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

Two plots from a recent observing system experiment, showing Meteosat water vapour channel data (bottom panel) and the corresponding simulated model output (top panel).

Such data will prove valuable when the 4D-var analysis is implemented as it is available every ½ hour.

(False colours – blue dry – red wet – white clouds)

See the article on page 2.

Editorial

ECMWF regularly reviews the impact that various observing system data has on the forecast. A recent series of such observing system experiments are reported in the first article (page 2).

In January 1997, there were severe floods in Greece due to torrential rainfalls over 12/13 January. The ECMWF forecast for that period is discussed on page 9.

Three more reports in the ECMWF Technical Memoranda series are summarised on pages 7,8 and 9.

The relentless march of computer technology improvements over the years has been, and still is, impressive. The article on page 12 highlights this; it reviews the changes that have occurred in the computer equipment at ECMWF over the past 17 years. As part of these ongoing changes in the Centre's computer equipment, the Fujitsu VPP700 will shortly be upgraded, the article on page 14 summarises the changes.

Finally, remember that ECMWF's telephone number is changing (see box on the inside front cover), this is due to an area code change for Reading.

Influence of observations on the operational ECMWF system

Introduction

The global meteorological observing system is extremely expensive and in the present economic climate conventional observations such as radiosondes are beginning to be severely reduced. At the same time improved satellite systems are becoming available. The operational observing network, which uses both conventional and satellite measurements, influences how accurately the initial atmospheric state can be prescribed and therefore to a large extent the resulting forecast accuracy. There is therefore an urgent need to investigate the importance of different observing systems on Numerical Weather Prediction performance.

In this work we quantify through Observing System Experiments (OSEs), the contribution made by the main ground-based and satellite-based operational systems to medium range forecasting. In an OSE the impact of a specified observing system is assessed by comparing extended data assimilation and regular forecasts based on the total operational system with those generated excluding the particular observing system under investigation. The value of a new or experimental observing system can be assessed in a similar manner.

In a previous study Kelly et. al. (1993) performed a series of OSEs using the then operational ECMWF system based on optimal interpolation (Lorenc; 1981). They used a baseline observing system consisting only of in-situ data, comprising radiosondes, aireps, synops, ships and buoys. OSEs were used to assess the impact of adding satellite derived temperature and humidity information (SATEMs) only, cloud motion winds (SATOBs) only, and SATEMs plus SATOBs to the baseline system. They found that adding SATEMs or SATOBs alone improved the forecast, but that adding both SATEMs and SATOBs gave less improvement than adding just one of those observing systems. These unsatisfactory results triggered an appraisal of the use of data in the O/I system and helped motivate the decision to develop a 3D-Var system which could analyse the Tiros Operational Vertical Sounder (TOVS) radiances directly, together with all other data. The purpose of this paper is to repeat and extend the earlier experiments, but this time using the 3D-Var assimilation system which became operational in early 1996. These new results, described in subsequent sections, demonstrate that the main observing systems are contributing in important ways to medium range forecasts in the Northern Hemisphere, in the Tropics, and in the Southern Hemisphere.

Experiment design

Two series of Observing System Experiments (OSEs) were run for periods in December 1996 and February 1997. These experiments used the operational system of the 15th May 1997, which includes a revision to the background cost function of the variational analysis (Andersson et. al. ; 1994, 1996), (Bouttier; 1997).

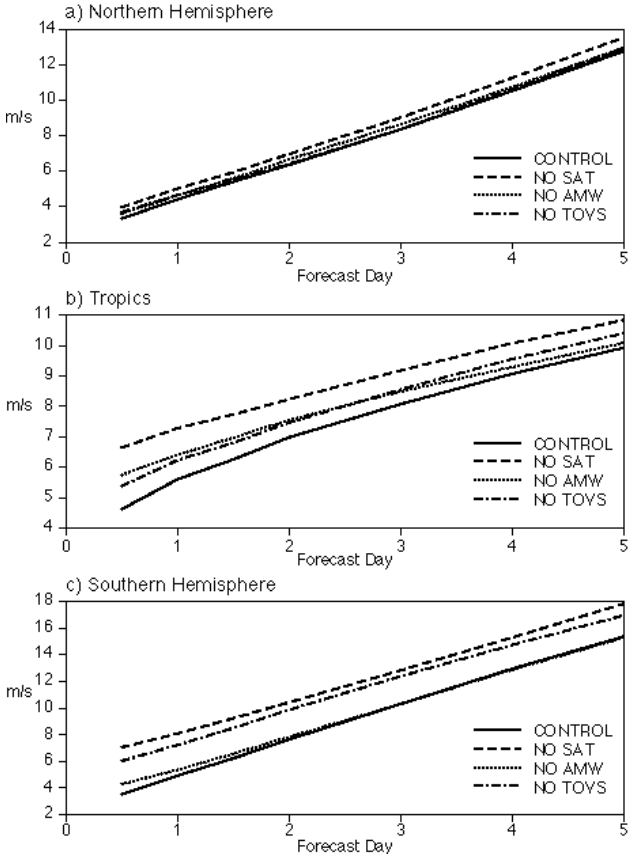

Fig. 1: *Satellite OSE 200hPa vector wind root mean square forecast for a) Northern Hemisphere, b) Tropics and c) Southern Hemisphere*

In the earlier experiments (Kelly et. al.; 1993) problems were revealed with the use of SATEMs in the Northern Hemisphere (Kelly and Pailleux; 1988) and as a consequence these satellite data were removed from the Northern Hemisphere and Tropics until the introduction of 1D-Var (Eyre et. al.; 1993). The interpretation of OSEs is not always straight-forward. The value of an observing system is most easily demonstrated when an energetic event seen only by one observing system occurs in the area being observed. A case study demonstrating this is included.

The new OSE experiments systematically removed the following observing systems from the full operational system:

- Satellite clear radiance data from the TOVS satellites (NOTOVS),
- Geostationary Atmospheric Motion Winds (AMWs) from cloud and water vapour (NOAMW),
- Radiosonde wind, temperature and humidity data (NORAOB),
- Aircraft winds and temperatures (NOAIREP),
- The combined removal of both TOVS and AMWs (NOSAT).

All of these experiments have been compared to the full operational system (CONTROL).

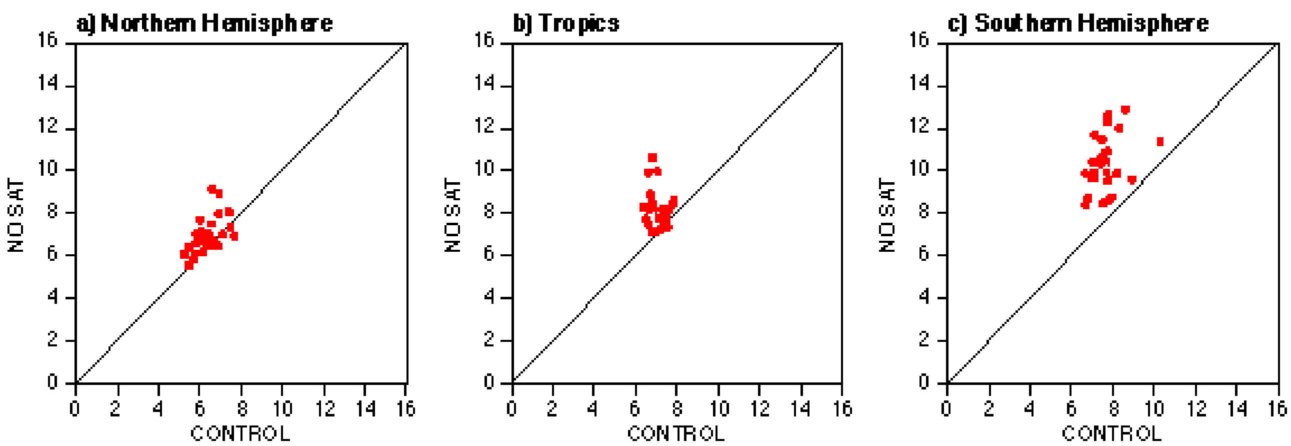

Fig. 2: *Satellite OSE (CONTROL v NOSAT) 200hPa vector wind root mean square 48 hour-forecast scatter plots for a) Northern Hemisphere, b) Tropics and c) Southern Hemisphere.*

The results are grouped into two sections, first the satellite measurements and second the conventional upper-air measurements. All verification statistics use the operational ECMWF analysis. Similar results were obtained for the two periods, 5th to 20th December 1996 and 1st to 14th February 1997, therefore the two experiments have been combined to give a 29 day set of forecast scores.

Results

Satellite OSEs

A series of four experiments have been included:

- CONTROL
- NOTOVS
- NOAMW
- NOSAT

The ERS scatterometer has been included in all experiments. Its impact will be discussed in a future article.

Forecast wind impact at 200 hPa

Aviation forecasts which use upper level winds are a major output product of NWP and the impact of satellite data on the upper level winds is illustrated. Figure 1 shows the overall impact of satellite data and it varies in the short range from ⅓ of a day in the Northern Hemisphere, to 1½ days in both the Tropics and the Southern Hemisphere. AMWs have most value in the Tropics but give a significant improvement in the Northern Hemisphere. TOVS have a significant impact in the Tropics and large impact in the Southern Hemisphere.

The scatter of these forecast wind rms errors for the NOSAT v CONTROL at day two is shown in Figure 2. Almost all forecasts are positive for the CONTROL. In all cases the CONTROL gives a smaller rms error and the largest deterioration in the NOSAT forecast varies from 2 m/s in the Northern Hemisphere to 5 m/s in the Tropics.

Forecast height impact at 500 hPa

Figure 3 shows the performance of the Operational ECMWF Forecast Model. The scores calculated are for 500hPa geopotential height anomaly correlation. The NOSAT experiment shows a negative impact in all areas. This is an important finding, as it is often suggested that it is not possible to demonstrate the impact of satellite data, especially in the Northern Hemisphere. This result comprises both the effect of AMWs and also of TOVS clear radiance data. In the case of AMWs there is a positive impact in the Northern hemisphere which must be largely influenced by the Tropics where the bulk of the data is found. TOVS data however has only a small impact in the Northern Hemisphere, its main impact being in the Southern Hemisphere where it improves the medium range forecast by up to 1½ days. In the presence of TOVS data the AMWs provide little additional benefit in the Southern Hemisphere. If TOVS data is not present then the AMWs have a positive impact.

Conventional upper air OSEs

The experiments included in this group are:

- CONTROL which uses all data,
- NORAOB which excludes radiosonde winds, temperatures and humidity data,
- NOAIREP which excludes aircraft wind and temperature observations.

For the purposes of comparison the NOAMW experiment will also be considered in this group in order to assess the relative importance of the AIREP and AMW wind observing systems.

Forecast wind impact at 200 hPa

The impact on the 200 hPa winds is shown in Figure 4. As with the satellite system experiments, all verification is based on the ECMWF operational analysis.

In the Northern Hemispheric the scores are dominated by the radiosondes, with their exclusion (NORAOB) the forecast accuracy is reduced by one day. The AMWs and AIREPs have a lesser effect and both have about equal weight, each degrading the forecast by ⅓ of a day if omitted.

In the Tropics, Radiosondes and AMWs have comparable impact of about ⅔ of a day. Removing the AIREPs has a smaller but still negative impact. Currently there are few Radiosondes in the Tropics and this number is being reduced. Such a reduction will adversely affect forecast skill.

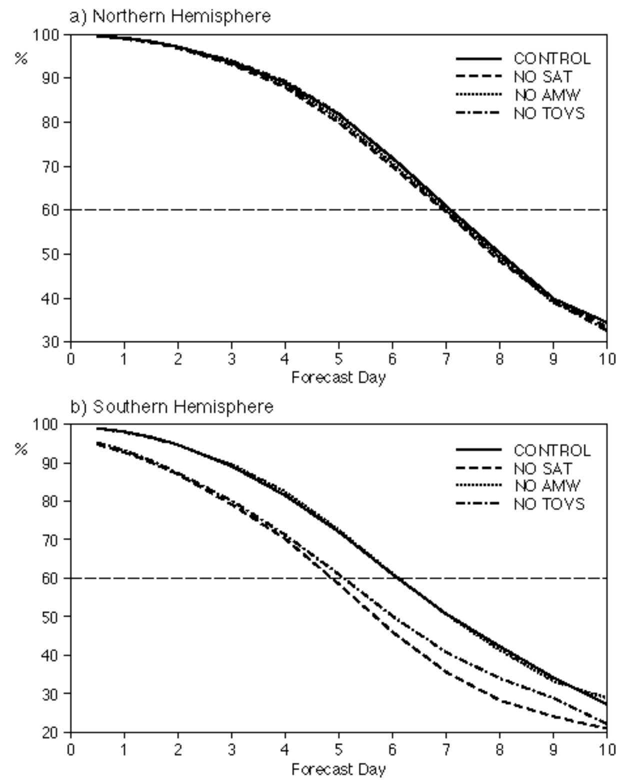

***Fig. 3:** Satellite OSE 500hPa anomaly correlation forecast for a) N. Hemisphere and b) S. Hemisphere.*

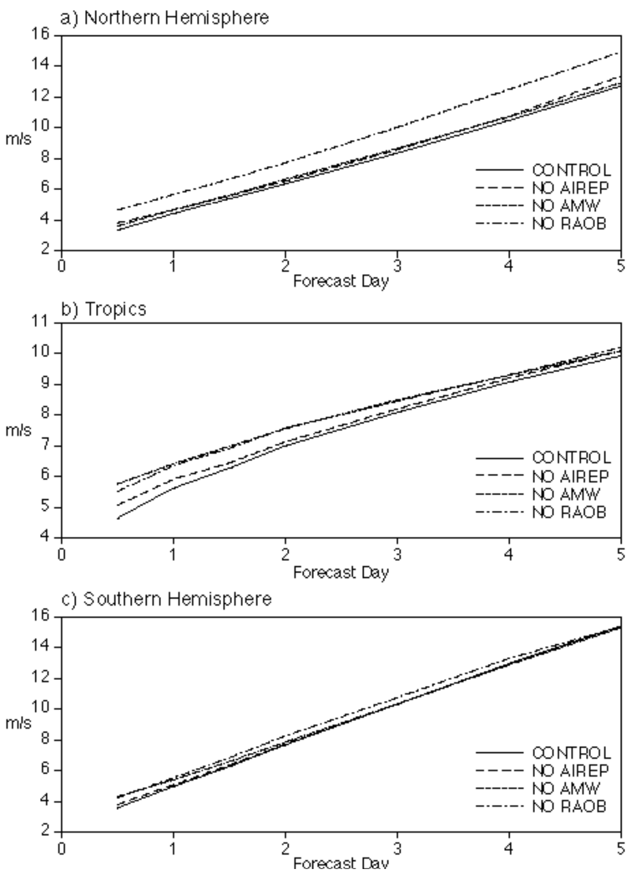

***Fig. 4:** Conventional upper air OSE 200hPa vector wind root mean square forecast for a) N. Hemisphere, b) Tropics and c) S. Hemisphere.*

As seen previously TOVS data dominates the forecast impact in the Southern Hemisphere. Removing them degrades the forecast skill by around 1½ days. Excluding the Radiosonde data in this region reduced the forecast skill about ⅓ of a day and AMWs and AIREPs have only a small impact.

Scatter plots are shown for NORAOB v CONTROL 48 hour forecasts (figure 5) and almost all points in the three regions show a negative impact due to the exclusion of Radiosondes. The largest radiosonde impact is in the Northern Hemisphere and has a magnitude of about 3 m/s.

Forecast height impact

Figure 6. shows the 500 hPa anomaly correlation scores In the Northern Hemisphere the forecast accuracy is dominated by the radiosondes without which the five day forecast skill is reduced by about 1 day. The Radiosondes have less impact in the Southern Hemisphere. The influence of AMWs and AIREPs is somewhat smaller but still positive.

Synoptic case

After the discussion of objective scores it is of interest to look at the synoptic impact. A case has been selected during the February experiment in which strong cyclogenesis occurred in the North Atlantic. The base date of this situation is 12 UTC on the 10th of February 1997. As discussed, the OSE experiments which have the most impact are NORAOB and NOSAT and forecasts from these have been compared with the control and the verifying operational analysis.

The 48 hour forecasts and the verifying analysis are shown in Figure 7. Large errors are clearly seen in both the NORAOB and NOSAT south of Newfoundland. Both of these forecasts fail to deepen the low pressure system which is captured well in the control. In order to capture this development both the radiosonde and satellite observations are required. Another region in which the NORAOB experiment is further degraded in comparison with the remaining experiments is in the complex low pressure system extending from the mid-Atlantic, over Scotland and to East of the Baltic. In this experiment the system is much weaker in intensity than the control or NOSAT. The low pressure system in this complex near Scotland is also poorly forecast.

Moving to the 96 hour forecasts, shown in Figure 8, the low pressure system south of Newfoundland at day two has moved north-west and deepened. The control forecast provides the best representation of this low pressure system. It deepens the low considerably more than the other experiments even though its central pressure is too high and its position has a two-degree error. The complex low pressure seen in the 48 hour forecast has now developed into four main low pressure systems. Large errors in the NORAOB forecast are evident providing a poor forecast for Europe. The NOSAT forecast is a little better but there is a large error in the system south of Ireland. In comparison the control forecast develops this low pressure system very well.

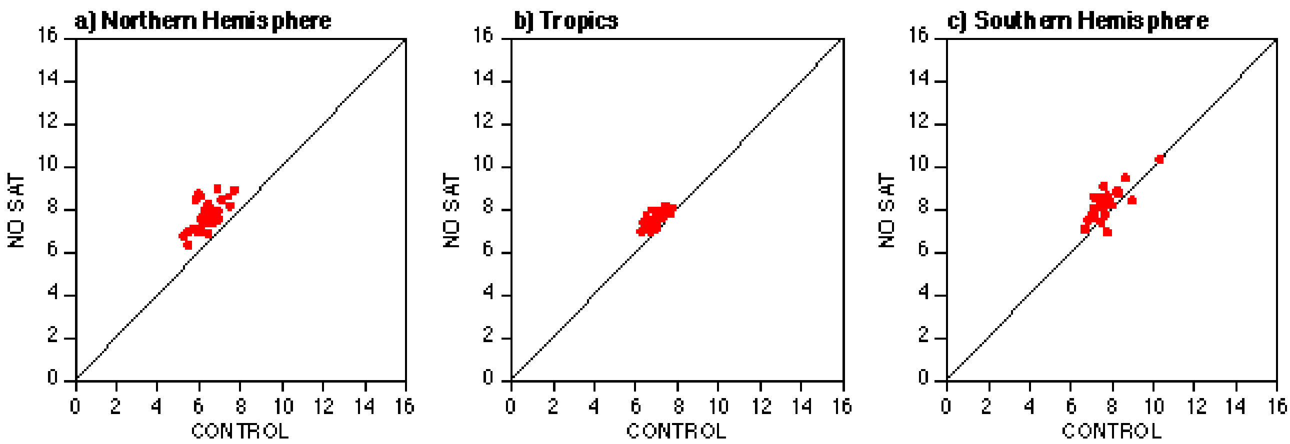

Fig. 5: Conventional upper air OSE *(*CONTROL *v* NORAOB*) 200hPa vector wind root mean square 48 hour-forecast scatter plots for a) Northern Hemisphere, b) Tropics and c) Southern Hemisphere.*

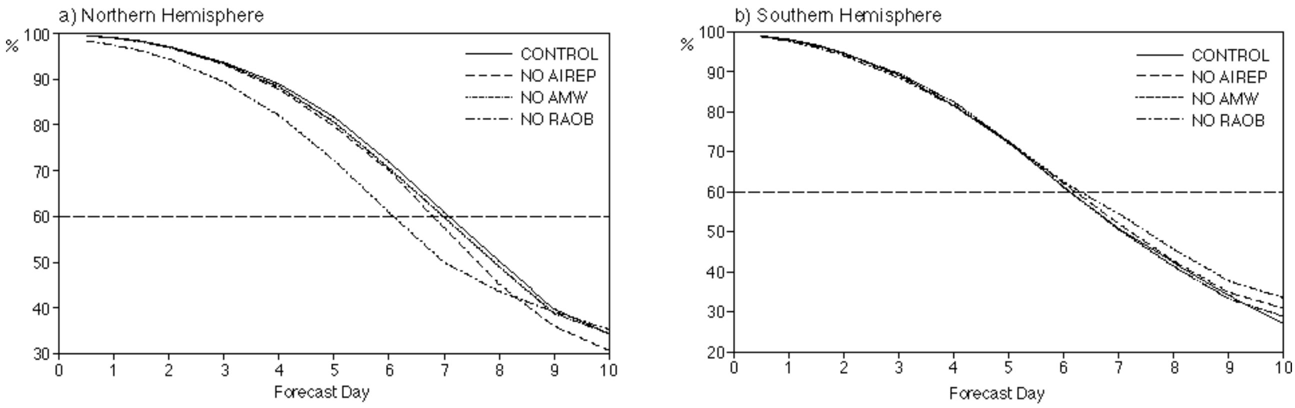

Fig. 6: Conventional upper air OSE *500hPa anomaly correlation forecast for a) Northern Hemisphere and b) Southern Hemisphere.*

Conclusions and recommendations.

In general the results obtained in this set of experiments are encouraging. The current operational 3D-Var system has been shown to benefit from the assimilation of both satellite data and conventional observations and broadly speaking its performance in each of the Northern Hemisphere, the Tropics and the Southern Hemisphere is satisfactory. It is clear that in some regions there is a degree of redundancy in the current observing system but this is necessary to provide coverage in the event of occasional failures in parts of the observing system. However the inclusion of each data type almost always improves the forecast system which was not invariably the case in the past.

The satellite OSEs show that AMWs have a positive impact particularly in the Northern Hemisphere and the Tropics. High space and time resolution AMWs are now being produced and their impact should be tested. Alternatively, the TOVS clear radiances provides benefits in the Southern Hemisphere and the Tropics but have only a small impact in the Northern Hemisphere. In this region the Radiosonde network is particularly important in improving forecast accuracy. However if the current trend to reduce this network continues it will soon result in a reduction in forecast accuracy. The large impact of TOVS in the Southern Hemisphere suggests that increasing the use of TOVS in the Northern Hemisphere could become a priority with the continuing decline of the Radiosonde network. Finally, TOVS data provided a large positive impact on the Tropical winds, a result which has not been reported from previous experiments.

References

Andersson E.,Haseler J., Undén P., Courtier P., Kelly G.,Vasiljevic D.,Brankovic C.,Cardina C., Gaffard C.,Hollingsworth A., Jakob C., Janssen P., Klinker E., Lanzinger A., Miller M.Rabier F., Simmons A.,Strauss B., Thépaut J-N., and Pedro Viterbo. 1996: The ECMWF implementation of three dimensional variational assimilation (3D-Var). Part III: Experimental results, accepted to *Q. J. R. Meteorol. Soc.*.

Andersson E., J. Pailleux, J-N. Thépaut, J.R. Eyre, A.P. McNally, G.A. Kelly and P.Courtier, 1994: Use of cloud-cleared radiances in three/four-dimensional variational data assimilation. *Q. J. R. Meteorol. Soc.*, **120**, 627-653.

Bouttier F.,Derber D.,and Fisher M. The 1997 revision of the J_b term in 3D/4D-Var. ECMWF Tech. Memo. No. 238.

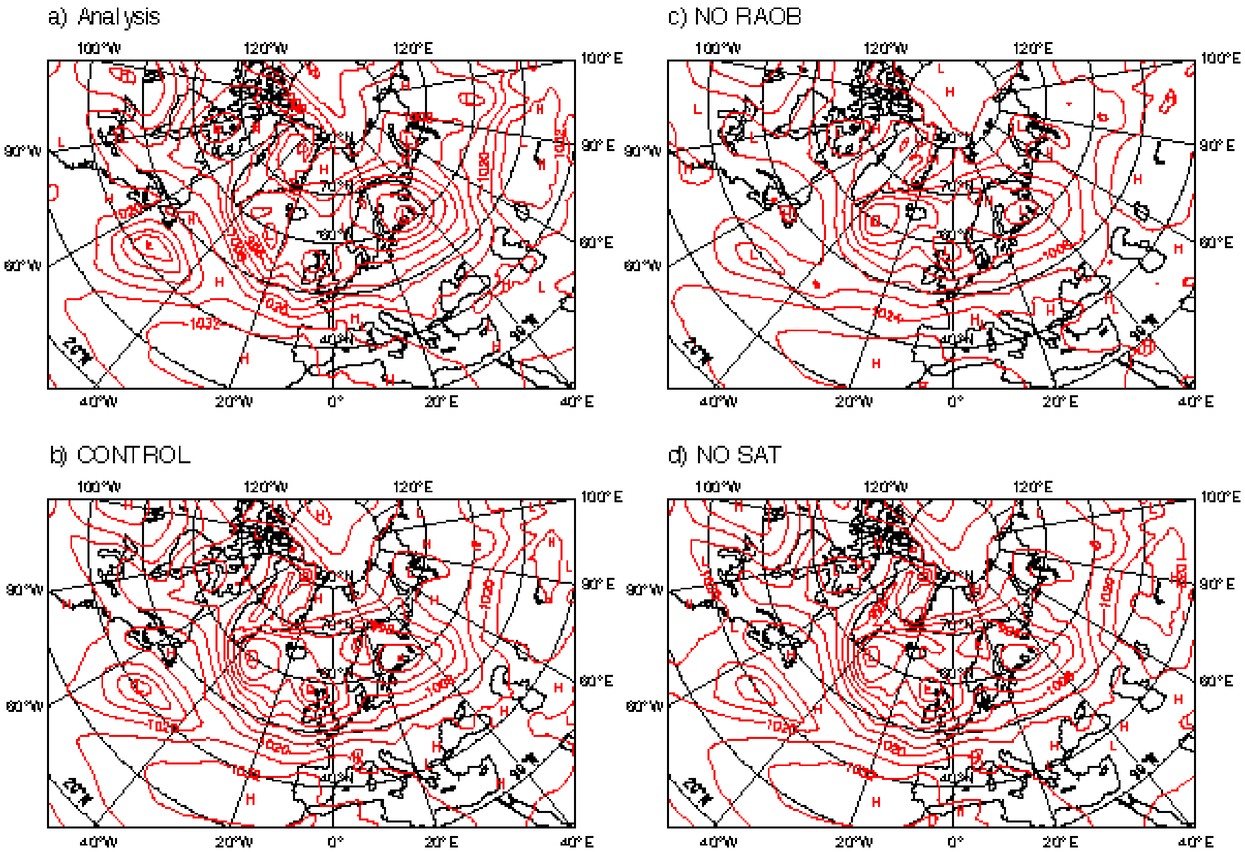

***Fig. 7 :** msl fields valid at 12 UTC 12/2/97 for a) analysis b) 48-hour forecast for control experiment, c) 48 hour forecast for 'NORAOB' experiment, d) 48 hour forecast for 'NOSAT' experiment.*

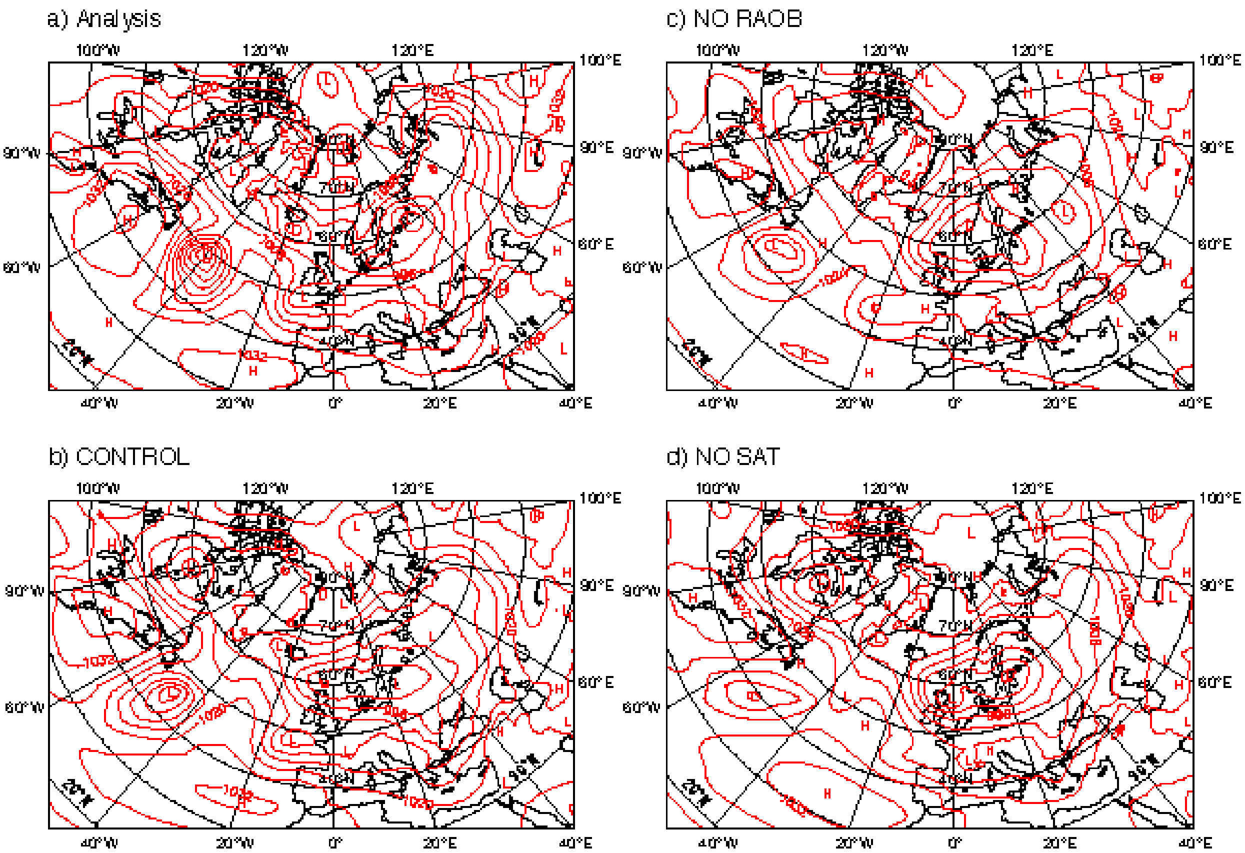

***Fig. 8 :** same as fig. 7 but for the 96-hour forecast valid at 12 UTC 14/2/97.*

Eyre, J.R., G.A. Kelly, A.P. McNally, E. Andersson and A. Persson, 1993: Assimilation of TOVS radiance information through one-dimensional variational analysis. *Q.J.R. Meteorol. Soc.*, **119**, 1427-1463.

Kelly, G.A. and Pailleux, J. 1988: Use of satellite vertical sounder data in the ecmwf analysis system, ECMWF Tech. Memo. No. 143.

Kelly, G.A., J. Pailleux,F. Rabier and Thépaut J-N. 1993: Observing System Experiments made with the ECMWF System. World Weather Watch Tech. Report.16. WMO/TD No. 594.

Lorenc, A.C., 1981: A global three-dimensional multivariate statistical interpolation scheme. *Mon. Wea. Rev.*, **109**, 701-721.

Graeme Kelly

Summary of ECMWF Technical Memorandum 223

ECMWF TOGA COARE data set

E. Klinker and A. Hollingsworth

The scientific objectives of the atmospheric component of the TOGA/COARE experiment include:

- Determination of the synoptic and mesoscale components of the large-scale slowly varying atmospheric circulation in the warm pool, including in particular the morphology of the most-convective stage of the 30-60 day mode and its sub-components.
- Determine the relationship of the phenomena described above to the heat, moisture, and momentum fluxes at the ocean atmosphere interface, in an attempt to relate the atmospheric phenomena of the western Pacific Ocean to transitions in the dynamics and thermodynamics of the upper-ocean structure.
- Determine the morphology of episodic westerly burst phenomena including their sources whether *in-situ* or remote.
- Determine the vertical heating distribution and life cycles associated with synoptic events and mesoscale convective cluster components, and compare these results to other regions in the tropics.

These issues play a key role in the formulation and validation of atmospheric models for weather prediction and for climate simulation. A special effort was made in the course of the experiment to provide a subset of the COARE observations to operational Centres via the WMO GTS. Those data were included in the real-time operational assimilation at ECMWF and other forecast centres.

In this memorandum we have used products from the operational assimilation and from additional model runs to calculate the three-dimensional diabatic forcing (with 6-hourly frequency) during the TOGA COARE experiment. Our main aim is to stimulate interest in comparisons of the field measurements with the model-based estimates from a state of the art assimilation system. After describing the assimilation system and the real-time TOGA/COARE data used in the assimilation we present the mean large-scale flow field during the major part of the Intensive Observation Period (IOP) December '92 to February '93, and the spatial distribution of the associated diabatic fields and surface fluxes.

For a period of three months including 15 November 1992 to 15 February 1993 additional short range forecasts have been run to provide a 3-dimensional picture of the diabatic forcing during TOGA/COARE. As the normal post-processing for the high resolution model (T213) provides only boundary fluxes from the output of the parameterization schemes, a lower horizontal resolution (T106) version of the model with extended diagnostic capabilities has been used to extract 3-dimensional diabatic and dynamic tendencies without stretching the computer requirements too much. Short range model integrations in this full diagnostic mode were run from the 6-hourly available analyses and the diabatic forcing for all major prognostic parameters (temperature, humidity and momentum) were post-processed.

The results show a large variability of the diabatic forcing in the warm pool area. Strong diabatic heating from the dominating process of cumulus convection occurs during the second half of December 1992, accompanied by comparatively strong westerly wind bursts of more than 10 m/s. A similar but less intensive event can be found at the beginning of February.

The vertical structure of the heat budget shows the dominance of the cumulus convection above the boundary layer with a maximum between 500 and 400 hPa. The melting level emerges as a sharp gradient of the heating between 600 and 700 hPa. The budget for the zonal momentum is again dominated by a process related to the cumulus convection. Large values of the so called cumulus friction are found in areas of large vertical wind shear and intense convective processes.

The above estimates of the components of the energy budget are generally consistent with current knowledge. However it is of interest to compare them with estimates made directly from in-situ data during the TOGA-COARE experiment. Preliminary results from observational studies seem to suggest that there is a fair amount of agreement between model results and observationally based estimates. One interesting difference has already emerged, which shows that the rather sharp gradient of the diabatic heating profile close to the melting level around 600 hPa is not found from observations. This could be an indication of the missing process of mesoscale downdraughts in the model that would distribute the cooling over a deeper layer below the melting level.

The rather incomplete use of TOGA COARE observational data in the real-time operational analysis requires a re-analysis for the full period. As the TOGA COARE IOP is part of the official ECMWF re-analysis period it is planned to rerun the diagnostic integrations from the re-analysis files that have a horizontal resolution of T106. Further improved analysis schemes will be tested in re-analysis experiments for shorter periods using the three-dimensional variational analysis.

Summary of ECMWF Technical Memorandum 229

Verification of the ECMWF wave forecasting system against buoy and altimeter data

Peter E.A.M. Janssen, Björn Hansen, Jean Bidlot

In this memorandum, we have reviewed the status of the ocean wave modelling at ECMWF. As part of the ECMWF operational task, global ocean wave forecasts are produced every day using the latest cycle of the WAM model (cycle 4). The wave model solves the wave energy balance equation for the wave spectrum and is forced by the ECMWF operational 10m winds. The initial conditions for the daily wave forecast are obtained from a 1 day assimilation starting from the previous analysis. The analysis procedure combines altimeter data from the ERS satellite with the wave first guess fields obtained with the analysed 10 metre winds.

Since its debut at ECMWF (1988), the wave forecasting system has been constantly upgraded. Firstly improvements have followed changes in the atmospheric model, for example the increased resolution to T213/31L in September 1991 with beneficial effects on the representation of surface wind features. Secondly, the wave model itself has been the object of constant improvements with the introduction of cycle 4 with modified physics for wind input and dissipation of wave energy (November 1991) and in July 1994, an increased resolution to 1.5 degree. Thirdly, assimilation of the newly available ERS-1 altimeter wave height data started in August 1993.

We have performed a systematic verification of all those combined changes on the quality of the wave analysis and compared the results to the few previous existing verifications. The quality of the wave forecasts has also been assessed for the first time.

Comparison of analysed wave data with independent moored buoy observations reveals the high quality of the wave analysis. Unfortunately, the geographical coverage of the buoy observations is rather limited; however, the first-guess wave field also seems to be of good quality as it follows from the verification against altimeter wave height data. From this comparison, it appears that considerable progress has been made with respect to wave model results from a decade ago. It is clear that the continued improvements of the atmospheric model analysis have had a positive impact on the quality of the wave first guess. Similarly, improvements in the wave model physics with respect to wind input and dissipation has contributed to a better wave model. Finally the inclusion of altimeter data was globally beneficial.

Regarding the wave forecast, we introduced comparison tools, such as the anomaly correlation, which have not been used before in wave modelling. The comparison of forecast wave heights with buoy observations has shown the slow deterioration of the quality of the wave forecast with time. It is suggested that the error growth is partly determined by a constant error (which we termed the error in swell) and an error caused by uncertainties in the local wind. Using the new comparison tools, we found by comparing the wave forecasts against the verifying analysis, that in the northern hemisphere we have a reasonable forecast up to day 5 while the limit of forecast skill in the southern hemisphere is between day 4 and day 5. It is fair to say that, since the wave forecast depends in a sensitive manner on the wind forecast, this verification confirms the high quality of ECMWF forecasts near the surface.

In spite of the overall good quality of wave forecast and analysis, we have pointed out a few problems which should require our attention (e.g. the poor quality of the wave analysis for the east coast of the US). Finally, we have concluded by mentioning the expected benefit to the wave analysis and forecasts from an improved atmospheric analysis with the advent of the variational analysis methods and new type of data.

Summary of ECMWF Technical Memorandum 232

Modification to the ECMWF WAM code

Jean Bidlot, Björn Hansen, Peter E.A.M. Janssen

This memorandum summarises the work which was performed for the migration of the wave model software to the new Fujitsu computer and the subsequent development. Indeed, with this new computer power at hand, the resolution of the global wave model was increased to a grid size of the order of 55 km (0.5 degree from 1.5 degree previously). Preliminary results show the benefit of the new high resolution grid in the vicinity of coastlines and in regions with sharp surface wind gradient.

The migration to a distributed memory computer such as the Fujitsu required the implementation of a new set of routines which were designed to carry out the necessary exchange of information from one processor to its neighbours using the message passing protocol. In this context, information is communicated between processors by sending and receiving messages. We opted to keep the one-dimensional array structures that were originally designed to favour a high vectorisation of the wave model computer code. For that reason, the global grid domain is divided into sub domains which are composed of a subset of latitude lines partitioned in such a way that the number of grid points is distributed as evenly as possible among the different processors. Except for the advection, all other wave physical processes are basically local and do not require information from neighbouring grid points. The upwind advection scheme used in the wave model only requires information from the 4 closest neighbouring grid points. Hence, the maximum length of the message that will be exchanged between the processors of two contiguous sub domains is determined by the number of grid points per latitude. It is then sufficient to exchange those messages before calling the advection routine and perform all computation on each processor independently from the other processors until the next advection call. Message passing is used again to recombine all fields before they are output to disk.

In regular spherical coordinates, the zonal grid spacing can become rather small as one approaches the poles. As a consequence, the size of the time step allowed for the wave model runs is severely limited for numerical stability reason by the cfl criterion. To alleviate this restriction, a new type of lat-lon grid was developed. It maintains a constant latitudinal increment but adjust the size of the longitudinal increment in such a way that the actual distance between grid points is almost constant. Such a grid is commonly known a reduced or irregular lat-lon grid. This grid was implemented in the wave model and required only a small modification of the code.

For operational production, it was decided to use the new reduced grid to increase the resolution to 0.5 degree at the equator. This means that the grid spacing is of the order of 55 km in both directions. We used a 5′ x 5′ global topographic data set to obtain the bathymetry for the high resolution grid. The key feature of this new grid is the better representation of coastlines and continental seas. The shallow water option was turned on to accommodate for those new shallow regions.

Preliminary results show that the wave model at the new high resolution captures more of the wind variability (and should therefore benefit from the future increased resolution of the atmospheric model). As a consequence, wave systems tend to have slightly sharper gradients in agreement with the wind speed distribution. The new bathymetry also has a positive impact for wave prediction in the North Sea and the coastal regions along the American coastlines, regions for which independent buoy observations are available for verification purpose.

The January 1997 Floods in Greece

Introduction

During 12 and 13 January 1997 torrential rainfalls over Greece led to widespread floods. At least six people were killed as a direct consequence of the floods and large damage was caused to property. In a 36-hour period from 11 January, 18 UTC, to 13 January, 06 UTC, more than 100 mm of rain were observed in central Greece, the southern mainland and the Pelopones peninsula (see figure 3). Precipitation amounts peaked over the northern Pelopones at Velos near Korinthos, where 344 mm in 36 hours were observed, more than half of which fell overnight between 11th and the 12th.

Synoptic situation

On 10 January a surface low was centred over the western Mediterranean while an upper level trough produced a cut-off low over the Gulf of Lions. During the following days this deep low moved slowly eastward, intensifying the southerly advection of warm and moist air masses over Greece. Further north-eastward progression of the system was prevented by the continental high extending over the northern Balcany. On its southern flank temperature gradients in a quasi-stationary front over Greece intensified.

In this situation, large scale lifting of the warm and moist air advected from the south, intrusion of initially cool and dry airmasses from the north-east, modified over the Black and Aegean Seas, and orographic convergence effects in many locations combined to cause the anomalous rainfall amounts described above.

Operational forecast performance

The cut-off low was forecast well in synoptic scale terms in the 3-day and also 5-day range (figure 1). The depth and positioning of the surface low and the distribution of lower tropospheric humidity over Greece and southern Italy were forecast well in the 3-day range, apart from a lack of humidity over the southern Aegean Sea and Crete (figure 2). At 120 hours the low was forecast slightly too far south and deeper than analysed. Values of high 700 hPa humidity are spread out more to the south, but the maximum over mainland Greece is missing.

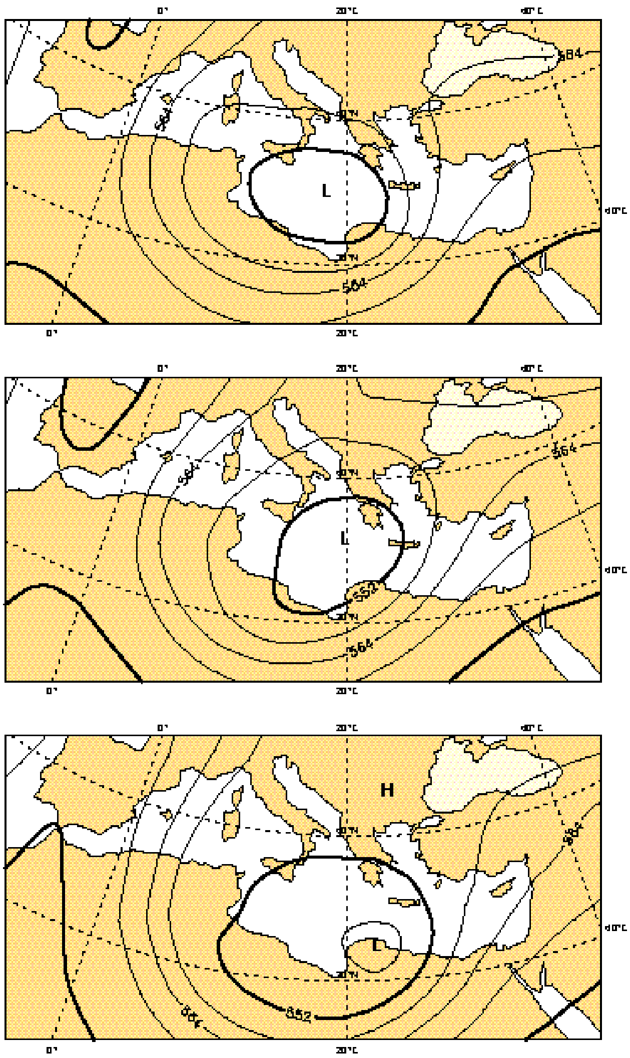

Fig 1: *500 hPa geopotential fields. Top: analysis from 12 Jan 1997, 12 UTC; middle: 72-hour forecast from 9 Jan 1997; bottom: 120-hour forecast from 7 Jan 1997.*

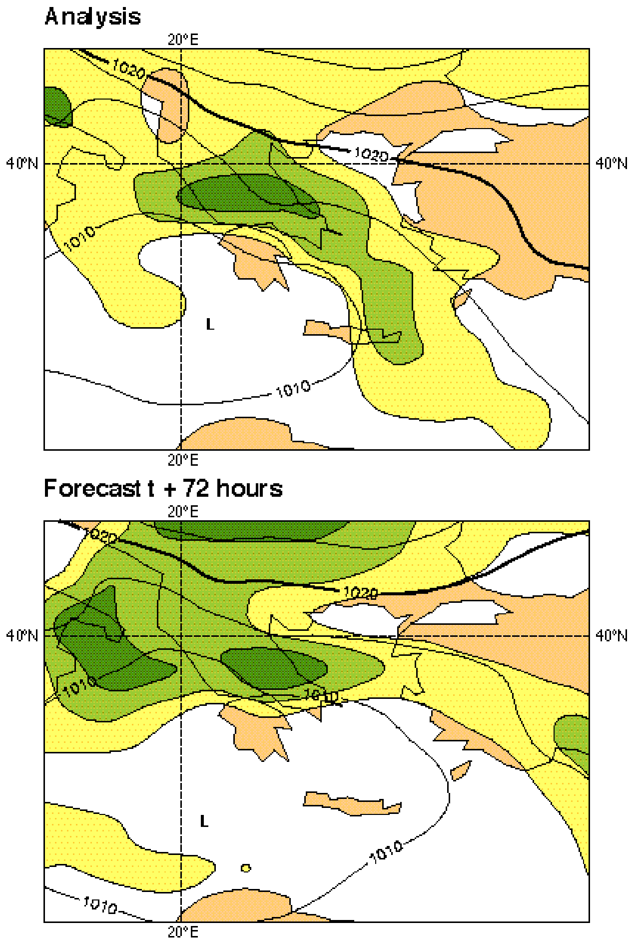

Fig 2: *700 hPa relative humidity (shading above 65%, 80% and 95%) and mean sea level pressure fields. Top: analysis from 2 Jan 1997, 12 UTC; bottom: 72-hour forecast from 9 Jan 1997.*

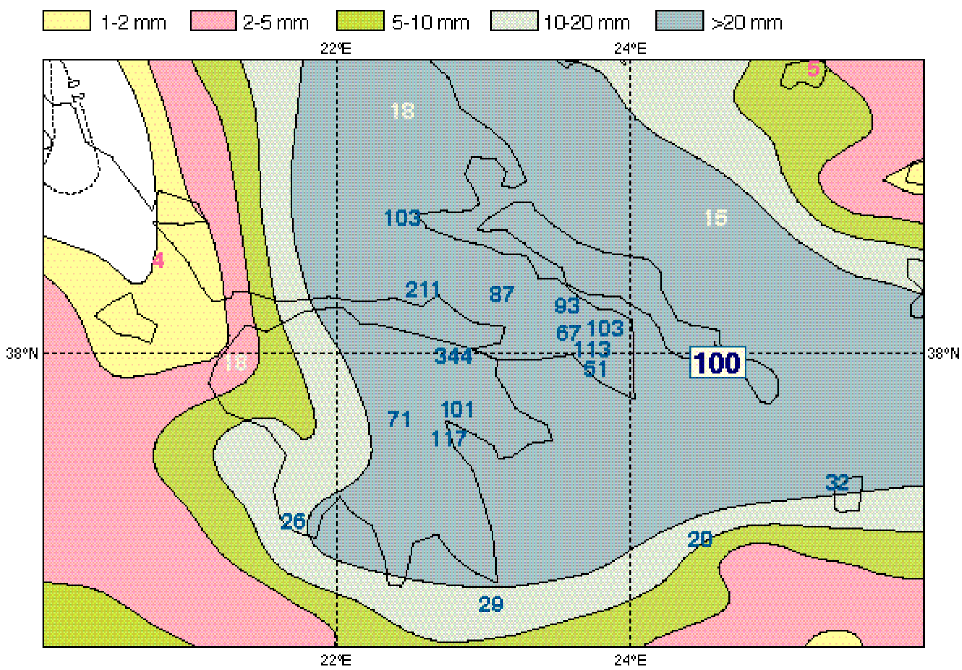

Fig 3: *Precipitation verification. Shaded: forecast of total precipitation from 9 Jan 1997 accumulated between steps t+54 and t+90 hours. The boxed number is a local field maximum. Small numbers: observations accumulated between 11 Jan, 18 UTC and 13 Jan, 06 UTC. (Partly obtained from Greek National Meteorological Centre).*

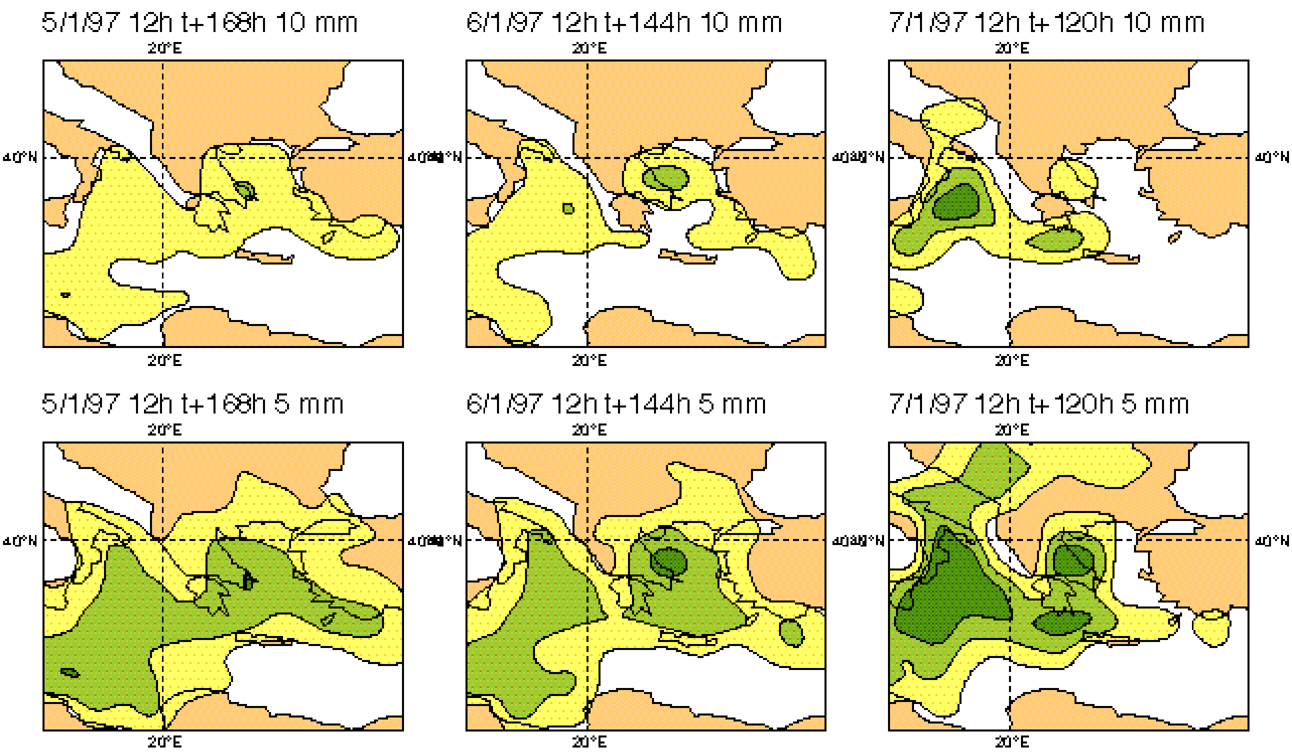

Fig 4: *Forecast probabilities of total precipitation exceeding 10 mm per day (top) and 5 mm per day (bottom), during day 7 (left), day 6 (middle) and day 5 (right) of the ensemble forecasts from 5, 6 and 7 Jan 1997, respectively. All valid for 11 Jan, 12 UTC to 12 Jan, 12 UTC.*

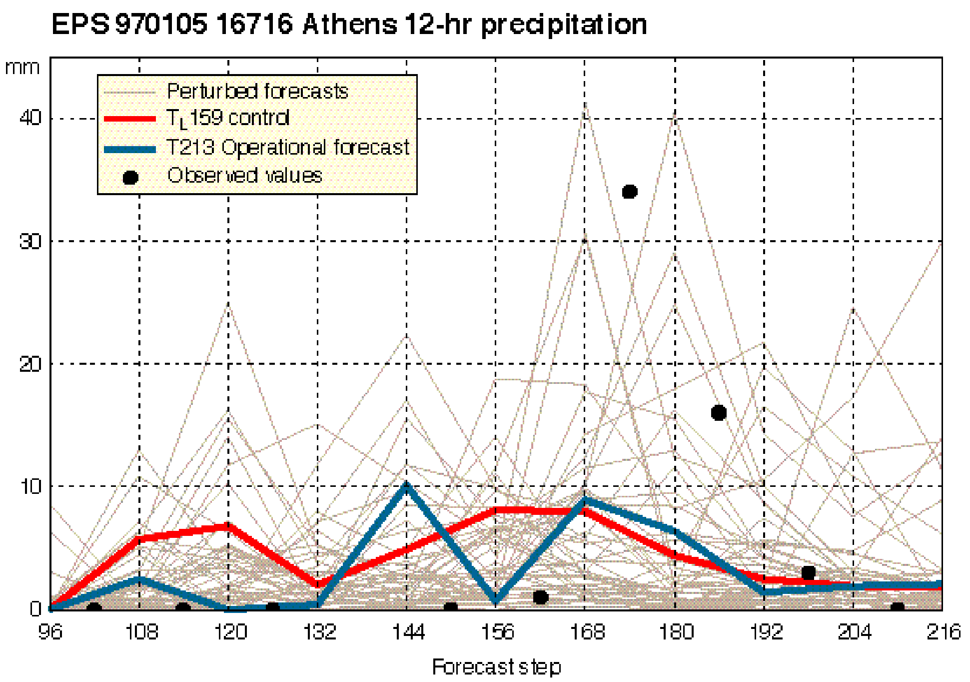

Fig 5: *Plume of 12-hourly precipitation from the ensemble forecast of 5 Jan 1997, at the grid point closest to Athens. Thin curves: perturbed forecasts (50); red solid: T_L159 con - trol forecast; blue dotted: T213 operational forecast; dots: observed values (06-18 UTC and 18-06 UTC, respectively)*

Figure 3 shows accumulated model precipitation in the early medium range, between t+54 and t+90 hours, valid in the 36-hour period of largest observed intensities. Although local extreme values cannot be represented the model gave good guidance to a major rainfall event. At shorter range, the distribution was forecast quite similar, however higher amounts were predicted.

Probability forecasts from the EPS

More than half of the ensemble members predicted 24-hour rainfall of more than 10 mm six and seven days in advance over a small area of eastern mainland Greece (figure 4). Five days in advance probabilities over this particular area were a bit smaller but more than 50% were forecast further south. For the 5 mm threshold more than 75% probabilities were predicted. Only few ensemble members forecast precipitation of more than 20 mm in 24 hours.

Twelve-hourly precipitation values at the grid point closest to the location of Athens from the T159 control, 50 perturbed forecasts and the T213 operational forecasts are shown in Fig 5, for the EPS from 5 Jan 1997. In the 6, 7 and 8-day range the operational and control forecasts produced amounts between 5 and 10 mm in 12 hours during the period of highest observed precipitation rates. Some ensemble members managed to forecast amounts close to observed values, highlighting the possibility of heavy rain a week in advance.

Acknowledgements

Many thanks to Mrs. Prezerakou and Mrs. Refene from the Hellenic National Meteorological Centre for their helpful comments on this case and for providing additional precipitation observation data for several locations in Greece.

Andreas Lanzinger

Computing at ECMWF - 1980 to 1997 - a personal viewpoint

I joined ECMWF in May 1980 as Head of Computer Division, but I am now leaving (30 June 1997) to join DWD as a Head of Department. In my 17 years here I have seen great changes in the computing environment. This article is a personal view of those changes, plus a light-hearted look at the future my successor may face.

When I joined ECMWF the computer equipment consisted of a single processor Cray 1-A, a Cyber 175 front-end, and a RegneCentralen RC8000 telecommunications system (see Fig. 1). Each of these provided a particular service, the Cray 1-A was the 'number-cruncher', the Cyber 175 provided both data storage and user access, while the RegneCentralen RC8000 handled the external network. This concept of separate servers for different tasks was then taken further over the years, as it allows great flexibility when it comes to upgrading. Each server can be replaced individually, without impacting the others. This means minimal interruption to the overall service. Thus, as I leave, the equipment is now three Fujitsu systems (46 processor VPP700, 9 processor VPP300, 4 processor VPP300), two data handling systems (IBM ES9000, IBM RS6000), a VAX cluster for telecommunications, three LANS (HIPPI, FDDI Ethernet), and an SGI/HP based service for direct user access (see figure 2).

Let us look at each of these areas in more detail. The number cruncher (or compute server) has expanded considerably over the 17 years. The Cray 1-A was a 160 Mflop (peak) system with 8 Mbytes of memory and 2.4 Gbytes of disk. Today's Fujitsu VPP700 is rated at just over 100 Gflops (peak), an increase of more than 600. Hence a program that took an hour of Cray 1-A CPU time to run, now takes less than 6 seconds. The total memory has risen from 8 to 2000 Mbytes per processor, and the disk capacity from 2.4 to 1000 Gbytes. However, the internal clock speed has changed very little, it was 12.5

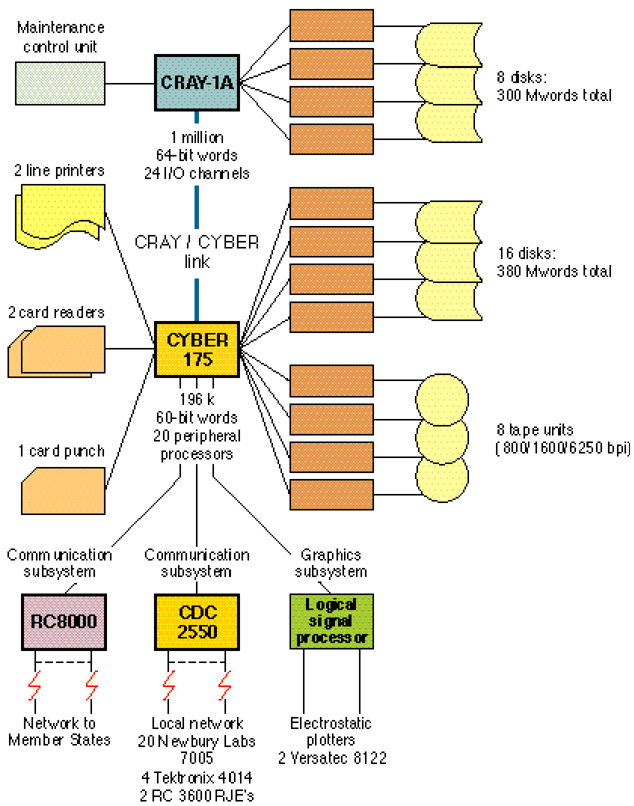

***Fig. 1:** Configuration diagram of ECMWF computer system 1980*

nano seconds on the Cray 1-A, today it is 7 nano seconds on the Fujitsu VPP700.

Even in 1980, it was apparent that future increases in computational speed would have to come primarily from sources other than clock speed improvements, in particular from multi-processor systems. Hence the Centre

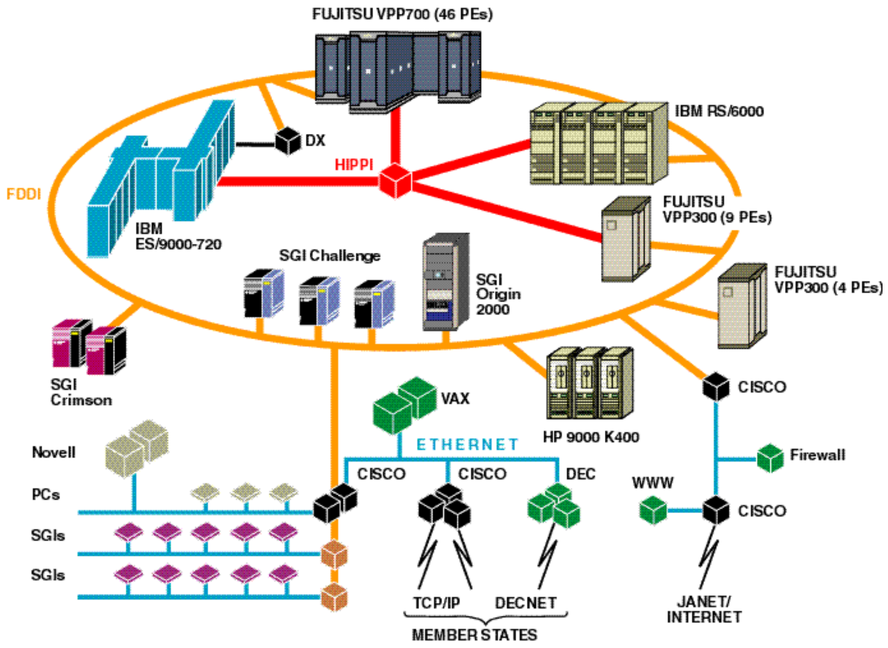

Fig. 2. *ECMWF computer configuration - 30 June 1997.*

invested a lot of effort into making its principal software multi-tasked. This has paid handsome dividends, as the Centre's computer server went through 2, 4, 8, 12, then 16 processors to today's 46 processor system (shortly to be 116 processors). By going multi-tasked the sustained computational rate for the Centre's main model has risen from 50 Mflops/sec in 1980 to 35 Gflops/sec today, an increase of a factor of 700!

A second major change occurred more recently. Up to the 16 processor Cray C90 (operational to September 1996), all machines had shared memory, thus all processors could directly access the whole of memory. However, by the mid 1980's it was apparent that distributed memory systems might be the way future machines would develop in order to sustain the growth in overall performance. Thus, again, the Centre expended much effort in making a distributed memory version of its main software, such that by 1995 it could satisfactorily benchmark both shared memory and distributed memory machines. The result was the present 46 processor Fujitsu VPP700, shortly to be increased to 116 processors.

Competition amongst vendors is a healthy thing, at least for the customers! In the early 1980's, Cray faced little real competition except from CDC/ETA. There were plenty of start-up companies that produced a working prototype or two, and some even sold a few production models. However, eventually all of them disappeared. For real competition, vendors with large resources needed to be tempted into the market. Also back in the early 1980's there was little knowledge in the meteorological field of how the large scale computer market was likely to develop, in particular the impact that parallel systems would possibly have. To address these issues, ECMWF began in 1984 a biennial series of workshops, under the general title of "ECMWF Workshop on the Use of Parallel Processors in Meteorology" (note the lack of a trendy acronym, with a title like that there was no obvious one that would fit!). These workshops brought together meteorologists, computer scientists, and vendors for a week long discussion. It certainly had the effect of raising the profile of the requirements of meteorologists, and of the particular problems they faced in getting the best performance for their type of work. In turn it gave the computer scientists a concrete problem to try their various algorithms on, while for the vendors it brought together a group of prospective customers with similar requirements. Out of all this, true competition emerged in the shape of Fujitsu, Hitachi, IBM and NEC.

When it came to acquiring a system, back in 1980 the main machines had been lease purchased. This led to the inevitable problem that the total capacity available for the period was constant, yet the demands on the service were growing steadily year by year. In addition, the Centre was left with the task of disposing of the old machine before acquiring a new one. To overcome this the concept of a service contract was adopted. Such a contract was not for a specific machine, but rather for the provision of a given level of performance, that level rising steadily year by year over the lifetime of the contract. The performance was measured by using the Centre's main forecast model as a yardstick. The successful vendor could then bring in, and remove, machines from his current range that met the required performance levels. This concept of service contracts has now been at the centre of ECMWF's budgeting policy for more than seven years, and has been very successful in ensuring the Centre has had adequate computing capacity to meet its operational and research needs.

Turning now to data storage, in 1980 it was provided by 1600, then 6250 bpi tapes. Users were allocated specific tapes, and they had to manage their own tape archives using those allocated tapes. It meant a lot of administrative work for each of the main user groups. To reduce this unnecessary work the idea of a data handling system was born, such a system taking data files (by name) and storing them as required on either disk or tape. The users would not know which disk or tape, the system would handle all that for them. The result was the first data handling system, built around IBM 4341 hardware.

At that time there was little commercially available software that could effectively handle the data volumes we estimated we would accumulate. However, a search of similar sites with large data archiving requirements revealed a product called CFS (Common File System) being developed at Los Alamos, USA. That product became the cornerstone of ECMWF's data handling system. It needed adapting to ECMWF's network connections of the time (RHF/LCN), and again later to Superlink, then TCP/IP. However, it has been an excellent product that has proved itself many times over. It is interesting to note that we originally selected it because it was already handling just under a million files at Los Alamos. At ECMWF we estimated it might eventually have to handle two to three million files. However, over the years this first system has expanded and grown so that today it stores 60 Tbytes in about 12 million files! But the end is in sight, looking to the future it is clear it cannot sustain the present growth rate for much longer, hence a new data handling system (IBM RS6000 based with ADSM software) is being brought into operation. This new system is destined to expand to at least 1 Pbyte (Peta = 10^{15}).

A crucial component of that first system in 1980 was the telecommunications network, by means of which the forecast was distributed to Member States, and Member State batch work was sent to the Centre. At that time computer networks were in their infancy, and as there were no standards for the necessary protocols the Centre had to invent its own. The star network thus created consisted of seven low speed lines (50 or 100 baud, dissemination only) plus four medium-speed lines (2400 or 4800 bps). All these were handled by the Regnecentralen RC8000 (also known as the NFEP - Network Front End Processor). Over the years the hardware changed to a VAX cluster, and the protocols from ECNET (ECMWF's own protocols), through DECnet, to TCP/IP. Line speeds are now 64 kbps minimum with a couple at 128 kbps (France, Austria), and one at 2 Mbps (UK). Networking for the individual user has gone from nothing (on the low-speed lines) to full interactive access over the Member State leased line network or international research networks.

Despite the tremendous increase, both in computational and storage capacities, the most visible sign of change for the individual user has probably been the method of accessing all that power. In 1980, most users were working with punched cards, a few code developers had alphanumeric Newbury Labs 7005 terminals. These few terminals accessed the then state of the art interactive system known as Intercom, running on the Cyber 175 under the NOS/BE operating system. Over the years, first the Newbury Labs terminals spread to all offices, then they were replaced by PCs (which offered office facilities such as word processing and electronic mail) which were, in turn, replaced by Unix based graphics workstations. It is interesting to note that the power on an individual desk today (typically 100 Mflops, 64 Mbytes memory, 2 Gbytes disk) equals or exceeds that shared by all the users of the first Cray 1-A system.

So, what can my successor look forward to? Assuming he/she stays 17 years, he/she will be writing this article in the year 2014. Based on a simple extrapolation of the growth rates for the past 17 years, he/she 'can look forward to' controlling a compute power of some 40 Tflops, plus an archive of three to four Ebytes (Exa = 10^{18}). The average office system will be as powerful as today's compute server, probably with (optional) voice input and output. The external network, which may well be a managed network covering both ECMWF's Member States and WMO's region VI, will be running at megabit speeds as a minimum. Finally, what language will the users be programming in? It will doubtless still be called Fortran, but what its actual structure will be like is anyone's guess!

Geerd-R. Hoffmann

Changes to the Fujitsu VPP700 configuration

As outlined in ECMWF Newsletter No. 74, the original agreement with Fujitsu envisaged that the VPP700 would be enhanced in 1998 to a system with approximately 240 processors (PEs), all secondary PEs having 0.5 Gbytes of memory each. Changes have now been agreed with Fujitsu that will alter both the configuration of the enhanced system and the timetable of its introduction. Given below is a general overview of these changes, together with a timetable for the introduction of the first of these changes.

In September 1997, the VPP700 system will be upgraded to a 116 processor system, consisting of:

- 1 primary PE — 2 Gb memory
- 10 IMPEs — 2 Gb memory each
- 105 secondary PEs — 1 Gb each
- 1.5 Tbyte disk space (RAID).

The proposed timetable of this change is as follows:

early July	Upgrade the operating system on the present VPP700 to the level required by the 116 PE system (UXP/V X97061)
15-30 August	Deliver and install a 70 PE VPP700 system, implementing on it level X97061 of the operating system
early September	Run a trial service on the 70 PE system for selected users, and make parallel runs of the operational forecast suite. Later move the operational suite to the 70 PE system
11 September	Switch off the current 46 PE VPP700 and move all PEs and disks to the new system. This will involve a shutdown of both machines for up to 10 hours.

It should be noted that:

(i) the new system will be binary compatible with the current one;

(ii) the main file system structure available on the current machine will be available on the new one;
(iii) each secondary PE on the 70 and 116 PE system will have 1 Gbyte of memory, as against 0.5 Gbytes in the original proposed upgrade and 2 Gbytes on the current system;
(iv) when the disks are moved to the new system on 11 September 1997, all data on those disks will be lost.

The end result will be a 116 PE VPP700 system that should be some 2.5 times as powerful as the present 46 PE machine.

There will not now be any changes to this configuration in 1998. However, in May 1999 additional capacity will be installed that will approximately double again the throughput capability.

Neil Storer

ECMWF Calendar 1997

25 Aug	ECMWF Holiday	
8 - 12 Sep	Seminar: Atmosphere-surface interaction	
29 Sep - 1 Oct	Scientific Advisory Committee	*26th*
7 - 8 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	

ECMWF Publications

EUMETSAT/ECMWF Fellowship Programme Research Report

No. 5 Tomassini, M, D Le Meur and R Saunders: Near-surface wind observations of hurricanes and their impact on ECMWF model analyses and forecasts, April 1997

Workshop Proceedings:

International TOVS Working Group. AReport on The Ninth International TOVS Study Conference, Igls, Austria, 20-26 February 1997.