager
* WorkShop - programming environment (also called CaseVision/WorkShop)
* INSIGHT - an online documentation tool, which provides access to documentation about the system (similar to the Sun Answerbook tool)
* EXPLORER - this is a data visualisation tool
* SHOWCASE - presentation graphics tool.

Further information about Insight, Explorer and ShowCase will be provided at a later date. For the moment, a brief introduction to WorkSpace and WorkShop follows. As with most GUI-based initiutive tools, it is much easier to understand how the tools work by trying them, rather than by reading about them.

WorkSpace

WorkSpace is an easy-to-use icon-based file manager. It has two types of windows - directory windows, which contain icons for all the files within the directory, and the WorkSpace window itself, which should be thought of as a work area in which to keep icons for frequently used files, directories and tools or applications. The icons used for files depend on the contents of the file, which makes it easier to find the files you are looking for.

Within directory windows Workspace provides the usual facilities for managing files - copying, moving, deleting, renaming files and using the mouse. It is possible to have many directory windows open simultaneously, which make it simple to move files about within directories by dragging them to the appropriate window. There is also a "Dumpster" for storing deleted files, but allowing the file to be restored if necessary; this is often called a WasteBasket or Trashcan on other systems.

Within the WorkSpace window, applications can be started by double clicking on the icon, or by dragging a file icon from a directory window and dropping it on the icon. (For example, if you have a printfile script which prints the file passed as argument to it and iconised in the window, dragging a file icon onto the printfile icon would cause the printfile script to execute taking the file dropped onto it as argument.)

WorkShop

In order to simplify the process of programming on the Indigos. the WorkShop programming environment is being installed. This supports C, C++ and Fortran programs, providing an intuitive Motif-based graphical user interface to various tools, in particular a source-based debugger, a static analyser, which can show the structure of a program in graphical form, a performance analyser, access to a user-specified editor, and also a build tool which interfaces to the Unix make tool.

The debug facilities are quite rich, augmenting the normal step, breakpoint, traps etc. with facilities such as watch points, array browser and graphical display of C structures.

SUN systems

Some of the services provided by the Silicon Graphics systems can be used on the Sun workstations, particularly running applications which use X Windows to display their output. In particular, software developed within the Centre will be able to be used in this fashion. However, a number of software application provided with the Silicon Graphics systems use a graphics library called GL to display their output; it is not possible to use these GL-based applications in conjunction with Sun workstations (or X-Terminals).

Performance

Perhaps the most exciting aspect of the new Indigos is their performance, which is among the highest available in workstations of this class. This performance will be an enabling technology, making it possible to use the Indigos for applications which are currently run on the Cray C90. In particular, since both MAGICS and a version of MARS will be available, it should be possible to run programs which retrieve data from MARS and then use MAGICS to produce plots of the data locally on the workstation, rather than submitting a job to the Cray.

The standard measures of performance of Unix workstations are the "SPECint92" and "SPECfp92" figures, which are quoted by the manufacturers for all Unix workstations. These figures are produced by running a set of 20 benchmark programs, which have been established by the "Standard Performance Evaluation Corporation" or SPEC. This is a non-profit corporation with representatives from all the leading manufacturers, which has put together a standardized set of relevant benchmarks. The latest set are the CINT92 and CFP92, comprising 6 integer and 14

floating-point based programs respectively. There are more floating point programs because floating-point benchmarks show substantially more variability in performance. The SPECint92 and SPECfp92 figures are computed by running each set of programs and taking the geometric mean of the "SPECratios" for each program. The SPECratio for a program is the time taken to run the program on the system being measured divided by the time taken to run the same test on a VAX 11-780. (This is called the SPEC Reference time, and is published in the SPEC Journal; the SPECint92 and SPECfp92 figures are also published in the SPEC journal.)

You may also see the "SPECmark" of a system being quoted - this is an old measure of performance which was basically calculated in the same way as the SPECint92 and SPECfp92 measures, but does not distinguish between integer and floating point capability. Although the Sun SPARCstations and the SGI Indigos have fairly similar integer and floating point performance, there are some systems which have floating point performance twice that of their integer performance, so the SPECmark can be distorted by this.

The SPECint92 measures for the SGI R4000 Indigos and the SUN IPC SPARCstations are 57.6 and 12.4 respectively, and the SPECfp92 measures are 60.3 for the Indigo and 11.4 for the SUN IPC SPARC.

However, it is stressed that these SPEC figures are to be used carefully, and that anyone buying a workstation should always measure the performance of the systems using their own programs. A portable version of the IFS model was successfully executed on the Indigos as part of the evaluation. A T21 resolution model with full physics was able to complete 5 timesteps in approximately 5 minutes. This would make it feasible to interactively debug such programs on the workstations before moving them onto the Cray for production.

Summary

The Centre is now moving into the final phase of its program to move to a predominantly workstation based environment. This phase promises to be very exciting, particularly in view of the opportunities created by moving to the Silicon Graphics environment.

- Richard Fisker

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

The following two articles are reprinted here by courtesy of SCD Computing News. The first appeared in the January/February 1992 issue, and the second in the March 1992 issue of that publication. Jeanne Adams chaired the International Programming Languages Committee of the ISO (International Standards Organisation), and chaired the ANSI Committee which developed Fortran 90. She also co-authored the book "The Fortran 90 Handbook: Complete ANSI/ISO Reference".

CF77 5.0 OFFERS FORTRAN 90 BIT MANIPULATION

This article explains the Fortran 90 features for problems in bit manipulation programming available on Cray's CF77 5.0 compiling system. CF77 5.0 contains both constant and edit descriptor facilities for bit data; it also contains intrinsic functions for bit manipulation as well as setting and testing bits.

In Fortran 90, an integer variable may be defined with a value by using binary, octal, decimal or hexadecimal number systems. Cray Fortran includes integer, octal, and hexadecimal constant definitions from Fortran 90, but not the binary constant definition. Cray uses two forms of the hexadecimal constant definition: the Fortran 90 form (Z'12AB') and Cray's own extension (X'12AB'). Fortran 90 limits these definitions to integer type, while Cray Fortran allows their use with variables of any default type, such as real or logical.

On input or output, both in Fortran 90 and Cray Fortran, the data edit descriptors for binary (Bw), octal (Ow), and hexadecimal (Zw) are available in addition to integer (Iw).

A test program

The program TESTBITS in Example 1 defines M within the octal number system and HEX within the hexadecimal number system. It then prints M and HEX using all four data edit descriptors - integer, binary, octal, and hexadecimal - for a variable of integer type. Two fields are summed, and the result is printed in INTEGER and then OCTAL.

Example 2 shows the output results for the program TESTBITS. Note that upper and lowercase letters may be used either in the program source or on output. No distinction between upper and lowercase is made in interpreting the meaning of the program. Upper and lowercase may be mixed in any convenient and readable way in Fortran.

Example 1. Binary, octal, and hexadecimal facilities

```
                PROGRAM TESTBITS
                INTEGER FIELD1, FIELD2
                INTEGER M, ANSWER, HEX
                M = O'0770'
                HEX = Z'002B1'
                print *, 'INTEGER, BINARY, OCTAL, HEXADECIMAL'
                WRITE (6, 102) M, M, M, M
                WRITE (6, 102) HEX, HEX, HEX, HEX
                PRINT *, 'integer, octal'
                FIELD1 = 28
                FIELD2 = 72
                ANSWER = FIELD1 + FIELD2
                write (6, '(I10, 3X, O10)') FIELD1, FIELD1
                write (6, '(i10, 3x, o10)') field2, field2
                WRITE (6, 100) ANSWER, ANSWER

C
 100            FORMAT (I10, 3X, O10)
 102            FORMAT (I10, 3X, B25, 3X, O10, 3X, Z10)
                END
```

Example 2. Output results from program TESTBITS

```
INTEGER          BINARY          OCTAL          HEXADECIMAL

 504           111111000        0000000770       00000001F8
 689          1010110001        0000001261       00000002B1

Integer      octal

   28        0000000034
   72        0000000110
  100        0000000144
```

Example 3. Bit functions

```
        PROGRAM BIT_FUNCTIONS
        IMPLICIT NONE
        DIMENSION ANSWER(20)
        INTEGER M, N, ANSWER, I
        LOGICAL BIT_TEST
C
        M = O'05'     !in binary, 0101
        N = O'03'     !in binary, 0011
C
101     FORMAT (2B20,10X,A)
        Answer(4) = NOT (M)              ! Logical Complement
        PRINT "(B20,3X,O20,8X,A)", M, ANSWER(4), "LOGICAL COMPLEMENT)
        ANSWER(1) = IOR (M, N)           ! INCLUSIVE OR
        ANSWER(2) = IAND (M, N)          ! LOGICAL AND
        ANSWER(3) = IEOR (M, N)          ! Exclusive OR
        PRINT "(A)","INCLUSIVE OR        LOGICAL AND     EXCLUSIVE OR"
        PRINT "(3B15)", M, M, M
        PRINT "(3B15)", N, N, N
        PRINT "(3B15)", (ANSWER (I), I=1,3)
C
        ANSWER(5) = ISHFT (M,2)
        PRINT "(2B20,10X,A)", M, ANSWER (5), "Shift, end off"
        ANSWER(6) = ISHFTC (N,3,8)
        Print "(2B20,10X,A)", N, ANSWER (6), "Circular Shift"
C
        ANSWER(7) = IBITS (M, 2, 2)
        PRINT 101, M, ANSWER (7), "CHOOSE A SUBFIELD"
        ANSWER (8) = IBSET (N, 8)
        Print 101, N, ANSWER (8), "Set a Bit"
        ANSWER (9) = IBCLR (N, 0)
        Print 101, N, ANSWER (9), "Clear a Bit"
        BIT_TEST = BTEST (M, 2)
        Print "(L40,10X,""Bit Test"")", BIT_TEST
C
        CALL MVBITS (M, 0, 5, N, 6)
        Print 101, M, N, "Move Bits from M to N"
        END
```

Example 4. Logical operations performed with bit functions

INCLUSIVE OR	LOGICAL AND	EXCLUSIVE OR	
0101	0101	0101	
0011	0011	0011	
0111	0001	0110	RESULTS

Bit functions

There are nine bit manipulation functions, one bit testing function (BTEST), and a bit transfer or copy subroutine (MVBITS). Example 3 is a simple demonstration of these bit functions.

The results of the test program BIT_FUNCTIONS in Example 3 assume that M = 0101 and N = 0011. The rightmost bit is bit 0; counting to the left, the next is bit 1, and so on. Given M = 0101, the logical complement using B64 as the edit descriptor is 111111111111...1010. The program in Example 3 uses octal for the complement because B64 is a very long constant in binary.

The output in Example 4 shows the logical operations that can be performed with bit functions (IOR, IAND, and IEOR).

There are two shift functions, an "end off" shift (ISHFT) and a circular shift (ISHFTC). An argument in the shift functions indicates the number of bits to shift: a left shift is positive, a right shift is negative, zero is no shift at all. In the following output, M was shifted left two positions; all bits in the word are shifted either right or left, depending on the sign of the second argument.

101	10100	Shift, end off

In the circular shift, only the rightmost eight bits were shifted three positions to the left. The zeros on the right were shifted into these positions from bits 61, 62, and 63. The rightmost bit position is bit position 0.

11	11000	Circular Shift

The function IBITS picked a subfield of length 2, beginning from bit position 2.

101	1	CHOOSE A SUBFIELD

There are two functions (IBSET and IBCLR) that clear or set a particular bit in the field.

11	100000011	Set a Bit
11	10	Clear a Bit

The bit inquiry function (BTEST) returns a logical value, depending on whether the specified bit position is 1 or 0.

T	Bit Test

Five bits are selected from M beginning at bit 0 (00101). These five bits are transferred to N beginning in bit position 6. The other bits in N are unchanged.

101	101000011	Move Bits from M to N

Not available yet on CF77 from Fortran 90 is the ability to create a data type of a specified number of bits using the KIND= parameter on the type statement.

- Jeanne Adams

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

FORTRAN 90 RECURSION: SEE IT NOW ON CF77 5.0

Recursion is a Fortran 90 feature included as an extension to FORTRAN 77 in the Cray compiling system CF77 5.0; it is also available on the Cray as a compiler option. Recursion is the ability of a procedure to call itself repetitively - that is, a way of repeating a process. Usually a recursive problem can be restated as an iterative algorithm.

Recursion is not a FORTRAN 77 standard feature because recursion requires stack storage to save copies of local variables in the subprogram to use on successive calls, and many FORTRAN 77 compilers use static storage. When stack processing is used (as is possible with CF77 5.0), additional CPU work is necessary for each subroutine call; but even with this extra overhead, recursive algorithms are sometimes more efficient than nonrecursive algorithms because they are simpler and more straightforward.

Fortran 90 and CF77 5.0 both offer direct and indirect recursion. However, to avoid any ambiguity that might occur, Cray Fortran does not support array-valued results in recursive functions.

Direct recursion is a program unit (either a subroutine or function) that invokes itself. An example of direct recursion is if FUNCTION FACTORIAL calls itself to compute the factorial. Direct recursion can always be detected by the compiler, because it is clear from the static program source of the program unit.

Indirect recursion is a program unit that invokes another program unit, which eventually results in another invocation of the first program unit. For example, indirect recursion occurs if SUBROUTINE A calls SUBROUTINE B, which invokes FUNCTION C, which calls SUBROUTINE A again. The compiler cannot always detect indirect recursion, since it may occur as the result of separate compilation of distinct files or by using already compiled libraries.

Example 1 shows the use of N factorial. Example 2 shows a simple recursive subroutine. This binary search requires that the elements of ARRAY must be in ascending order.

Example 1. The use of N factorial

```
PROGRAM MAIN
 INTEGER FACTORIAL, I
 DO I = 1, 5
        WRITE (6,*) 'FACTORIAL(',I,') =', FACTORIAL (I)
 END DO
END

RECURSIVE INTEGER FUNCTION FACTORIAL (N)
 INTEGER N
 IF((N .EQ. 1) .OR. (N .EQ. 0)) THEN
        FACTORIAL = 1
 ELSE
        FACTORIAL = FACTORIAL(N - 1) * n
 END IF
END
```

Example 2. A recursive subroutine

```
RECURSIVE SUBROUTINE SEARCH (MATCH, VALUE, ARRAY, LOWER, UPPER)
 INTEGER MID, MATCH, VALUE, LOWER, UPPER, ARRAY (UPPER)
 MID = INT( (LOWER+UPPER)/2)
 IF (MID .LT. LOWER .OR. MID .GT. UPPER) THEN
        PRINT "(A, I5)", "ERROR OCCURRED", MID
        STOP                    !or RETURN
 ELSE IF (VALUE .LT. ARRAY (MID) ) THEN
        UPPER = MID - 1
        CALL SEARCH (MATCH, VALUE, ARRAY, LOWER, UPPER)
 ELSE IF (VALUE .GT. ARRAY (MID) ) THEN
        LOWER = MID + 1
        CALL SEARCH (MATCH, VALUE, ARRAY, LOWER, UPPER)
 ELSE
        MATCH = MID
 END IF
END
```

Syntax for recursive functions and subroutines in CF77 5.0 and Fortran 90

Recursive functions

The syntax for recursive functions in CF77 5.0 can be stated as follows, using square brackets for optional items:

[type-spec] RECURSIVE FUNCTION function-name ([d [, d] ...])

RECURSIVE [type-spec] FUNCTION function-name ([d [, d] ...])

where type-spec is one of the following: INTEGER, REAL, DOUBLE PRECISION, COMPLEX, LOGICAL, or CHARACTER [*len], and d is the dummy argument name.

Fortran 90 has a RESULT clause in the syntax for a recursive function in addition to the prefix of RECURSIVE. The use of the result clause simplifies some of the repetitive processes in that the result clause name and function name are different local variables. Optional items are enclosed in square brackets. The syntax for a recursive function in Fortran 90 is:

prefix FUNCTION function-name ([d [, d] ...]) [RESULT(result-name)]

where prefix is [type-spec] RECURSIVE or RECURSIVE [type-spec], function-name is a name (not result-name), d is a dummy argument name, and result-name is a name different from function name.

Recursive subroutines

The syntax for recursive subroutines is the same in Fortran 90 as in Cray Fortran:

RECURSIVE SUBROUTINE subroutine-name [([d [, d] ...])]

where d is a dummy argument name, or * (for an alternate return).

- Jeanne Adams

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

CFS LEVEL 61 FEATURES

CFS level 61 was introduced at ECMWF in November 1992 as a replacement for the previous level 59 CFS. The major new features that are of interest to the users are improved Tape Management and Wildcard Delete.

Tape Management

Prior to this release of CFS, the space on cartridge volumes had been controlled by means of bitmaps, which imposed a limit on the amount of data that could be stored on a cartridge. This was fine for the 18 track cartridges with maximum data capacity of 200 Megabytes thus far used at ECMWF. However, in order to accommodate the new higher density cartridge technologies which include longer, thinner tapes, data compression, 36 track recording and Enhanced Cartridge format it has become necessary to give up using bitmaps.

Instead of bitmaps, CFS now uses the block-ids used by the cartridge controllers to define the start and end position of each file on a cartridge. These block-ids are kept for each file in the CFS Directory and are used to:

* rapidly position at the start of each file being read or written;

* calculate the amount of free space remaining at the end of a cartridge each time a new file is written.

This feature has already enables us to gain the following benefits:

* a marked improvement in cartridge positioning times, particularly when writing a file at the end of a cartridge holding many small files;

* the use of 25% longer cartridges with data compression, increasing the average cartridge capacity from 190 Megabytes to 330 Megabytes;

* the ability to completely fill all cartridges, even in cases where a set of highly compressible files may result in up to a GigaByte of data on a cartridge.

Level 61 CFS will also support the upgrade of the cartridge drives in the Storage Technology silos to support 36 track and Enhanced Format cartridges. This will eventually result in cartridges holding more than 1 GigaByte of CFS data.

At present it has been possible to increase the maximum CFS file size to 240 Megabytes, the maximum we can guarantee to get on a single cartridge.

Wildcard Delete

Wildcard Delete enables a user to delete all the files in a CFS subtree with a single command. In order to protect users from inadvertent or over-zealous use of this feature, there is also an automatic delayed delete feature whereby any files deleted are temporarily renamed and marked as dying when listed by the user. The delay in deleting these files defaults to one day, but this can, optionally, be increased by the user. As part of this feature there is also a new RESCUE command to recover a dying file.

The following article gives further details of the effect of the new features in CFS on UNIX-ECFILE.

- Tony Stanford, Jean-Luc Pepin

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

ECFILE - ECMWF'S DATA ARCHIVAL SYSTEM

Introduction

In 1988, ECFILE was introduced on the Cray X-MP, which was running the COS operating system. ECFILE was loosely modelled on the features of ACQUIRE and DISPOSE, which were COS commands to retrieve and store data on the front-end systems connected to the Cray machine. ACQUIRE would first look to see if the data already existed on the Cray and, if not, would fetch the data from the front-end and save it on the Cray's disks. In 1987, Cray released a software product called SUPERLINK. Superlink was in fact two pieces of software: one component ran on the Cray and the other ran under MVS on the IBM. These two components communicated with each other using the Superlink protocol, which was based on OSI protocols, over FEIs. FEIs are Front-End Interface boxes, used to connect Cray channels to IBM channels. Over the years ECFILE has developed through several stages:

Stage 1:

Initially, ECFILE did not transfer the data itself to/from CFS, instead it used the File Transfer Access Mechanism of Superlink (FTAM). ECFILE only passed information about the data to the IBM. This procedure suffered from the need to stage the file out of CFS onto disk on the IBM and from there to the Cray disks, or vice versa. As this was slow and restricted by the amount of IBM disk space, stage 2 was produced.

Stage 2:

An ECMWF-defined protocol called MIRE was developed to allow data to be transferred directly between the Cray and CFS. This speeded up the transfers considerably and used fewer resources on the IBM.

Stage 3:

With the introduction of the Cray Y/MP8 a major overhaul of the user interface to ECFILE was carried out. This was because this machine ran UNICOS and not COS. The underlying parts of ECFILE were also modified so that about 50% of the code was original (stage 2), written in Fortran and 50% of the code was new, written in C. It was decided not to rewrite the code completely as there was a very short time in which to migrate from COS to UNICOS, and ECFILE was required at the earliest possible time to enable the rest of the migration, which needed to access CFS.

Stage 4:

Once ECMWF started to acquire other UNIX platforms, such as workstations, it was decided to provide ECFILE on these. This gave the opportunity to produce a "generic UNIX" ECFILE. ECFILE is now written completely in C, and there is a single source, with only a few conditional compilation statements, depending on which UNIX system it is being compiled upon. It uses TCP/IP protocols with FTP for the data transport and runs on all of ECMWF's UNIX systems (UNICOS, SunOS, IRIX, EP/IX, AIX, ULTRIX).

Usage

The MAN page for ecfile can be seen on any of the UNIX systems by typing:

```
man ecfile
```

The synopsis is reproduced below:

```
NAME
        ecfile - ECMWF's file archival and retrieval system

SYNOPSIS
        ecfile  [-d disable-string] [-e enable-string] [-g group]
                [-h existing-file-handling] [-i informative-text] [-k keep-delay]
                [-n new-name] [-o list-options] [-p CFS_path(s)] [-r release-date]
                [-s secret-password(s)] [-t type-of-file] [-v validation-entry(s)]
                [-x exit-status-file] [-V]
                CFS-function [UNIX-filename(s)]
```

New features

Some new features in CFS are supported with the "generic UNIX" ecfile, as are some local UNIX ecfile features. These are as follows:

the "-k keep-delay" option

This may be used on a "delete" request. Its meaning is:
Delete the file in CFS, but keep it available such that a "rescue" request can "undelete" the file within the "keep-delay" (days) deadline. For example:

```
ecfile -p /xyz/my_file -k 2 delete          # Whoops a mistake

ecfile -p /xyz/my_file rescue               # OK "undeleted"
```

the "-x exit-status-file" option

Using this option will cause ECFILE to write into the specified file, a record containing 3 blank-separated fields. These fields are:

ECFILE	exit status (0=OK, e.g. 9 = CFS error)
CFS	exit status (0=OK, e.g. 2 = CFS node already exists)
FTP	exit status (0=OK, e.g. 530 = LOGIN failed).

These fields could be used in scripts to take corrective action in case of an ECFILE failure, e.g. in the case where the "-d a" option is used to prevent ECFILE from returning a non-zero status after the "set -e" Bourne Shell command has been used to abort the job on detecting a non-zero exit status.

e.g.

```
# Add a directory, but don't abort if it already exists
# --------------------------------------------------------
set -e
ecfile -d a -x $TMPDIR/exit_status add mydir
set -- 'cat $TMPDIR/exit_status          # You could use 'cut' instead
if[$1 -ne 0]; then
   if[1 -ne 9]; then                     # '9' = 'CFS error'
        exit $1
   elif[$2 -ne 2]; then                  # '2' = 'node already exists'
        exit $1
   fi
fi
```

There is a script on the Cray in file /usr/local/bin/ecfcheck that is designed to check the exit status file. The above example could be rewritten as:

```
# Add a directory, but don't abort if it already exists
# ......................................................................
set -e
ecfile -d a -x $TMPDIR/exit_status add mydir
ecfcheck $TMPDIR/exit_status 9 2
```

the "-V" option

This just causes ECFILE to send to stderr information regarding which version of ECFILE is being used, e.g.

ECFILE version TCP/IP 1.0, compiled on Mar 12 1993 at 12:49:03.

the "audit" function

The command:

ecfile -p /xyz audit

will produce a listing of all the files under /xyz in the CFS directory. This listing is produced once every 24 hours in the early hours of the morning, so it will only reflect the contents of the directory as they were at that time.

the "rescue" function

The command:

ecfile -p /xyz/my_file rescue

will rescue a file deleted with the "-k keep-delay" option. This will effectively "undelete" a file, provided that the "keep-delay" time limit has not been exceeded.

CFS "wildcard" characters

These special characters, which are:

*, ?, [], and ^

can be used to specify CFS pathnames with the *list, delete* and *rescue* functions of ECFILE. Their use is modelled after their implementation in the UNIX shells, i.e.

- `*` will match any string (including the NULL string)
- `?` will match any single character (excluding the NULL character)
- `[...]` enclosing a sequence of characters matches any one of those characters. Enclosing two characters separated by a minus sign matches any character lexically within the range of the 2, using the ASCII collating sequence
- `^` functions as a NEGATIVE or NOT function.

N.B. This feature should be used with care.

Deleting files using wildcard characters is VERY dangerous. It is very easy, using this facility, to delete huge numbers of files accidentally. To help alleviate such problems, CFS has been set up to cause wildcard deletes to have a "keep-delay" of 1 day. This means that if you do make a mistake on the wildcard delete, you have 1 day in which to reverse the deletes by using the "rescue" function. This ONLY applies to wildcard deletes, NOT normal ones; e.g.

```
ecfile -p /xyz/* delete          # Whoops my mistake
ecfile -p /xyz/* rescue          # Got them back.
```

Incompatibilities

The new ECFILE is compatible with the old, Superlink-based, ECFILE in all but two respects:

(1) The "list" function produces its output in lower-case, whereas the old ECFILE produced its output in upper-case;

(2) the old ECFILE would allow multiple files to be specified on a single command. This little-used feature is not supported with the new ECFILE;
e.g.

```
ecfile get file1 file2 file3 file4    # Not supported
```

However, this can still be done using the Shell's iteration structures,
e.g.

```
#! /bin/sh
for file in file1 file2 file3 file4
do
        ecfile get $file
done

#! /bin/csh
foreach file (file1 file2 file3 file4)
        ecfile get $file
end
```

- Neil Storer

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

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 293). 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
265	Lost UNICOS outputs submitted via RJE or VAX Microfiche changes
266	Reminders on how to import/export magnetic tapes
268	Changes to WMO FM 92 GRIB

No.	**Still Valid Article**
270	Changes to the Meteogram system
271	New ECFILE features on UNICOS
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
286	Improving the performance of "model" jobs on the Cray Y/MP8 at ECMWF (pre-allocating disk space)

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

TABLE OF TAC REPRESENTATIVES, MEMBER STATE COMPUTING REPRESENTATIVES AND METEOROLOGICAL CONTACT POINTS

Member State	TAC Representative	Computing Representative	Met. Contact Point
Belgium	Dr W Struylaert	Mrs L Frappez	Dr J Nemeghaire
Denmark	Dr A M Jørgensen	Mr P Henning	Mr G R Larsen
Germany	Dr B Barg	Dr B Barg	Dr Rüge
Spain	Mr T Garcia-Meras	Mr J Juega	Mr R Font Blasco
France	Mr J Goas	Mr D Birman	Mr J Goas
Greece	Mr D Katsimardos/	Mr I Iakovou/	Dr N Prezerakos/
	Dr G Sakellarides	Mr G Konstantinidis	Mrs M Refene
Ireland	Mr J Logue	Mr L Campbell	Mr T Sheridan
Italy	Dr M Capaldo	Dr S Pasquini	Dr M Conte
Yugoslavia*			
Netherlands	Mr S Kruizinga	Mr S Kruizinga	Mr G Haytink
Norway	Mr K Björheim	Ms R Rudsar	Mr P Evensen
Austria	Dr G Wihl	Dr G Wihl	Dr H Gmoser
Portugal	Mrs M Leitao	Mr C M Fernandes	Mrs I Barros Ferreira
Switzerland	Mr M Haug	Mr B Bachofner	Mr M Schönbächler
Finland	Dr M Alestalo	Mr T Hopeakoski	Mr P Nurmi
Sweden	Mr G Bergh	Mr S Orrhagen	Mr O Åkesson
Turkey	Mr M Örmeci	Mr M Örmeci	Mr M Örmeci
United Kingd.	Dr R Wiley	Dr A Dickinson	Mr C R Flood

* At its 37th Session (December 1992) the Council decided that the telecommunications link between ECMWF and Belgrade would be terminated with immediate effect, and that henceforth no ECMWF documents would be sent to the Federal Republic of Yugoslavia (Serbia and Montenegro).

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

ECMWF CALENDAR 1993

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
22 - 24 September	**Scientific Advisory Committee**, 21st session
4 - 6 October	**Technical Advisory Committee**, 18th session
6 - 8 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. 190:	A microwave radiative transfer model
Technical Memorandum No. 191:	Report on the seventh meeting of Member State Computer Representatives, 28 - 30 September 1992
Technical Memorandum No. 192:	Impact of a simple vertical diffusion scheme and of the optimisation time interval on optimal unstable structures
Technical Report No. 69:	A preliminary study of the impact of C-band scatterometer wind data on global scale numerical weather prediction

Forecast and Verifications Charts up to end January 1993

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

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.

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
Console fax number	-	+44 734 499 840		
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	-	John Hennessy (acting)	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