ECMWF Newsletter

Number 69 - Spring 1995

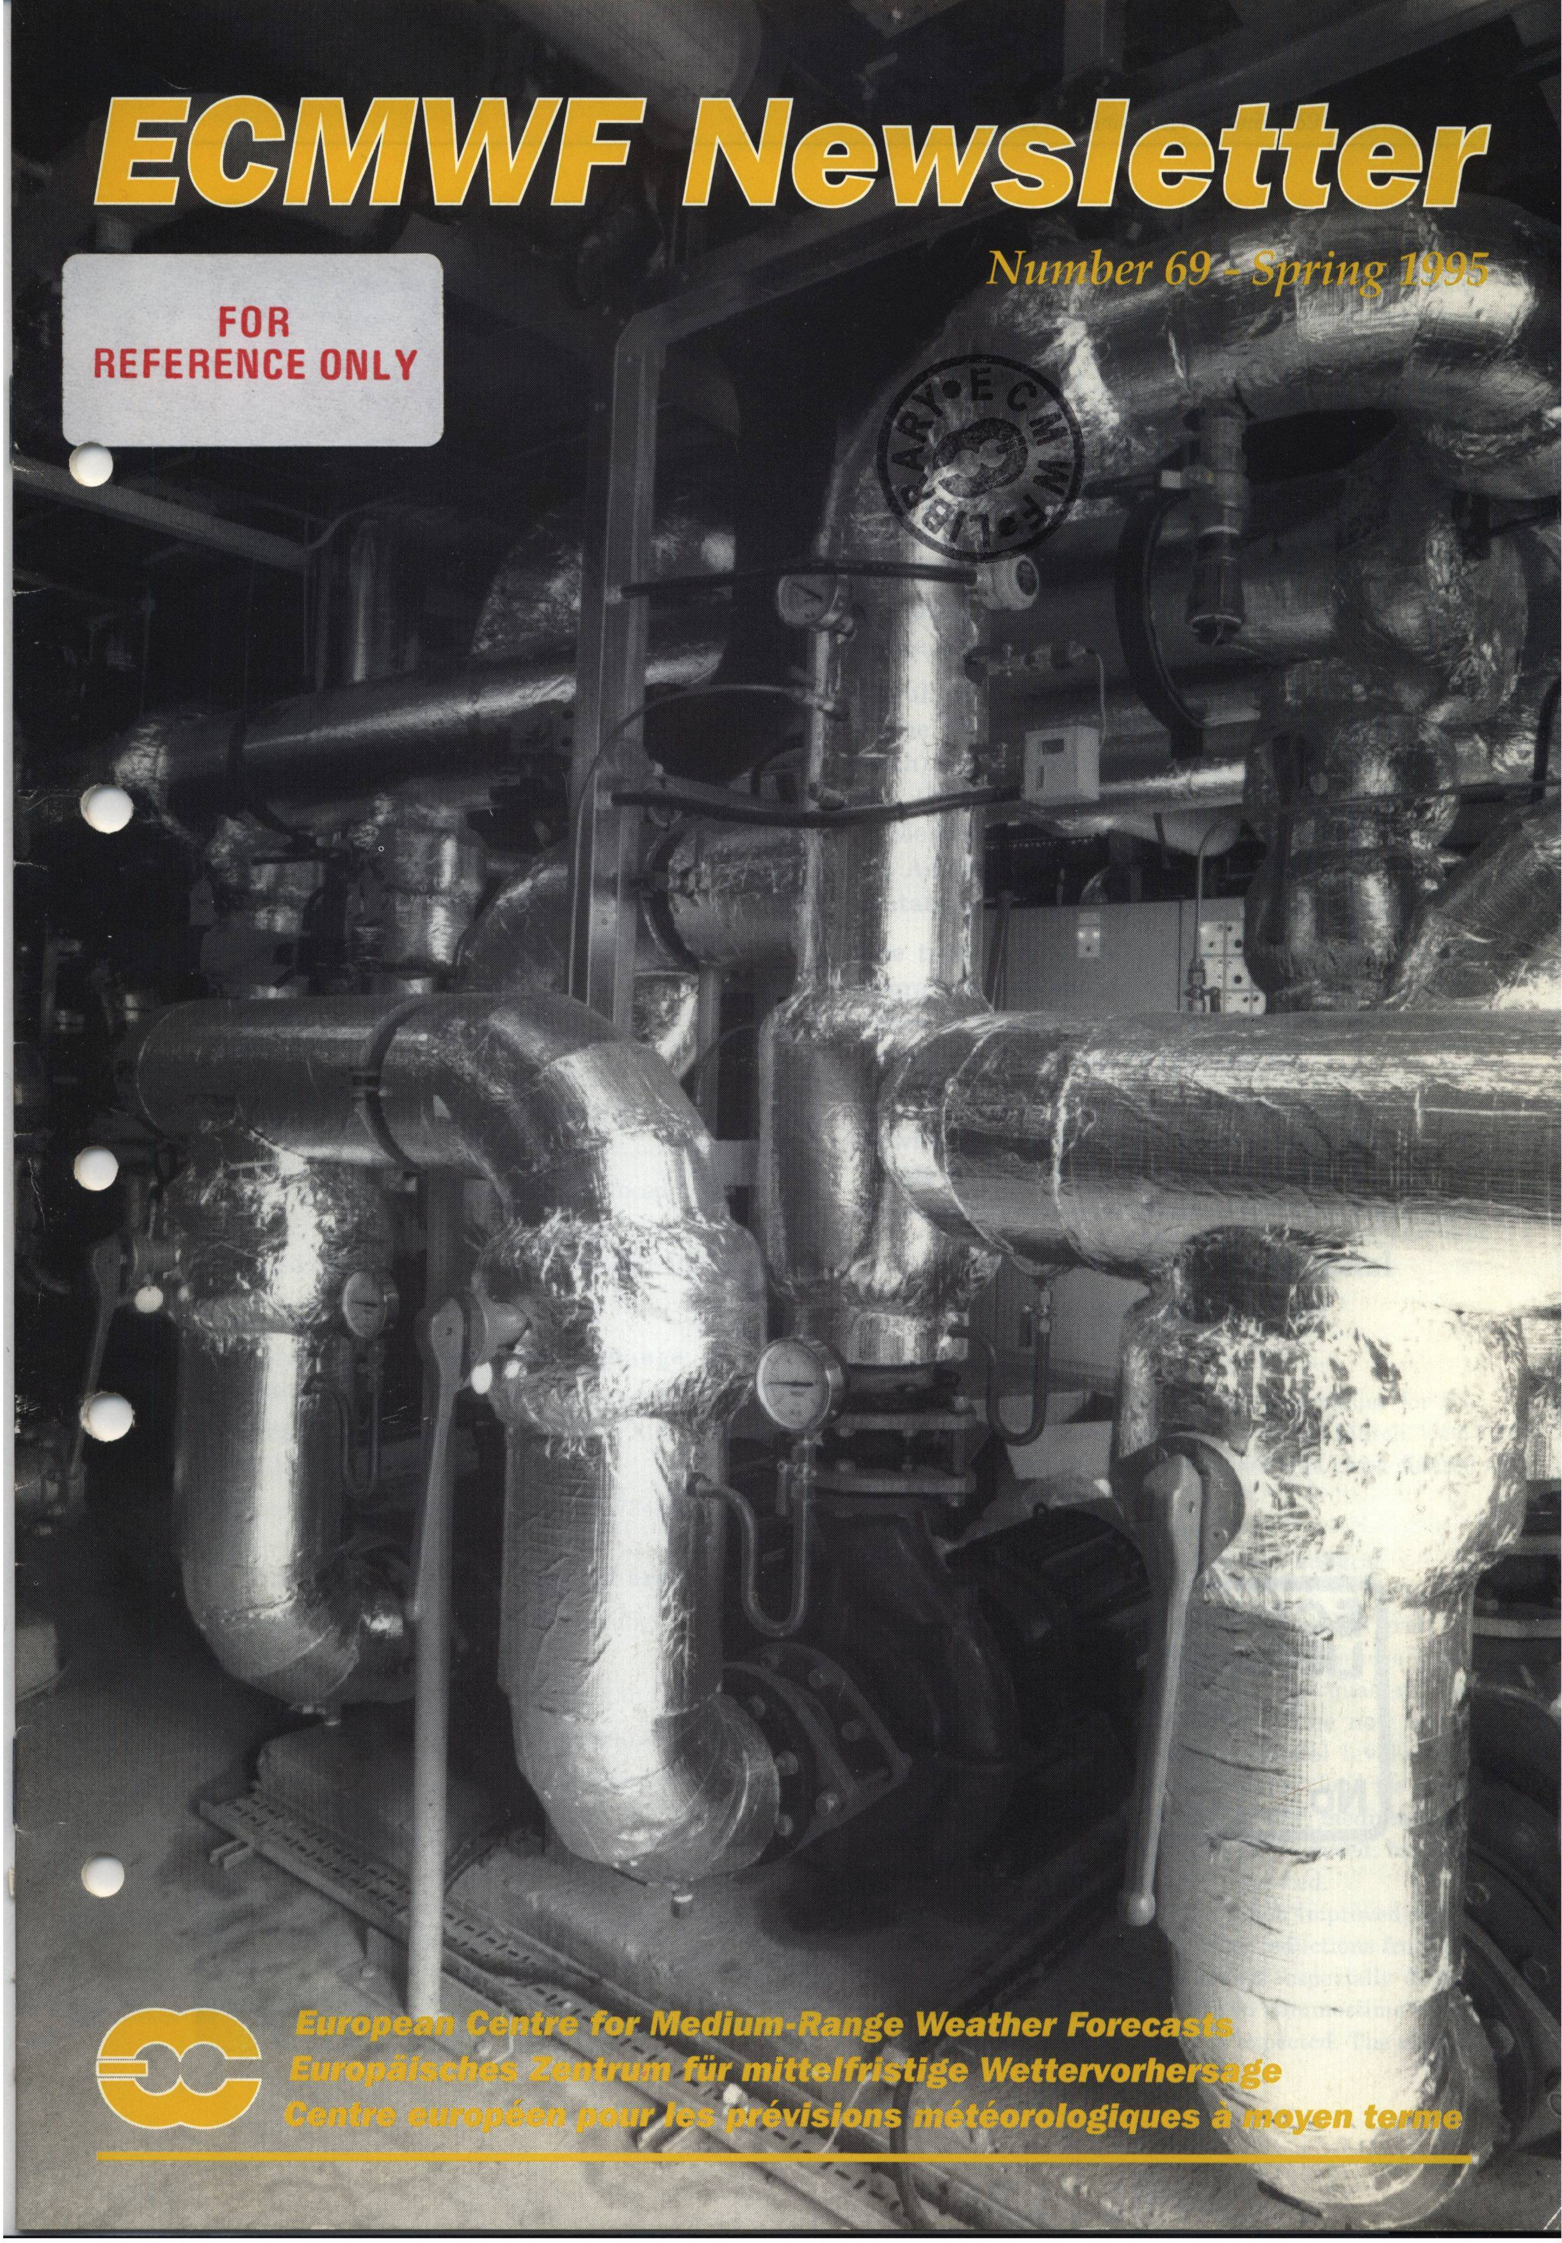

European Centre for Medium-Range Weather Forecasts
Europäisches Zentrum für mittelfristige Wettervorhersage
Centre européen pour les prévisions météorologiques à moyen terme

In this issue

METEOROLOGICAL

COMPUTING

GENERAL

This Newsletter is edited and produced by User Support.

The next issue will appear in Summer 1995.

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, Berkshire RG2 9AX,
England.

Telex: 847908 ECMWF G, Fax: (01734) 869450

Telephone Number: Change of area code

Please note that the telephone number for ECMWF is now

National	01734 499000
International	+44 1734 499000

Cover

Pipework for the upgraded chilled water system (see article on page 14)

Editorial

This edition of the Newsletter, after documenting the recent changes to the operational forecasting system, continues with a description of the effect of initial values of soil water on the quality of the summer forecasts, and a scheme for resetting these initial values to avoid model drift.

This is followed by a discussion of past, present and future high-performance computing requirements to carry out all the operations necessary for the production, and enhancement, of medium-range weather forecasts.

Regularly-recurring items are the announcements of the ECMWF Annual Seminar, and of the next (fifth) Workshop on Meteorological Operational Systems. Other routine articles include the updated lists of Member States' representatives on the Technical Advisory Committee and Computing Representatives and Meteorological Contact Points.

The trials and tribulations of implementing an upgrade to the chilled water system are also described in this issue, giving a good idea of what goes on behind the scenes in the Centre!

Readers will of course see that this edition of the Newsletter has been presented in a new format, which we hope will prove 'user-friendly'.

Changes to the Operational Forecasting System

Recent changes

A revision to the definition of the tropopause for the humidity analysis was introduced on 28 March 1995. A new version of the forecast model was implemented on 4 April, including:

- new prognostic cloud scheme;
- smoothed mean orography and new subgrid orography parametrization;
- modified quasi-regular (reduced) Gaussian grid.

Specific humidity on model levels and pressure levels was changed from spherical harmonic to quasi-regular Gaussian grid field. In addition, three new quasi-regular Gaussian grid fields on model levels were introduced: cloud liquid water content, cloud ice water content and cloud cover. Fields of low, medium, high and total cloud cover continue to be produced, but convective cloud cover is no longer produced.

This set of changes leads to a much improved representation of cloud cover, with large reductions from the previous systematic negative biases, especially during the day. A significant reduction in summertime warm bias of two-metre temperature is expected. The change should also improve the precipitation forecast, with smoother precipitation patterns and reduced spin-up problems in the early range of the forecast.

An improvement of the RMS and anomaly correlation scores in the medium-range over Europe, reflected in the overall synoptic pattern, was noted during the experimentation, particularly in summer.

In addition, the parametrization of the ocean surface albedo has been changed to better represent the variation of reflectivity with solar zenith angle.

Planned changes

A 3-d variational analysis system and a high resolution wave prediction model for the Mediterranean and Baltic Sea will be implemented. *Bernard Strauss*

Initial values of soil water and the quality of summer forecasts

Introduction

Despite the sensitivity of medium-range weather forecasts to initial conditions of soil water (see Garratt, 1993, for a review), there are at present very few methods to define the soil water content in data assimilation systems. Soil water values are neither observed nor exchanged regularly, making the direct use of observations impossible. Almost all schemes for initialisation of soil water are based on finding the equilibrium value of soil moisture, given *climatological* estimates of sensible and latent heat fluxes, and radiative fluxes at the surface (see e.g. Mintz and Serafini, 1992). They are therefore inappropriate for use with data assimilation schemes, where the goal is to find a soil water field representing an adequate balance of *real-time* estimates of the above fluxes. The error of short-range forecasts of summer-time near-surface temperature and dewpoint, when compared to the plentiful SYNOP observations, is normally a good indicator of the quality of the soil water field; in broad terms, too warm and too dry near surface atmospheric model states during daytime are associated with too dry values for the soil wetness. Mahfouf (1991) developed an optimal interpolation scheme for initialisation of soil water, relating the analysis increments of soil moisture to short-range forecast errors of near-surface temperature and dew point.

A new surface/boundary layer scheme was put into operation at ECMWF in August 1993 (Viterbo and Beljaars, 1995, see article in Newsletter no. 63). The model represents time scales ranging from the diurnal cycle to the seasonal cycle, with free drainage for moisture and zero heat flux for temperature as lower boundary conditions. Initial values for soil water and soil temperature were set to the first guess value, reflecting the values of precipitation, evaporation, sensible heat and radiation fields in very short term forecasts. Because there is no relaxation to a climatology there was some concern that the model could drift into its own climate state, if the above fluxes showed significant systematic deficiencies. Preliminary testing however had demonstrated that the surface model has negligible biases, when forced with the correct (observed) surface fluxes for a handful of selected locations corresponding to intensive experimentation campaigns. This article documents the model drift that effectively occurred in June 1994 and the quick solution to it implemented in July 1994, followed by a description of a very simple, but physically based, scheme for initialisation of soil moisture, based on the ideas of Mahfouf (1991). This scheme was implemented in December 1994.

The model drift in late spring and early summer of 1994

Figure 1 shows the weekly errors (forecast minus observation) in day two of the forecast for 2m temperature, specific humidity, and cloud cover, averaged for the European area, for the period January to September 1994: each curve represents one synoptic time, top curves show the standard deviation, bottom curves show the bias. The cold daytime bias in the winter reduces in early spring, but by mid-April a warm bias starts to build, reaching a maximum of almost 3 K in the last week of May, and staying this high throughout June. This is accompanied by the error in the humidity, where the model evolves from an unbiased winter and early spring, to a dry bias in late spring and early summer. In contrast to the temperature bias, the dry humidity bias continues to increase throughout June, reaching a maximum of -1.8 g kg^{-1} in the last week of June. The cloud bias is systematically negative, with little seasonal variation; a bias of -1.5 octas is typical of daytime values (out of an observed value of around 4 octas).

As a ‘quick therapy’ for the low level warm and dry model bias, the soil moisture was initialised to field capacity in the beginning of July, and then allowed to evolve normally in the assimilation cycle.The sudden reduction of the dry warm bias in the beginning of July coincides with the one-off wetting of the soil (3 Jul 94). The fact that the July and August errors never reach values similar to May and June shows that the wetting was an effective ‘fix’ for the problem.

The serious drift in the model surface parameters in May and June 1994 affected the quality of the forecasts at all ranges: June 1994 was a month of relatively poor scores (when compared to other centres). The forecasts showed large warm biases in the boundary layer, with a tendency to develop spurious thermal lows in ‘stagnant’ synoptic situations, and a positive height bias at upper levels. Figure 2 shows the 2-day forecast 50 kPa geopotential height bias, for DWD and ECMWF, over Europe (a similar picture was seen for the Northern Hemisphere). It is clear that the ECMWF positive bias in June, larger than the DWD values, is reduced in July

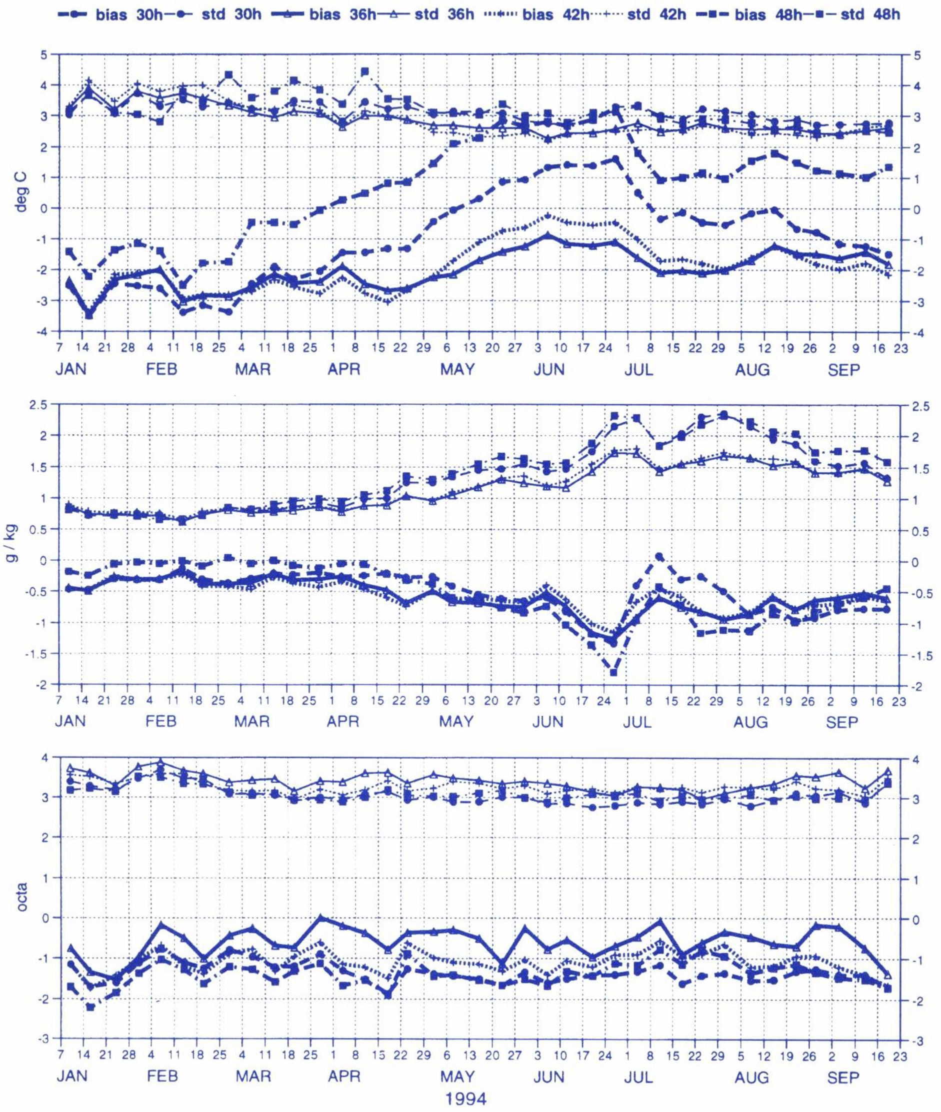

Fig. 1 Weekly European comparison with observations of day 2 forecast of weather elements for: a) 2m temperature; b) specific humidity; c) cloudiness. In each panel, the top curves represent the standard deviation, the bottom curves show the bias. Each curve corresponds to one synoptic time.

1994 following the one-off wetting of the soil. It seems plausible that the above drift is caused mainly by deficiencies in the cloud cover. The negative bias in cloud cover translates into excessive shortwave radiation at the surface, of the order of 100 Wm^{-2}; the excessive energy available at the surface causes too much evaporation early in the spring, drying the surface reservoir too rapidly, too early in the year. The new cloud scheme (Tiedtke, 1993) was implemented in April 1995 and reduces the cloud bias significantly.

The surface model formulation was last modified in 1993. Two 5-year runs at T63L31 resolution were performed, with Cy47 (the pre-August 1993 surface/boundary layer scheme, Blondin, 1991) and Cy48 (the new surface/boundary layer scheme, Viterbo and Beljaars, 1995). Cy48 has clearly less summer rainfall than Cy47 (which is itself already too dry). However, it was felt that the short-term forecasts would keep the atmosphere in a realistic state, preventing a serious drift in data assimilation. Moreover, the superiority of Cy48 over Cy47 in simulating the US floods of July 1993 (Beljaars et al., 1995) was an encouraging signal for the new surface scheme. The very large sensitivity of the 1994 summer scores to the state of the surface was therefore somewhat unexpected. Although the catalyst for the dry bias is clear (too much radiation at the surface caused by insufficient cloud cover), the reasons for its exaggerated magnitude lie in a precipitation/evaporation positive feedback caused by:

(i) a too small surface water reservoir (currently 15 cm of water), making the soil moisture too sensitive to a positive bias in the radiative forcing;

(ii) summer (convective) precipitation too tightly connected to local evaporation enhancing the positive feedback especially in areas of comparatively large recirculation of moisture in summer;

(iii) anomalies in the circulation for June 1994 specially sensitive to the warm/dry bias.

The results documented above provided the motivation for the development of the simple method for initialisation of soil water described hereafter.

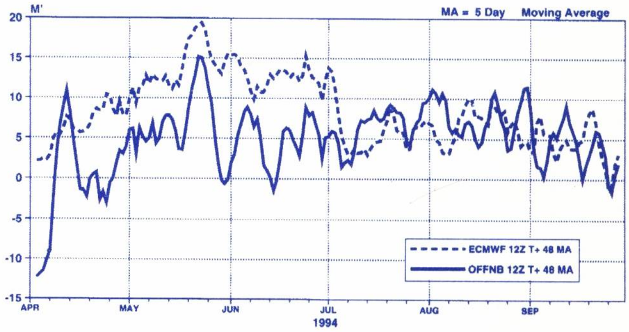

Fig. 2 Bias of the 50 kPa geopotential over Europe, for day 2 ECMWF (dashed curve) and DWD (solid curve) forecast.

A simple method for initialisation of soil moisture

A large model bias such as the one documented above, is normally reflected in the analysis cycle, where a very short term forecast (first guess) is compared to observations: a warm and dry model bias is compatible with negative temperature increments (analysis minus first guess) and positive humidity increments. The analysis increments for May 1994 exhibit the model problem in low levels. Inspection of the mean May analysis increment for lowest model level specific humidity, at 12 UTC (not shown here), shows large positive increments over land, in Europe, Africa, Northern Australia, and India, with maxima of around 2.5 g kg^{-1} in Europe. The same pattern, but with reduced amplitude, can be seen up to model level 26 (around 85 kPa). Negative analysis increments for temperature at low levels can be seen in the same areas, ranging from -0.5 to -1 K in Europe (not shown). For other synoptic types, the pattern of dry/warm first-guess bias has maxima around local afternoon. It is clear that the analysis is trying to moisten and cool the boundary layer, to compensate for the model bias, suggesting that the low level humidity increments can be used to identify the areas where the bias is present.

Assuming that the low-level humidity (q) bias responds instantaneously to a soil moisture (θ) bias in the root layer, then we can write

$$(\theta_a-\theta_g) = C_v D \Delta t (q_a-q_g) \qquad (1)$$

where Δt = 6 hr, and the subscripts g and a stand for first-guess and analysis values, respectively, D is a

Northern Hemisphere *1/6/94–20/6/94 (20 cases)*

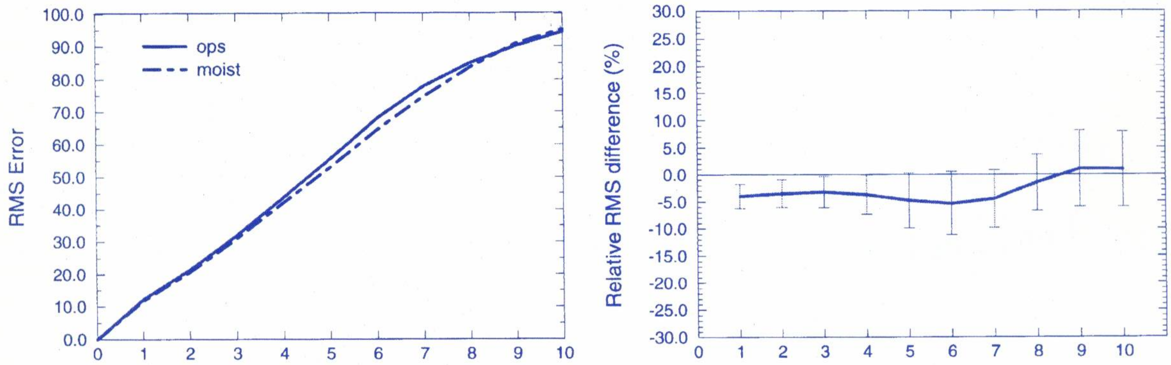

Europe *1/6/94–20/6/94 (20 cases)*

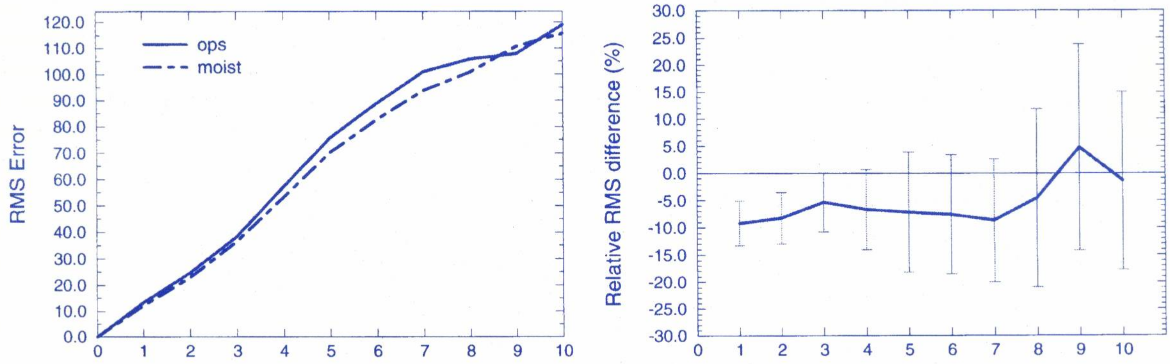

North America *1/6/94–20/6/94 (20 cases)*

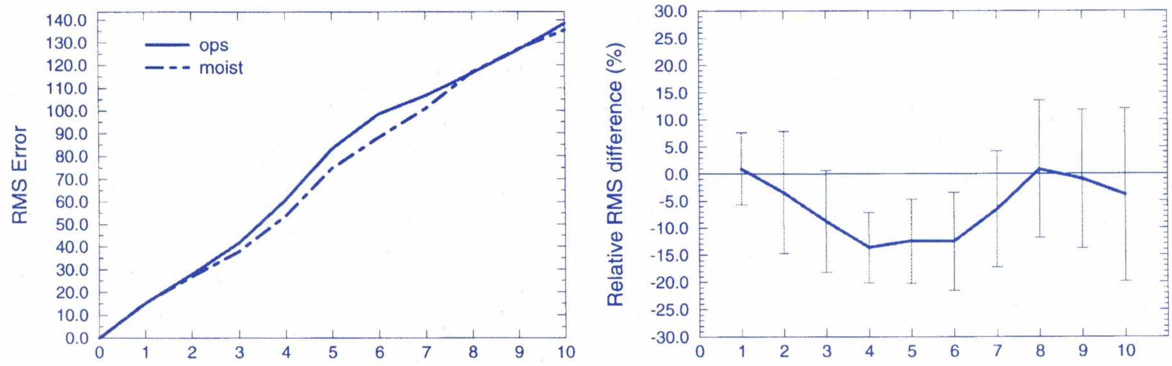

Fig. 3 Mean root mean square of geopotential height for control (ops) and test (moist), and mean difference and confidence interval at 95% level, 1-20 June 1994, for 50 kPa Northern Hemisphere and Europe, and 20 kPa North America.

(global) constant, $D\Delta t$ represents the ratio between a unit increment of soil moisture and a unit increment of specific humidity for a fully vegetated area, and θ represents the vertically integrated value of soil moisture. At each analysis cycle the above increment, weighted by the vegetation fraction C_v, is added to the first guess soil water; the weighting ensures that the scheme is not active over deserts. As further precautions, no increments are produced in the presence of snow, and the increments are limited in such a way that the field capacity and the permanent wilting point thresholds are never crossed. The root integrated soil water increments are distributed in each of the three soil layers following the model root extraction. Only the humidity increments in the lowest level of the model are used, because the other boundary layer increment fields have the same structure.

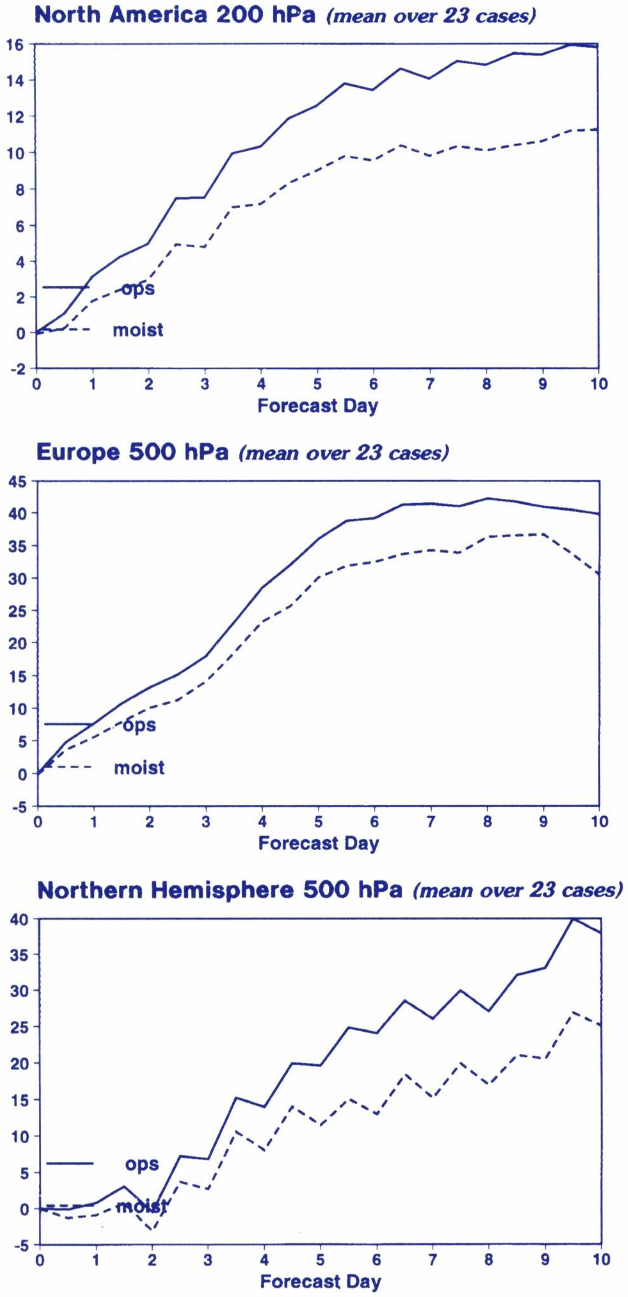

Fig. 4 Same as Fig. 3, but for the mean bias.

The above method can be interpreted as a nudging of the soil water in the analysis, proportional to the magnitude of the boundary layer humidity increment. The value of the nudging (relaxation) time constant used for the summer pre-operational test described in the next section and subsequently used for operational implementation corresponds to a moisture analysis increment of 3 $\mathrm{gkg^{-1}}$ filling 0.15 m of water in the soil in 1.5 days. This constant is discussed further below.

In order to assess the impact of the scheme on the poor performance of the forecast in June, an analysis experiment at T213L31 resolution has been run, for the period 15 May - 3 July 1994. Ten-day forecasts were run on 1-20 June and on 1-3 July. Figure 3 shows the differences in root mean square error for the geopotential at 50 kPa, Europe and Northern Hemisphere, and 20 kPa, North America. The scheme performed better at almost all ranges, and the improvement was significant at the 95% confidence level at least up to day two. The area where the impact was the largest of all, North America at 20 kPa, has a significant positive impact up to day six. The large impact comes mainly from a reduction in the bias, as can be seen in Figure 4: rerunning the same period with the new cloud scheme (see below) removes the bias that can still be seen in the 'moist' analysis. As should probably be expected, the impact over the Southern Hemisphere was small (not shown).

The impact on 2m temperature and specific humidity is shown for Europe in Figure 5 (verification against synop observations). The scheme managed to bring the daytime temperature error over Europe down from around +3 K to about +1 K, and to keep it at that level even during the period when the drying-out of the soil in operations was at its maximum, after 20 June. A similar behaviour is seen over Africa (not shown). However, the night-time cold bias was increased to about -2K over Europe. Overall, the behaviour of T and q was brought close to what it was after the resetting of the model soil to field capacity on 4 July, with no negative impact on the short-range forecast of precipitation.

Extended tests of the proposed scheme

Early results for May-August 1979 from the ECMWF Reanalysis Project showed large humidity analysis increments in the boundary layer, although smaller than the operational increments for the corresponding months of 1994. A rerun with the initialisation of soil moisture was started from April 1979, and continued until December. The benefit of the initial value scheme was clearly seen through reduced boundary layer increments, showing that the soil moisture increments contribute towards creating a more balanced atmospheric state. The change in the soil initial conditions affects the model surface fluxes in a dramatic way. Comparison of mean monthly values of the Bowen ratio (H/LE,

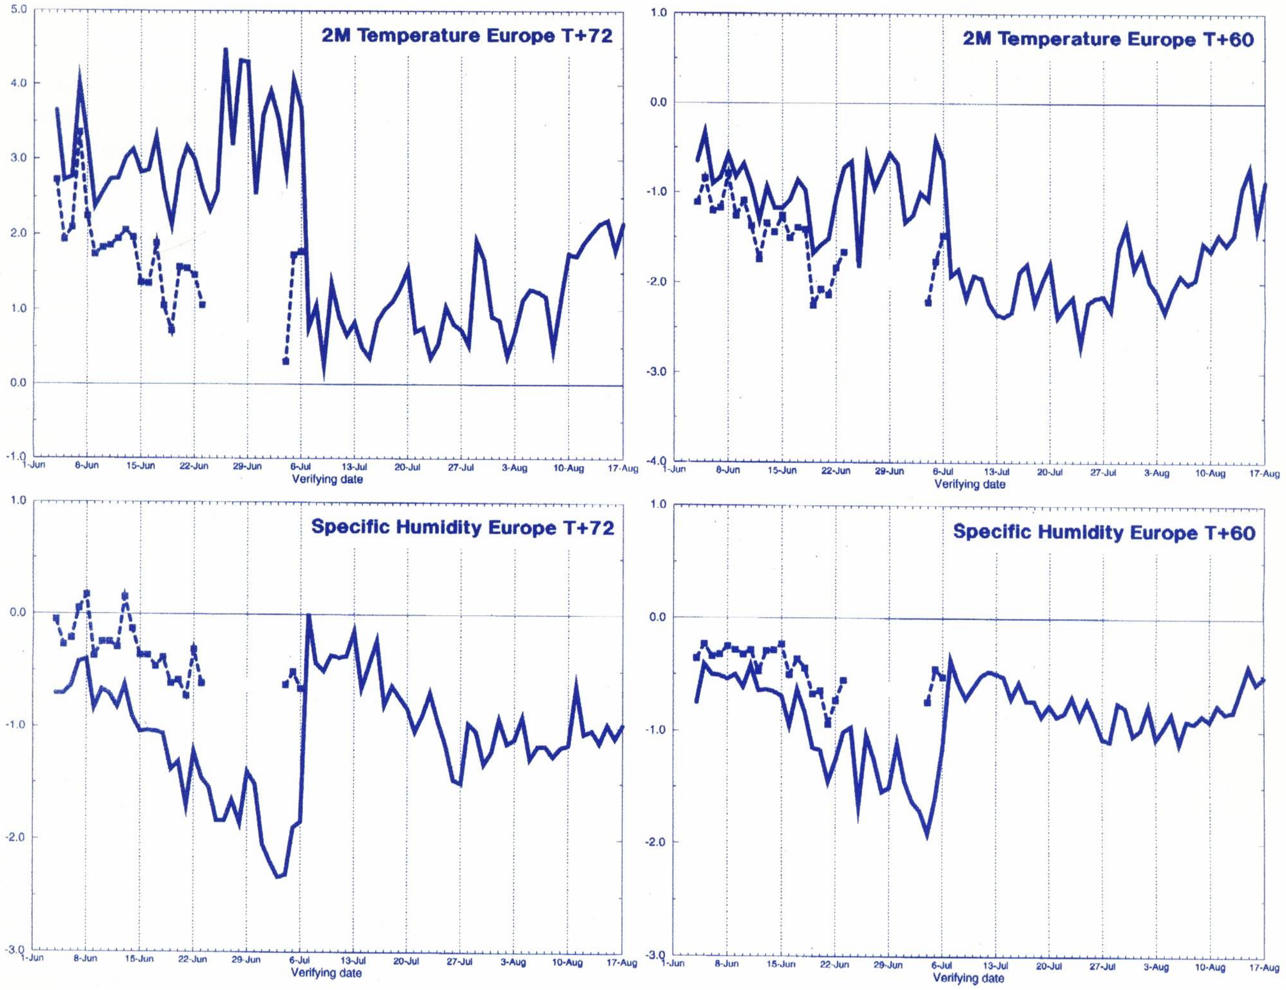

Fig. 5 Daily mean forecast error of 2 metre temperature and specific humidity verified against synop observations, T+60 and T+72, Europe. Thin line: operations, solid line with squares: test suite.

where H and LE are, respectively, the sensible and latent heat fluxes averaged over the day 1 forecast) over land for the summer of 1979, with and without the soil moisture initialisation are significantly different: global values of the Bowen ratio range from 0.6 to 0.9, without soil water initialisation, and from 0.4 to 0.6, with soil water initialisation, corresponding to larger values of latent heat and smaller values of sensible heat.

Similar changes to the fluxes were seen in the pre-operational test described above. Figure 6 presents a zonal average over all land points of the mean sensible and latent heat fluxes for all day two forecasts between 940601 and 940620. The increase in latent heat in the Northern Hemisphere is compensated almost exactly by the decrease in sensible heat, showing that the surface radiation (and the cloudiness) are not significantly affected. There is very little change in the winter hemisphere and in high latitudes, as expected. A similar diagram for precipitation shows 15% more convective precipitation at day two in the forecast for the pre-operational test, in latitudes corresponding to the maximum difference between the fluxes. Figure 7, showing the mean hydrological and energy budget at the surface for the same period, shows that this increase in precipitation corresponds to reduced spin-up of precipitation, suggesting a more balanced initial state. Note that a slow increase of evaporation during the forecast in operations is replaced by a decrease in the pre-operational test. Rerunning the same period with the new cloud scheme (see below) removes the spin-up in precipitation that can still be seen in the 'moist' analysis.

Because there is a time scale implicitly involved in equation (1), the soil moisture evolution in the analysis can be thought of as the combined effect of the surface model and its atmospheric forcing in short-term forecasts, and the additional information from observations suitably filtered in time. Note that the filtering properties only exist because equation (1) is applied to obtain positive (wetting) and negative (drying) increments. In the absence of rain, during summer, the main atmospheric forcing to the surface is a drying effect of evaporation. In order to have a net wetting effect, the amount of water added by the increments, in case of a dry bias, has to compensate for the evaporation loss during the next 6-hour forecast. If evaporation is larger than the increment added, the soil water content will decrease from one analysis time to the next. Because the wetting increments occur mainly in the daytime, and be-

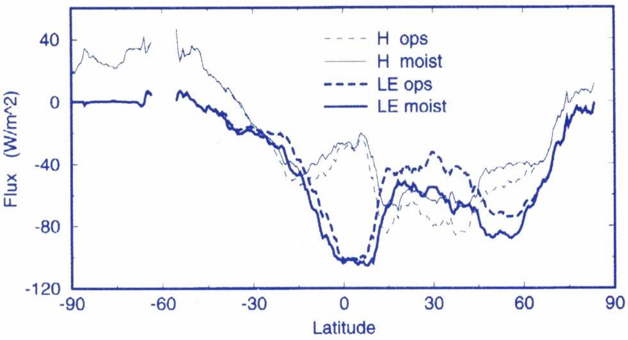

*Fig. 6 Latent heat (thick curves) and sensible heat (thin curves) zonally averaged for land points, for all day 2 forecasts between **940601** and **940620**, for the moist experiment (solid) and operations (dashed).*

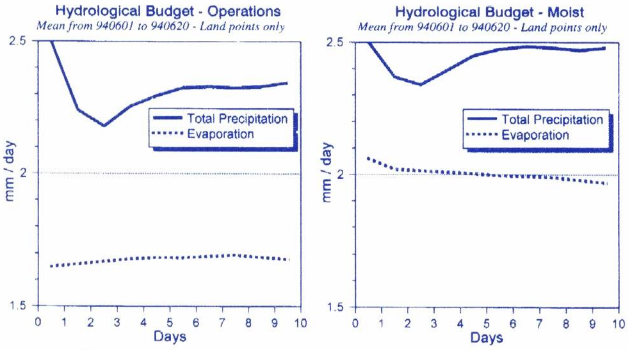

*Fig. 7 Hydrological budget of the land surface for the moist experiment (right panel) and operations (left panel), averaged for all forecasts between **940601** and **940620**.*

cause of the drying effect of evaporation, the interpretation of the 1.5 days referred to earlier as a time constant has to be considered with caution. This time scale is *much shorter* than the effective time scale of the net change (difference of the values between two consecutive days) to the soil water content. A way of verifying if the filtering properties of the nudging scheme are adequate, is to check if there is a diurnal cycle in the root integrated soil water in the analysis. Note that, although the soil water in the top soil layer should have a diurnal cycle, the root integrated value should be constant. Inspection of the 00 UTC and 12 UTC values of root integrated soil water, averaged over the pre-operational test period, reveals no evidence of a diurnal cycle.

The relaxation constant D should be large enough to prevent the drying drift, but not so large as to imply a diurnal oscillation of root integrated soil moisture. Ideally, it should also be dependent on cloudiness and other factors: essentially we want increments in soil moisture only in situations where the dry bias can be due mainly to deficiencies in the surface model or its forcing. A proper estimation of D for each situation is in fact a feature of the Mahfouf (1991) assimilation scheme, with some 'quality control criteria' to be met. In the simple method described here we can only use a global constant value for D.

In order to test the sensitivity of the surface and boundary layer to the physical parameterization, a number of 5-month T63L31 experiments were run, with initial conditions 930401, but alternative parametrization schemes. An 85 kPa warm bias of up to 8 K over the Eurasian area for the control experiment is reduced to around 4 K in an experiment with the prognostic cloud scheme (Tiedtke, 1993) recently introduced in operations, as a result of the increase of cloud cover of the new cloud scheme. A number of changes in the parametrization schemes were tried, and sensitivity was found to the closure of the convection scheme (Nordeng, 1994), and to an alternative formulation of the stress function for transpiration, albeit to a much lesser extent than the new cloud scheme.

Conclusions

A very simple way of initialising the soil water content has been presented, and extensive experimentation, concentrating on the May / June 1994 period shows:

(i) a much reduced summer warm/dry bias;
(ii) better 2 metre temperature and humidity during daytime;
(iii) a large positive impact in terms of summer objective scores;
(iv) only a small impact on the winter hemisphere and other periods;
(v) a small negative impact on nighttime temperatures;
(vi) no over-prediction of precipitation as a result of the increased soil wetness.

The method presented here has been operational at ECMWF since December 1994, and is used by the ECMWF Reanalysis Project. As discussed above, the value of D proposed for operational use is deliberately a 'cautious' choice. Even without the implementation of the method, the new cloud scheme would reduce the dry/warm bias, but control through the analysis system is still desirable, albeit with a smaller value of D. The period of June 1994 will be used in future experiments aimed at understanding better the nature of the model precipitation-evaporation feedback described above. Future experimentation with changes in the closure of the convection scheme or the size of the soil moisture reservoir will be analyzed in this context. Given the greater sensitivity of the summer scores to initial conditions of soil moisture, work towards a better method of defining those, along the lines of Mahfouf (1991) and Bouttier et al. (1993), is already under way.

Pedro Viterbo

References

Beljaars, A.C.M., P. Viterbo, M.J. Miller, and A. Betts, 1995: The anomalous rainfall over the USA during July 1993: Sensitivity to land surface parametrization. *ECMWF Tech. Mem. 214.*

Blondin, C., 1991: Parametrization of land-surface processes in numerical weather prediction. *Land surface evaporation: Measurement and parametrization,* T.J. Schmugge and J.C. André, Eds., Springer, 31-54.

Bouttier, F., J.-F. Mahfouf, and J. Noilhan, 1993: Sequential assimilation of soil moisture from atmospheric low-level parameters. Part I: Sensitivity and calibration studies. *J. Appl. Meteor.*, **32**, 1335-1351.

Garratt, J.R., 1993: Sensitivity of climate simulations to land-surface and atmospheric boundary layer treatments - A review. *J. Climate*, **6**, 419-449.

Mahfouf, J.-F., 1991: Analysis of soil moisture from near-surface parameters: A feasibility study. *J. Appl. Meteor.*, **30**, 1534-1547.

Mintz, Y., and Y.V.Serafini, 1992: A global monthly climatology of soil moisture and water balance. *Clim. Dynamics*, **8**, 13-27.

Nordeng, T.E., 1994: Extended versions of the convective parametrization scheme at ECMWF and their impact on the mean and transient activity in the tropics. *ECMWF Tech. Mem. 206.*

Tiedtke, M., 1993: Representation of clouds in large-scale models. *Mon. Wea. Rev.*, **121**, 3040-3061.

Viterbo, P., and A.C.M. Beljaars, 1995: A new land surface parametrization scheme in the ECMWF model and its validation. *ECMWF Tech. Rep. 76.*

High-performance computing requirements for medium-range weather forecasting

Introduction

Until recently, the computational requirement for numerical weather prediction has been determined principally by the need to model the complex behaviour of the atmosphere with as fine a resolution as possible, in order to produce a realistic simulation of the evolution of weather systems within an elapsed time of at most a few hours. Use of finer resolution has been made possible both by substantial increases in computational power and by development of more efficient numerical techniques. Much improved and more comprehensive representations of the dominant physical processes have also been introduced. New types of observation have become available and better ways have been developed for processing observations to define the starting state for the atmospheric model. The resulting increases in accuracy have been dramatic for both short-range forecasts (from a few hours to a few days ahead) and medium-range forecasts (from a few days to a week or two ahead). For example, the level of accuracy reached on average by three-day forecasts in the pioneering medium-range prediction experiments reported by Miyakoda et al.[1] in 1972 was reached typically by five-day forecasts in 1979/80 during the first winter of operational forecasting by ECMWF. Currently, average seven-day forecasts come close to reaching this level.

These developments have brought us to a point at which refinement of the model is no longer the predominant reason for requiring higher computing performance, at least for medium-range prediction. Predictability studies such as that by Simmons et al.[2] suggest that better estimation of the initial state offers the principal path to more accurate individual (deterministic) forecasts, although there is still scope for benefit from model improvement. The most promising method for better exploitation of observational data, particularly newer types, is the computationally demanding method of four-dimensional variational data assimilation (Courtier et al.[3], and references). Moreover, the use of deterministic forecast information in the later medium range is hindered by variations in forecast quality which reflect both uncertainty in initial conditions and the varying degrees of predictability of different flow regimes. This has led to development of ensemble (probabilistic) prediction systems based upon multiple lower-resolution integrations from perturbed initial states (e.g. Molteni et al.[4]).

It is thus an appropriate time to look back at the past use of computing for numerical weather prediction, to summarize present use, and to look forward a few years to examine the computational implications of the development of variational data assimilation and ensemble prediction. Our present and future views will take the ECMWF forecasting system as an example, although interest in variational assimilation and ensemble prediction is by no means limited to ECMWF.

A centre such as ECMWF that has both operational and research functions, and which provides a remote computing service, has many requirements of its high-performance computing system. Important items include reliability of hardware and software, functionality of operating system, efficiency of running a mixed workload, effectiveness of compilers and debugging tools, memory and I/O resources. These will not be discussed further here. Rather, we shall consider only basic computational performance, as judged either by the time taken to execute a model forecast, or by the floating-point computation rate achieved by the model.

The past

The origin of numerical weather prediction can be traced back to at least 1904, when Vilhelm Bjerknes[5] set out the basic principles of the computation of atmospheric evolution. The conclusion of his article included the following words (taken from a translation of the original German by a later pioneer of numerical modelling, Yale Mintz): "It may be possible one day, perhaps, to utilize a method of this kind as the basis for a daily practical weather service. But however that may be, the fundamental scientific study of atmospheric processes sooner or later has to follow a method based on the laws of mechanics and physics. And thereby we will arrive, necessarily, at a method of the kind outlined here."

Bjerknes' visionary work "exercised a considerable influence" on the later pioneering practical studies by L.F. Richardson[6], published in 1922 in a book entitled "Weather Prediction by Numerical Process." Richardson set out with remarkable detail and completeness a scheme for computing the evolution of the atmosphere. The scheme was "complicated because the atmosphere is complicated." His hand-worked example used initial conditions derived in part from isobaric charts which Bjerknes himself had prepared. Richardson used slide rule and log tables to produce a six-hour forecast at five different heights at the centre of an initial horizontal array of nine 'momentum' and sixteen 'pressure' points. The computation time was six weeks on "a heap of hay in a cold rest billet" while he served in an ambulance unit during World War I. Lynch[7] has recently re-examined that first numerical forecast.

The penultimate chapter of Richardson's book was confidently entitled 'Some remaining problems.' It began with the words: "The two great outstanding difficulties are those connected with the completeness necessary in the initial observations and with the elaborateness of the subsequent process of computing." Richardson could not foresee the electronic computer, but asking to "play with a fantasy" he set out a scheme for massively parallel computation. He envisaged a "forecast-factory" of 64,000 human computers. Groups of these computers were assigned to each atmospheric column (one grid-square, five levels in the vertical). Each group communicated results optically to neighbouring groups responsible for the columns to north, south, east and west, and to a central team of clerks responsible for dissemination of the forecast. A central coordinator kept the calculations in step.

It required the more or less simultaneous development of the electronic computer and of rational simplifications of the governing equations for Charney, Fjörtoft and von Neumann[8] to complete the first successful numerical forecast experiments some thirty years later. Their model predicted the evolution of the atmospheric flow at just a single level, and at an array of 19x16 grid-points covering North America and neighbouring waters. The 736km grid interval was significantly larger than anticipated by Richardson. Each 24-hour forecast took about 24 hours to complete, and involved about 250,000 multiplications and divisions. The authors concluded: "With a larger capacity machine, such as is now being built ... (it) is estimated that the total computation time with a grid of twice the Eniac-grid's density will be about ½ hour, so that one has reason to hope that Richardson's dream(1922) of advancing the calculation faster than the weather may soon be realized."

The hope indeed became reality soon afterwards, and was followed in subsequent years by major developments in technique, computing power and the accuracy of results. Shuman[9] has given a review of more than three decades of operational numerical weather prediction in the USA, and illustrated the dramatic improvements in 36-hour forecasts that occurred between 1955 and 1988. Comparable developments took place in Europe and elsewhere.

Routine forecasting began at ECMWF in 1979. The first operational model used finite-differences and covered the globe with a 1.875° (192x96) grid. In the vertical, atmospheric variables were represented at 15 levels and soil variables at 3 levels. A ten-day forecast took between four and five hours to produce on a single-processor CRAY-1 vector computer. The effective computation rate was about 40 million floating-point operations each second (40Mflops) including operational post-processing overheads, and about 50Mflops without overheads. A model using the spectral transform technique was introduced with T63 truncation in 1983, and multi-processing was introduced when the horizontal resolution was increased to T106 in 1985 on a two-processor CRAY X-MP machine. The computation rate of this model exceeded 300Mflops on a four-processor X-MP in 1986, at which time vertical resolution was increased to 19 levels. The Gflops range was reached when an eight-processor CRAY Y-MP system was installed in 1990. Simmons et al.[10] have reviewed the numerical formulations used operationally over this period, and computational aspects of the multi-tasking spectral model have been reported by Simmons and Dent[11].

The present

Ritchie et al.[12] have described the numerical formulation of the current operational ECMWF forecast model and presented some computational details of its performance on the Centre's 16-processor CRAY Y-MP/C90 computer. The most recent major operational change was in 1991 when horizontal resolution was increased to T213, vertical resolution was increased to 31 levels, and a semi-Lagrangian treatment of advection was introduced. At the same time, the computational grid was changed in the east-west to be quasi-regular in distance rather than longitude. The grid-spacing is now about 62.5km, and the globe is spanned by more than 134,000 points at each level. The forecast proceeds to ten days in time-steps of 15 minutes.

The latest Integrated Forecast System (IFS) model code can produce a ten-day forecast in about 75 minutes when run without overheads using 16 processors. The computation rate approaches 6 Gflops, more than one hundred times the speed achieved fifteen years earlier, implying close to a doubling of speed every two years on average. The main operational forecast is multi-tasked over 14 processors, and runs with post-processing jobs and some other general work, typically completing in around two hours. Considerable effort has recently been devoted to producing a flexible, portable version of this model code which can be run on distributed-memory machines. This version can use efficiently as many as 3000 processors in parallel at the operational T213, 31-level resolution. Details, including results obtained on various types of machine, are presented in papers by Barros et al. and Dent et al.[13].

The computational cost of assimilating observations into the model to provide the starting point for the forecast cannot be neglected. Part of this cost is that of the short-range forecasts that are run to provide the background information for the actual analysis of observations. Currently, analyses are carried out for the main observation times of 00, 06, 12 and 18UTC, and a nine-hour forecast is needed from each of these times to provide the background for the analysis of the observations made in the six-hour time period centred around the following main observation time. The ten-day medium-range forecast is run daily from the 12UTC analysis, and the first nine hours of it doubles as the background forecast for the next 18UTC analysis. The background forecasts are relatively expensive as they each include initialization costs. They require a net CP time of about 20% of that needed by the ten-day forecast (apportioning the cost of the first nine hours of the 12UTC forecast equally between the background and the medium-range forecast.) The analyses have used the optimal interpolation method described by Lorenc[14] since the start of operational forecasting. They typically use a little more CP time than is needed for the background forecasts. Adding in the costs of pre-processing observations and post-processing results, the total computational cost of the 24-hour data assimilation amounts to not much less than 50% of the cost of the ten-day forecast.

Only the 12 UTC analysis is time critical, as observations are received soon enough for earlier analyses and the background forecasts to be completed well before the cut-off time for data receipt for the 12 UTC analysis and forecast. Parts of the analysis are multi-tasked over eight processors in operations. Other work runs in parallel. The elapsed time for the analysis is typically around a third of its CP time.

The future

There has been a vast increase in observations of the atmosphere since L. F. Richardson wrote: "The two great outstanding difficulties are those connected with the completeness necessary in the initial observations and with the elaborateness of the subsequent process of computing." Nevertheless, uncertain knowledge of the initial state remains an outstanding problem for numerical weather prediction. Additional or more accurate observations offer one possible way to alleviate the problem, but progress is also possible by other means. One is to develop better methods of exploiting existing observations to produce more accurate initial states. Another is to quantify forecast uncertainty through development of new forecasting techniques. At ECMWF, the first is being followed through development of a four-dimensional variational data assimilation system and the second through development of an ensemble prediction system. The computational costs of these two approaches dominate the Centre's requirement for computational power in excess of that provided by its CRAY Y-MP/C90 computer. In this section we discuss the requirements that are foreseen for the coming five years or so.

Variational data assimilation

Four-dimensional variational data assimilation involves determining values of the model state $\underline{x}$ at time t_1 that minimize a scalar function $J(\underline{x})$, given a 'background' estimate $\underline{x}_b$ of $\underline{x}$ and a set of observations $\underline{o}$ made in the time interval $t_1 \leq t \leq t_2$. The dimension of $\underline{x}$ is of the order of $6 \text{x} 10^6$ for a model with T213, 31-level resolution. J may be written:

$$J = J_o + J_b + J_c \text{ where,}$$

J_o is a measure of the difference between observed values and corresponding model estimates of these values (appropriate functions of $\underline{x}$),
J_b is a measure of the deviation of $\underline{x}$ from $\underline{x}_b$, and
J_c represents any additional constraints.
J_o is given by

$$J_O = \tfrac{1}{2}(\underline{H}(\underline{x}) - \underline{o})^t \underline{\underline{O}}^{-1} (\underline{H}(\underline{x}) - \underline{o}) \text{ where,}$$

$\underline{\underline{O}}$ denotes the covariance of observation errors (including effects of representativeness), and
$\underline{H}(\underline{x})$ maps model variables at initial time t_1 to variables observed at time t.

Evaluation of $\underline{H}(\underline{x})$ in general involves a forecast from t_1 to t, and a transformation from the model variables to the observed variable, for example from model temperatures and humidities to radiance as measured by satellite.

Efficient solution of this problem requires knowledge of the gradient of J with respect to $\underline{x}$, at each iteration of the algorithm used for the minimization. Direct evaluation would be prohibitively expensive given the dimension of $\underline{x}$. However, Le Dimet and Talagrand[15] showed that adjoint methods can be used to good effect. The J_o term is the most expensive computationally, and its gradient is computed in the following way. First, the full model is integrated forward in time from t_1 to t_2, starting from the latest estimate of $\underline{x}$, and storing the forecast state at each time-step. Then, the adjoint of the model equations linearized about the forecast states stored during the forward integration is integrated backwards from t_2 to t_1. The adjoint model includes forcing by a term which depends on the differences between the observations and model estimates based on the stored forecast states. The backward integration starts with all fields set to zero, and the result at time t_1 gives the required components of the gradient of J_o with respect to $\underline{x}$.

If the linearized model equations are written in general form $\underline{x}(t) = \underline{\underline{R}}\underline{x}(t_1)$ then the adjoint, $\underline{\underline{R}}^*$, of the 'resolvent' operator $\underline{\underline{R}}$ is such that, for any $y(t)$,

$$[\underline{y}(t), \underline{\underline{R}}\, \underline{x}(t_1)] = [\underline{\underline{R}}^* \underline{y}(t), \underline{x}(t_1)]$$

where [,] is an appropriate norm (typically based on the total energy). Talagrand and Courtier[16] give an explicit example of the adjoint of the barotropic vorticity equation. In general, calculation of the adjoint of the linearized model equations involves the evaluation of more terms than in the original non-linear model; for cost estimates we assume the increase in calculation is by a factor of two. One iteration using observations over a one-day period is thus approximately equal in cost to a three-day model forecast.

It can then readily be seen that if 100 iterations are required to achieve sufficient convergence of the minimization, then the cost of carrying out the data assimilation will be about equal to the cost of making a 300-day forecast, or about 30 times the cost of making the medium-range forecast.

Courtier et al.[3] therefore sought to reduce further the cost of the assimilation. They proposed an "incremental" variational method. This uses a lower-resolution model (with simpler physical parametrizations) to compute increments to the high-resolution initial state that reduce the deviation from observations of the subsequent high-resolution forecast. The scheme comprises outer and inner iterative loops. One outer iteration entails forward integration of the full high-resolution model, storing deviations from observations, and forward integration of the lower-resolution model, storing fields at each time-step (the "trajectory"). Each inner iteration involves forward integration of the lower-resolution model linearized about the trajectory, saving the perturbations at observation points, followed by backward integration of the adjoint of the linearized lower-resolution model, forced by the net deviation from observations.

This incremental approach has been shown to work in practice, and to bring further significant computational savings. It has the added advantage of reducing the memory needed for storage of the trajectory. It is used also in the three-dimensional variational assimilation scheme which is expected soon to replace the optimal interpolation method used for operational assimilation at ECMWF. This variational scheme is in fact a special case of the four-dimensional scheme outlined above. The background states $\mathbf{x}_b$ are obtained from forecasts from preceding analysis times as in the optimal interpolation scheme, and the observations made in a time window centred on time t_1 are applied at t_1. In this case the operator $\underline{H}(x)$ involves simply the mapping of model to observed variables.

Although the model integrations are the dominant computational component of four-dimensional variational assimilation, other components cannot be neglected when it comes to considering implementations on highly parallel computer systems. If only the model components are efficiently parallelized, then the cost of the other components may significantly influence the elapsed time of the assimilation, just as a single non-vectorized loop can noticeably slow the performance of the current highly-vectorized model code. Isaksen and Barros[17] discuss the parallelization of variational data assimilation.

Ensemble prediction

An experimental Ensemble Prediction System (EPS) has been run daily at ECMWF since May 1994, following successful results in trials run three days per week starting in December 1992. The principal aims of this approach are:

- provision, in advance, of an estimate of the skill of the ten-day high-resolution- deterministic forecast;
- provision of possible alternative evolutions of the atmospheric circulation pattern;
- provision of probabilities of the occurrence of specific weather events.

The system entails execution of a set of numerical forecasts using a lower-resolution model, starting from suitably perturbed initial conditions. T63 horizontal resolution and 19-level vertical resolution are used in the present experimental system. A control forecast is run starting from a truncated and interpolated form of the high-resolution operational analysis, and a further 32 forecasts are run from perturbed initial conditions. Execution of these forecasts is the most computationally demanding part of the EPS and is ideal for effective use of parallel processing. The perturbed forecasts are currently run in groups of four on the Centre's CRAY Y-MP/C90 computer, and the full set typically requires just a little longer to complete than the single high-resolution ten-day forecast.

The initial perturbations are chosen to be those which give the most rapid growth away from the control forecast during the early forecast range when linear theory may still apply. This time interval, $t_1 \leq t \leq t_2$., is currently chosen to be two days. Knowing the forecast from the unperturbed analysis, $x = x_0(t)$, we consider perturbed forecasts, $x = x_0(t) + x'$. The 'tangent-linear' model gives $\underline{x}'(t_2) = \underline{\underline{R}}\ \underline{x}'(t_2)$. We seek structures $x'(t_1)$ which give large growth.

Specifically, we seek large $$\frac{[\underline{\underline{R}}\,\underline{x}'(t_1)\,,\ \underline{\underline{R}}\,\underline{x}'(t_1))}{[\underline{x}'(t_1)\,,\ \underline{x}'(t_1)]}$$

or equivalently, large $$\frac{[\underline{x}'(t_1)\,,\ \underline{\underline{R}}^{*}\,\underline{\underline{R}}\,\underline{x}'(t_1)]}{[\underline{x}'(t_1)\,,\ \underline{x}'(t_1)]}$$

where $\underline{\underline{R}}*$ is the adjoint of $\underline{\underline{R}}$. The required structures are the eigenvectors of $\underline{\underline{R}}*\underline{\underline{R}}$ with large eigenvalues. A discussion of these 'singular-vector' calculations and of the resulting structures is given by Buizza and Palmer[18]. Molteni et al[4] give details of how the singular vectors are combined and scaled to produce initial perturbations for the EPS that have magnitudes consistent with estimates of analysis error.

In practice, $\underline{\underline{R}}*\underline{\underline{R}}x'(t_1)$ is determined for given $x'(t_1)$ by forward integration of the tangent linear model followed by backward integration of the adjoint model. A Lanczos algorithm is used to compute a comprehensive set of singular vectors, and it requires repeated evaluations of $\underline{\underline{R}}*\underline{\underline{R}}$. The cost is not substantial in the current routine configuration in which a coarse T21 truncation of the model is used. However, significantly faster growth has been found in sample calculations carried out at T42 resolution, and when adjoint methods are used to calculate the structure of analysis error most responsible for major forecast error (Rabier et al.[19]) it is found that a significant component is captured by T42 resolution but not T21. Some 70 or so iterations are necessary in the Lanczos algorithm to obtain a suitable set of singular vectors, and the cost is thus equivalent to that of several hundred days of forward integration of the forecast model. At T42, 19-level resolution the calculation requires a CP time which is about 30% of that used by the high-resolution ten-day forecast.

Notwithstanding the cost of calculation of the initial perturbations, the size of the ensemble and resolution of the forecast model are seen as the main factors determining the computational cost of the EPS for the immediate future. Given an increase in computing power, a balance has to be found between a larger ensemble and higher resolution. The number of amplifying singular vectors suggests that benefit may be gained from use of an ensemble size larger than 32, but the future ensemble size is likely to depend also on progress in the construction of the initial perturbations. An increase in horizontal resolution above T63 would provide better indications of possible extreme events, an improved spread of forecasts, and less discrepancy between the deterministic and ensemble forecasts associated with differences in model representations of orography. An increase to T106 resolution, keeping ensemble size unchanged, would require close to a factor of five increase in computational performance.

Estimation of analysis error covariance

A reliable estimate of the error covariance of the analysis is required for:

- improved definition of the background cost function J_b in variational data assimilation;
- preconditioning to reduce the number of iterations needed in the minimization in variational data assimilation;
- more realistic initial perturbations for ensemble prediction.

The determination of dynamic, flow-dependent estimates of analysis error is a developing research area. The basic approach of Kalman filtering is well established theoretically, but the computational requirement renders a full implementation intractable. However, Courtier[20] has set out ideas for progress in this field, and initial tests of these ideas have yielded promising results. Computationally, the approach involves repeated low-resolution integrations, based in part on a Lanczos algorithm (for which integrations have to be carried out in sequence) and in part on randomization (for which integrations can be carried out in parallel.)

Model refinement

Refinement of the model will be an important ongoing activity. Some changes will be primarily in support of the development of the data assimilation system. Examples include a possible increase in stratospheric resolution to improve the use of data in this region, and a possible introduction of ozone as a prognostic variable so that ozone-sensitive satellite measurements can be used in four-dimensional variational assimilation to improve lower stratospheric wind analyses. The production of adjoints of the physical parametrization routines has also to be addressed for the development of the variational assimilation. General development of the physical parametrizations will continue, to improve forecasts both of the large-scale atmospheric circulation and of specific weather elements such as surface temperature and precipitation. Some improvements may be implemented without a substantial increase in model cost, but others, for example relating to the frequency at which radiation is computed, may increase the cost more significantly. An increase of vertical resolution in the planetary boundary layer is a further possibility. However, there may be a further gain in efficiency from introduction of a two-time-level version of the semi-Lagrangian advection scheme. Overall, whilst there will be increases in the cost of the model, these are likely to be small compared with the requirements of variational data assimilation and ensemble prediction. This assumes no major increase in horizontal resolution for the deterministic forecast, although such an increase may become justified if there is sufficient improvement in the accuracy of initial conditions or demand for more detailed weather products.

The computational requirement

An indication has been given of the current computational requirements of the operational analysis, the high-resolution medium-range forecast, and the daily experimental ensemble prediction system. The distribution of computation between these applications is approximately:

Data assimilation	*20%*
Main ten-day forecast	*40%*
EPS	*40%*

Taking into account that the data assimilation cost includes that of the background forecasts, almost 90% of the computation is accounted for by model integrations, just over half of these at the high resolution of the main medium-range and background forecasts.

Estimates have been made of the minimum computation likely to be needed for the operational implementation of four-dimensional variational assimilation (4DVAR). The requirement is for an increase by a factor of at least five over that currently measured for the forecast model. This should also enable a significant improvement of the EPS. Although details of the initial implementation will depend on experimentation yet to be carried out, a possible distribution of computing power for an enhanced operational forecasting system in 1997 is:

4DVAR	*50%*
Main ten-day forecast	*10%*
EPS	*40%*

The cost of this forecasting system will be dominated by the cost of running the model code in either forward or adjoint mode. A more relevant breakdown from the viewpoint of computational performance is in terms of the model versions to be run:

Integration of full T213 model	*20%*
Integration of near-adiabatic T106 model	*40%*
Parallel integrations of EPS model with resolution in the range T63 - T106	*40%*

This assumes that the inner loops of the incremental variational assimilation are run at the target resolution of T106, initially with only the simplest physical parametrization.

A further increase by a factor of at least two is required for significant further development of the forecasting system, in particular the planned improvement of the estimation of the analysis error covariance. As we look further ahead the configuration of the operational forecasting system becomes more uncertain, but the computing for the 1999 system may be distributed among model versions in the following ranges:

Integration of high-resolution forecast model	*10-30%*
Integration of intermediate-resolution model in inner loop of 4DVAR	*30-50%*
Parallel integrations of intermediate-resolution model in EPS	*30-50%*
Partly parallel integrations of lower-resolution model(s) to determine estimates of analysis error and EPS perturbations	*10-30%*

Conclusions

There is evidently still a significant requirement for higher-performance computing for medium-range weather forecasting. Although this requirement is becoming less dominated by the cost of running a single high-resolution forecast model, the cost of the future forecasting system remains dominated by the cost of running the forecast-model code. Although other components of the forecasting system are less demanding of computational performance, their parallelization cannot be neglected if the system is to be run on a massively parallel machine.

High performance has to be attainable across a range of model resolutions, and for adjoint as well as direct versions of the model. This poses more of a challenge than simply running the highest possible resolution, as high resolution lends itself to efficient integration either on a few powerful processors where long vector lengths are needed for efficiency, or on micro-processor systems where large numbers of processors are needed to provide the necessary power. The multiple lower-resolution forecasts of the ensemble prediction system can, however, be integrated efficiently in parallel if memory is sufficient.

The higher performance has to be affordable as well as deliverable across a range of model versions. Doubling the computer power applied to numerical weather prediction does not lead generally to a halving of forecast error, but more typically to rather small improvements in several aspects of the forecast. The larger numerical forecasting centres have traditionally used (and by implication been able to afford) "top-of-the-range" machines, or configurations not substantially below the maximum. This is becoming no longer the case, with scalable distributed-memory machines promising peak performance at the Tflops level and beyond, but at a price. The effort being devoted to cutting the cost of four-dimensional variational data assimilation is indicative of concern to make the most effective use of the affordable power.

Adrian Simmons

References

[1] K. Miyakoda, G.D. Hembree, R.F. Strickler and I. Shulman, *Mon. Wea. Rev.* **100** (1972) 836-855.

[2] A.J. Simmons, R. Mureau and T. Petroliagis, *Quart. J. Roy. Met. Soc.* **121** (1995) to appear.

[3] P. Courtier, J.-N. Thépaut and A. Hollingsworth, *Quart. J. Roy. Met. Soc.* **120** (1994) 1367-1387.

[4] F. Molteni, R. Buizza, T.N. Palmer and T. Petroliagis, *Quart. J. Roy. Met. Soc.* **121** (1995) to appear.

[5] V. Bjerknes, *Met. Z.* **21** (1904) 1-7.

[6] L.F. Richardson, *Weather Prediction by Numerical Process* (Cambridge University Press, 1922, republished by Dover, 1965) 236pp.

[7] P. Lynch, in *The Life Cycles of Extratropical Cyclones, Vol I*, ed. S. Grønås and M.A. Shapiro (ISBN 82-419-0144-5, 1994).

[8] J.G. Charney, R. Fjörtoft and J. von Neumann, *Tellus* **2** (1950) 237-254.

[9] F.G. Shuman, *Wea. Forecasting* **4** (1989) 286-296.

[10] A.J. Simmons, D.M. Burridge, M. Jarraud, C. Girard and W. Wergen, *Met. Atmos. Phys.* **40** (1989) 28-60.

[11] A.J. Simmons and D. Dent, *Comp. Phys. Rep.* **11** (1989) 165-194.

[12] H. Ritchie, C. Temperton, A. Simmons, M. Hortal, T. Davies, D. Dent and M. Hamrud, *Mon. Wea. Rev.* **123** (1995) 489-514.

[13] Proceedings of Sixth ECMWF Workshop on Use of Parallel Processors in Meteorology, 21-25 November 1994.

[14] A.C. Lorenc, *Mon. Wea. Rev.* **109** (1981) 701-721.

[15] F.-X. Le Dimet and O. Talagrand, *Tellus* **40A** (1986) 97-110.

[16] O. Talagrand and P. Courtier, *Quart. J. Roy. Met. Soc.* **113** (1987) 1311-1328.

[17] Proceedings of Sixth ECMWF Workshop on Use of Parallel Processors in Meteorology, 21-25 Nov 1994

[18] R. Buizza and T.N. Palmer, *J. Atmos. Sci.* **52** (1995) to appear.

[19] F. Rabier, E. Klinker, P.Courtier and A. Hollingsworth, *submitted to Quart. J. Roy. Met. Soc.*.

[20] P. Courtier, in *Proceedings of 1993 ECMWF Seminar on Developments in the Use of Satellite Data in Numerical Weather Prediction,* 189-207.

Upgrade to the chilled water system

Background

The Centre's original chilled water system was installed 18 years ago when ECMWF Headquarters were built. In 1990, in readiness for the installation of the CRAY Y-MP8, the capacity of the system was increased from the original 1,100 kW to 1,475 kW by the addition of an extra chiller and by upgrading the pumps. This upgrade increased the system pressure in the pipework to around 6.5 bar compared to the maximum design pressure of 10 bar.

The original pipework was mostly plastic (ABS). The ABS pipework had installation advantages in that it took less time and required lower levels of skill to install. However, the system was designed without effective isolation valves so that making modifications or repairs to the system required extensive downtime as chemical weld curing times were of the order of 12 - 20 hours. During the period up to 1993 a small number of leaks in the system clearly demonstrated where the problems of maintenance were!

During the parallel run of the CRAY X-MP4 and the Y-MP8 in 1990, and of the CRAY Y-MP8 and CRAY C90 in 1992, it was clear that the chilled water system would not have sufficient capacity to handle a parallel run when the C90 was replaced. In 1992, it was necessary to switch off equipment to allow the essential computer service to continue.

Specifications

Early in 1993 a study was undertaken to see how best to satisfy the Centre's cooling requirements for the 15-20 year period from 1995. The requirements were:

- to upgrade the maximum capacity of the system to around 2,500 kW;
- to replace the ABS pipework system which was prone to problems by a pipework system which would facilitate expansion and repair;
- that the new system should be as energy efficient as possible; and
- that the work should require the minimum amount of downtime for the computer service.

During 1993 a consulting company was employed to propose how the requirements might best be met. It was clear that the most demanding requirement was the ability to upgrade the system without significant downtime to the computer service. Difficulty arose because the capacity and energy efficiency requirements meant that the main pipework had to be increased in diameter from 200 mm to 300 mm. The lack of space in the plant rooms meant that a very complex three-dimensional geometry problem had to be solved.

Eventually a possible, though expensive, method was devised using temporary pipework to maintain the existing service during installation of the new pipework. The plan then was to keep the old chillers and cooling towers, adding to them when necessary to increase the capacity. Under this plan the capacity of the old system would be reduced, to allow new pipework to be installed. This capacity reduction was acceptable, providing it could be done during the winter period when cooling needs were at their lowest.

At that point the condition of the old chillers and cooling towers was closely examined and it became clear that time had taken its toll. Expensive refurbishment would be required, at a cost comparable to that of replacing the equipment. Replacement, however, would take a total of around 13 months - a period in excess of the time available.

In the spring of 1994 the Centre's staff re-examined the problem and came up with the following outline plan which is now being implemented:

- two of the original five chillers and one cooling tower will be removed and the cooling tower compound extended to make space for new pipework and replacement chillers;
- new packaged air cooled chillers with a capacity of 1,200 kW plus a further 600 kW standby will replace the old water cooled chillers;
- a much improved control and monitoring system will be installed;
- the new chillers will initially be connected to the old pipework, allowing the old chillers to remain in place for a 30-day parallel run; and
- the electrical distribution arrangements for the ancillary services will be modernised.

The amount of external noise to be produced was also seen to be a problem since the Centre had earlier received complaints about excessive noise from its old diesel standby generators. (This problem was resolved when the Eurodiesel system was installed.) A noise survey was undertaken VERY early one morning in August 1994 and the specification was altered to set maximum noise limits. The equipment subsequently selected has special noise suppression facilities.

The Centre issued Invitations To Tender in summer 1994 for the upgrading of the chilled water pipework and for replacement chillers. Contracts were awarded to a local Reading company *(G.H. Marshall & Co).*

The work

Work began immediately on signature of contracts in December 1994, and on 17 December 1994 the first service shutdown took place when two chillers and one cooling tower were disconnected so that they could be removed to provide space for the new pipework and pumps.

Detailed planning went on during December to find realistic routes for the new pipes and electrical supplies. The challenge was to install the new pipework in a space already full of pipes and equipment. At the end of the operation, when the old pipes and equipment have been removed, anyone looking at the mass of pipes and equipment squeezed into one end of the plant room will wonder what the problem was, since there will be plenty of apparently unused empty space!

Early on 12 January 1995, a temporary pipe connection made on 17 December broke due to ineffective fastening in an enclosed space. Around 1,200 litres of antifreeze were lost down the surface water drain and an unscheduled shutdown of a couple of hours took place. The immediate problem was rapidly fixed but there was the potential problem of the antifreeze polluting the local environment. The UK National Rivers Authority Pollution Department were contacted and their staff undertook tests to see whether damage had been done. Fortunately, the leak had been a slow one over a number of hours, and the antifreeze had been diluted by rainwater in the drain so that no measurable pollution had occurred.

Installation work has continued during the first half of 1995, with various minor problems being encountered and solved. These included:

- a chilled water buffer tank, 2,400 mm high, 2,200 mm diameter and weighing 1,500 kg had to be installed in an area with only a few centimetres clearance. Skill and brute force exerted over a period of several days finally got the tank positioned.
- The chillers (supplied by Carrier) were subject to last minute delays due partly to a customs drugs search at the ferry port and partly due to a failing in communications.
- The building works to construct the new chiller compound were affected by the very wet winter in the UK. The new chillers were delivered in mid-February and on 18 April they were turned on, taking the Centre's entire cooling load. On 4 May the CRAY T3D and Y-MP2E were moved to new Computer Hall circuits and on 18 May the IBM ES9000-720 was moved to the new circuit. On 20 May a further shutdown took place at which time the CRAY C90 and the air conditioning equipment were moved to the new Computer Hall circuit and work could begin to remove the remaining old equipment. The project is scheduled to be completed by the end of June 1995.

Conclusion

At the time of writing the new equipment has only been in service for a few weeks. However, it is already clear that substantial energy savings will be made with the new system and that the new control and monitoring system will provide a valuable management tool for the future.

The foundation has been laid to bear future increases in capacity to between 2,500 kW and 3,000 kW cooling, depending on requirements. When the result of the invitation to tender for the C90 replacement is known, plans will be made to add chilling capacity, if necessary, in readiness for its installation in 1996.

Peter Gray, Michael O'Brien

Table of TAC Representatives, Member State Computing Representatives and meteorological contact points

Member State	TAC Representative	Comp. Representative	Met. Contact Point
Belgium	Dr W Struylaert	Mrs L Frappez	Dr J Nemeghaire
Denmark	Dr P Aakjæer	Mr N Olsen	Mr G R Larsen
Germany	Dr B Barg	Dr B Barg	Dr Rüge
Spain	Mr T Garcia-Meras	Mr E Monreal Franco	Mr R Font Blasco
France	Mr J Goas	Mr D Birman	Mr J Goas
Greece	Mr D Katsimardos/ Dr G Sakellarides	Dr G Sakellarides	Dr N Prezerakos/ Mrs M Refene
Ireland	Mr J Logue	Mr L Campbell	Mr T Sheridan
Italy	Dr S Pasquini	Dr G Tarantino	Dr G Maresca
Yugoslavia*			
Netherlands	Mr S Kruizinga	Mr S Kruizinga	Mr G Haytink
Norway	Mr K Björheim	Mrs 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	Mrs C Raeber	Mr M Schönbächler
Finland	Dr M Alestalo	Mr T Hopeakoski	Mr P Nurmi
Sweden	Mr L Moen	Mr S Orrhagen	Mr O Åkesson
Turkey	Mr M Örmeci	Mr M Örmeci	Mr M Örmeci
United Kingdom	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).*

Fifth workshop on meteorological operational systems

Introduction

The planned biennial Workshop on Meteorological Operational Systems, to be held at ECMWF 13-17 November 1995, will be the fifth in the series.

The workshop will review the state of the art of meteorological operational systems and address future trends in the use of medium-range forecast products, data management and meteorological applications on UNIX workstations.

Use and interpretation of medium-range forecast guidance

The session will address the problems and solutions related to the use of numerical guidance in medium-range weather forecasting. With the current developments in numerical modelling, the medium-range forecaster is faced with yet a new challenge: to prepare weather forecasts based on the output from one or several high resolution global models as well as the output from ensemble prediction systems, which can be exploited through appropriate operational procedures to estimate in advance the uncertainty in the numerical forecast guidance.

Operational centres will present their approaches to medium-range weather forecasting and report on their experiences with a combined use of output from high resolution deterministic models and from ensemble prediction systems. The concept of probability forecasts and tailored products for certain categories of end users may also be addressed and will be further discussed in a working group.

Operational data management systems

The ever increasing use in meteorological applications of UNIX systems and associated software tools and utilities to provide data storage and access methods means that for users the distinction between databases and archives is becoming unclear. Operational database and archive systems will be reviewed showing how use of standard software utilities and languages facilitates the development and portability of such systems. Special attention will be given to ease of use and speed of access to data, as well as data transformations performed by retrieval systems on behalf of the users. The current role of commercial databases in meteorological applications will also be reviewed.

Meteorological UNIX workstation applications

Much progress has been seen in this area since the previous workshop in 1993. Some meteorological UNIX workstation applications have become available and others are being developed. Current and planned meteorological workstation applications will be presented and demonstrated at the exhibition during the workshop.

De facto standards are being used in some areas, e.g. the X-Window System and Motif, whereas elsewhere different solutions are developed. A working group will discuss areas where no standardisation has emerged including different user interface paradigms, operational aspects, data access methods, handling of repetitive tasks and 2D/3D presentation techniques.

Horst Böttger

ECMWF Annual Seminar: Predictability

4-8 September 1995

The ECMWF 1995 seminar will be a pedagogical presentation of both theoretical and practical issues for the study of the predictability of the atmosphere. With an opening lecture by Professor E.N. Lorenz, the seminar will cover topics such as the predictability of weather regimes and their transitions, ensemble prediction schemes, and predictability of the coupled ocean-atmosphere system. Specific questions such as similarities and differences between breeding vectors and singular vectors will also be addressed.

The Seminar forms part of our educational programme and is aimed at young, mainly post graduate/ post doctoral scientists in the ECMWF Member States. Posters and application forms are mailed to the national meteorological services and major universities in the ECMWF Member States.

Further information can be obtained from e.kooij@ecmwf.int.

ECMWF Publications

Technical Memoranda:

No. 211 A spurious mode in the "Lorenz" arrangement of ϕ and T which does not exist in the 'Charney-Phillips' arrangement.

No. 213 Treatment of the Coriolis terms in semi-Lagrangian spectral models.

No. 214 The anomalous rainfall over the USA during July 1993: Sensitivity to land surface parametrization and soil moisture anomalies.

European Space Agency Contract Report:

The ECMWF Contribution to the Characterisation, Interpretation, Calibration and Validation of ERS-1 Scatterometer Backscatter Measurements and Winds, and their use in Numerical Weather Prediction Models.

Executive summary:

European Space Agency Contract Report: The ECMWF Contribution to the Characterisation, Interpretation, Calibration and Validation of ERS-1 Scatterometer Backscatter Measurements and Winds, and their use in Numerical Weather Prediction Models.

Forecast and Verification Charts to March 1995

ECMWF Calendar 1995

26 April - 23 June	Meteorological training course: **Met 1** (26 April - 5 May): *Numerical methods, adiabatic formulation of models* **Met 2** (9 - 19 May): *Parametrization of diabatic processes* **Met 3** (22 May - 2 June): *Data assimilation and use of satellite data* **Met 4** (5 - 9 June): *General circulation, systematic model errors and predictability* **Met 5** (12 - 23 June): *Use & interpretation of ECMWF products*
4 - 5 July	Council, 42nd session
28 August	*ECMWF holiday*
4 - 8 September	Seminar: Predictability
11 - 13 September	Workshop: Non-linear aspects of data assimilation
25 - 27 September	Scientific Advisory Committee, 24th session
11 - 13 October	Technical Advisory Committee, 22nd session
17 - 18 October	Finance Committee, 55th session
24 - 27 October	Workshop: Surface fluxes (WGNE)
30 Oct - 3 Nov	WGNE meeting
6 - 8 November	Workshop: Semi-Lagrangian methods
13 - 17 November	Workshop: Meteorological Operational Systems
4 - 5 December	Council, 43rd session
25 - 27 December	*ECMWF holiday*

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.

	No	Date	Page
GENERAL			
ECMWF publications - range of	26	June 84	16
COMPUTING			
Data handling system	57	Mar 92	16
CRAY T3D			
- *system overview*	67	Autumn 94	17
- *programming*	68	Winter 94/5	29
ECMWF Documentation			
- *Current Computer Bulletins*	64	Dec 93	46
- *On line*	64	Dec 93	44
(Exabyte) 8mm data cartridge service	67	Autumn 94	36
Fortran 90 features in CF77	60	Dec 92	23
&	61	Mar 93	29
&	62	June 93	26
&	63	Sept 93	33
&	65	Spring 94	20
&	66	Summer 94	34
MAGICS - the ECMWF graphics package	62	June 93	15
Massively parallel computing - ECMWF's current investigations	61	Mar 93	15
METVIEW	68	Winter 94/5	9
Multitasking ECMWF spectral model	60	Dec 92	3
Networks			
- *ECMWF's internal network*	54	June 91	20
- *New LANs at ECMWF*	59	Sept 92	20
POSIX shell under UNICOS 8	66	Summer 94	17
Supervisor Monitor Scheduler (SMS)	59	Sept 92	13
Telecommunications			
- *The ECNET system*	54	June 91	16
- *Digital links to Member States*	57	Mar 92	19
Workstations at ECMWF	55	Sept 91	7
&	61	Mar 93	21
Y-MP/C90:			
- *An introduction*	57	Mar 92	10
- *I/O optimisation techniques*	64	Dec 93	21
METEOROLOGY			
Comparison between SSM/I and ECMWF total precipitable water	57	Mar 92	3

	No	Date	Page
Data acquisition - ECMWF's new system	46	June 89	21
Development of the operational 31-level T213 version of the ECMWF forecast model	56	Dec 91	3
Envelope orography - discussion of its effects	33	June 86	2
ECMWF Analysis			
- *New version of analysis system*	35	Sept 86	16
- *Divergent structure functions*	42	June 88	2
- *Revised use of satellite data*	39	Sept 87	4
- *The variational analysis scheme - main features and some early results*	62	June 93	5
- *Use of TOVS satellite data at ECMWF: a new approach*	61	Mar 93	3
ECMWF Preprocessing - new scheme	43	Sept 88	3
ECMWF Re-analysis (ERA) Project - plans and current status	64	Dec 93	3
Ensemble prediction	58	June 92	5
Ensemble prediction system -			
- *expert meeting*	63	Sept 93	18
- *status and plans*	65	Spring 94	3
ERS-1 mission	54	June 91	8
Evaporation from tropical oceans	51	Sept 90	3
Forecast model			
- *Cloud cover scheme*	29	Mar 85	14
- *Convection - parametrization of*	43	Sept 88	6
- *Increased resolution - studies of*	38	June 87	10
- *Initial conditions - the spin-up problem*	39	Sept 87	7
- *New surface/boundary layer formulation*	63	Sept 93	3
- *Parametrization of gravity wave drag*	35	Sept 86	10
- *Revisions to physics*	46	June 89	3
- *Revision of the clear-sky and cloud radiative properties*	61	Mar 93	3

	No	Date	Page
Global forecast experiment at T213 resolution	41	Mar 88	3
Good prediction of a severe storm over southern Sweden	50	June 90	10
GORBUSH - a storm in the Mediterranean Sea	53	Mar 91	4
MARS - the ECMWF meteorological archival and retrieval system	32 & 33	Dec 85 Mar 86	15 12
Meteorological applications at ECMWF utilizing EMPRESS	64	Dec 93	11
Minimum temperature forecasts at the Regional Meteorological Service of the Emilia Romagna region (N. Italy) by the application of the Kalman filter technique	60	Dec 92	9
Monte Carlo forecast	49	Mar 90	2
Ocean Wave Forecasting in the Mediterranean Sea	68	Winter 94/5	3
Performance of the ECMWF model in tropical cyclone track forecasting over the western north Pacific during 1990-1991	58	June 92	16
Potential vorticity maps at ECMWF	50	June 90	3
Recent verification of 2m temperature and cloud cover over Europe	54	June 91	3
Skill forecasting - experimental system	40	Dec 87	7
Surface wind observations from the ERS scatterometers	66	Summer 94	3
Systematic errors - investigation of, by relaxation experiments	31	Sept 85	9
Use of reduced Gaussian grids in spectral models	52	Dec 90	3

Useful names and 'phone numbers within ECMWF

	Room*	Ext.**
Director		
David Burridge	OB 202	2001
Deputy Director and Head of Research Department		
Anthony Hollingsworth	OB 119A	2005
Head of Operations Department		
Massimo Capaldo	OB 010A	2003
Advisory:		
Available 9-12, 14-17 Monday to Friday		2801
Other methods of quick contact: Telefax (+44 1734 869450) VMS MAIL addressed to ADVISORY Internet mail addressed to Advisory@ecmwf.int		
Registration		
Project Identifiers - Pam Prior	OB 225	2384
User Identifiers - Call Desk	CB 022A	2315
Computer Operations		
Console - Shift Leaders	CB Hall	2803
Console fax number +44 1734 499 840		
Fault reporting - Call Desk	CB 022A	2303
Service queries - Call Desk	CB 022A	2303
Tape Requests - Tape Librarian	CB 022A	2315
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 Section 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			
Computer Co-ordinator			
David Dent	OB 123	2702	

---

* CB - Computer Block
OB - Office Block

** The ECMWF telephone number is READING (01734) 499000, international +44 1734 499000. Direct dial is available to (01734) 499 + last three digits of individual extension number, e.g. the Director's direct number is (01734) 499001.

DEC mail

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

Internet

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