In this issue

European Centre for Medium-Range Weather Forecasts

Shinfield Park, Reading, Berkshire RG2 9AX, UK

Fax: . . . . +44 118 986 9450

Telephone: National. 0118 949 9000

International . . . . +44 118 949 9000

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

Cover

The front page shows tracks of tropical cyclones in the western North Pacific and South China Sea during 1998, together with a satellite visual image of typhoon Babs.

Editorial

The article on page 2 by Queenie Lam (Hong Kong Observatory, China) describes how improvements in the Centre's model during the past decade have been reflected in better forecasts of tropical cyclones in the western North Pacific and the South China Sea. She states that the greater accuracy of the model forecasts, particularly since 1996, has meant that they are now regarded by forecasters as being reliable sources of information that are making significant contributions towards improving the prediction of tropical-cyclone positions up to 48 hours ahead, and beyond.

On page 8 Craig Dorman (US Office of Naval Research Europe) provides an external critique of the Centre's activities. He concludes that the Centre's success and high international reputation is a consequence of its reluctance to grow too large, and its policy of cycling the most talented European meteorologists between itself and home institutions in Member States.

Increasing demands on the ECFS file management system have necessitated a redesign to cope with current and future file storage requirements, as described by Tony Stanford on page 10. He points out that the ECFS is intended primarily to be an archival storage system, and that users need to use other methods of storing frequently accessed files. Changes in coding conventions within the Integrated Forecasting System that use several Fortran 90 features have been introduced to enable the code to be more portable, and to encourage better and safer coding styles; these are outlined by David Dent and Mats Hamrud on page 11.

Changes to the Operational Forecasting System

Recent changes

On 22 October 1999 the reference model's vertical resolution was increased from 50 to 60 levels, most of the extra resolution being in the planetary boundary layer. The EPS vertical resolution has ben increased in a consistent way from 31 to 40 levels.

These resolution changes have been coupled with several modifications to the atmospheric model (Cycle 21r4):

- extensive changes in the cloud and convection schemes (mainly a new parametrization of the effects of cloud overlap on precipitation/evaporation);
- a new post-processing of the 10 m winds aimed at a better representation of weather-station environments;
- a revised evaluation of the background-error cost function used in 4D-Var (J_b);

- a revised scheme for the correction of satellite (A)MSU radiance biases;
- an active assimilation of SSM/I wind speed through a 1D-Var assimilation of SSM/I radiances;
- a bugfix of the humidity computation from SYNOP and TEMP observations below 0° C.

New global orography and subgrid orography fields have been introduced, based on a new high resolution dataset. The changes are mostly isolated ones except for the Greenland massif and Australia, where more substantial changes in the basic orography are evident.

The changes have been proved to have had a positive impact on the error scores (Z500 in the extratropics, the upper-level winds and temperature in the tropics). The frequency of occurrence of moderate to strong 10m winds and precipitation were also improved.

Planned changes

The next upgrade in 2000 will be related to the data assimilation system, i.e. 12-hour cycling of 4D-Var, possibly with increased resolution of the minimized increments.

François Lalaurette

Recent performance of the ECMWF model in forecasting the tracks of tropical cyclones over the western North Pacific and the South China Sea

The Hong Kong Observatory (HKO) has been receiving ECMWF model forecast products based on 12 UTC analyses in GRID format via the Global Telecommunication System (GTS) since 1986. Since early 1999 enhanced-resolution datasets in GRIB format have been retrieved by the HKO via Offenbach and Beijing. These data are available at 2.5° (2.5° resolution at 24-hour intervals up to 120 hours ahead.)

The HKO has been using the ECMWF model for more than a decade to aid operational forecasting in the short–to–medium range. As early as in 1986, the low-pressure centres on surface prognoses and the cyclonic circulation centres on 1000 hPa and 850 hPa forecast charts were used to forecast the movement of tropical cyclones (TCs). Chan et al. (1989) assessed the performance of the ECMWF model in predicting the movement of a particularly difficult TC in 1986 and suggested some rules in using ECMWF model for TC track forecasting at that time.

The Director of the HKO from 1965 to 1981, G. J. Bell, found that subjective forecasts issued by the warning centres were hard pressed to beat forecasts obtained from a simple combination of persistence and climatology *(Bell, 1962)*. Bell (1979) recognised that conventional synoptic/statistical type 24-hour forecasts were at the limits of development. He drew attention to the need to look to models for improved forecasts for periods of 48 hours or more, though the arrangements for the dissemination of observations and the methods of objective analysis of TCs were inadequate at that time to generate useful forecasts for 48 hours ahead. A decade after Bell, Lam (1992) reviewed the techniques and performance of TC forecasting by subjective and objective means from the perspective of the HKO's operations. Models at that time were still not able to beat simple methods like persistence in 24-hour TC position forecasts, and model performance had not yet shown significant impact on the quality of forecasts up to 48 hours and beyond.

Over the subsequent years, the ECMWF model has been evolving. Front-line forecasters have observed significant improvements in the forecasting of TCs by the ECMWF model in recent years. It is now of interest to review again the model performance in TC forecasting, and to see whether recent advances in numerical modelling have led to significant impacts on operational TC forecasting. This note reports on the recent performance of the ECMWF model in TC track prediction over the western North Pacific and the South China Sea.

Verification of tropical cyclone forecasts

The model forecast positions of TCs in 1991-1998 have been verified against the "best track" analysed by the HKO over the Hong Kong area of responsibility for issuing TC warnings for shipping, that is 10–30°N, 105–125°E. This verification area was chosen to enable easy comparison with the HKO subjective forecasts. The subjective forecasts during the period 1975–1998 were also verified against the best-track positions to see the long-term trend, and to compare with the model performance during the latter part of the period. In this intercomparison, the forecast error has been defined as the great-circle distance between the forecast position and the best-track position of the TC. The best-track positions of the TCs are determined after the event with the benefit of all available data and the wisdom of hindsight.

The positions of the minimum surface pressure, as determined from ECMWF prognostic data, were treated as the TC centres. In addition, the positions of the maximum 850 hPa relative-vorticity centres were also used to track the TCs in the most recent year of 1998 for comparison with the position errors associated with tracking using the minimum surface pressure centres. In view of the vertical alignment of centres at different levels in mature TCs, predictions based on 850 hPa vorticity centres should not be too much different from those based on the surface pressure centres. One advantage of using the 850 hPa vorticity field is that such centres are often easier to locate in weak systems with circulations not be adequately represented in the surface pressure field. Because of this, the 850 hPa vorticity field generally gave better continuity in the course of tropical cyclone evolution *(Lam and Lai, 1994)*.

When using the 850 hPa relative-vorticity field to track the TCs, a threshold of 50 x 10^{-6} s^{-1} was adopted to identify and track the TC positions on the prognostic charts through an automated algorithm. The extracted positions of the TCs were further checked manually. The forecasts were verified even if the TC had not formed at the model initial time; this indicated to some extent the model's ability to forecast TC genesis.

Tropical cyclones over the western North Pacific and the South China Sea are classified into four categories, namely tropical depressions (TDs), tropical storms (TSs), severe tropical storms (STSs) and typhoons (Ts) according to their maximum sustained surface winds near their centres (Table 1).

Class	Maximum sustained surface winds near TC centres (knots)
Tropical Depressions	22 – 33
Tropical Storms	34 – 47
Severe Tropical Storms	48 – 63
Typhoons	Greater than 63

Table 1: The classification of tropical cyclones over the western North Pacific and the South China Sea.

More detailed analysis of the 1998 dataset was carried out with the verification area expanded to cover the western North Pacific, that is 0–45°N, 100–140°E. The tracks of the TCs in 1998 are shown in Figure 1. A total of 20 TCs occurred, of which 19 TCs were included in the verification dataset. TS Nichole was not verified because its maximum 850 hPa vorticity never reached the threshold during its relatively short life span.

Figure 2 shows the trend of annual mean errors of model forecasts compared with those of the HKO operational subjective forecasts up to 48 hours ahead over the area 10–30°N, 105–125°E. Looking at the results over the past two decades, there are some indications of better subjective forecasts during the 1990s than in previous years. However, the 24-hour forecast errors remained about 200 km while 48-hour forecast errors were around 400 km. Similar results were found in an earlier study by Lam (1992) based on 1961–1990 data.

The relatively large error in the 48-hour subjective forecasts for the year 1998 was contributed mostly by two TCs, namely Vicki and Zeb. In the early stage of Vicki, there was another vortex to its west over the South China Sea, and the divergent scenarios forecast by different global models puzzled forecasters. In the case of Zeb, its acceleration after recurving over the western North Pacific was difficult to predict.

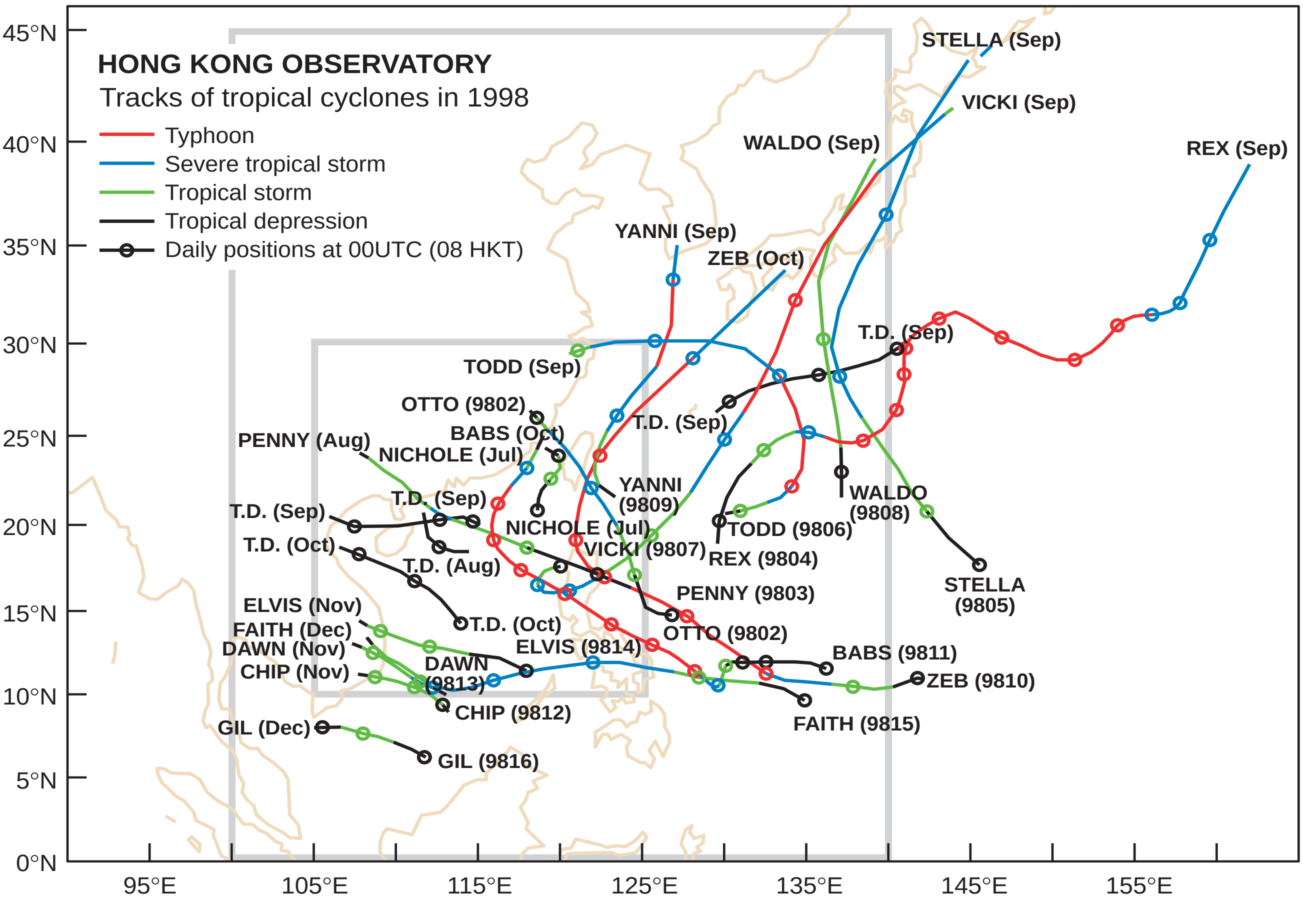

Figure 1: Tracks of tropical cyclones over the western North Pacific and the South China Sea in 1998. The two verification areas are marked out with thick, grey lines.

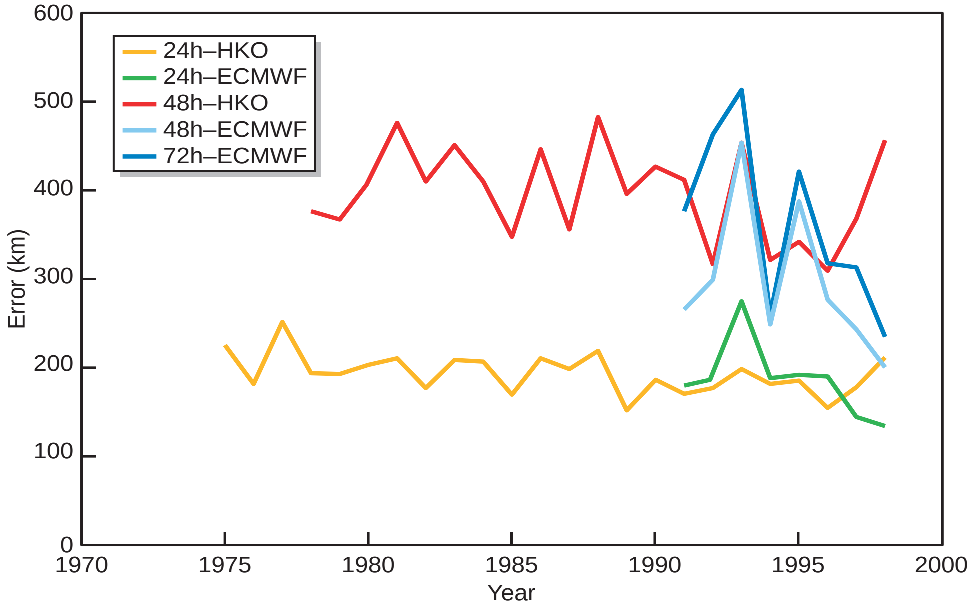

Figure 2: Annual mean errors of the model forecasts and the HKO subjective forecasts for TCs over the verification area 10–30°N, 105–125°E.

In contrast to the situation regarding subjective forecasts, the model forecast position errors for all forecast ranges showed a long-term decreasing trend (Figure 2), particularly since 1996. In the early 1990s (1991–1993), the model mean forecast errors for 24-hour, 48-hour and 72-hour positions were around 220 km, 340 km and 450 km respectively. In 1998 they were reduced to around 140 km, 200 km and 240 km, respectively. Figure 2 indicates that the improvement in ECMWF model forecasts of TC positions was most noticeable during the last three years.

The mean position errors and the standard deviations of the model and subjective forecasts between 1991 and 1998 are given in Table 2 for comparison. It is seen that the model performance at 24 hours was comparable with the subjective forecasts, but the 72-hour model forecast error was about the same as the 48-hour subjective forecast error during the period. The HKO carried out experimental 72-hour subjective forecasts in 1976 and the mean error was 563 km; the accuracy of the model forecasts for 72 hours was much better than this in the 1990s.

Forecast hour	ECMWF model forecast Mean error (Standard deviation)	HKO subjective forecast Mean error (Standard deviation)
24	187 km (42 km)	183 km (17 km)
48	299 km (82 km)	372 km (61 km)
72	364 km (99 km)	–

Table 2: Forecast position errors – mean and standard deviation – for TCs over the verification area 10–30°N, 105–125°E during 1991-1998.

As a measure of tropical cyclone forecast skill, forecast errors are usually compared with those obtained by the simple climatology-persistence (CLIPER) method *(Bell, 1962; Neumann, 1972)*. The skill score is defined as (CLIPER error minus forecast error)/(CLIPER error), expressed as a percentage. A positive skill score means that the model is more skilful than CLIPER. Figure 3 shows the skill score of operational subjective forecasts by the HKO and of predictions by the ECMWF model in forecasting TC positions in the region 10-30°N, 105-125°E. The CLIPER method often outperformed the operational subjective 48-hour forecasts before 1990. However, the situation seems to have reversed in the last few years. On the other hand, the model forecasts have been more accurate than CLIPER for all forecast ranges since 1994. A significant increase in model skill against CLIPER has been observed, particularly for the T+72 hour forecasts, over the past few years. Putting these results together, it would seem that the subjective forecasts have benefited from the much improved model forecasts over the years. That the 48-hour subjective forecasts have not improved as rapidly as the model forecasts is an issue to be addressed by the HKO meteorological service.

Looking more closely at the dataset in 1998, the model mean forecast position errors using the minimum surface pressure centres and using the maximum 850 hPa relative-vorticity centres to track the TCs are given in Table 3 for comparison.

Forecast hour	Mean error (standard deviation) [number of cases] for surface pressure	Mean error (standard deviation) [number of cases] for 850 hPa vorticity
0	102 km (59 km) [35]	87 km (52 km) [35]
24	136 km (80 km) [35]	131 km (67 km) [35]
48	203 km (124 km) [32]	207 km (124 km) [32]
72	237 km (136 km) [29]	246 km (126 km) [29]

Table 3: The ECMWF model forecast position errors for TCs in 1998 over the verification area 10–30°N, 105–125°E. Cases with maximum 850 hPa relative vorticity (≥ 50 x 10^{-6} s^{-1}) are included in the comparison study.

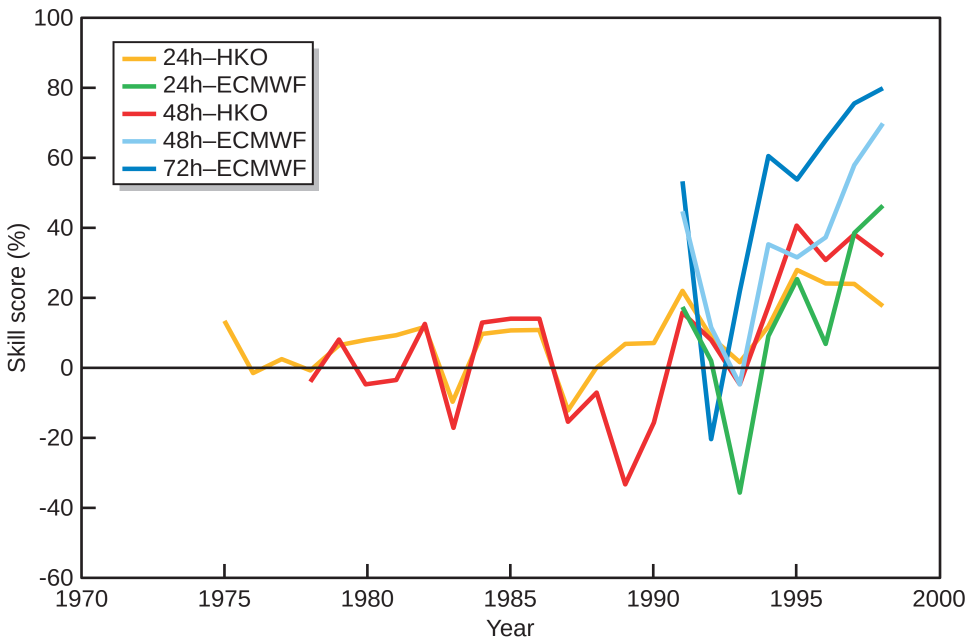

Figure 3: Operational forecasting skill and model skill in tropical cyclone position forecasts relative to climatology–persistence for TCs over the verification area 10–30°N, 105–125°E.

In operational forecasting, TCs formed over the western North Pacific are also monitored by the HKO forecasters as they may enter the South China Sea and affect Hong Kong. In view of this, it is also of interest to see the model performance for TCs over this region. Table 4 summarises the model mean position errors over the area 0–45°N, 100–140°E, covering the South China Sea and much of the western North Pacific.

Forecast hour	Mean error (standard deviation) [number of cases] for surface pressure	Mean error (standard deviation) [number of cases] for 850 hPa vorticity
0	94 km (55 km) [64]	86 km (55 km) [65]
24	140 km (89 km) [65]	133 km (81 km) [65]
48	232 km (168 km) [58]	229 km (165 km) [59]
72	308 km (241 km) [51]	311 km (238 km) [51]

Table 4: The ECMWF model forecast position errors for TCs in 1998 over 0–45°N, 100–140°E. Cases with maximum 850 hPa relative vorticity ($50 \times 10^{-6}\ s^{-1}$) are included in the comparison study.

From Tables 3 and 4, it can be seen that there is not much difference in the errors of the TC position forecasts between verifications using the centre positions deduced from the surface pressure field or from the 850 hPa vorticity field. With the inclusion of the western North Pacific in the verification area, the model forecast errors were larger than those over the South China Sea only. The larger differences for the longer forecast ranges (around 14% for 48-hour forecasts and 30% for 72-hour forecasts) arose because recurving TCs over the western North Pacific tended to accelerate after recurving, bringing about large errors.

Figures 4(a)–(d) shows the spatial scatter of the model analysed and forecast positions using 850 hPa vorticity field relative to the best-track positions during 1998. It should be noted that the analysed and forecast TC positions were interpolated from model output with the relatively coarse 2.5° resolution that is available on the GTS. It therefore gave rise to rather large position errors at T+0 hours. These errors would be probably less if the full resolution model data were used. Figure 4(a) shows that the maximum model analysed position error was around 200 km. However, this error could be as large as 400 km in the early 1990s (Lam, 1992). Initial positions of TCs have been more accurately depicted on model analysis in recent years. This may be one of the reasons contributing to the overall improvement in model forecast TC track.

In the context of operational tropical cyclone forecasting, warning the public and the marine community about tropical cyclones has always been the ultimate goal. For warning purposes, initial detection, determination of position and intensity, forecasting movement, intensification and the point of landfall, as well as the timing and magnitude of the weather impact, are all important issues for an operational warning centre. Typhoon Babs is discussed below to illustrate the issues of concern to the HKO in tropical cyclone situations, and to show how model forecasts have helped in the warning services.

Figure 5 shows the observed and successive ECMWF forecast tracks for Babs. The life history of typhoon Babs was the second longest of the western North Pacific TCs during 1998. It posed a potential threat to Hong Kong after crossing Luzon and moving towards the south China coast. Babs slowed down and recurved within a distance of 300 km to the south-east of Hong Kong at around 12 UTC 25 October. The ECMWF model forecasts provided useful guidance in forecasting the recurvature of Babs to the east of Hong Kong, starting from the model run with the initial time of 12 UTC 24 October. This was an important piece of information as the effect of a TC to Hong

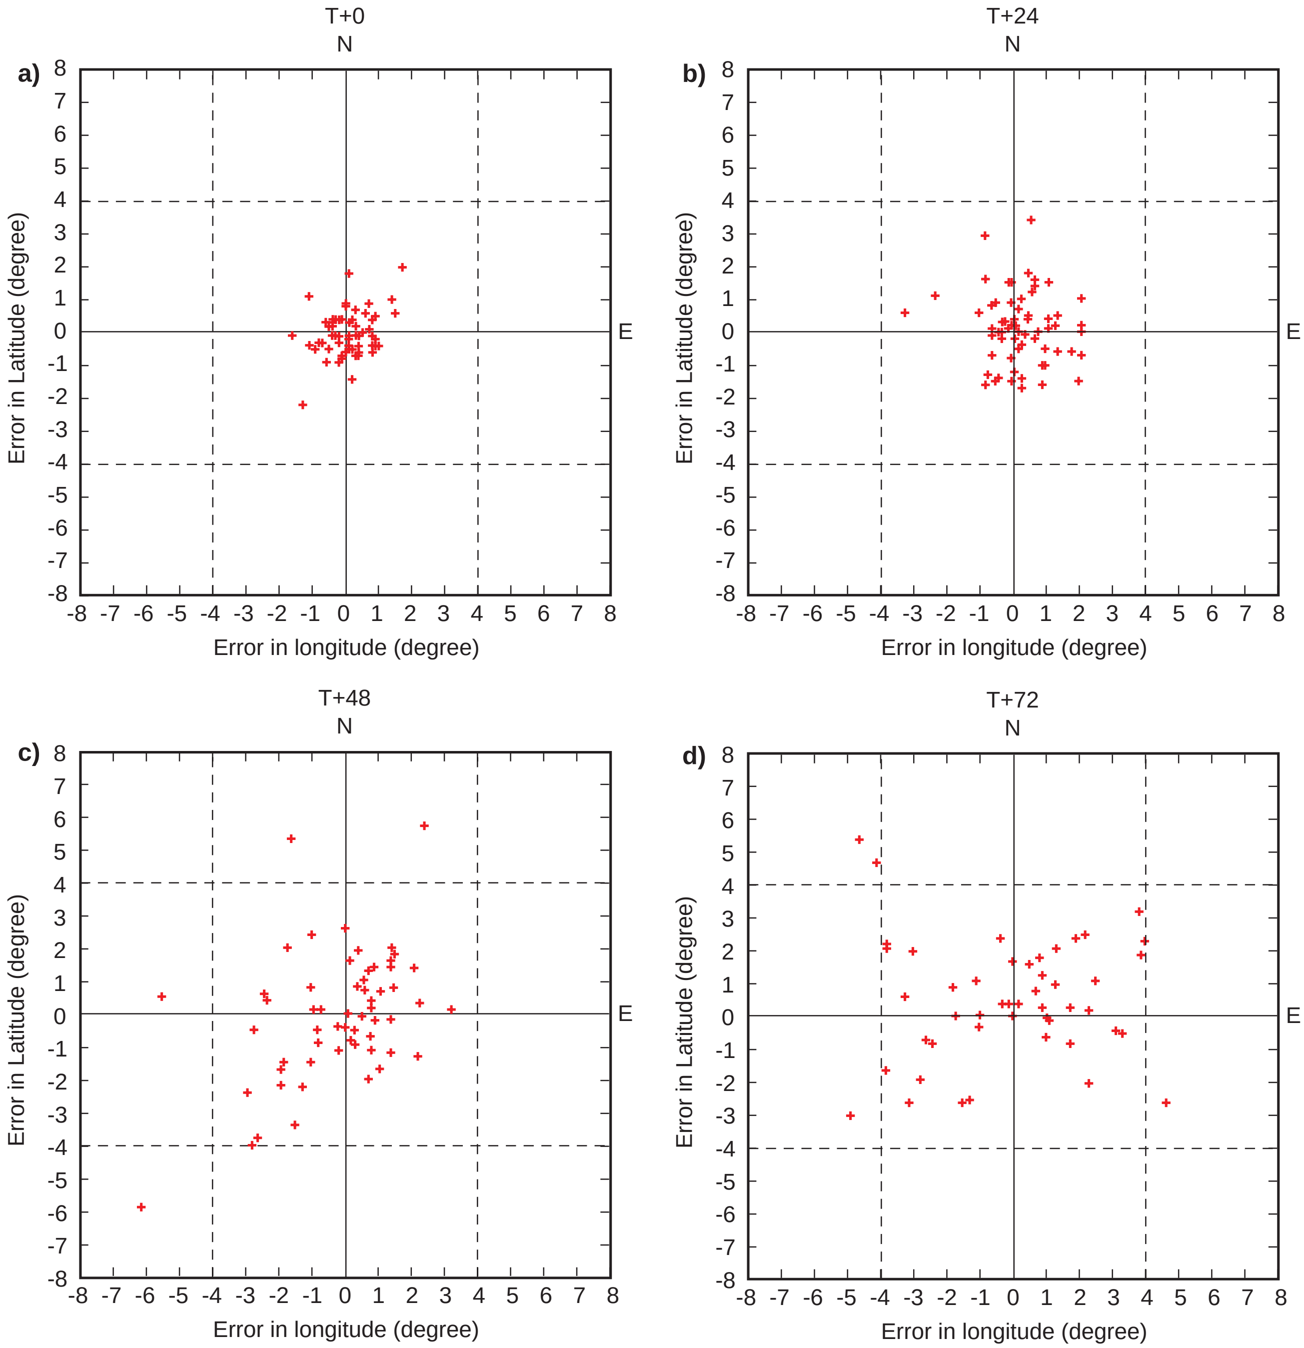

Figure 4: Scatter diagram of forecast position errors of the 850 hPa vorticity fields with respect to the HKO best-track analysed position for (a) T+0, (b) T+24, (d) T+48 and (d) T+72 hours for TCs during 1998 over 0–45°N, 100–140°E.

Kong depends on whether the TC will make landfall to the east or west and, of course, on the distance of closest approach and the strength of the TC. Hong Kong is less susceptible to the damaging effects of TCs if they make landfalls to the east of Hong Kong, owing to the sheltering of the urban areas from the northerlies by mountains. In the case of Babs, the mean forecast errors for the ECMWF model were 93 km, 165 km, 248 km and 350 km, respectively, for the T+0, T+24, T+48 and T+72 hour forecasts; these figures were based on centres deduced from the 850 hPa vorticity field using around ten data points for each forecast hour.

Overall assessment

The recent performance of the ECMWF model in forecasting tropical cyclones over the western North Pacific and the South China Sea has been reviewed using the 2.5° resolution data that is available on the GTS for the period 1991–1998. Since 1994, model forecasts of TC positions have been more accurate than those made from the simple climatology-persistence method over 24-hour to 72-hour forecast ranges. Improvements in model TC track forecasting have been most noticeable since 1996. However, subjective forecasts for the area 10–30°N, 105–125°E have not shown as much improvement during the past two

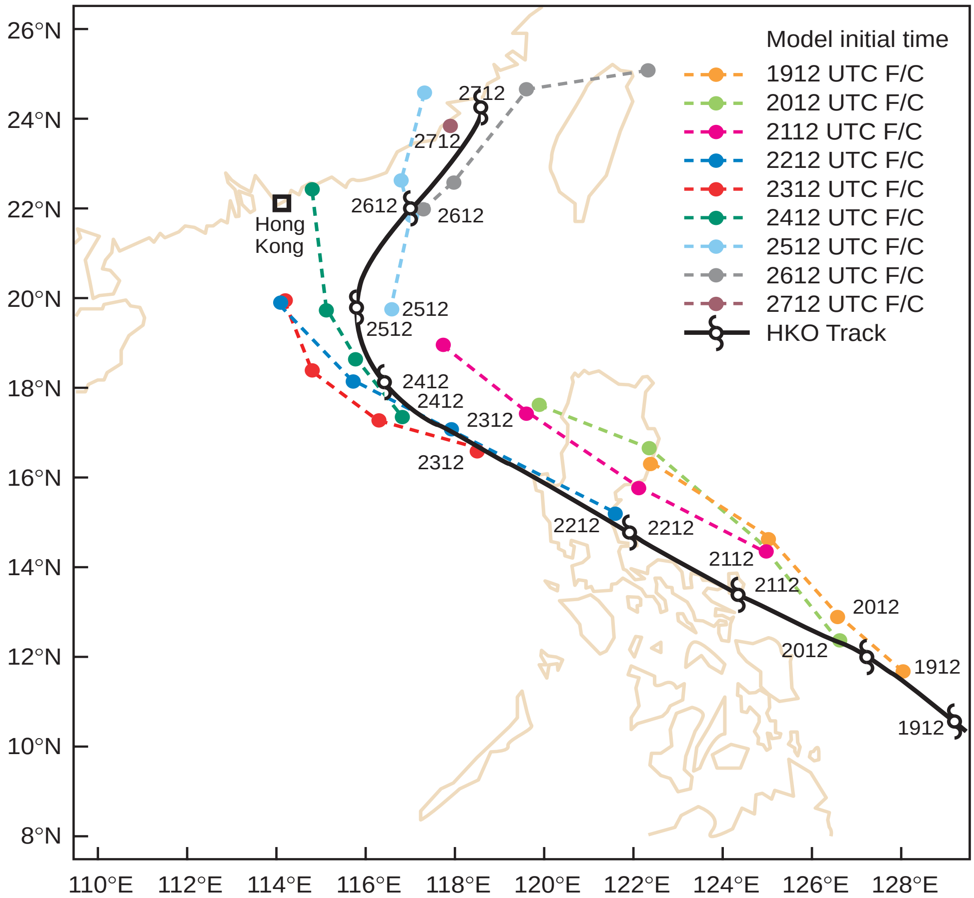

Figure 5: Forecast tracks for typhoon Babs (9811) with different model initial times (dashed lines) based on centres deduced from the 850 hPa vorticity field during the period 19–27 October 1998. The HKO analysed best track is marked as a solid line. Time marks are given in DDHH format, where DD stands for the day and HH for the time in UTC.

decades or so, with mean errors staying in the region of 200 km and 400 km for 24-hour and 48-hour forecasts, respectively. On the model side, the mean errors for 24-hour, 48-hour and 72-hour forecasts were reduced to around 140 km, 200 km and 240 km, respectively, during 1998. However, over the larger verification area of 0–45°N, 100–140°E, the 48-hour and 72-hour forecast errors were larger by around 14% and 30%, respectively, and their standard deviations were much larger.

It is pleasing to note that numerical modelling has become a reliable means of improving TC position forecasting. The verification results presented here should motivate operational forecasters to make better use of model outputs. Accurate forecasting of tropical cyclone motion helps to reduce the associated natural disasters through improved meteorological warning services. This is a most fitting achievement to mark the end of the International Decade for Natural Disaster Reduction (IDNDR). Two decades after Bell's (1979) wise words, numerical weather prediction has, at long last, come of age and is making significant impacts on forecasts of tropical cyclone positions up to 48 hours and beyond.

Acknowledgments

The author would like to thank Mr. C.Y. Lam and Mrs. Hilda Lam for their constructive comments. Preparation of the tropical cyclone dataset by C. K. Chow, C. H. Au and their colleagues is appreciated. ❑

References

Bell G.J., 1962. 'Predicting the movement of tropical cyclones in the region of the China Sea'. Pp. 195–198 in *Proceedings of the inter-regional seminar on tropical cyclones* 18–31 January, Tokyo, Japan

Bell G.J., 1979. Operational forecasting of tropical cyclones: Past, present and future. *Austr. Meteorol. Mag.*, **27**, 249-258

Chan J.C.L. and **H. Lam**, 1989. Performance of the ECMWF model in predicting the movement of typhoon Wayne (1986). *Weather & Forecasting*, **4**, 234-245

Lai S.T., 1995. 'Application of ECMWF 850 hPa vorticity fields to tropical cyclone forecasting in Hong Kong'. Pp. 16-26 in Tropical cyclone programme Report No. TCP–39, *WMO/TD No. 752*, World Meteorological Organisation, Geneva

Lam C.C. and **S.T. Lai**, 1994. 'Use of ECMWF 850 hPa vorticity fields in the forecasting of tropical cyclones and intense lows in June–July 1994'. Pp. 137-156 in *Proceedings of 9th Guangdong-Hong Kong–Macau joint seminar on hazardous weather*, Hong Kong Observatory

Lam C.Y., 1992. 'Operational tropical cyclone forecasting from the perspective of a small weather service'. Pp. 530-541 in *Proceedings of ICSU/WMO international symposium on tropical cyclone disasters*, Peking University Press

Neumann C.J., 1972. 'An alternative to the HURRAN tropical cyclone forecasting system'. *NOAA Tech. Memo.*, NWS SR–62, National Centres for Environmental Prediction, Washington, USA

Queenie C.C. Lam, (Hong Kong Observatory, China)

ECMWF status and plans: a view from the USA

The European Centre for Medium-Range Weather Forecasts (ECMWF) is located outside Reading, England. Its excellent WWW homepage http://www.ecmwf.int/ describes its organisation and products. The following is a brief introduction, and the author's impressions from his visit.

Discussions about the need for a European meteorological centre began under COST (European Co-operation in Science and Technology) in 1967 at about the same time that NCAR was established in the US. After considerable discussion, the European nations decided to focus on operational aspects - medium-range forecasts for distribution to the meteorological services of the 18 Member States (ECMWF also has co-operative agreements with WMO, EUMETSAT, ACMAD, Iceland, Hungary, Slovenia, and Croatia), and the associated research and development of numerical methods.

The stimuli for operational collaboration were the computational intensity and associated costs of mid-range (out to 10 day) forecasts, the benefits of pooling the best talent and then rotating them back to home offices, and an analysis that projected about a 20:1 benefit to cost ratio if the then quality of two-day forecasts could be achieved at six days (this level of improvement in forecasting has effectively been achieved; cost-benefit studies have been discussed but, except for a study of the economic benefit to the UK of hosting ECMWF that the Centre carried out in 1995, they are still in the too-hard box).

The establishing Convention was signed in 1973, and the first forecasts provided in 1979. ECMWF's ruling Council has up to two representatives per Member State; usually one from the nation's met office and one from the responsible ministry (one exception is the UK; the UK Meteorological Office comes under the MoD, and both Council members have so far been from the UKMO). Contributions to the annual budget (£21M in 1998) are determined by GNP (updated every three months by OECD). Thus the UK, France, Italy and Germany are the big members, followed by Spain; Germany's share increased considerably when the calculations switched to post-unification GNP.

The Council normally works well by consensus but votes, when needed, are based on the double–two–thirds rule (⅔ of the funding, plus ⅔ of the members). The budget provides for the facilities, computers, a dedicated telecommunications net (64 Kbps rates to each member as part of the basic subscription; wider bandwidth – used by the UK and France – is paid for by the user nation), and 145 staff positions plus a number of consultancies. ECMWF can also accept external contracts (e.g. NSF, JMA and the EC are contributing to the cost of the reanalysis).

All staff positions are filled by nationals from the member nations, but visiting scientists from other nations are accepted (e.g. Ron Gelaro from NRL Monterey spent a TDY period there and Jim Doyle, also from Monterey, did nine months work on the coupled wind-wave model).

The organisation is simple - a Research Department (38 staff) under Dr Tony Hollingsworth (Ireland) with Model and Data Divisions; an Operations Department (76 slots) under Dominique Marbouty (France) with Computer and Meteorological Divisions; and an Administration Department. The three Departments report to the Director, Dr. Dave Burridge (UK). This has been basically the organisation from the start. ECMWF is small enough that management is quite informal and the staff all know each other; thus there is no need to change style to resolve bureaucratic issues.

We chatted over lunch about the stability and success of the structure. My host Dr Austin Woods cited "Grooming, Gossip and the Evolution of Language" by Prof. Robert Dunbar of Liverpool University (published by Faber) which contends, based on the correlation of brain size and group size in primates, that 150 is about the natural size for a human group. At any rate this has worked exceedingly well for ECMWF for some 20 years, and they want not to grow.

At issue, of course, is tapping into the best European talent – their ability to draw from the 18 Member States is envied in many circles in the US – and, from my brief exposure, although there are national differences the quality and diversity of the staff, as well as the rotation process, are rightfully envied. The multinational structure combined with the links to the European meteorological services provides them with a distinct competitive advantage. The US has not as effectively rationalised its efforts in mid-range (or seasonal) forecasting. We still have enough money that we haven't yet been forced to join forces across organisational boundaries (e.g. NOAA, Navy, and NASA) to the detriment of out nation skill.

One concern, although it hasn't impacted yet, is the potential consequence of changing policies on data rights. Satellite, radiosonde, etc. data are the lifeblood of any forecasting group, and restrictions or excessive subscription costs could severely impact the product. As indicators of change my hosts cited ECOMET (an "Economic Interest Grouping" (private) established to commercialise the products of the European weather services) and EUMETNET (an umbrella organisation of weather services to foster collaborative projects and potentially the embryo of a European Meteorological Organisation - there are 25K meteorologists in Europe servicing the 400M population compared with some 5–10K in the US for its 280M population; some degree of rationalisation and job sharing will eventually occur). I have heard of similar concerns concerning data management at NESDIS and the Navy in the US, and the US policy (free access to government data) is certainly at odds with the trends in the rest of the world.

ECMWF's main product is its daily medium-range (10-day) forecast, distributed to the Member State met. services. They also produce lower-resolution ensemble predictions (currently 50 integrations of the prediction model) that are starting to be used in innovative ways by some of the larger met. services. Research associated with these efforts focuses on exploitation of satellite data for initial conditions, data assimilation techniques, physics, and numerical methods to enhance speed, resolution and accuracy.

One initially unappreciated fallout of the accuracy of ECMWF analyses and predictions is their ability to precisely and rapidly spot anomalies in data. The information is routinely fed back to the collectors and has resulted in significantly improved quality, and also in the re-calibration of some of the space-based sensors on ERS-1 and -2.

Another extremely valuable product is the ECMWF's archive of data and forecasts. This archive is widely used throughout the world, and the collaboration with the user meteorologists (and climatologists, of course) has beneficial feedback to ECMWF's research. The 15-year reanalysis (1979–93) using the baseline 1995 operational system (2.5°, T106L31) is of tremendous value, and is already being exploited in the ECMWF's own seasonal prediction work.

Of particular interest to the Navy are the ECMWF's wave predictions using the WAM model. Operational forecasts out to 10 days started in 1992; the global resolution is now 1.5°. They assimilate ESR-1 SAR data and verify against observations from ERS-2 and buoys. The WAM 1st guess (before assimilation) standard deviation versus the ERS-1 wave heights is now about 40 cm compared with a metre or so in the days of SEASAT. Although the Navy also uses WAM, the ECMWF's forecasts are more accurate across a wider range of sea states, perhaps because the ECMWF uses the wind field for forcing, while the Navy uses calculated stresses. This is in consonance with the overall ECMWF philosophy of predicting and assimilating directly observed data, rather than values calculated from the data.

Other recent initiatives

The 1979-1993 reanalysis using the 1995 operation system has been completed, many enquiries concerning the data and results are being received and work is under way on ERA-40. This is an internationally funded and supported effort, and the resultant consistent dataset generated with a stable modern assimilation system should be of tremendous value to researchers world-wide.

The seasonal forecasting project continues. Assessment of atmospheric predictability on seasonal time scales, given sea surface temperature and soil moisture, has produced encouraging results. ENSO impact is now widely appreciated; at issue are the nature and location of other drivers, and the predictability in areas such as Europe that are less immediately impacted by ENSO. Initial results have indicated predictability, at least in cases of strong interannual changes and large SST anomalies. The meteorological services have conducted similar studies with the same boundary conditions to test the influence of different models. The second part of this project has been the development and operation of a coupled ocean-atmosphere model. The German HOPE is the basis for the ocean model, and the French OASIS is used as the coupling interface.

Personal comments

Like most workers in the field I have heard about ECMWF and taken advantage of its available products for many years. This was my first opportunity to visit, and the reality more than lived up to my expectations. Perhaps more impressive to me, even than the well-known excellence of the products and the research, is the stability and the quality of the management. Most research institutions have gone through, or shortly will go through, major perturbations because of changes in national research policies and reductions in funding. While ECMWF has not been immune to such influences, its international structure and the commitment of its leadership to spurn growth and stick to producing top quality research, plus a product that is most highly valued by its members, have minimised disruptions. ECMWF was well designed at the start, is well situated in an environment conducive to top quality research, has been and remains well managed and staffed, and knows where it's headed. I can't say that about many institutions I've visited over the last several years.

ECMWF's selection of the Fujitsu VPP represents a significant change for those who care about supercomputers. ECMWF was an obvious key target for the Japanese because it is a highly respected and multinational world leader.

ECMWF research promises consistent and continued progress in medium-range forecasting. Consistent accuracy at greater than 80% levels at five days is in itself astounding, at least to me, and the combination of new machines, new data, new algorithms, new assimilation techniques and ensemble forecasts should press the limits of predictability. Over 25 years this field retains every element of its excitement and promise. The fully coupled wind-wave model became available in 1998; the wave forecasts were vastly improved and are seeing much wider application, with commensurate feedback on the research.

ECMWF's venture into seasonal forecasts is likewise exciting. While the initial justification for the seasonal research was improvements in the medium-range forecasting capability (and vice versa), The potential for effective operational seasonal forecasts in Europe and Africa has immense ramifications for economics and quality of life.

ECMWF is obviously in the best position to exploit its own reanalyses, which will become the world standard. Others' use of the reanalyses will only enhance and improve their own research, at negligible cost. Also, ECMWF's entry on the scene of the coupled ocean-atmosphere models gives tough competition to US researchers. Their move into this work at the time of the reanalysis and the acquisition of a significant increase in computing power was prescient. ❑

(This is an updated version of an article originally published in March 1996 by Dr Craig Dorman, Chief Scientist and Technical Director, US Office of Naval Research Europe. With the kind permission of Dr Dorman, changes have been made to bring the technical information concerning the computer system etc. up to date. The original article can be read in http://www.ehis.navy.mil:80/dnews2.html. The opinions and assessments are solely those of the author and do not necessarily reflect official ONREUR, Navy, or US Government positions. ONR Europe is on the World Wide Web. Our home page contains information about European activities, conferences, and other newsletters. Find us at: http://www.ehis.navy.mil/).

Craig Dorman (US Office of Naval Research Europe)

The ECFS file management system

The ECFS file management system was developed as part of the Data Handling System (DHS) project to provide an archive facility for private user files. The initial ECFS-1 system was introduced in October 1996 to support the introduction of the first Fujitsu VPP systems and replace the CFS-based ECFILE system. The ECFS provides a UNIX-like interface to files stored on file systems managed by IBM's Hierarchical Storage Manager (HSM), which manages a very large disk cache and uses the ADSM system to archive files to Magstar tapes within two large IBM robots.

The ECFS started off with one RS6000 server managing two HSM file systems: **ecfs_shared** for those parts of the operational and research data that could not be accommodated in MARS; and **ecfs_users** for the private user data. It soon became apparent that a single HSM file system on the IBM AIX system could not be easily managed once it had grown to much beyond 750,000 files.

The ECFS-2 system was developed and introduced at the beginning of 1998 to enable the ECFS to support overflow file systems, and manage file systems spread over multiple DHS servers. This has been achieved by having an ECFS server and database on each DHS server supporting the ECFS. The database has an entry for each directory indicating which file system(s) should be referenced for files in that directory and to which file system new files in that directory should be written.

ECFS-2 now supports 8 HSM file systems spread across two RS-6000 servers named **dartagnan** and **constance**. This has entailed migrating over a million files from the original two HSM systems into new file systems dedicated to operational, reanalysis, seasonal and Member States' data.

We have nevertheless still been struggling to keep the original **ecfs_users** and **ecfs_shared** file systems below a million files. At the beginning of 1999 **ecfs_users2** was established as an overflow file system such that all new directories created by internal ECFS users are put on **ecfs_users2**. So far, 30,000 directories and 225,000 files have been stored on **ecfs_users2**. This growth is now being accelerated by having the most active users write all new files to this file system.

The ECFS now holds 3 million files and 30 Terabytes of data, and typically handles up to 18,000 save/retrieve operations per day transferring up to 200 Gigabytes of data with a peak hourly transaction rate of 1000 save/retrieve operations and 20 Gigabytes of data. This load is evenly split between saves and retrieves; this contrasts with the original aim of developing the ECFS as an archival storage manager from which files would be retrieved only occasionally. Consequently, the current ECFS can cope with the expected growth of data, but would find it difficult to manage the corresponding growth in the number of files stored and retrieved.

The next big migration entails establishing a new **ecfs_rdx** file system and migrating the 385,000 files and 25,000 directories belonging to RDX which are in the **ecfs_shared** file system. However, such migrations barely cope with the growth, and are not a long-term solution.

We are therefore enhancing the ECFS to reduce the number of files in each HSM file system by creating 'tar' files of suitable candidate files within the ECFS directories. This will be done in a way that is completely transparent to the user; all ECFS commands will work in exactly the same way whether they reference files within a tar'd directory or not. A command making a first reference to an ECFS tar'd directory may be delayed while that tar file is copied from tape onto disk cache, but any further references to that ECFS directory will experience no delay.

During 2000 it is planned to develop ECFS-3 to enable a single intelligent ECFS server to manage ECFS client transaction processes across multiple DHS systems which may be a mixture of disparate file-system technologies.

Strategies for Reducing ECFS Growth

The ECFS is being used in a manner for which it was not designed. It was designed to be an ARCHIVAL system, to store reasonable numbers of relatively large files and libraries and to provide occasional retrievals as necessary.

Unfortunately, the ECFS has become regarded as an extension to the file systems on the Centre's machines, and as a facility for providing an extended secure SCRATCH file-system capability. The resulting storage and retrieval of such a large number of files, many of which are very small, causes many operational problems and can tie up valuable resources for long periods of time.

Analysis of daily activity has shown that some users retrieve the same file from the ECFS many times a day. This again is very wasteful of the ECFS resources. It would be much better to keep the original file on-line on a system where it can be easily accessed when it is needed and to keep an archive copy in the ECFS for those (hopefully) rare occasions when the file has been removed from the local disks.

By following these simple recommendations you can help to ensure that the ECFS will provide a fast and effective service to all users, and at the same time almost certainly improve the turnaround time for your own jobs.

a) Rather than save many small files in the ECFS, let them accumulate on a secure local disk system and then 'tar' or 'cpio' them into a file that can be archived into the ECFS (or use the 'ecfsdir' utility provided for this purpose).

b) If you continually use a file but it is too large to save in your HOME file system, save it on a 'temporary' file system and keep a copy of it in the ECFS. Then modify your scripts to only retrieve it from the ECFS if it has been removed (using the ecp -n option) or is out-of-date (using the ecp -t option). Temporary file systems that can be used in this way include:

$TEMP and $LTEMP	on the VPPs
$SCRATCH	on the SGI servers

c) If you want to save a file in the ECFS that is only required for a few weeks then save it into the ECFS 'temporary' file system 'ectmp', rather than in 'ec' (files within 'ectmp' are automatically deleted 90 days after the date of last modification).

d) If you are planning a new project which requires a large amount of computer resources, in particular ECFS resources, first contact User Support for advice. User Support will then consider your requirements, help in the optimisation of your work and advise on the best use of the resources at your disposal. When appropriate for really large accumulations of ECFS data, the ECFS group will become involved in the planning, and may provide special facilities or resources.

e) Try to structure your directories such that there are between 20 and 200 files in each directory; the management of HSM file systems becomes very inefficient when they have a very large number of directories with few files. Also, once a directory has much more than 200 files, the time to list and manage that directory becomes unreasonably long.

It is not necessarily the amount of data in the ECFS that causes operational problems, it is rather the number of files, how those files are organised within the file system, and the frequency with which those files are accessed. Please try to help us improve the overall service by making more efficient use of the ECFS. ❑

Tony Stanford

Fortran Developments in IFS

The Integrated Forecasting System (IFS) has existed as a jointly developed application for more than 10 years (for a review of its history, see Newsletter No 75, Spring 1997). Over that period the content of the implementation language has changed considerably. Initially, Fortran 77 was used, but with a liberal addition of extensions as supplied by the CRAY compiler of that time; these were principally used to provide dynamic memory management so that the IFS could be coded independently of the chosen resolution.

In IFS cycle 14 (1995), these non-portable features were removed by converting the relevant code to conform to the Fortran 90 standard. Primarily, these were POINTERS (Fortran 90 style) and ALLOCATABLE arrays. 'Automatic' arrays and namelists that were CRAY extensions became part of the standard along with 'bit intrinsics'.

Since then, several more steps have been taken to introduce gradually more Fortran 90 features, where it was felt that these would result in better and safer coding styles. For example, in Cycle 18, MODULES, and the associated USE statement to identify which MODULE elements were referenced in the code, replaced COMMON blocks.

This article outlines the latest step in this process, introduced in Cycle 21 and released in March of this year. As usual, all developments are the result of collaborative work with Météo-France, who have contributed significantly to the effort required. All of the changes described were made through the use of *awk* and *perl* scripts operating on the Cycle 20 source code.

New Fortran features

The principal new Fortran 90 features that have been incorporated into Cycle 21 are:

- IMPLICIT NONE
- Explicit KIND values
- Free source format

IMPLICIT NONE

This statement is added to every subroutine and disables the normal Fortran typing conventions. Thus, ALL variables must have an explicit type declaration. Although this introduces an initial overhead for the programmer when writing a new routine, it is generally accepted to be the best working practice. Generally, there are fewer run-time errors due, for example, to mis-typing variable names, since such problems are trapped at compile time.

Explicit KIND Value

This is the Fortran 90 method of specifying the precision of variables and replaces, for example, REAL*8 as a method for specifying double precision. The new declaration is:

```
REAL (KIND=JPRB) :: TEMPERATURE (NLON , NLEV)
```

The parameter JPRB is defined by use of an intrinsic function:

```
INTEGER,PARAMETER :: JPRB = SELECTED_REAL_KIND
(13,300)
```

Which defines a kind value (available on the platform being used) of a real data type with decimal precision of at least 13 digits and a decimal exponent of at least 300. On machines that have been built to follow the IEEE floating-point format, this will result in 8 byte (double precision) real numbers.

To make the coding simpler, the IFS has some defined MACROS: e.g.

```
#define REAL_B     REAL(KIND = JPRB)
```

so that a declaration for an array becomes:

```
REAL_B ::    TEMPERATURE (.....)
```

Note that the macros defined are NOT part of the Fortran 90 language. However, the ability of compiling systems to expand user-provided macros is fairly widely available and does not seem to harm portability.

Constants defined in the code can also be subject to a KIND value:

```
PI = 3.14159265_JPRB
```

The statement

```
PI = 3.14159265
```

will cause the constant to be represented at the default precision (4 byte single precision on the FUJITSU).

Prior to the introduction of KIND in Cycle 21 of the IFS, nearly all declarations were for ‘single precision’ real numbers and we relied on a compiler feature (auto double) to convert these into double precision. Since not all Fortran 90 compilers support this option, the code is now more portable. Additionally, it is now possible to mix single and double precision variables inside the same code, leading to possibilities for saving memory.

The use of macros has been extended further by providing a number of commonly used constants:

```
#define _ONE_     1.0_JPRB
```

which can be used in a statement thus:

```
REAL_B :: X , Z
Z = X + _ONE_
```

To be consistent with the appropriate variable kind, it is important that the generic form of all intrinsic functions is used. Thus, for example, the function names EXP,LOG and SQRT are used regardless of their argument KIND.

Free source format

Free source format relaxes the fixed column requirements of traditional Fortran:

- Lines may start in any column
- A maximum line length of 132 characters is allowed
- Continuation lines are indicated by terminating a line with &
- Comments can occur anywhere in a line after !
- Blanks are significant

To aid readability, the IFS documentation standard now specifies that:

- Non-indented lines start in column 1
- Continued lines also commence with &
- Indentation by 2 spaces is used for DO, IF and WHILE statements

Additional Fortran 90 features

The use of array syntax in IFS has so far been limited to the initialisation of arrays and other simple statements. There is some concern that compilers in general do not yet create as efficient code from array syntax as from traditional DO loop syntax. Additionally, there are some limitations that are significant. For example, in ‘physics’ type conditional code, it is common to have nested IF..THEN..ELSE constructs. The equivalent array syntax structure is WHERE..ELSEWHERE which cannot be nested. However, that restriction is removed in the Fortran 95 standard.

A large number of array intrinsic functions have been added (for example MATMUL). While it is good to use these for portability reasons, there are again performance concerns since the code generated may not be competitive with the optimised library routines currently in use. ❑

David Dent, Mats Hamrud

ECMWF Publications

Technical Memoranda

No.277 **Morcrette, J-J.**: On the effects of the temporal and spatial sampling of radiation fields on the ECMWF forecasts and analyses *June 1999*

No.278 **Morcrette**, **J-J.** and **C. Jakob**: The response of the ECMWF model to changes in cloud overlap assumption *April 1999*.

No.279 **Buizza**, **R.**, **M. Miller** and **T.N. Palmer**: Stochastic representation of model uncertainties in the ECMWF Ensemble Prediction System *April 1999*.

No.280 **Buizza**, **R.**, **J. Barkmeijer**, **T.N. Palmer**, and **D.S. Richardson**: Current status and future developments of the ECMWF Ensemble Prediction System *May 1999*.

No.283 **Temperton**, **C.**, **M. Hortal** and **A. Simmons**: A two-time-level semi-Lagrangian global spectral model *June 1999*.

No.285 **Cullen, M.J.P.**: On the accuracy of the semi-geostrophic approximation *July 1999*.

No.286 **Buizza**, **R.** and **A. Montani**: Targeting observations using singular vectors *August 1999*.

ECMWF Proceedings

Recent developments in numerical methods for atmospheric modelling, *7-11 September 1998*

Workshop on diagnosis of data assimilation systems, *2-4 November 1998*

Other reports

Application and verification of ECMWF products in Member States and Co-operating States, Report 1999

Seasonal forecasting users meeting, 17-18 June 1999

ECMWF Global Data Monitoring Report July 1999