In this issue

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, RG2 9AX, UK

Fax: (0118) 986 9450

Telephone: National 0118 949 9000

International . . . . +44 118 949 9000

Internet URL . . . . http://www.ecmwf.int

Cover

"De Elfstedentocht 1997" at Sloten - a major skating event in the Netherlands - see page 6 for how ECMWF's EPS forecasts assisted in this year's event.

Editorial

ECMWF has been running operationally an Ensemble Prediction System (EPS) since 1 May 1994. In December 1996, two major improvements were made to the EPS system, the model resolution was changed from T63 with 19 levels to T_L159 with 31 levels, and the number of members in the ensemble was increased by 32 to 50. An overview of these changes is described on page 2.

Already this enhanced EPS has been put to good use in the high profile skating race "Elfstedentocht" held in the Netherlands. This race can only be held whenever the ice on the canals and lakes is thick enough. The involvement of the EPS to predict when the weather conditions would be suitable is graphically described on page 6, resulting in the race being held on 4 January 1997.

Anew global Ocean Wave model with a spatial resolution of 55 km was introduced into operation in December. This has been a major step forward compared to the previous model and was made possible by the introduction into service of the Fujitsu VPP700. The article on page 9 has the details.

Two further ECMWF Technical Memoranda (numbers 219 and 220) are summarised on pages 12 and 13.

Each year talks are given at ECMWF of the current status and plans for the computing service. Asummary of the most recent of these talks is on page 14.

In the last Newsletter an article covered the history, current situation and future plans of ECMWF's involvement in distributed memory computer systems. One of these systems is now installed and operational at ECMWF, a Fujitsu VPP700. The article on page 17 describes this system in some detail.

ECMWF produces a range of publications, a summary of the major publication types is given on page 21. In addition, the article on page 20 details the recent changes to ECMWF computer user documentation.

Changes to the operational forecasting system

Recent changes

- A high resolution version of the optional project global wave forecast model was introduced on 5 December 1996. The new model has a resolution of approximately 50km using a quasi-regular latitude/longitude grid with 0.5 degree intervals between latitude rows and varying numbers of points along latitudes.

- Model cycle 15R7 was introduced on 10 December 1996 in both the T213 L31 10-day forecast and the EPS. The main features of the new cycle are:
  - (i) a two-time level semi-Langrangian scheme
  - (ii) advection of cloud variables
  - (iii) changes in the snow albedo, expected to reduce the cold bias in 2m temperature in spring
  - (iv) change to the humidity analysis to eliminate unrealistic drying of the extra-tropical lower stratosphere
  - (v) change to the use of SSM/I sea ice data to eliminate erroneous sea points over permanent ice
  - (vi) minor change to the model orography and slight adjustments to other fixed surface fields.
- On the same date, the EPS was enhanced to run at T_L159 resolution (linear grid) with 50 ensemble members and the control.The model physics are computed on the quasi-regular Gaussian N80 grid.
- Model cycle 15R8 was introduced on 21 January 1997. This cycle corrects an error affecting the degree of balance between the mass and wind analyses.

Brian Norris

The enhanced high resolution Ensemble Prediction System

The new EPS

From 10 December 1996, the operational Ensemble Prediction System was changed in two ways aimed at providing greater skill and usefulness.

(i) The model on which the EPS is run, has been upgraded from T63 resolution with 19 levels to T_L159 truncation with 31 levels (linear grid, time step 2700 sec).

(ii) The number of perturbed ensemble members has been increased from 32 to 50. A control run is also made at T_L159L31.

Experiments have indicated that increases in both resolution and ensemble size were desirable for the next generation EPS. The chosen configuration appears to be a satisfactory compromise incorporating both requirements. A scientific paper documenting results from these experiments has been submitted to the Quarterly Journal of the Royal Meteorological Society (Buizza, R., Petroliagis, T., Palmer, T.N., Barkmeijer, J., Hamrud, M., Hollingsworth, A., Simmons, A., and Wedi, N., 1997. Impact of model resolution and ensemble size on the performance of an ensemble prediction system. *Q.J.R. Meteorol. Soc.*, submitted on 21 January 1997).

This enhanced configuration became operationally feasible on the VPP700 due to the combination of two advances in the model numerics. First, the introduction of a two-time-level semi-Lagrangian integration scheme has significantly increased the computational efficiency of the model. Second, the use of a linear grid enables the accuracy of a T_L159 model to be obtained at a cost similar to that of the original (quadratic grid) T106 model.

There are several advantages to run the EPS with a higher resolution model:

(i) Since the eddy kinetic energy of the T_L159 resolution model is larger and perturbations amplify more realistically, the model will have a higher and more climatologically correct synoptic variability. This will be noticed both with small scale intense baroclinic developments and large scale atmospheric features, like occurrence of blockings or large scale changes in the flow patterns.

(ii) The great similarity between the T_L159 and the operational T213 models will ensure that the deterministic forecast normally lies within the cluster.

(iii) The resulting larger spread should reduce the number of occasions when the ensemble does not cover the verifying analysis. The increase in number of ensemble perturbations should also support this.

(iv) The high T_L159 resolution allows improvements in modelling of physical processes like convection and moisture fluxes. Surface weather parameters like 2 m temperatures, 10 m winds and precipitation amounts should be much more realistic than with the T63 resolution model.

(v) A more detailed land-sea mask allows for regional climatological differences, e.g. the difference near coastal areas. It will also reduce interpolation errors when data are retrieved (Fig.1).

(vi) The relation between ensemble spread and the skill of the Control, an important statistic in assessing the quality of an ensemble, should be further improved.There has been a decrease in the number of undesired cases when the ensemble has small spread and the Control turns out to be bad.

Precipitation probabilities

Preliminary verification results show that an overall improvement has been achieved. Fig. 2 shows reliability diagrams for 72-96 hour forecast rainfall for Europe during two early winter periods 1995-96 and 1996-97. Figs. 2a) and c) show the verifications for probabilities for >1 mm/24h.The verification line is in both cases close to the 45-degree line which indicates a perfect correspondence between forecast probabilities and observed frequencies, while the later verification period with the new system yields slightly better correspondence.

The improvement is more clear for rainfall >10 mm/24h (Figs. 2b) and d)). For high probabilities 80-100%, with the old system, rain occurred only around 50%, with the new system it is around 70%.

Case study

On 28 October 1996, a violent storm affected Ireland, Wales and England (fig.3). The storm had its origin as a tropical cyclone LILI over the Caribbean (and will be called low L). The ECMWF D+5 forecast from 23 October placed “low L” northwest of Portugal instead of over Ireland.

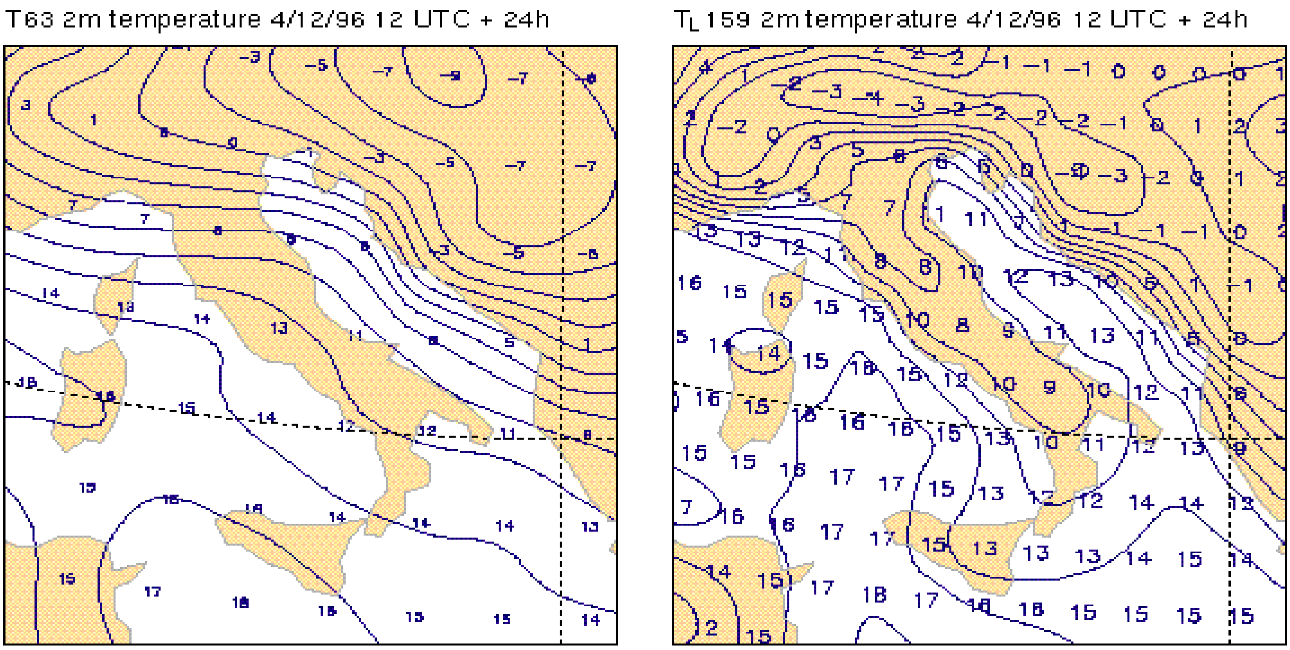

***Fig.1:** Examples of the 2m temperature forecasts from the old T63 system and the new T_L159 system.*

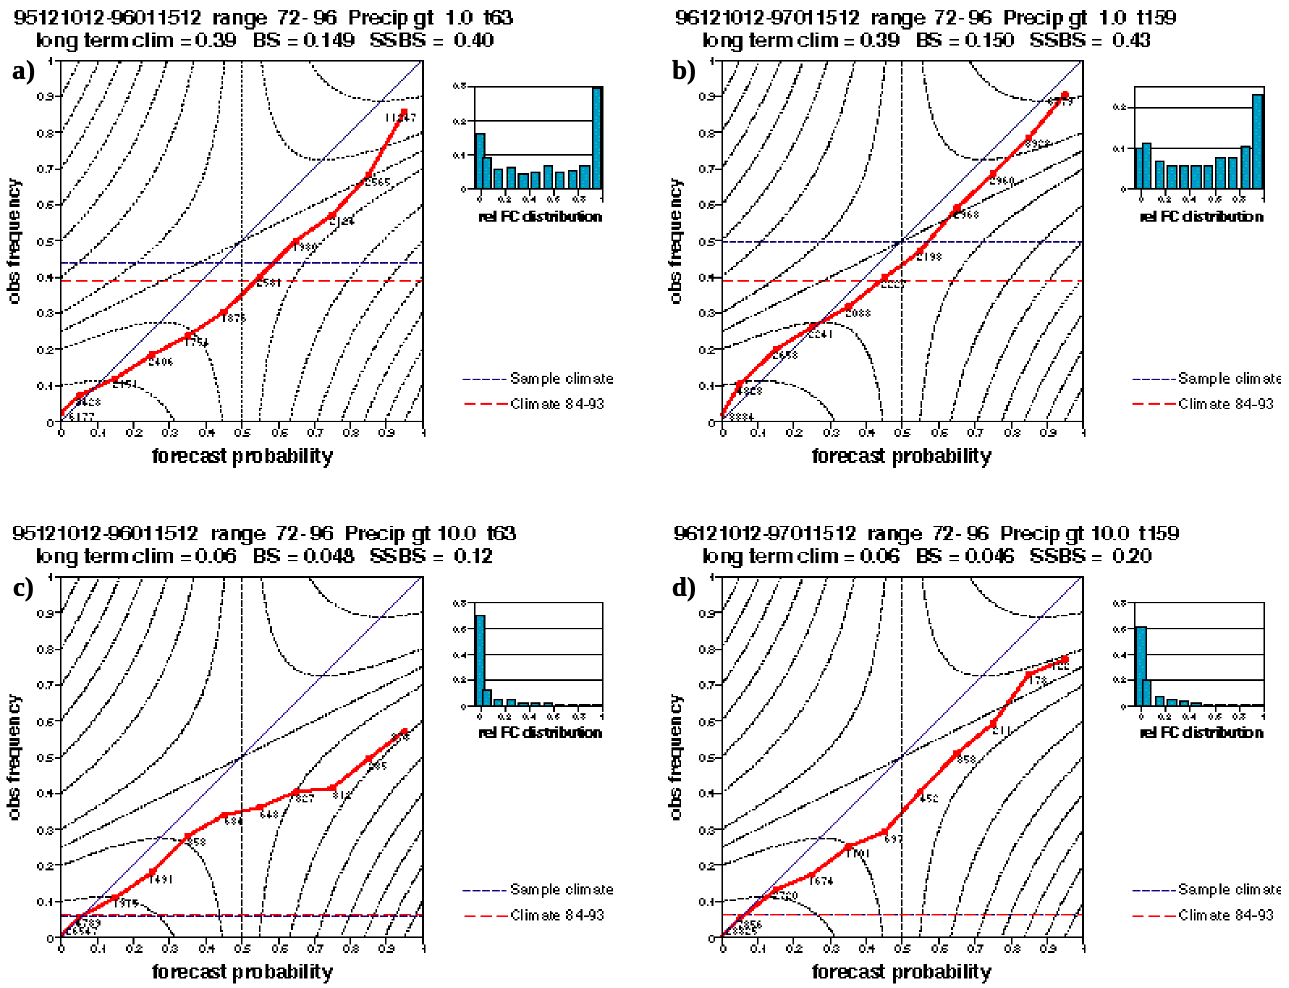

***Fig.2:** Reliability diagrams for 72-96 hour forecast rainfall for Europe 10 December 1995-15 January 1996 (a and c), for the same period 1996-97 (b and d). a) and c) show the verifications for probabilities for >1 mm/24h, c) and d) for >10 mm/24h. The 45-degree line indicates a perfect correspondence between forecast probabilities and observed frequencies.*

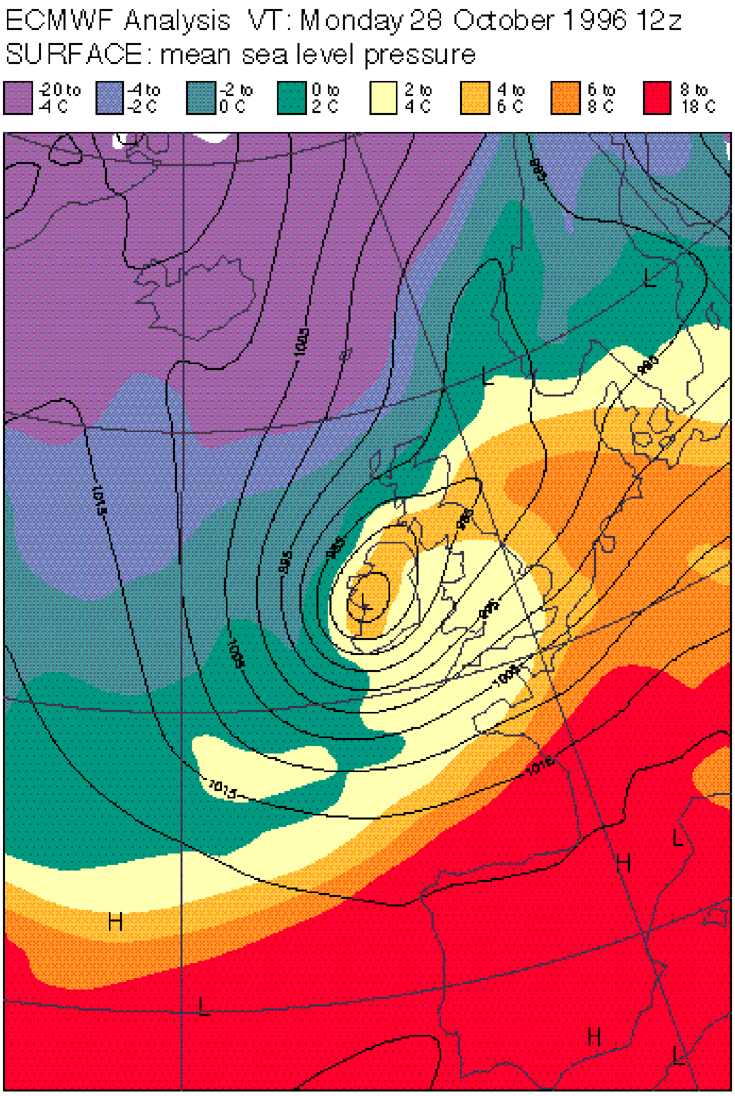

Fig.3: *[left] The ECMWF analysis 28 October 1996 12 UTC when “Low L”, the reminiscences of ex-tropical cyclone “LILI” passed over Ireland, Wales and southern England causing structural damage and fatal accidents.*

Fig.4: *[below] The ECMWF operational D+5 forecasts from 23 October 1996:*
a) the T63 control forecast where “Low L” only remained as a sharp trough NW of Spain;
b) the positions of forecast depressions below 995 hPa in the 33 ensemble members, numbers indicate the respective member. Most of the forecasted lows referred to another cyclone over the Norwegian Sea.

Fig. 5: *[facing page, top] a) The T_L159 forecast based on the same analysis as in Fig.3. The “Low L” is now in the unperturbed forecast, although in the wrong position;*
b) around 30 members put “Low L” near the unperturbed forecast or further to the west, but 10-15 clustered in the right position around Ireland; the positions northeast of the British Isles refer to another low

Fig. 6: *[facing page, bottom] Three typical examples of the 10-15 members which not only were forecast in the right area, but also displayed a synoptic structure which was close to or identical to the extra-tropical cyclone in the verification (upper row); The three “best” members from the T63 ensemble had their positions quite correct but the synoptic structure resembles a more normal extra-tropical cyclone (lower row).*

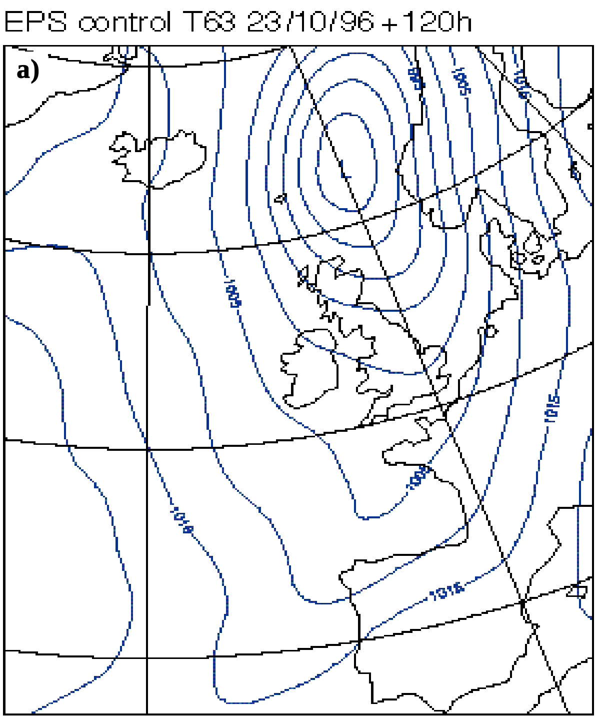

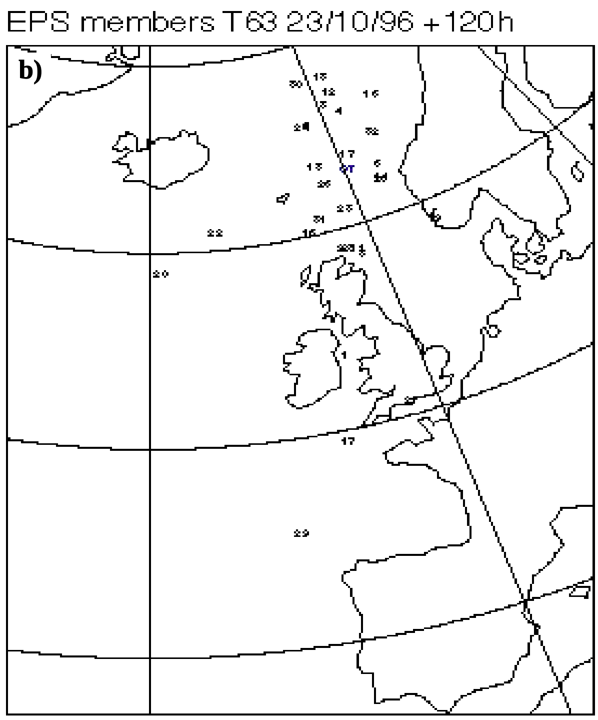

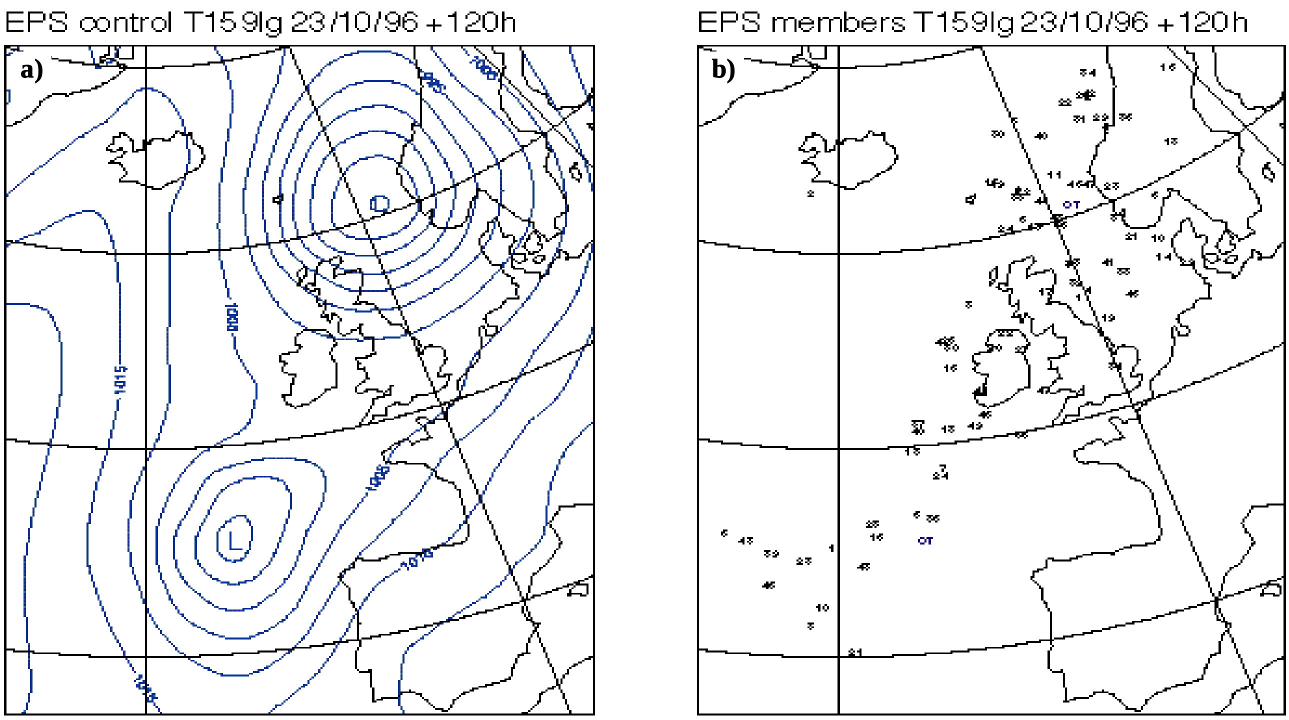
EPS control T159lg 23/10/96 +120h
a)
EPS members T159lg 23/10/96 +120h
b)

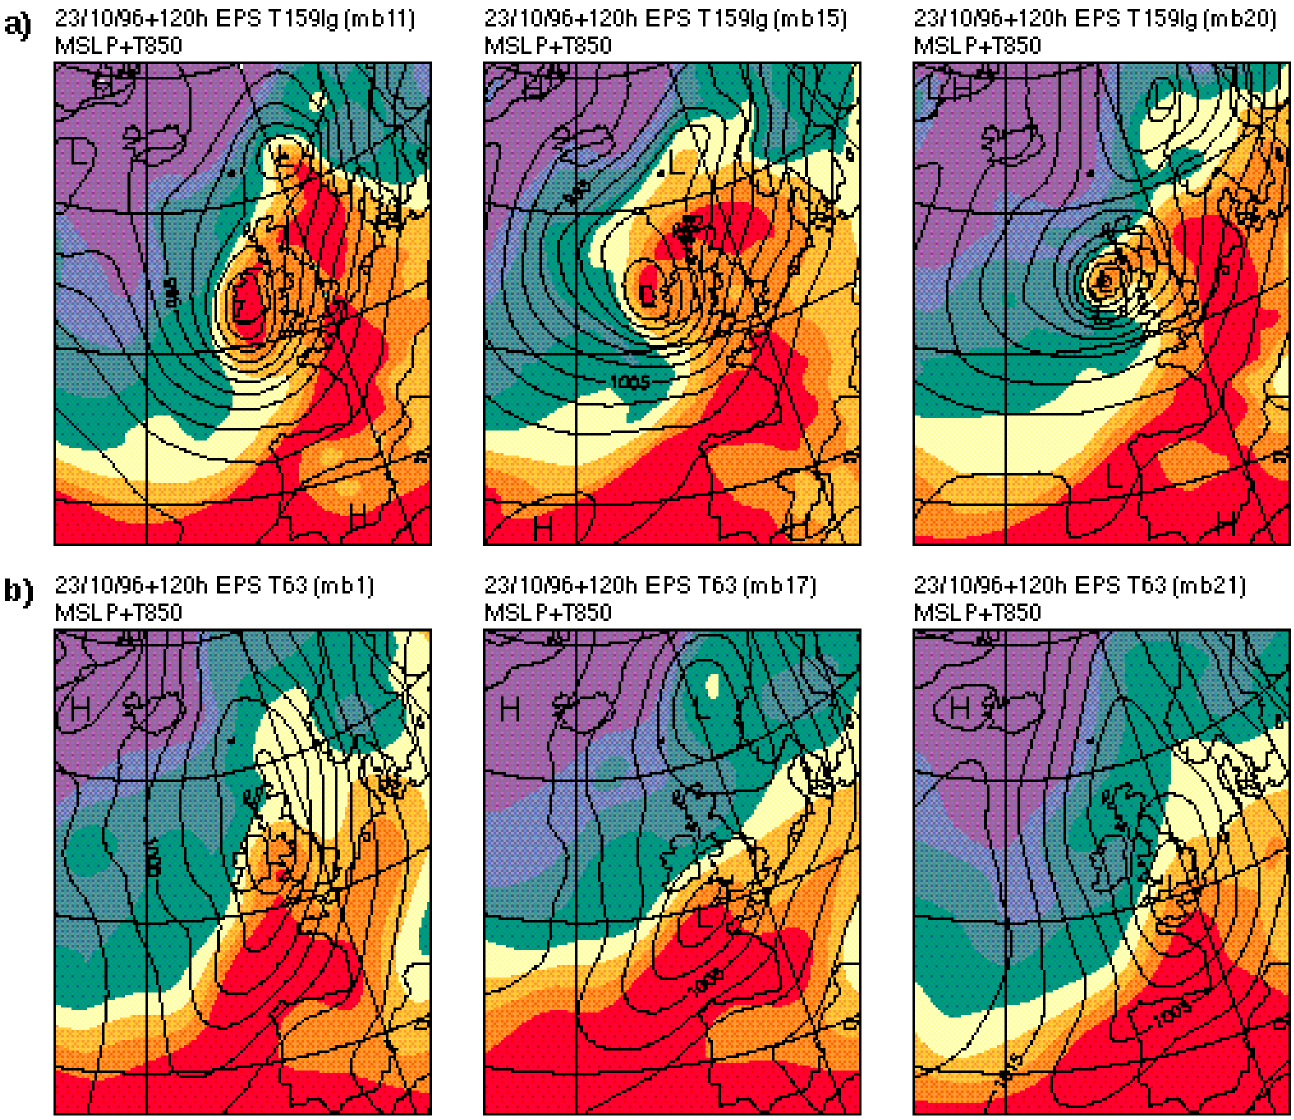
a)
23/10/96+120h EPS T159lg (mb11)
MSLP+T850
23/10/96+120h EPS T159lg (mb15)
MSLP+T850
23/10/96+120h EPS T159lg (mb20)
MSLP+T850
b)
23/10/96+120h EPS T63 (mb1)
MSLP+T850
23/10/96+120h EPS T63 (mb17)
MSLP+T850
23/10/96+120h EPS T63 (mb21)
MSLP+T850

The T63-based EPS did not give any useful guidance: the Control had a trough north of Portugal and only five members developed a low, three of them in the approximate correct position (Fig.4).

The T_L159-based EPS was quite different. The Control was similar to the T213: it managed to maintain "low L", but placed it off the European coast, as did most of the 50 members. However, 10-12 members forecast "low L" at or close to the observed position (Fig.5). The shape and intensity of "low L" in these T_L159-based forecasts were in good agreement with the observed structures (Fig.6 upper row), in contrast to the three best members from the T63 system (Fig.6 lower row).

Summary

(i) Initial evaluations of the new system indicate that the EPS should be able to improve the detection of extreme weather events.

(ii) The new EPS should provide the forecasters with improved guidance to assess probable alternative forecast solutions.

(iii) The surface weather parameters (temperature, wind and precipitation) have become more realistic both physically and geographically.

Anders Persson, Clive Temperton

Skating on EPS

The ECMWF ensemble system EPS assisted in the organization of a major skating event in the Netherlands. In this contribution we would like to explain how EPS was used, but in order to appreciate its significance, the reader should be aware of the impact of such a skating event on the population of the Netherlands. So, before we give the account of the use of EPS, we will first give a short overview of the history of the "Elfstedentocht".

Skating Tours

Every few years in the Netherlands, a winter is sufficiently cold to make it possible to skate on canals and lakes. We have many shallow canals and lakes (not much deeper than 2 metres), which freeze over very rapidly and there is usually enough wind to blow the snow off the ice. Such conditions may last for a few weeks and the astonished foreign visitor will find, on those days, almost the entire population out skating, making the country side look pretty much like a historic painting of Breughel. This is all good fun, despite the occasional concussed brain. However, every now and then the situation runs a bit out of hand! That is when the frost is so severe that it is actually possible to skate from one village to the next, and people start to organise tours. Several such tours exist, but by far the most famous is held in the province of Friesland, and completes a 200km loop through 11 cities. It is therefore called, very imaginatively, the Eleven Cities Tour, or "Elfstedentocht", a tour which starts in the early morning at 6am in total darkness, and finishes for most people in the late evening, again in darkness. The deadline is set at 12 midnight!

Who wants to skate 200km, partly in darkness, and in cold and windy conditions, you may wonder? Well, the answer is: many . The number of participants has, in fact, to be limited by the Tour Organization to 20,000(!) because more would be too dangerous, as the ice might wear out too soon (or might simply give way).Apart from all these worries the Organizing Commission also has to deal with a crowd of about half a million spectators. This is a major responsibility, and the decision to go ahead with the Tour is not a light one. So, there is a great deal of stress involved, also because the period of frost is rarely long enough.Most winters it is touch and go, with the result that the Tour has often been held on the day before thaw sets in.

The Tours have, therefore, been very few and far between.The first was held in 1909, and, until this year, only 14 had been held. There was one very famous Tour in 1963 which many had to abandon with frozen toes or eyelids. After this heroic Tour of 1963, however, a myth was born, which grew to incredible proportions, because after 1963, we had to wait more than 20 years for the next Tour. But then, in 1985, after 22 years, there was a Tour, and, immediately in the year after, in 1986, there was another one. (We will spare you the opinions of the "climate experts" who believed in those years that the Tour could never be held again).

After these two Tours we had to wait a long time again. Last year was close. Unfortunately the Tour Organization hesitated just a bit too long, the thaw came and chances were lost. The reader will, by now, understand that criticism was heaped upon the poor Tour Organizers by the entire Dutch population. Criticism also came from the weather forecasting world because the forecasts had been completely ignored. In fact, the policy of the Tour Organization had always been one of extreme caution: the ice had to have a minimum thickness of about 20 cm. Only when that simple condition was satisfied would the tour be announced for two days later. After last winter's experience, contacts were established between the Tour Organization and KNMI , and we became the official Tour weather advisers.

Forecasting assistance

In summer and autumn 1996, KNMI and the Tour Organization met and agreed on the collaboration. First of all it was decided that contacts would run through one

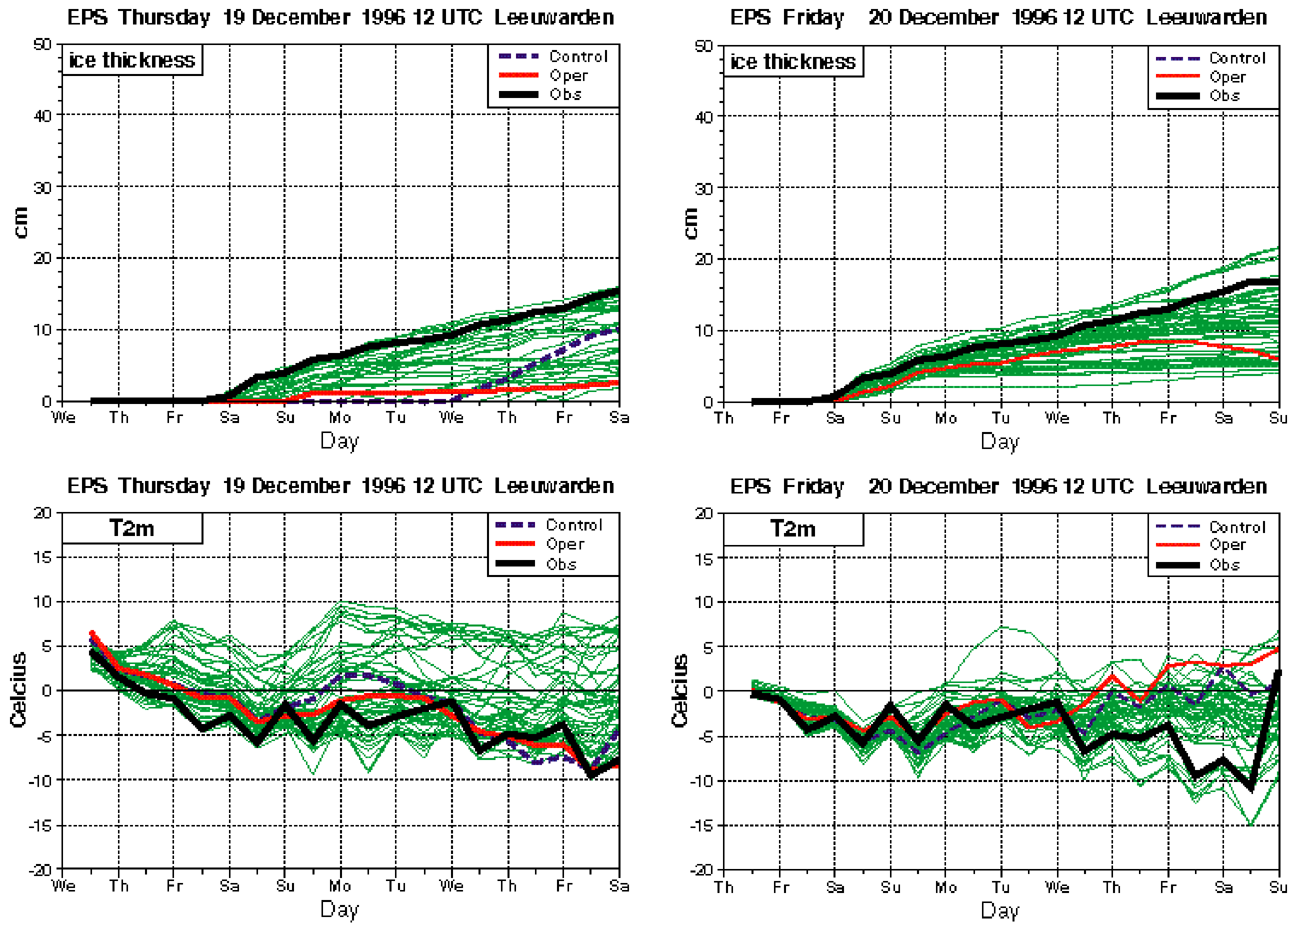

Fig.1: *The predicted ice thickness (upper panel) and 2 m temperature (lower panel) at 00 and 12 UTC, based on the 19 and 20 December 1996 ECMWF ensemble forecast for the Frisian capital Leeuwarden.*
The forecast based on the T213 forecast is in red, the T_L159 in blue and the 50 ensemble members in green; the observed temperature is indicated by the black line. The observed ice thickness is calculated by the ice model, but with the observed temperatures of the past days. The operational forecast from 20 December is forecasting considerably milder weather of the late medium range. Note that several members of the ensemble were still forecasting the extremely cold weather. The forecast from 19 December illustrates that ensemble averaging would, in this case, not have worked.
(The date in the diagram refers to the date of availability for the forecaster, which is one day after the analysis on which the ECMWF forecast was based).

representative of the Tour Commission (a non-meteorologist, but someone with an engineering background), while one of us (HvD) would be the fixed forecasting contact point at KNMI.

There were several “novelties” for the Elfstedentocht Organization this year.

- The forecast would, officially, be incorporated in the decision process
- The 10-day forecast, based on EPS, would be used to get a feel for trends to anticipate logistics decisions.
- Two types of EPS output were made available:
  - (1) Clusters for 500hPa height fields for +120 hr, +144, ... + 240 hr forecasts. These clusters are generated, at KNMI, for each forecast range separately.
  - (2) The “plume” of the 2 metre temperature. This plume is constructed from the ASCII file which ECMWF makes available to the Member States in a quasi-operational mode.
- An ice formation model was coupled to the EPS plume of the surface parameters. In this model (**de Bruin and Wessels, J** of *Appl Met*, 1988) the physical process of formation and growth of ice is simulated. Relevant parameters and processes in the model are: 2 metre temperature, wind speed and precipitation (these parameters were derived directly from EPS), radiation, exchange of latent and sensible heat between the ice and both the atmosphere and the water, and diffusion of heat through the ice layer (ice formation occurs at the bottom of the ice layer). In the daily routine run, EPS-like ice formation scenarios were calculated for a layer of undisturbed water (no water flow, no anomalous heat sources) of a typical depth of 2 metres.
- Internet : The EPS output and the maps of the operational 10-day forecast were displayed on a protected Internet Site. The meteorological representative in the Organization had a lap-top PC at his disposal, and had (through a password) access to all information.

What happened?

At the start of the winter we did not really expect to get immediately involved in making ice forecasts, but almost as soon as the new T_L159 ensemble (!) had been introduced (11 December 1996) the model indicated, in the medium-range, severe frost for the Netherlands. A little suspiciously we noticed wide bifurcations in the T2m plume of the new ensemble between 14 and 20 December, but we also saw the operational model flip flopping a few times between very cold and very warm solutions, entirely consistent with the bifurcations. See plume of 19 December 1996 (analysis 961218) and plume of 20 December 1996 (analysis 961219) [fig. 1]. Frost came, eventually, on 20 December, but, still, we all went on our holiday breaks in the firm belief that it was not going to last. It surely couldn't, could it? However, maximum temperatures would never reach values above zero for another month (with the exception of 29 December), and when we came back from our holidays after Xmas we realised, a little to our mortification, that the Tour was near.

The daily press conferences by the Tour Organization started on New Year's Eve. There was heavy, daily, contact between our Dedicated Forecaster and the Tour Organization, during which period EPS and the ice model were frequently consulted for medium-range trends and developments. The Internet contact proved to be particularly useful: the KNMI forecaster in De Bilt and the Organization in Friesland could discuss forecasts as soon as they were available, whilst looking at the same maps and graphs, in full colour display. On Thursday 2 January, finally, the announcement came, live on television (!), that the Eleven City Tour would be held on the following Saturday, 4 January.

In figure 2 the EPS forecast available at this day is shown. The bottom panel shows the 2m temperature plume in the Frisian Capital of Leeuwarden, the top panel the predicted ice thicknesses. All ensemble members are shown in green, the T213 forecast is in red, and the T_L159 Control in blue (note that the forecasts are based on the analysis of 1 January). There was a marked upward temperature trend during the weekend which could worsen ice conditions, and chances were that patches which were not frozen over yet, would not improve. The conclusion, at the time, was that there was a small probability that temperatures would reach zero, and it was therefore decided to organize the Tour on the Saturday. An alternative option for Monday 6 January was abandoned. It should also be mentioned that the model had been forecasting thaw for this weekend in several preceding forecasts. Now, we know that the operational forecast did not verify (heavy black curve).

However we feel (and more importantly, the Tour Organization feels) that the decision was correct. Verified over December and January EPS showed very good skill, in fact the 2 January forecast was an exception. In particular the surface parameters have improved a great deal over the last year. The RMS error in the T2m by the operational T213 forecast and the Control T_L159 forecast varied from degrees in the short range to about 5 degrees in the medium range, whilst the bias was less than 0.5 degrees in the first 6 days of the forecast, and about +2 degrees between day 6 and 10. The latter mainly manifested itself in the frost period. The same applies to the ice thickness forecasts. The distribution of the ensemble members with respect to the analysis, as calculated in Talagrand diagrams, showed only a small bias. A more thorough documentation of the verification is in preparation.

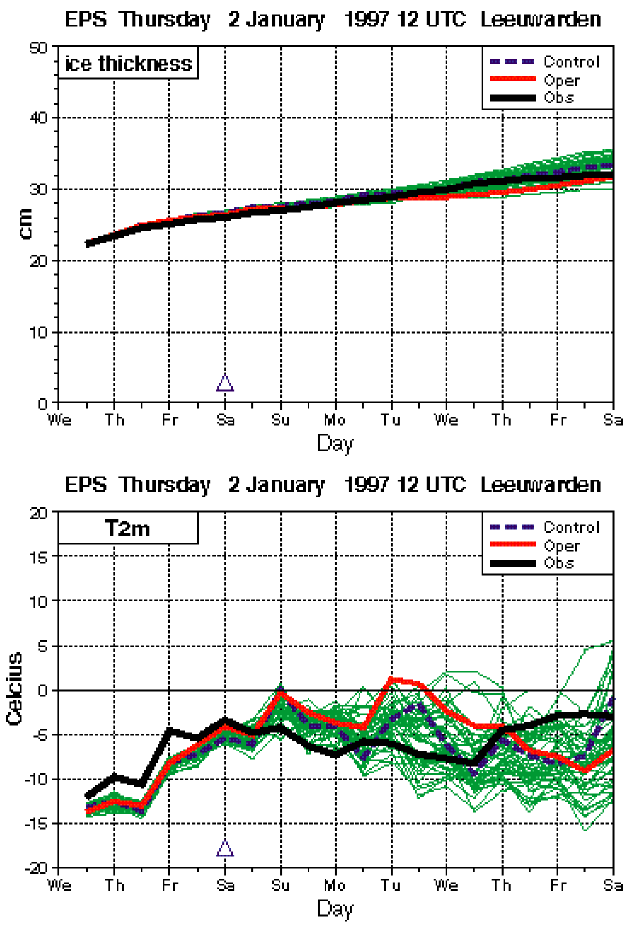

Fig. 2: *as for fig.1, based on the analysis from 1 January 1997. The T213, the Control and most ensemble members indicate temperatures going up to the melting point, which made the Tour organizers decide not to take any chances , but to go for the Saturday .*
Note that the spread in the "ice plume"is small. Ice does not grow as fast when it has formed a thick layer

The 15th Elfsteden Tour went ahead, and was a great success. The small group of racers (about 100) started at 5.30am, in darkness, and the winner completed the 200km in nearly 7 hours. The 20,000 tour skaters got going between 6am and 10.30am. The first arrived at about 2pm , but the bulk arrived somewhere between 6pm and 12pm.About 10 million people (out of the available 14 million) watched it, live, on television,which was more than for any of the matches of the Dutch team during the European Soccer Championships.

Robert Mureau (KNMI Research Dept.), Herman Wessels (KNMI Development Dept.), Henk van Dorp (KNMI Forecasting Dept.)

Operational implementation of the high resolution Ocean Wave model

Introduction

On 5 December 1996 the Centre introduced a new, global Ocean Wave model with a spatial resolution of 55 km. Compared to the previous operational wave model (with a resolution of 1.5 degrees) this may be regarded as a major step forwards, which was made possible by the introduction of the new Fujitsu VPP700 computer in September 1996. In this article, some steps are briefly described that allowed high-resolution wave modelling, while also first verification results of the month of December 1996 are presented. Finally, near-future developments are given.

High-resolution wave modelling

The Centre has been doing operational wave forecasting since June 1992. The model was developed by the WAM group and solves the so-called energy balance equation for the wave spectrum, where the relevant processes such as generation of waves by wind, dissipation by white-capping and nonlinear transfer are modelled in an explicit way (Komen et al, 1994). Initially, a modest spatial resolution of the global wave model of 3 degrees was chosen, while the two dimensional wave spectrum had 25 frequencies and 12 directions. This first implementation performed satisfactorily in many respects except that extreme events were underpredicted. This changed to some extent with the introduction of the 1.5 degree version of the wave model,although the underprediction of extreme events still occurred occasionally. The reason for this improvement when going to higher resolution is that more details and variability of the driving wind fields are included. It is well-known that wave model results, in particular during extreme events, are sensitive to the variability of the wind because the wind-wave system acts as a rectifier. Therefore, additional improvements in predicting extreme events are expected with the introduction of the present operational wave model which runs on a spatial grid which (almost) matches the grid of the atmospheric model.

Another reason for going to higher resolution is that coasts such as the east coast of the United States and the North Sea are not well represented on the coarse, 1.5 degree grid. However, in order to describe wave system near coasts properly, shallow water effects are important as well. Therefore, the shallow water option in the WAM model was switched on.

Furthermore, even on the open oceans there are benefits from going to higher resolution, since for example frontal systems have more details than may be described on a 1.5 degree grid. An example of the presence of fronts in the wave height field for the North Atlantic area (including North Sea, Baltic and Mediterranean) is shown in Fig. 1.

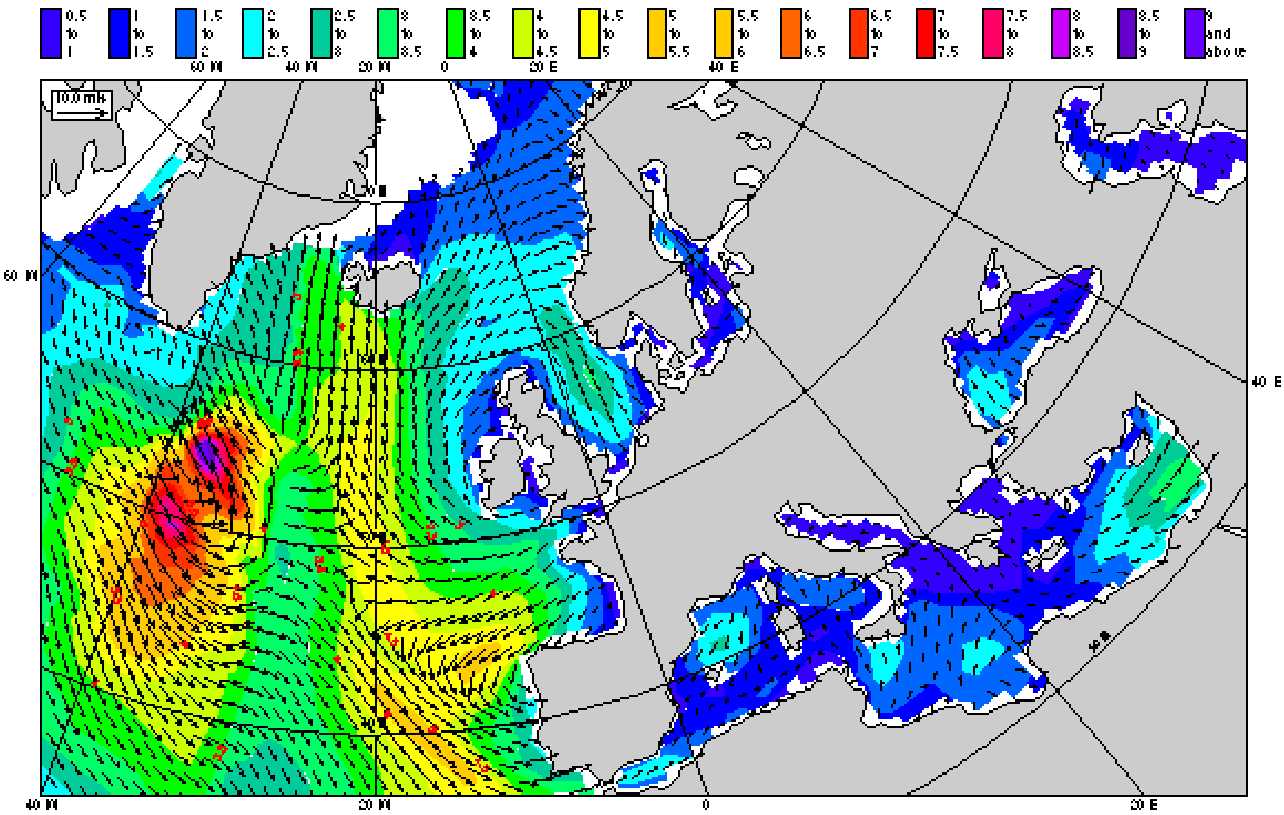

Fig. 1: *Analysed wave height (in m) and wave direction for the North Atlantic and European waters.*

Here, contours of constant wave height are plotted together with the wave direction which is presented by arrows. In addition, it is seen that the present high-resolution wave model gives a fair representation of the wave field in coastal areas and even in the Baltic and Mediterranean.

In order to be able to do high-resolution wave modelling a number of developments were necessary (for more details see **Bidlot et al** , 1997). First of all, the numerics of the WAM model is limited by the well-known CFL criterion which imposes an upper bound on the size of the integration time step of the advection scheme, and the upper bound gets smaller for decreasing spatial step. In regular spherical coordinates the increments are constant in longitude and latitude, so that near the poles the zonal spatial step may become rather small. To circumvent this problem an irregular lat-lon grid was introduced. It still has a constant latitude increment but adjusts the size of the longitude increment in such a way that the actual distance between grid points is almost constant. In this way the CFL criterion becomes less restrictive near the poles. When compared to the regular grid, additional advantages of the irregular grid option are a reduction in the number of computations per time step and a reduction in memory because the latter option requires 25% fewer grid points to cover the globe.

The resulting version of the WAM model was nevertheless demanding in CPU usage, so that a message passing version of the model was required in order to make optimal use of the available computing power on the Fujitsu. The numerical code which is submitted to the parallel computer must be modified to contain instructions that will tell each processing element (PE) what to do. Moreover since the Fujitsu is a distributed memory computer, information contained on one PE is not automatically shared with the others and must be shared explicitly if there is a need for it. The information exchange is actually done in the message passing context, whereby one PE can send a message which is received by one or more other PE's. In its very basic implementation, the message is nothing more than a one-dimensional array of a given type containing values that are needed by the other PE('s) plus all the necessary information about the sender and the receiver. Each PE now has the task to solve the energy balance equation on a subdomain of the globe. The physics source terms, such as wind input, are local and require no exchange of information with other PE's. However, advection is nonlocal, and, therefore, before doing the advection the PE waits until it has received the necessary information from the others.

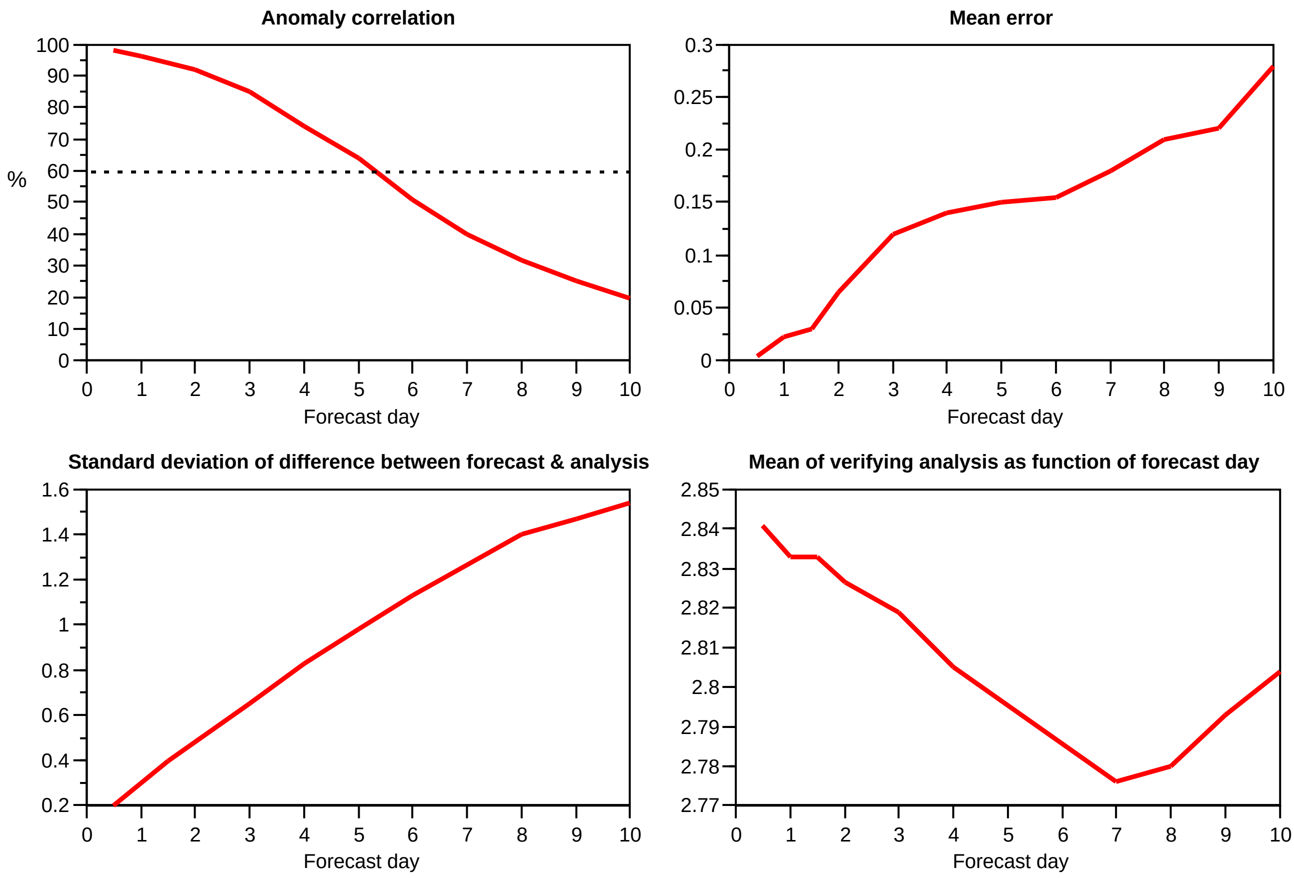

***Fig. 2:** Forecast verification against analysis for the Northern Hemisphere. Top left panel shows the anomaly correlation, the top right panel shows the mean error, the bottom left panel shows the standard deviation of the difference between forecast and analysis, while the bottom right panel shows the mean of the verifying analysis as function of forecast day.*

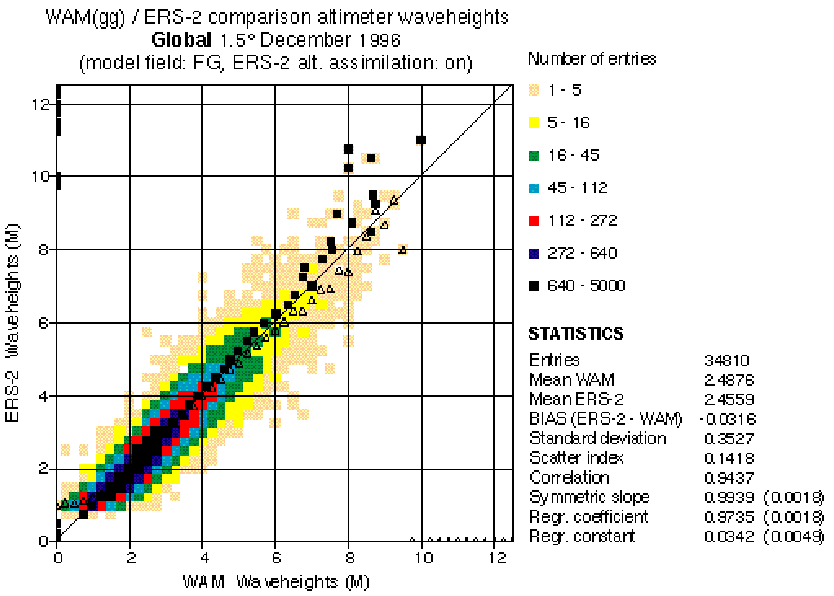

Fig. 3: *Comparison of first-guess wave height and ERS-2 altimeter wave height for the globe.*

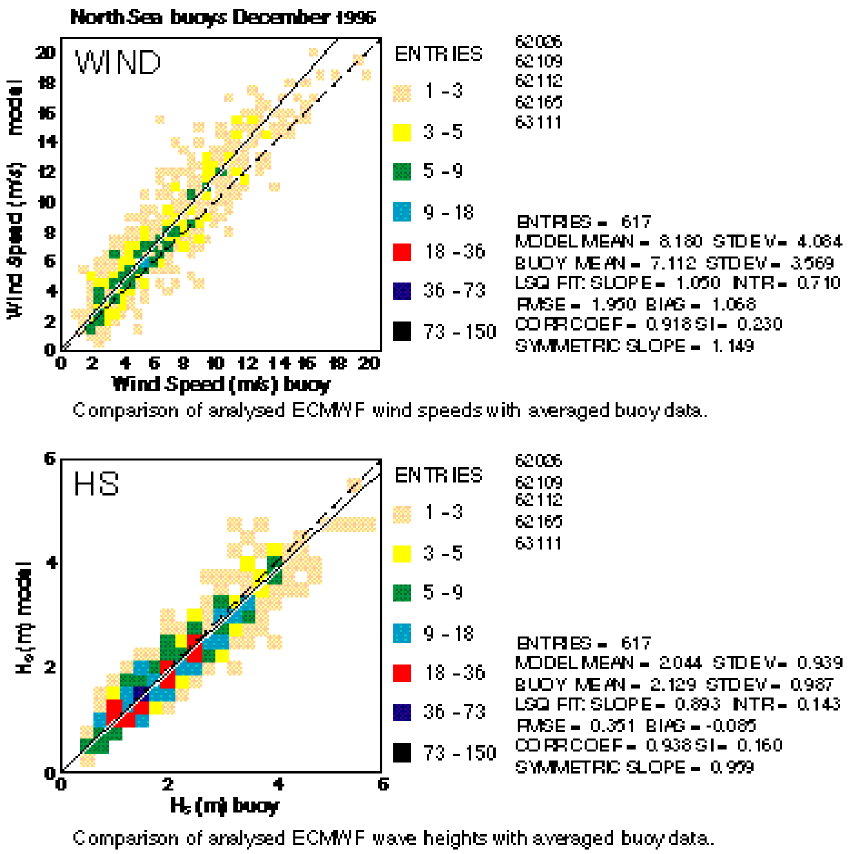

Fig. 4: *Verification of analysed wave height against buoy wave height for a number of buoys in the North Sea.*

In this way an efficient code was obtained. On 6 processors a one day analysis and 10 day forecast with the present high resolution wave model is performed in 35 minutes elapse time. On a single processor such a task would take more then 3 hours so that a speed up is achieved of about 5.5.

First verification results

During a parallel run of about one month the performance of the high resolution wave model was extensively compared with the then operational 1.5 degree model. The forecast performance in areas such as the Northern and Southern Hemisphere was found to be very similar, while in coastal areas improvements were noted.

As an illustration the first wave height verification scores for the month of December are presented. Fig. 2 shows forecast scores for the Northern Hemisphere. From the plot of anomaly correlation the forecast skill is about 5 days which is a result that is typical for this area since the last two years. Also typical is the growth of the mean error during the forecast which may amount to more then 10% of the mean wave height.

In Fig.3 the first guess wave height is compared with Altimeter data from ERS-2. The verification statistics show a good performance of the new version of the wave model, since the scatter index (=normalised standard deviation of error) is only 14% while the is almost no bias. Finally, in order to show the good performance of the high resolution model in coastal areas, the verification of analysed wave height against buoy data for the North Sea is shown in Fig.4. Although there is a slight under estimation by the model of about 4% the scatter index is only 16% confirming the above claim.

Future prospects

The introduction of the high resolution version of the WAM model seems to be successful. In particular we have noted some improvements in coastal areas. Further gains are to be expected with the introduction of the T319 version of the atmospheric model, because this will result in higher variability of the surface wind.

Since the wave model is now running on almost the same grid as the atmospheric model, we are ready to implement a coupled atmosphere, ocean-wave model in the near future. The coupling of the two models allows a better representation of the air-sea momentum transfer, which in case of early cyclogenesis will give an extra slowing down of the air flow (Janssen, 1994). Since the wind decreases somewhat the result is a lower wave height during the forecast. Therefore the expectation is that this will reduce the error growth in wave height during the forecast as seen in Fig.2.

References

Janssen, PAEM 1994, Weather and Wave Prediction over the Oceans, ECMWF Newsletter, 67,3-15

Bidlot,J , B Hansen and PAEM Janssen , 1997, Modifications to the ECMWF WAM code, ECMWF Tech Memo, to be published.

Komen, GJ, L Cavaleri, M Donelan, K Hasselmann, S Hasselmann and PAEM Janssen , 1994. Dynamics and modelling of ocean waves, Cambridge University Press.

Peter Janssen

Summary of ECMWF Technical Memorandum 219

The impact of increased resolution on predictability studies with singular vectors

R Buizza, R Gelaro*, F Molteni, T N Palmer

The dominant singular vectors of the tangent propagator of the ECMWF numerical weather prediction model are an essential component of the ECMWF Ensemble Prediction System (EPS). These singular vectors describe the principal finite-time linear instabilities of the northern hemisphere extratropical atmospheric circulation.

The impact of increasing the horizontal resolution of the tangent model from T21 to T42 on three different types of initial perturbation, which make use of these singular vectors, is considered. The increase of resolution allows the possibility of describing more accurately instabilities with an upscale cascade of energy from sub-synoptic to synoptic scales. Two of the perturbations are referred to as the pseudo-inverse and sensitivity vectors. These are both diagnostic, and involve estimating from the short-range forecast error, the component of initial error in the unstable subspace. The third type of perturbation is used to construct the set of initial states for the ensemble prediction scheme. Linear and non-linear integrations are described using these different types of perturbation.

All the results point to the conclusion that the higher resolution calculations lead to more accurate results. Specifically, it is shown that significant amounts of energy in the pseudo-inverse and sensitivity vectors reside in scales smaller than T21, and that the skill of the EPS is improved when generating the initial perturbations using T42 singular vectors. This is found to be especially true in a number of cases when the control forecast was particularly poor.

The benefit of the high resolution singular vectors was not felt uniformly over all cases studied. Specifically, the impact of the T42 singular vectors was largest for cases of low predictability in which both the sensitivity pattern and the pseudo-inverse pattern had relatively large amplitude. As a consequence, the predictability of synoptic-scale disturbances in the atmosphere on timescales of a few days is likely to be determined by errors in the initial state on sub-synoptic scales.

In addition to these conclusions, it is proved that the analysis of the relationship between the EPS and the pseudo-inverse perturbations can be used to define the initial amplitude of the EPS perturbations.

As a result of this study, together with a more extensive set of quasi-operational tests (not reported here), the real-time configuration of the ECMWF Ensemble Prediction System was modified (on 14 March 1995) to calculate T42 singular vectors. At the same time, the optimisation interval was increased to 48 hours to be consistent with routine real-time sensitivity studies. As a result, the amplitude of the ensemble perturbations was set to a value smaller than the “T42” configuration studied here (by a factor $\sqrt{2/3}$).

* Permanent address: Naval Research Laboratory, Monterey, California, 93943,USA.

Summary of ECMWF Technical Memorandum 220

Estimating the Covariance Matrices of Analysis and Forecast Error in Variational Data Assimilation

M Fisher and P Courtier

Data assimilation attempts to combine observed and predicted values of the state variables of a system in an optimal way to produce the best possible estimate of the true state. This requires knowledge of the corresponding covariance matrices of observational and short-term forecast errors. The former matrix is relatively easy to specify, since it depends primarily on the error characteristics of measuring instruments. Estimation of the covariance matrix of short-term prediction error is more difficult.

For a linear system, the Kalman filter (Kalman, 1960) produces the “Best Linear Unbiased Estimate” of the state of the system at the final time of an observing period. A natural generalization to weakly nonlinear systems is provided by the extended Kalman filter (see Ghil and Malanotte-Rizzoli, 1991).

The Kalman filter may be divided into an analysis step and a forecast step. In the analysis step, the state of the system is estimated as a combination of observed and forecast values, and the covariance matrix for the forecast errors is modified to give the covariance matrix of analysis error. In the forecast step, both the analyzed state and the covariance matrix of analysis error are propagated forward in time to give the forecast state and covariance matrix for the next analysis cycle.

In current numerical weather prediction systems, the dimensions of the covariance matrices of analysis and prediction error are of the order of $10^6 \times 10^6$. This is far too large for the matrices to be represented explicitly, or for the equations which govern their evolution to be solved exactly. A number of simplifying approximations must be made.

It is well known (e.g. Rabier et al., 1992) that, for the linear case and a perfect model, the Kalman filter and four dimensional variational data assimilation (4D-Var) are equivalent in that they produce the same estimate of the state of a system at the final time of an assimilation period. However, in 4d-Var the covariance matrices of analysis and forecast error are not explicitly known.In a cycling analysis system, it is necessary to determine an approximation for the covariance matrix of forecast error, since this provides the covariance matrix of background error for the next analysis cycle. It is also useful for diagnostic purposes to be able to estimate the covariance matrix of analysis error.

The mathematical tools provided by the introduction of adjoint techniques to data assimilation, together with the mathematical clarity of the variational formulation of the data assimilation problem, provide the opportunity for significantly improving on currently used methods for approximating the covariance matrices of analysis and short-term forecast errors. ECMWF Technical Memorandum 220 discusses how approximate solutions of the Kalman filter's analysis and forecast equations for the covariance matrix may be produced in the context of variational data assimilation. The emphasis is on methods which are sufficiently computationally efficient to be of practical use in an operational system.

Three methods for estimating the covariance matrix of analysis error in incremental variational data assimilation are presented. All three methods are based on the equivalence between the inverse of the Hessian matrix of the cost function in variational data assimilation, and the analysis error covariance matrix.

Rabier and Courtier (1992) showed that the Hessian matrix is equal to the covariance matrix of a particular set of random gradients. The randomization method uses this equivalence to produce an approximation to the Hessian which may be inverted to give an approximation to the covariance matrix of analysis error.

The Lanczos method applies a Lanczos algorithm (Lanczos 1950) to determine the leading eigenvectors of the Hessian matrix. Again, this produces an approximation to the Hessian which is easily inverted to give an approximation to the covariance matrix of analysis error. The Lanczos algorithm may be combined with conjugate gradient minimization to give an algorithm which simultaneously minimizes the cost function and determines an approximate covariance matrix of analysis error.

The BFGS method depends on the approximation to the inverse of the Hessian matrix which is generated during quasi-Newton minimization. As for the Lanczos method, the approximate analysis error covariance matrix is generated at no extra cost during the minimization.

All three methods are shown to produce reasonable estimates of the variance and covariance structure of analysis error in a one-dimensional, univariate implementation of 3D-Var. The most promising method (the Lanczos algorithm) is evaluated in the IFS 3D-Var analysis system. It is shown that the main features of the horizontal pattern of analysis error variance can be determined with relatively few vectors. The method also provides a good first guess and a preconditioner which can be used to accelerate the convergence of the minimization in a subsequent analysis. However, a significant proportion of the difference between the variances of background and analysis error remains unexplained even if more that 50 vectors are used.

Two approximations to the background error covariance matrix are demonstrated. The approximations allow diagnosis of the true standard deviations of background error for a multivariate analysis system in which the standard deviations of background error are not explicitly known.

Finally, the paper considers how the various approximations to the covariance matrix of analysis error may be propagated in time. Since the approximate analysis error covariance matrices presented in the paper all consist of low-rank corrections applied to the covariance matrix of background error, temporal propagation of the approximations may be achieved by propagating the corrections as well as an approximation to the background error covariance matrix. This can be done at the cost of a few integrations of the tangent linear model. An alternative is to approximate the forecast error covariance matrix using a few leading eigenvectors. These vectors may be calculated without approximation by solving a generalized eigenvalue problem. Unfortunately, existing algorithms for solving such problems are computationally expensive. A computationally more efficient alternative solves an ordinary eigenvalue problem, but this requires an approximation of the covariance matrix of analysis error.

References

Ghil M.and P . Malanotte-Rizzoli , 1991.Data assimilation in meteorology and oceanography. *Adv. Geophys.*, **33**, pp 141-266.

Kalman R. , 1960. A new approach to linear filtering and prediction problems. Trans. ASME, Ser. D, *J. Basic Eng.* **82**, pp 35-45.

Lanczos C , 1950. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat.Bur. Standards*, **45**, pp 255-282.

Rabier F . and P . Courtier , 1992. Four-dimensional assimilation in the presence of baroclinic instability. *Quart. J. Roy. Met. Soc.*, **118**, pp 649-672.

Rabier F., P. Courtier , J. P ailleux, O. Talagrand, J-N. Thépaut and D . Vasiljevic , 1992. Comparison of four-dimensional variational assimilation with simplified sequential assimilation. Proc. workshop on variational assimilation with special emphasis on three dimensional aspects, 9-12 November 1992.

ECMWF's computer status and plans

(This article is based on talks given by G.-R. Hoffmann at ECMWF on 7 October and 13 December 1996)

Overview

The computer configuration in 1995 was based on Cray, IBM ES/9000 and SGI/SUN systems. This year has seen a lot of change. The Cray systems have been replaced by two Fujitsu VPP series machines, three IBM RS6000 systems have been added as the first part of the new data handling system, and three HP 9000 servers have been installed as a high availability server cluster. Each of these will be discussed in more detail later.

Over the 18 years the Centre has had Cray equipment the availability of that equipment has in general improved year by year, starting from 45 hours mean time between unscheduled interrupts in 1980, to 437 hours in 1996. The Fujitsu equipment has started at a reasonable level of availability considering both the hardware and software are new. The mean time between unscheduled interrupts is currently 120 hours overall (1200 hours hardware, 125 hours software).

Looking at the forecast delays we were doing well this year up to September. The forecast suite was moved to the VPP700 system in late September, and as is inevitable with any new system, delays were then experienced as the system was tuned. Even so this year looks as though it will be the second or third best ever.

Turning now to Computer Division personnel, there have been a few changes recently. Following the sudden death of Peter Gray, Claus Hilberg has taken over Computer Operations Section temporarily and Neil Storer has thus become Acting Head of Systems Section. Two posts are currently being filled.There are a few new consultants:

Jim Almond	*working on a FECIT (Fujitsu European Centre for Information Technology) project concerning seamless High Performance Computing*
Mike Connally	*working in the new Data Handling System area supporting AIX etc.*
Ken Kagoshima	*from Japan working on systems support for the VPP machines*
Umberto Modigliani	*working on general User Support duties*
George Mozdzynski	*working on a European Union funded project looking at extensions to HPF (High Performance Fortran) which would remove the need to write IFS code using message passing explicitly.*

Computer servers

As said before, the Cray systems (C90, T3D) were generally stable towards the end of their lifetime, although the C90 continued to experience problems at roughly three months intervals. The availability at 166.7 hours per week (averaged over 12 months) was excellent for such a system, while the user utilisation at 93.1% was the best the Centre has ever achieved. The job load increased towards the end of its lifetime as some users attempted to complete their work and so avoid migrating to the Fujitsu systems, even so there were no major queue delays.

The reliability of the T3D was very good, 167.4 hours per week availability over the past 12 months, with very little downtime. Its main workload was the Ensemble Prediction System (EPS) which it ran continuously to capacity at nights and weekends, with a more general load during the working day.

Turning now to the Fujitsu systems, the timetable set up at the beginning was well met by Fujitsu Ltd. The preliminary system (a 16 processor VPP300) arrived in March and was accepted soon after. The secondary

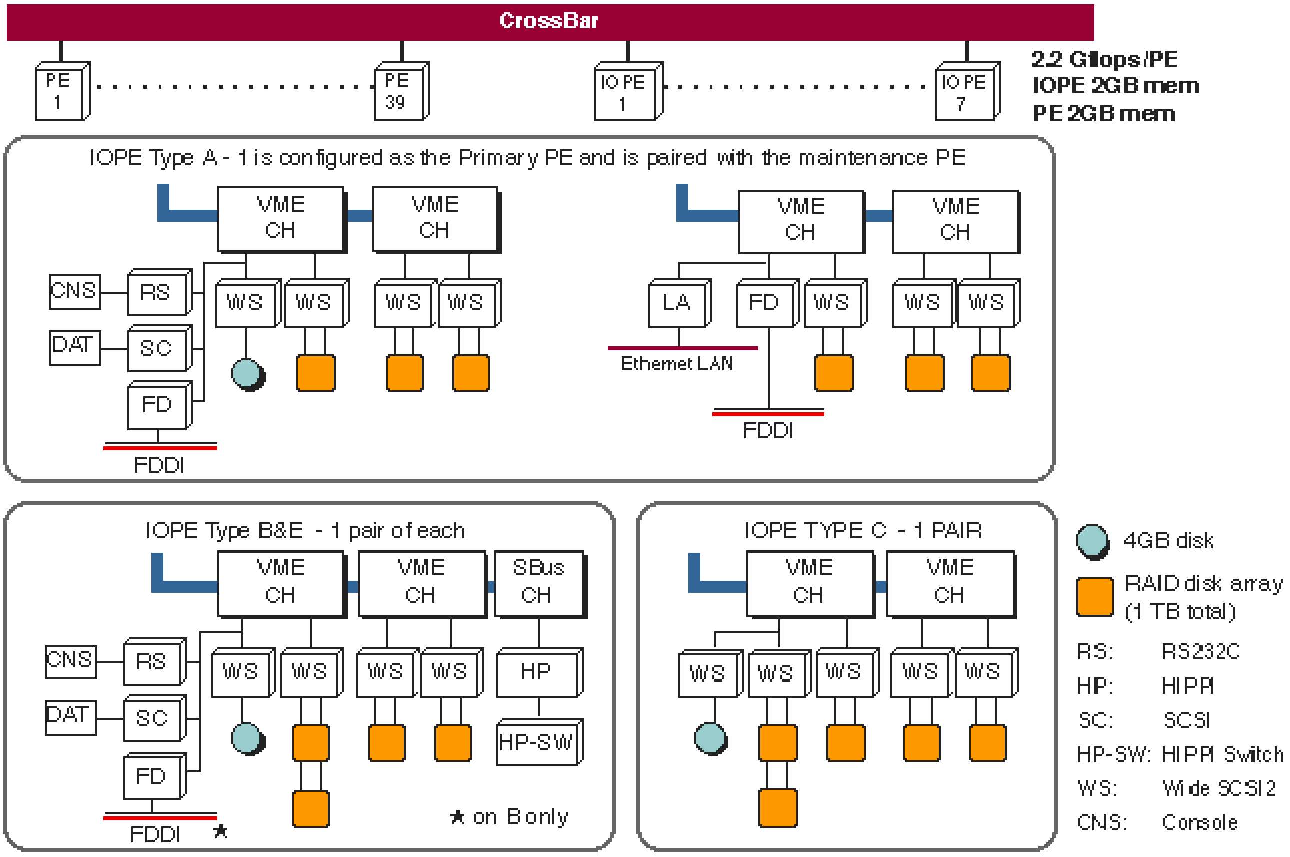

Fig. 1: *Secondary system*

system (a 46 processor VPP700 - see Fig. 1) arrived in June and passed its initial acceptance on 2 July. The configuration of these systems are:

VPP700 46 processors (PEs), 7 ns cycle time per PE, 2 Gbytes memory per PE, 1 primary PE, 39 secondary PEs, 1 Tbyte RAID disks

VPP300 16 PEs, 10 ns cycle time per PE, 2 Gbytes memory per PE, 1 primary PE, 15 secondary PEs, 200 Gbytes RAID disks.

To date, the reliability of both Fujitsu systems has been excellent, with each machine showing only 0.6 hours per week (average) lost due to hardware or software faults. The CPU utilisation depends critically on which PE you look at.The primary PE (PE0) is heavily used, typically 50% going on system time. This is an area of concern that we are currently discussing with Fujitsu, because once the system time on PE0 reaches 100%, then all other machine scheduling etc. halts. The 16 processors currently used by the forecast are averaging 50% utilisation,and the processors used by general user work are also 50% loaded.This is roughly in line with expectation.

Initial experience has shown that on compute bound problems that vectorise well these systems are extremely efficient. However, extensive I/O can degrade their performance. Fujitsu is well aware of the problem and is putting a lot of effort into this area.Already the I/O performance has been improved by relocating some file systems and tackling specific bottlenecks. An action plan has been put forward by Fujitsu that should lead to considerable improvements over the next six months. However, bearing in mind the system software is very new it is running well.

The current plan is that enhancements to the VPP700 will be done in two phases in 1998, to 153 processors in March of that year, and finally to 240 processors in September. ECMWF now possesses the most powerful system in Europe, the VPP700 being rated at just over 100 Gflops peak rate.

Within the Computer Hall we have currently moved all the tapes out of the Tape Library so that we can reinforce the floor under that area. Once that has been completed, the area where the Crays were installed will be done. This follows the floor strengthening we had to do before we could install the VPP700. The original floor design was based on wrong loading assumptions which did not permit the Centre to install the new systems without strengthening the floor.

Other equipment

The CFS based archive on the IBM ES9000 has reached 51 Tbytes held in 10 million files and is still growing strongly.

Our VAX based telecommunications system will shortly only be handling telecommunications because the data pre-processing will move to the HP system. The old ECMWF defined protocols are no longer used, and only two Member States (Portugal, Turkey) rely exclusively on DECnet protocols. All other Member States use either TCP/IP exclusively, or a mixture of TCP/IP and DECnet.

The server systems are much as before, except the HP cluster has taken over the handling of the home file system (from the CD4860 which has now gone) plus some functions from the SUN based servers. An SGI Origin server has been delivered at the end of the year and will be brought into operation early in 1997. This will allow us to redistribute some of the existing servers, providing more capacity for both Centre staff and Member State users.

The availability of the current servers is very good. For example, ecgate1 has been fully (100%) available over the past 12 months. Usage is rising quickly, sometimes we see peaks of 100% usage for short periods. Some of this load increase is due to Member State MARS jobs being migrated from the Cray system.

We have recently taken delivery of some SGI Indigo 2 workstations. They have passed acceptance and have been installed in user offices.

The HIPPI based internal network now connects the Fujitsu and IBM (old and new) systems plus one of the workstation servers. Two Gigarouters control this network. The FDDI network has expanded to three rings. However, despite (or perhaps because of) the LAN complexity it all continues to work very well. Breakdowns are rare, the users are happy with the service they get.

The wide area network now runs at 64 Kbps to all Member States, except Austria and France (128 Kbps) and the UK (2 Mbps). Data dissemination continues to increase, currently averaging 1600 Mbytes per day total (was 900 Mbytes a year ago). The incoming data is largely static at around 100 Mbytes per day.

Changing now to discuss the work on parallel systems

- the RAPS work continues, with the next workshop to be held at GMD, Germany, on 12/13 May 1997
- the PPPE project (tools environment) completed in June, resulting in the Centre gaining experience in the use of parallel tools
- the HPF+ project (as mentioned earlier) started January 1996.

Plans

The major items in the 1996 annual plan have either been completed, or are well on the way to completion.

The IBM RS6000 based new data handling system contract was signed in July 1995, the equipment for phase 1 arriving in October 1995.There are two further phases, the relevant equipment being scheduled for 1997 and 1998 respectively, to enhance the system to a level where it can cope with the full load from the final VPP700 configuration.

The phase 1 configuration consists of three servers (each with 4 PowerPC 604 processors and 1 Gbytes memory),375 Gbytes of disk, and an IBM 3494/L10 tape robot system with 16 tape drives (10 Gbytes per tape). The software is AIX plus ADSM. Because of problems passing its initial acceptance, IBM have now provided an additional 8 tape drives (on top of the 8 original drives), an extra 400 Gbytes of disk (to enable RAID 1 running instead of RAID 5 as originally proposed), plus some tape partitioning features to be delivered in December. Various stages in the provisional acceptance have been successfully passed, final acceptance should be completed before the end of 1996.

Looking ahead to 1997 the major items planned are:

- install phase 2 of the new DHS, consisting of another system and another tape robot, plus additional disks and memory for the present systems;
- a full MARS service on the new DHS will be introduced. Currently it is being developed and tested. This new MARS service should be more flexible than the present one. The old MARS should become read only around the middle of 1997. This will then allow one year for the existing MARS data to be copied to the new system;
- phase 3 of the new DHS has to be agreed,and the necessary contract amendments put before the committees and Council;
- install a 4 PE Fujitsu VPP700 test system,and reduce the VPP300 to a 9 PE machine. The VPP300 will stay on site as a gift to ECMWF from Fujitsu for seasonal forecasting and re-analysis work;
- better facilities for the Member State network, in particular the RMDCN proposal where it is proposed to merge the ECMWF WAN with that of the WMO GTS Region VI. Technically, it looks quite feasible, the political aspects may be more difficult to resolve. For Member States it would mean 128 Kbps links minimum, with possibly higher speeds for some larger ones;
- introduce optional access via the Internet to ECMWF for external users;
- it is recognised that full access to ECMWF's computer systems is hindered by the (relatively) low speed links. Hence the DAWN project which will hopefully provide a proof of concept of using the high speed (34 Mbps and higher) European research networks to give much better access from Member States to the ECMWF systems. The first links, based on a 34 Mbps backbone should be available next year, with higher speeds 155 Mbps coming a couple of years later;
- a new accounting system for our Administration Department, the current system being based on 15 year old technology.

Beyond that into the 1998-2000 time frame plans become more vague, but should include:

- two upgrades to the Fujitsu system in 1998;
- installation of phase 3 of the new DHS and de-installing the current IBM ES9000 (1 July 1998);
- an ITT for the continued supply of workstations and servers beyond the end of the current three year contract;
- upgrade of the LAN in 1999 to be ready for Gigabit per second transfer speeds required from 2000 onwards;

- possible replacement of the VAX VMS system which may well not be supported by DEC at that time. This will require the full co-operation of the Member States who, we hope, will all have moved to TCP/IP by then;
- the contract with Fujitsu finishes at the end of the year 2000, hence an ITT may be required in 1999 followed by a parallel run in 2000. In 1998, Council will need to decide what computing capacity the Centre requires in 2001 and beyond, and how the Centre could acquire it.

Summary

Overall, 1996 has seen a lot of changes, which have been successfully carried out without too much upheaval.The Cray systems have been replaced with Fujitsu systems, and all codes migrated. The first phase of the DHS project has been brought into service. The high availability server cluster has been installed and has taken over several major services.

All in all a good year.

Andrew Lea

The VPP700/46 System at ECMWF

In the last ECMWF Newsletter No. 73, Autumn 1996) there was an article on “Distributed Memory Systems at ECMWF”. The current article will describe in some detail one of these systems; namely the Fujitsu VPP700/46, the centrepiece of ECMWF’s operational computer system.

Description of the System

The VPP700/46 is a distributed memory, vector parallel supercomputer, with a theoretical peak performance of just over 100 GFLOPs. The main components of this system are:

- 46 independent 2.2 GFLOPs vector processing elements (PEs), plus a spare PE. Each PE has a memory storage unit (MSU) holding 2 GBYTEs of synchronous dynamic random-access memory (SDRAM);
- A non-blocking 2-dimensional bi-directional crossbar switch, connecting each PE to all the other 45 PEs.

Figure 1 is a schematic diagram of the crossbar network. The crossbar switch is bi-directional, which means that a PE can be sending data to another PE while simultaneously receiving data from the same, or a different, PE. The speed of the data transfer in each direction is 0.57 GB/s, making 1.1 GB/s in total. The crossbar is also non-blocking, which means that (2 times) 45 pairs of connections can be transferring data simultaneously making the maximum total bandwidth of the switch nearly 50 GB/s.

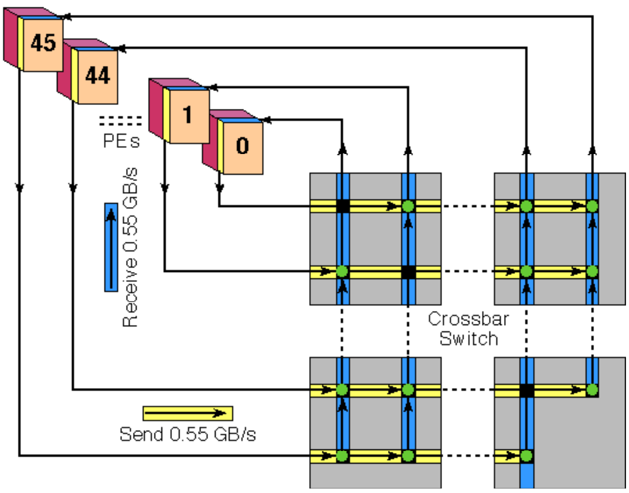

Fig. 1: *Schematic diagram of the crossbar network.*

Each of the 46 PEs comprises:

- A scalar unit (SU);
- A vector unit (VU);
- A memory storage unit (MSU);
- A data transfer unit (DTU) to connect the PE to the crossbar switch.

Seven of these PEs have I/O devices connected, e.g. RAID disks etc. Five of these also have network connections, other than the crossbar, such as HIPPI and/or FDDI. One of the seven is called the Primary-PE (P-PE) and the other 6 are called IMPEs (IPL Master-PEs).The remaining 39 PEs are called Secondary-PEs (S-PEs).

Figure 2 is a schematic diagram of a PE. Figure 3 is a schematic diagram of the VPP700/46 system as a whole with only a single line in one dimension of the 2-D crossbar network shown.

The P-PE is the “controlling” PE of the system.This is where subsystems such as the NQS daemons and Partition Manager execute.

The PEs are logically grouped together into what is known as “IPL groups”. IPL is an acronym for “Initial Program Load”.The groups are defined in a special configuration file on the P-PE. IPL group 0 comprises 5 S-PEs plus the P-PE, while the IPL groups 1 to 6 comprise 5 or 6 S-PEs plus an IMPE.

Starting the VPP700 from “cold” first requires the P-PE to be IPLd (or booted). Once the P-PE is up then commands are executed on it to concurrently IPL the IMPEs and S-PEs belonging to IPL group 0. As each IMPE comes up, it IPLs the S-PEs belonging to its IPL group. In this way the PEs are booted in parallel. Each S-PE receives the Unix kernel and tunable parameters from its IMPE. It also uses this IMPE for other “system” files when it is up and running. Certain “system” files however, such as the NQS spool directory and so on, only reside on the P-PE.

Tables

The SU, VU, MSU and DTU of each PE are detailed in the tables below:

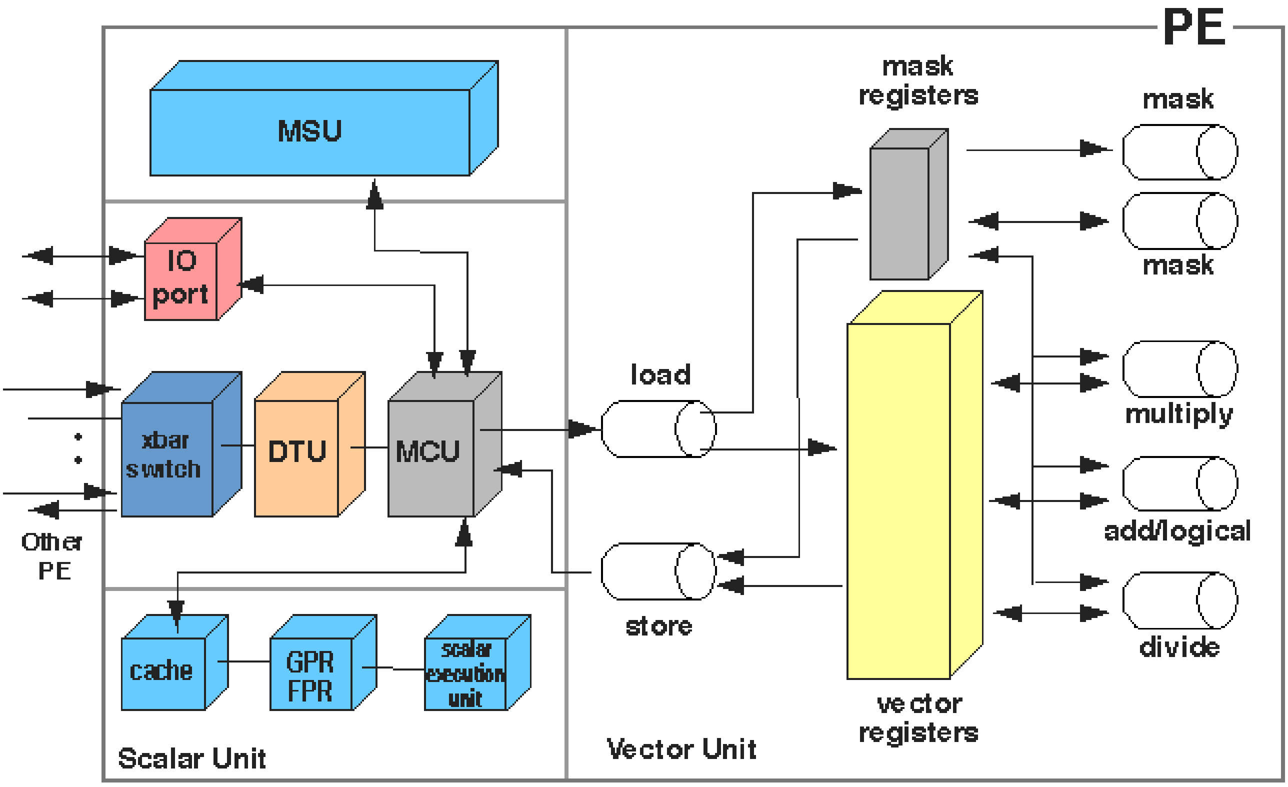

GPR: General Purpose register
FPR: Floating Point register

Fig. 2: *Schematic diagram of a Processing Element.*

Scalar Unit (SU)

Property	Value
Cycle time (τ)	7 ns
Floating-point registers	32
Load or Store rate	8 Bytes per τ
Functional units	Add (latency 6τ) Multiply (latency 6τ) Divide (latency 36τ)
Integer size	32 bits
Cache (Direct)	32 KB (Instruction) 32 KB (Data) 64 Bytes/cache-line cache-miss delay = 38τ

Vector Unit (VU)

Property	Value	
Instruction issue rate	1 per 4τ	
Startup latency	10τ	
Number of Load pipes rate (in 8 byte *words*)	1 8 per τ	2C2
Number of Store pipes rate (in 8 byte *words*)	1 8 per τ	
+ rate (flops per τ)	8	3C2
* rate (flops per τ)	8	
/ rate (flops per τ)	8/7	
Vector register size	256 x 64 (8 Byte words)	

Concurrency:

(i) 3C2 means that of the 3 functional units, only 2 may operate concurrently;

(ii) the load and store pipes can operate concurrently with any of the arithmetic pipes;

(iii) the Add and Multiply pipes can operate concurrently, each one producing 8 (64-bit) floating-point results every 7 nanoseconds clock period, hence the theoretical peak performance of the vector unit is:

$$2 \times 8 / (7 \times 10^{-9}) = 2.2 \text{ GFLOPs}$$

Memory Storage Unit & Data Transfer Unit (MSU & DTU)

Property	Value
SDRAM chip size	16 Mbits
Capacity	2 GB
Interleave (banks)	512
DTU bandwidth (GB/s)	2 x 0.57 (bi-directional)
Parallel jobs per PE	2
DTU <=> MSU (GB/s)	2 x 1.1

I/O Considerations

Each of the PEs runs a complete version of Unix SVR4 called UXP/V. There is no microkernel, or “single-system” image, as there is on some other MPP systems. This has several consequences, several of which have do to with I/O.

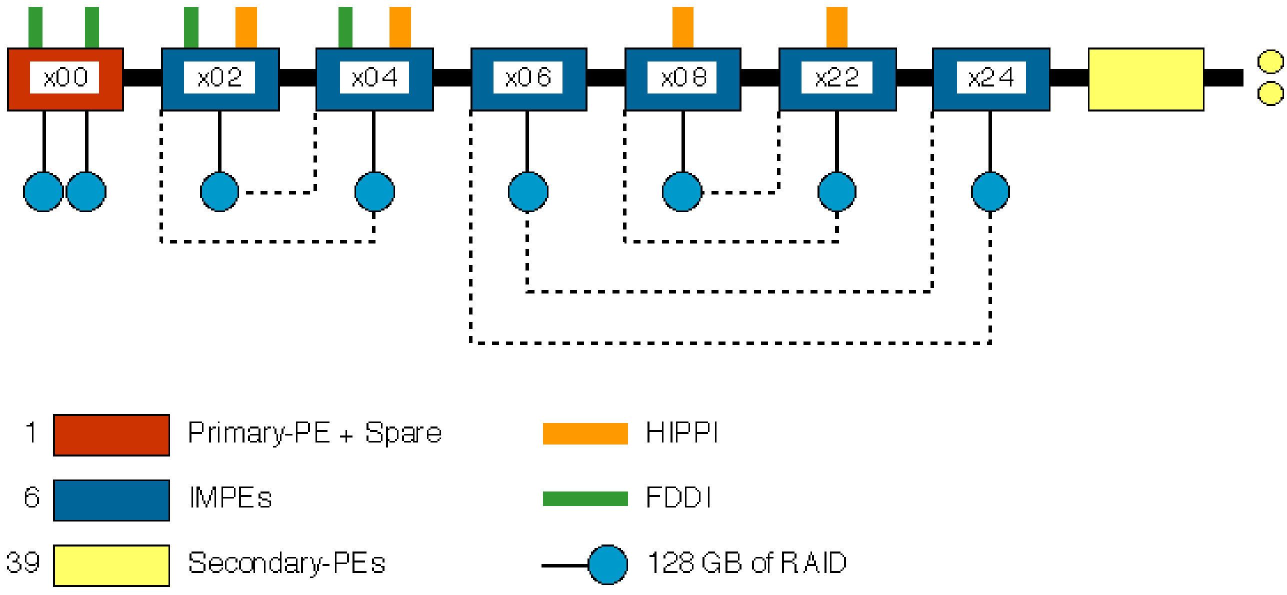

Fig. 3: *Schematic diagram of the VPP700/46 system as a whole with only a single line in one dimension of the 2-D crossbar network shown.*

The P-PE and IMPEs have RAID disks attached, and are collectively known as IOPEs, S-PEs have no attached disks. Filesystems can span multiple disks, but they cannot span IOPEs. PEs "share" files using NFS, which operates over the crossbar network. There is a special, high performance, low overhead version, called PXNFS which is used automatically for large I/O transfers.

The directory structure on each PE is set up in such a way as to hide the fact that, in general, the filesystems do not exist on the PE on which a job is running. For example, if there is a filesystem "/work" on an IOPE, then it can be accessed from any PE, including that IOPE, by the name "/vpp700/work". The user does not need to know on which IOPE the filesystem actually resides. Of course, if many jobs accessed that filesystem concurrently, or if even only 1 job accessed the filesystem but in a very inefficient way, then that IOPE could get overloaded handling all the (PX)NFS calls. If the IOPE in question is the P-PE,which is used by all the jobs in the machine, it can have dire consequences for the whole system.

The main key to using the VPP700 system efficiently is to reduce the amount of I/O done via the crossbar network. One way of achieving this is to do I/O to a memory-resident filesystem on the PE on which the program is running. If the actual amount of I/O cannot be reduced, then the next best thing is to perform the I/O efficiently (using a small number of large requests) and use files which have been pre-allocated on a VFL (Very Fast and Large) filesystem, rather than use a regular UNIX style one. This can easily be done for Fortran programs, simply by increasing the size of the run-time library buffers via shell environment variables. Jobs which are totally I/O-bound can, if necessary, be run on the IOPE where the data resides, though in many cases, programs doing efficient I/O to VFL filesystems using PXNFS over the crossbar actually cause less overhead than the same program running on the IOPE.

Usage of the System

The VPP700/46 was delivered in June 1996, Provisional Acceptance starting on 1 July.

Initially usage of the system was light, while applications and jobs were migrated from the Cray C90 to the VPP700. Later on, when the workload began to build up on the machine we found that for long periods of time the P-PE would be saturated, spending almost 100% of its time executing operating system code. This affected all of the PEs on the VPP700, not just the P-PE. By August, when the workload increased, due to large amounts of testing of the operational suite on the VPP700, it became clear that something had to be done to improve the I/O performance of the machine, in particular to reduce the load on the P-PE.

In September, Fujitsu supplied some modifications to NFS to reduce the P-PE overhead. In October, analysts from Fujitsu and ECMWF came up with some suggested changes to the placement of files and filesystems, which also reduced the P-PE load. The analysts from Met. Applications and the staff from the Research Division made some application code changes to improve I/O which dramatically cut down the overheads. With the termination of the Cray service at the end of September, the Computer Operations Section could concentrate much more on the VPP700 service, and came up with improved ways of scheduling the workload on the machine. Between November and January, Fujitsu supplied several modifications to UXP/V, which cut down on the overheads even more.

The result of all these changes is that the P-PE overhead has been reduced to an average of about 50%. Occasionally it peaks at 80%, but it rarely reaches the 100% mark. The overhead on the IMPEs has also been reduced by these changes. Much more work is now being processed on the VPP700 than could ever have been achieved before these modifications. The peak daily usage of the machine has improved from about 45-50% which occurred during August-September to over 70% today. That is equivalent to 3.5 times the Cray C90 and, at this early stage in the life of the system, is an excellent achievement by all concerned.

It should be noted that the operating system, UXP/V, is a new operating system,ECMWF being one of its first users. Its stability at this stage in its life is extraordinarily good, with a current meantime between interrupts of around 125 hours. The modifications described above are the first steps of a longer term plan that Fujitsu intend to implement in order to improve the performance of the operating system. Modifications will be made on a regular basis throughout 1997, which will allow us to make even better use of the VPP resources over the coming year.

Concluding Remarks

The first half of this article concentrated on the physical hardware aspects of the VPP700/46.The second half described how I/O considerations were the key to using the system efficiently, and how changes to the software and filesystem configurations have enabled us to get much more out of the system than was possible earlier on. Other software aspects of the VPP700 system will be elaborated upon in an article in a later Newsletter.

Neil Storer

Changes to ECMWF's computer user documentation

Following the termination of the Cray service and the introduction of the Fujitsu service, there have been some changes to the documentation supplied to ECMWF's computer users.

Manuals

All Cray manuals have been replaced with the following set of Fujitsu manuals:

UXP/V User's Guide	*Second edition, June 96*
UXP/V ANSI C & Programming Support Tools	*J2U5-0160-01EN*
UXP/V Network Queuing System Handbook	*Second edition, June 96*
UXP/V C Language User's guide	*J2U5-0110-01EN*
UXP/V C/VP User's Guide	*J2U5-0120-01EN*
UXP/V Analyzer User's guide	*J2U5-0130-01EN*
UXP/V PVM User's Guide	*J2U5-0140-01EN*
UXP/V Programmer's Guide: System Services & Application Packaging Tools	*Second edition, June 96*
SSLII User's Guide	**99SP4020E1*
SSLII Extended Capabilities User's Guide 2	*J2X0-1360-01EN*
SSLII Extended Capabilities User's Guide	**99SP4070E2*
SSLII/VPP User's Guide	*J2X0-1370-01EN*
UXP/V Fortran Syntax Translator Handbook	*94AR0010E1*
UXP/V C++ Language System Handbook	*J2X0-0670J-01EN*
UXP/V Fortran90/VP User's Guide	*J2U5-0050-01EN*
UXP/V Fortran90 Messages	*J2U5-0060-01EN*
FJ VP Programming Handbook	**99SP0080E1*
UXP/VFortran90/VPP User's Guide	*J2U5-0080-01EN*
VPP Workbench	*J2S1-0660-01EN*
VPP MPI User's Guide	
VPP PARMACS User's Guide	
MPI:A Message-Passing Interface Standard	
PARMACS V6.0 Specification	
MPTOOLS-Message Passing Library Toolset	

Computer Bulletins

As stated in previous Newsletters (e.g. No. 64) the set of printed Computer Bulletins is being reduced, in favour of on-line documentation. Again, largely following the termination of the Cray service, the following ECMWF Computer Bulletins have been withdrawn:

B1.0/3, B2.2/1, B2.8/1, B2.8/2, B3.3/1 to B3.3/5, B5.2/14, B8.5/1.

The current list of valid Computer Bulletins is therefore:

0.1/1	ECMWF Computer Division management and personnel list
0.2./3	Security
1.0/2	An introduction to the VAX interactive system
1.0/5	User authentication - Passwords and SecurID cards
1.2/1	Accounting *(to be updated)*
1.5/1	Advisory and visitor services
1.7/1	Migration from Cray
3.4/2	Integrated electronic mail services
3.4/3	INTERNET
5.2/5	The ECMWF meteogram system
5.2/8	Reference manual for MAGICS
5.2/9	User's guide for MAGICS
5.2/10	Pocket guide for MAGICS
6.0/1	Software libraries available at ECMWF
6.7/2	MARS user guide
8.2/1	Supporting incoming/outgoing magnetic tapes at ECMWF
8.3/1	ECFILE concepts

Andrew Lea

Range of ECMWF's publications

ECMWF has an obligation to keep Member States informed of its activities. Also, as part of the wider international community and with a global responsibility in the field of numerical weather prediction,it is important that non-Member States are kept informed of the Centre's operational and research activities.

One method the Centre does all this is through its publications. These fall into four broad categories:

- Regular publications
- Scientific and technical publications
- Handbooks, catalogues and guides
- On-line information.

Regular publications:

ECMWF Newsletter - *published quarterly*

The Newsletter contains articles of a general meteorological and computing nature. It is intended for all meteorologists and computer users who make use of ECMWF products and facilities, or who are interested in ECMWF's activities. Included in the Newsletter is information about:

- research and operational activities
- future plans
- educational and publicity material.

Reference Charts *- published monthly*

Charts showing mean sea level pressure and 850 hPa temperature, and charts of 500 hPa height and temperature, are printed for each operational forecast during a particular month.There are 10 charts per page with the following formats:

- analysis and previous forecasts from up to 10 days all verifying at the time of the analysis (verification charts);
- analysis and forecast out to day 10 starting from that analysis (forecast charts).

These reference charts are published for operational and research staff.

Scientific and technical publications:

Seminar and workshop proceedings

Every autumn, a seminar is arranged at the Centre to review the current state of some aspect of meteorology which is relevant to the work of the Centre and the National Meteorological Services of the Member States.

Each year, several workshops are held to discuss specialised topics which are associated with the current research and development activities at the Centre. The published proceedings contain the papers presented at the workshop as well as a report of the discussions and recommendations of the working groups.

The seminar and workshop proceedings are mainly of interest to people actively engaged in research and development related to the topics chosen.

Technical Reports

These are reports of scientific work carried out at the Centre in the fields of meteorology, numerical analysis and computer applications. They are reviewed internally and contain material which is intended for eventual publication in scientific journals. The Technical Reports have a wide external distribution to universities and meteorological services all over the world.

Technical Memoranda

This series contains unrefereed reports of scientific and operational matters. Distribution is limited to the Member States.

Research Reports

This series describes the work carried out in the joint ECMWF/EUMETSAT Fellowship programme. They are distributed to Member States only.

Lecture Notes

The Lecture Notes for the ECMWF Training Course fall into three subject areas:

- dynamical meteorology and numerical methods
- analysis, initialisation and the adiabatic formulation of models
- physical parametrization and the inclusion of orography in numerical models.

They are intended for Training Course participants.

Handbooks, catalogues and guides:

Meteorological Bulletins

These are a series of individual bulletins covering the design and use of the ECMWF forecast suite and its products. They often contain reference material not found elsewhere and include detailed manuals describing the scientific basis of the Centre's analysis and forecast system. Three bulletins should be mentioned specifically:

- the User Guide to ECMWF Products gives a non-mathematical description of the operational forecast systems, an outline of the product range, the characteristics of the products and guidance on their use;
- the Dissemination of ECMWF Products to Member States gives the technical details for requesting real-time dissemination products;
- the MARS User Guide is a detailed reference manual of the Centre's data archive.

Meteorological Bulletins are intended for meteorologists using the ECMWF forecast suite and its products.

Computer Bulletins

The Computer Bulletins detail the formal rules and working conventions needed to use the ECMWF computer system,with reference material not covered in the manufacturers' manuals. They are intended for all computer users and are distributed to Member States and the Centre's staff. It is planned to eventually replace them with on-line documents (see below).

Computer News Sheets

These contain computer information of immediate but short term interest, and are available to all computer users.

Graphics documentation

The Centre's principal graphics visualisation package is called METVIEW. Documentation for it consists of a multi-part manual (four sections of which have been published to date) plus on-line help.

Catalogue of ECMWF Products, and Guide to the Catalogue of ECMWF Products

The Catalogue outlines the real-time meteorological and oceanographic products available from the ECMWF forecasting system, the conditions applied to their supply, the tariffs and licence arrangements.

It is intended for use by national meteorological services, service providers and end users.

The Catalogue is available in hard copy and via the Internet (see http://www.ecmwf.int/).

The Guide outlines administrative procedures, principles for calculating tariffs and practical determination of product price. It includes the "Rules governing the distribution and dissemination of real-time products".

On-line information

ECMWF provides an Internet server. Open to all, it is accessible via http://www.ecmwf.int.As well as giving an overview of the Centre it provides lists of each of the main publications available. In addition, a limited set of forecast charts are published daily via this server.

As stated before, it is the intention to replace much of the paper-based computer documentation for registered users with on-line versions.

Distribution

In order to make the publication widely accessible, the Centre prefers to have libraries on its regular mailing list,rather than individuals. However, in exceptional circumstances, individuals can be added to the mailing list for a limited number of publications. Normally, a request for a particular item is met without further formality.

All requests for publications should be made in writing to the ECMWF Library.

Andrew Lea

Member State computer resource allocations 1997

Member State	Fujitsu (kunits)	Data (Gbytes)	Member State	Fujitsu (kunits)	Data (Gbytes)
Belgium	31.0	61	Norway	24.1	48
Denmark	25.3	50	Austria	28.8	57
Germany	151.2	298	Portugal	22.1	43
Spain	50.6	100	Switzerland	33.3	66
France	104.2	205	Finland	22.3	44
Greece	22.4	44	Sweden	29.8	59
Ireland	19.0	37	Turkey	26.8	53
Italy	88.2	174	United Kingdom	84.3	166
Yugoslavia*	20.3	40	Special projects	91.3	180
Netherlands	38.0	75	**Total**	**913.0**	**1800**

* In accordance with UN Security Council Resolution 757 (1992) of 30 May 1992, the Council instructed the Director to suspend the telecommunications connection to Belgrade with immediate effect. This took place on 5 June 1992. As a consequence no operational products are disseminated to Belgrade and access to the Centre's computer system is not available to Belgrade.

ECMWF Calendar 1997

Meteorological Training Course (7 April-18 June)

7 - 16 Apr	**Met1**	Numerical methods
17 - 25 Apr	**Met2**	Data assimilation
28 Apr - 2 May	**Met3**	Predictability
6 - 16 May	**Met4**	Parametrization
27 May - 5 Jun and 9 - 18 Jun *(two sessions)*	**Met5**	Use and interpretation of ECMWF products

5 May	ECMWF Holiday	
12 - 13 May	Security Representatives' meeting	*5th*
23 - 26 May	ECMWF Holiday	
3 - 4 Jun	Technical Advisory Committee	*24th*
2 - 3 Jul	Council	*46th*
25 Aug	ECMWF Holiday	
8 - 12 Sep	Seminar:Atmosphere-surface interaction	
29 Sep - 1 Oct	Scientific Advisory Committee	*26th*
6 - 7 Oct	Computing Representatives' meeting	
8 - 10 Oct	Technical Advisory Committee	*25th*
13 Oct	Policy Advisory Committee	*8th*
15 - 16 Oct	Finance Committee	*58th*
20 - 22 Oct	Workshop:Predictability	
10 - 12 Nov	Workshop:Orography	
17 - 21 Nov	Workshop on Meteorological Operational Systems	*6th*
2 - 3 Dec	Council	*47th*
24 - 26 Dec	ECMWF Holiday	

Special Project allocations 1997

Member State		Institution	Project title	Allocations 1997 Fujitsu units	Allocations 1997 Data storage Gbytes
Continuation Projects					
Denmark	**1**	DMI (Christensen)	Regionalization of climate simulations	4,000	4.2
France	**2**	CNET/CETP (Eymard)	Determination of ocean surface heat fluxes using satellite data and the ECMWF model	37	1.3
	3	L.A.M.P.,Aubière (Cautenet)	Chemistry, cloud & radiation interactions in a meteorological model	93	2
Germany	**4**	Institute for Geophysics and Meteorology (Speth)	Interpretation and calculation of energy budgets	93	4.2
	5	MPI, Hamburg (Roeckner)	Modelling the earth's radiation budget and evaluation against ERBE data	3,733	7
	6	MPI, Hamburg (Bengtsson)	Numerical experimentation with a coupled ocean/atmosphere model	8,000	7
	7	MPI, Hamburg (Bengtsson)	Simulation and validation of the hydrological cycle	8,000	7
	8	University of Munich (Wirth/Egger)	The behaviour of cutoff cyclones in ECMWF analysis: impact of diabatic processes on their development and decay	70	1
	9	FU, Berlin (Fischer/Thoss)	Comparison of the ECMWF cloud scheme with multi-spectral satellite data in the Baltic Sea	70	5
	10	GKSS, Geesthacht (Rockel)	Energy and water cycle components in regional forecasts, remote sensing and field experiments	47	0.2
	11	DLR, Oberpfaffenhofen/ UnivMunich (Hoinka/Egger)	The climatology of the global tropopause	14	4.2
Italy	**12**	Istituto per lo Studio della Dinamica delle Grandi Masse, Venezia (Cavaleri)	Testing and applications of a third generation wave model in the Mediterranean Sea	140	1.5
	13	Univ Bologna (Rizzi)	TOVS lB radiances and model simulations	467	7
Netherlands	**14**	KNMI, De Bilt (Siegmund)	Stratosphere-troposphere exchange	187	4
	15	KNMI (van Velthoven)	Chemistry and transport studies with a 3D off-line tracer model	1,400	4.2
	16	KNMI (Komen)	North Sea wave climate	933	5
	17	KNMI (Drijfhout)	Agulhas	467	0
	18	KNMI (Komen)	Validation of re-analysed A/S fluxes	933	5
	19	KNMI (v Meijgaard)	Simulation of the Antartic climate model with a regional atmospheric model	4,000	4.2
	20	KNMI (Siebesma)	Large Eddy Simulation (LES) of boundary layer clouds	4,667	7
	21	KNMI (Opsteegh)	Short term,regional probabilistic forecasting using IFS	2,800	7
Norway	**22**	Geophysical Institute, Univ. of Bergen (Grønås/Kvamstø)	Cloud parametrization in general circulation models	93	0.2
	23	Univ Oslo (Isaksen)	Ozone as a climate gas	467	4.2
	24	Univ Oslo (Iversen)	Climatic impacts of anthropogenic aerosols	560	4.2
	25	DNMI (Røed)	The role of mesoscale features in the heat transport in the Nordic seas	4,000	4.2
Sweden	**26**	SMHI, Norrköping (Källén)	The HlRLAM 3 project	3,500	5
UK	**27**	Univ Reading (Hoskins)	Routine back trajectories	1,167	4.2
	28	Univ Oxford (Sutton)	Tropical seasonal cycle in the coupled atmosphere-ocean system	4,000	4.2
New projects					
Austria	**29**	Univ Vienna (Kromp-Kolb)	Vertical ozone transports in the Alps	467	4.2
	30	Univ Vienna (Ehrendorfer)	Covariance evolution and forecast skill	397	4
	31	Univ Vienna (Hantel)	Estimating the global mean sub-gridscale energy conversion term	163	2

Special Project allocations 1997 (continued)

Member State		Institution	Project title	Allocations 1997 Fujitsu units	Data storage Gbytes
France	**32**	CERFACS (Thual)	Universal software for data assimilation: variational method for global ocean	1,120	7
Italy	**33**	CINECA,Bologna (Molteni)	Influence of springtime land-surface conditions on the Asian summer monsoon	8,000	7
Netherlands	**34**	KNMI (Kelder)	The relation between satellite ozone observation errors & dynamical errors in the ECMWF model	1,400	7
	35	KNMI (v d Hurk)	EOS-execution phase	700	7
	36	KNMI (Burgers)	OGCM mixed-layer modules	350	0
Turkey	**37**	Wea Forec Dept (Yildirim)	Atmospheric long term range dispersion model evaluation	1,867	1
	38	Wea Forec Dept (Yildirim)	Numerical study of small-case disturbances formed in Mediterranean Sea and their impact on Turkish provinces	2,333	1
UK	**39**	CGAM,Univ Reading (Slingo)	Predictability experiments for the Asian summer monsoon	8,000	7
	40	Br Antartic Survey, Cambridge (Turner and Leonard)	Assessment of ECMWF forecasts over the high latitude areas of the Southern Hemisphere	47	0
			Total	**78,782**	**161.4**
			Reserve (to be allocated by ECMWF)	**12,518**	**18.6**
			Overall total	**91,300**	**180**

ECMWF Publications

Technical Memoranda

No.229 Janssen, PAEM, B Hansen and J Bidlot: Verification of the ECMWF Wave Forecasting System against buoy and altimeter data, July 1996

No.230 Lea,A (Compiler): Report on the ninth meeting of Member State Computing Representatives, 7-8 October 1996, October 1996

No.231 Janssen, P A E M, H Wallbrink,C J Calkoen, D van Halsema, W A Oost and P Snoeij: The VIERS-1 scatterometer model, December 1996

No.232 Bidlot, J., B. Hansen, P. Janssen: Modifications to the ECMWF WAM code, January 1997 (in preparation)

Workshop Proceedings

Workshop on Semi-Lagrangian methods, 6-8 November 1995

ECMWF/EUMETSAT Fellowship Research Reports

No.2 Ottenbacher, A., M. Tomassini, K. Holmlund and J. Schmetz: Low-level cloud motion winds from METEOSAT high-resolution visible imagery, November 1996

No.3 Vesperini, M.: Humidity in the ECMWF model: monitoring of operational analyses and forecasts using SSM/I observations, November 1996

No.4 Fernandez, P., G. Kelly and R. Saunders: Use of SSM/I ice concentration data to improve the ECMWF SST analysis, December 1996

ECMWF Forecast and Verification Charts to 31 December 1996