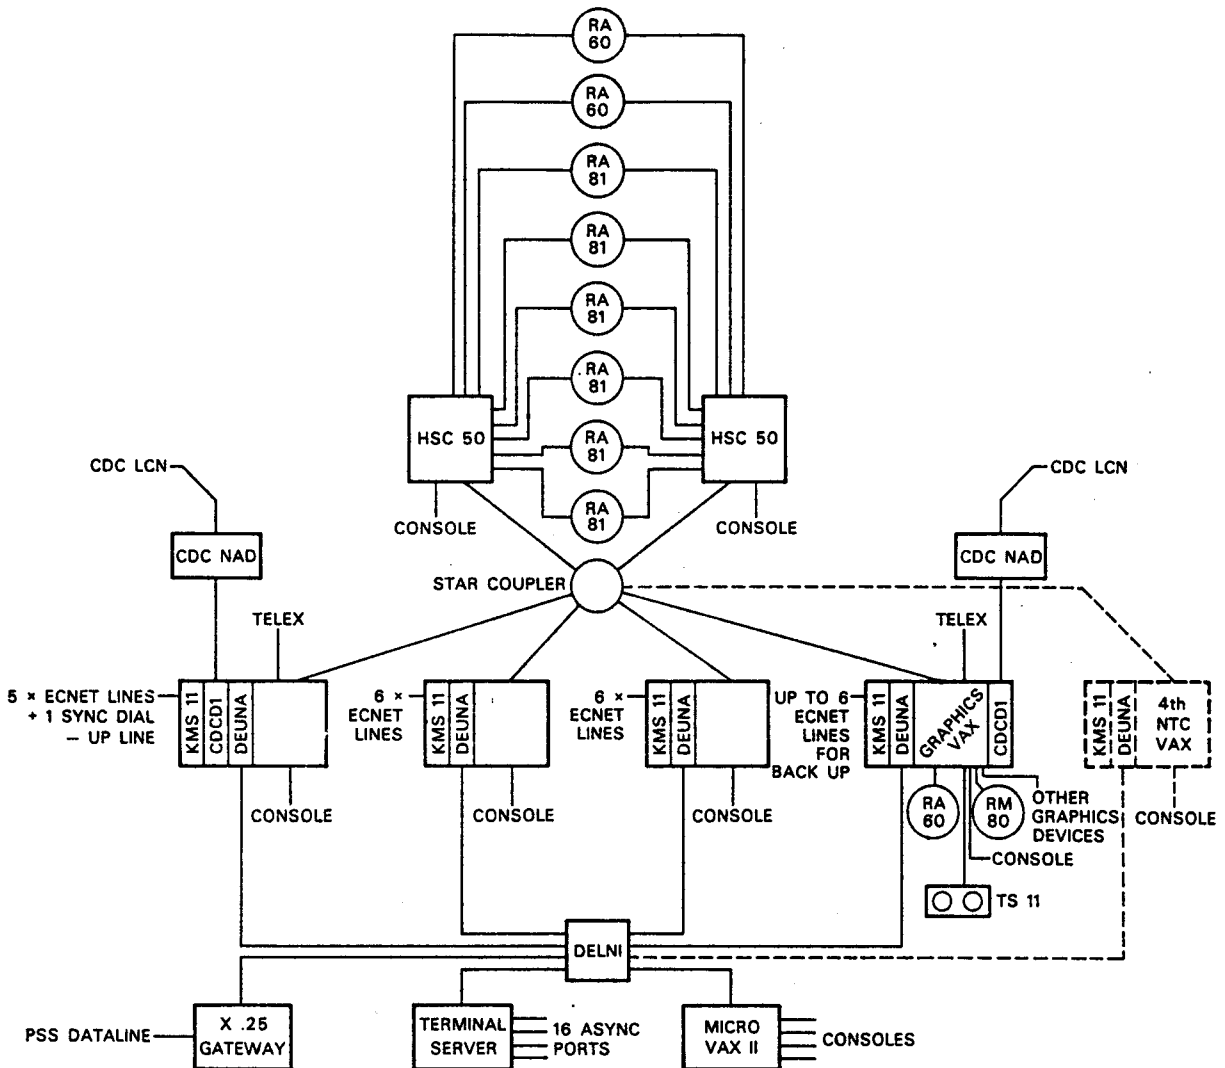




ECMWF NEWSLETTER

Shinfield Park, Reading, Berkshire RG2 9AX, England. Tel: U.K. (0734) 876000, Int. (44 734) 876000, Telex: 847908

Number 31 - September 1985



IN THIS ISSUE

METEOROLOGICAL

Changes to the ECMWF operational forecasting system	2
The introduction of and initial experience with the high resolution T106 forecast model	3
Investigation of systematic errors by relaxation experiments	9

COMPUTING

The new ECMWF telecommunications system	13
---	----

COMPUTER USER INFORMATION

Ideas for reducing the number of magnetic tapes used	17
How to use multifile tapes	19
Still valid news sheets	24

GENERAL

ECMWF publications	25
Calendar of events at ECMWF	25
Index of still valid newsletter articles	26

COVER: Configuration of the new telecommunications system

This Newsletter is edited and produced by User Support.

The next issue will appear in December 1985. --

The Centre is pleased to announce the completion of negotiations for the acquisition of a new telecommunications system. The purchase contract was signed on 11 July and work towards the availability of a replacement system in 1986 is progressing well. Details of the configuration of the new system are given on page 13.

Following the introduction into operation of the T106 version of the operational forecast model on 1 May, the last issue of the Newsletter (No. 30, June 1985) gave a necessarily brief summary of its initial performance. Several months have now passed and additional studies have been undertaken. Their findings are described in an article on page 3.

* * * * *

CHANGES TO THE ECMWF OPERATIONAL FORECASTING SYSTEM

Recent changes

(i) After the successful implementation of the high resolution forecast model (T106) on 1 May 1985, the performance of the analysis and forecasting system was closely monitored during the subsequent months. Verification results for the free atmosphere for May and June 1985 were in agreement with the performance of the model during the development and implementation phases. May 1985 gave the best result for that month since the beginning of operational forecasting, when using the anomaly correlation of height at 500 mb as a measure of the quality of the products. However, some problems with near-surface weather parameters (e.g. the experimental products: temperature at 2 m and precipitation amount) were identified. Minor modifications to the convection scheme and the treatment of the evapotranspiration were introduced on 18 July 1985 resulting in a reduction of near surface temperature biases and the over-forecasting of precipitation.

Forecasts during June, but especially in July, may have been adversely effected by the use of erroneous sea surface temperatures. The error was corrected on 31 July 1985.

Further information on the performance of the new forecast model is given in the article on page 3 of this Newsletter.

(ii) The day 3 forecast of the 850 and 200 mb wind field was added to the ECMWF GTS products on 1 August 1985. The ECMWF grid code products currently available on the GTS are:

for the northern (20-90°N) and southern (20-90°S) hemisphere:

mean sea level pressure and 500 hPa height analyses and forecasts for 24, 48, 72, 96 and 120 hours ahead;

for the tropical belt (35°S-35°N):

850 and 200 hPa wind vector analyses and forecasts for 24, 48 and 72 hours ahead.

Planned changes

No further major changes are planned for this coming quarter.

Horst Böttger

* * * * *

THE INTRODUCTION OF AND INITIAL EXPERIENCE WITH THE HIGH RESOLUTION
T106 FORECAST MODEL

The new model

On 1 May 1985 the Centre's high resolution (T106) version of the spectral forecast model was brought into operation. This followed a period of development and testing on the Centre's Cray X-MP/22 computer, the capacity of which is used to the full by this model. Several improvements and enhancements of the numerical methods brought about a highly efficient version of the model, so that it now takes no longer to run than the T63 model on the Cray 1. In parallel to the enhanced resolution, a revised scheme for cumulus convection, including shallow convection, a modified cloud scheme and an altered orography were developed. After obtaining encouraging test results, these modifications were introduced together with the higher resolution. The extensive tests carried out on the impact of resolution and choice of orography also provided a much deeper insight into the relative effects of these changes.

Subsequent testing under quasi-operational conditions during a parallel run of both old and new models confirmed the effects of these changes on the mean behaviour of the model's 10 day forecasts, which became more realistic, and systematic deficiencies in mid-tropospheric temperatures and tropical circulation were reduced. The expected improvements in scores for Northern Hemisphere extratropical forecasts, however, were not fully realised during the parallel run. In the light of results of objective verification during May, when the highest ever 500 mb height anomaly correlation scores for the Northern Hemisphere in this month were found, the modest improvements seen during the parallel run may be linked to the particular synoptic situation prevailing during that time.

One aspect of the new model was expected to be particularly beneficial in terms of the practical usefulness of the forecast products: the higher resolution brought about a more realistic representation of both the land-sea mask and orography and frontal systems became sharper and better defined. This, in turn, would lead to an increased usefulness of the so-called direct weather parameters, such as 2 m temperature, rainfall, 10 m wind and cloudiness. Synoptic assessments of experimental forecasts had indicated substantial benefits in these areas.

Initial operational experience with the new model

Both the objective verification results and the subjective impression of synoptic behaviour during May and, to a lesser extent also during June and July, proved that the new model's forecasts of the free atmosphere were of very high quality, as can be seen from the daily objective verification scores in Fig. 1. It should be noted that the improvement compared to May and June 1984 may be partly attributable to other changes introduced in the meantime, such as the new radiation scheme, and to the interannual variability of predictability. It soon became apparent, though, that the modified convection scheme was over-active in producing rainfall, particularly over Europe. Linked to this problem was the effect that too much of the available radiation was used to evaporate the surface moisture, leading to a marked negative bias in the 2 m temperature in daytime, in particular when dry, sunny conditions were observed.




Fig. 1: Comparison between 1984 and 1985 of daily RMS errors at 500 mb for the northern hemisphere

This problem was identified both within ECMWF and by some Member States. On 18 July 1985 small modifications were made to both the convection scheme, so that a threshold of the moisture content for an atmospheric layer is now required for precipitation, and to the treatment of evapotranspiration, where an "efficiency factor", depending on temperature, was introduced. These two modifications are seen as a temporary solution and further research is being carried out in this field. A reduction of the biases both in rainfall and temperature was found after the introduction of these modifications. Over-forecasting of precipitation was eliminated for the northern part of central Europe, and the 2 m temperature bias was much reduced there and in Fennoscandia.

The reduction of the positive bias in predicted rainfall has not been achieved at the risk of omitting important precipitation events. When the alpine and northern Mediterranean area was affected by disastrous storms and floods during the first week of August, predicted rainfall amounts out to 3 or 4 days ahead ranged from 50-90 mm in 24 hours. An example of these predictions is given in Fig. 2, in which observed and predicted rainfall for central Europe on 7 August 1985 are compared.

Unfortunately, a combination of unforeseen circumstances led to the rejection of sea surface temperature analyses received from NMC Washington by the Centre's data assimilation system. For the period 31 May to 30 July 1985 the analysed values were kept constant. At the end of the period, this led to discrepancies of up to 10°C between observed and "analysed" temperatures over some sea areas.

At this stage, two forecast experiments using climatological sea-surface values for July were performed to evaluate the impact of this error on the forecasts during the period. It appears that the main effect of the colder sea surface was an intensification of the oceanic anticyclones. This affected mainly the lower levels, where RMS error and anomaly correlation coefficients of height fields were adversely influenced. Fig. 3a shows the differences at D+8 for 1000 mb height between the operational forecast from 23 July 1985 and an experimental forecast with climatological sea-surface temperatures. A comparison with the mean D+10 errors for all July forecasts (Fig. 3b) indicates the systematic nature of the forecast error introduced by the incorrect sea-surface temperatures. The extension of the anticyclone from the eastern Atlantic to well over the British Isles led to a northward shift of the typical cyclone tracks, and the strengthening of the anticyclone over the western Pacific would have had a similar effect in the area of Japan. Following the correction of this error on 31 July 1985 there appears to have been noticeable improvement in the daily verification results for 3-day forecasts, as can be seen from Fig. 4 for the anomaly correlation at 500 mb for the northern hemisphere.

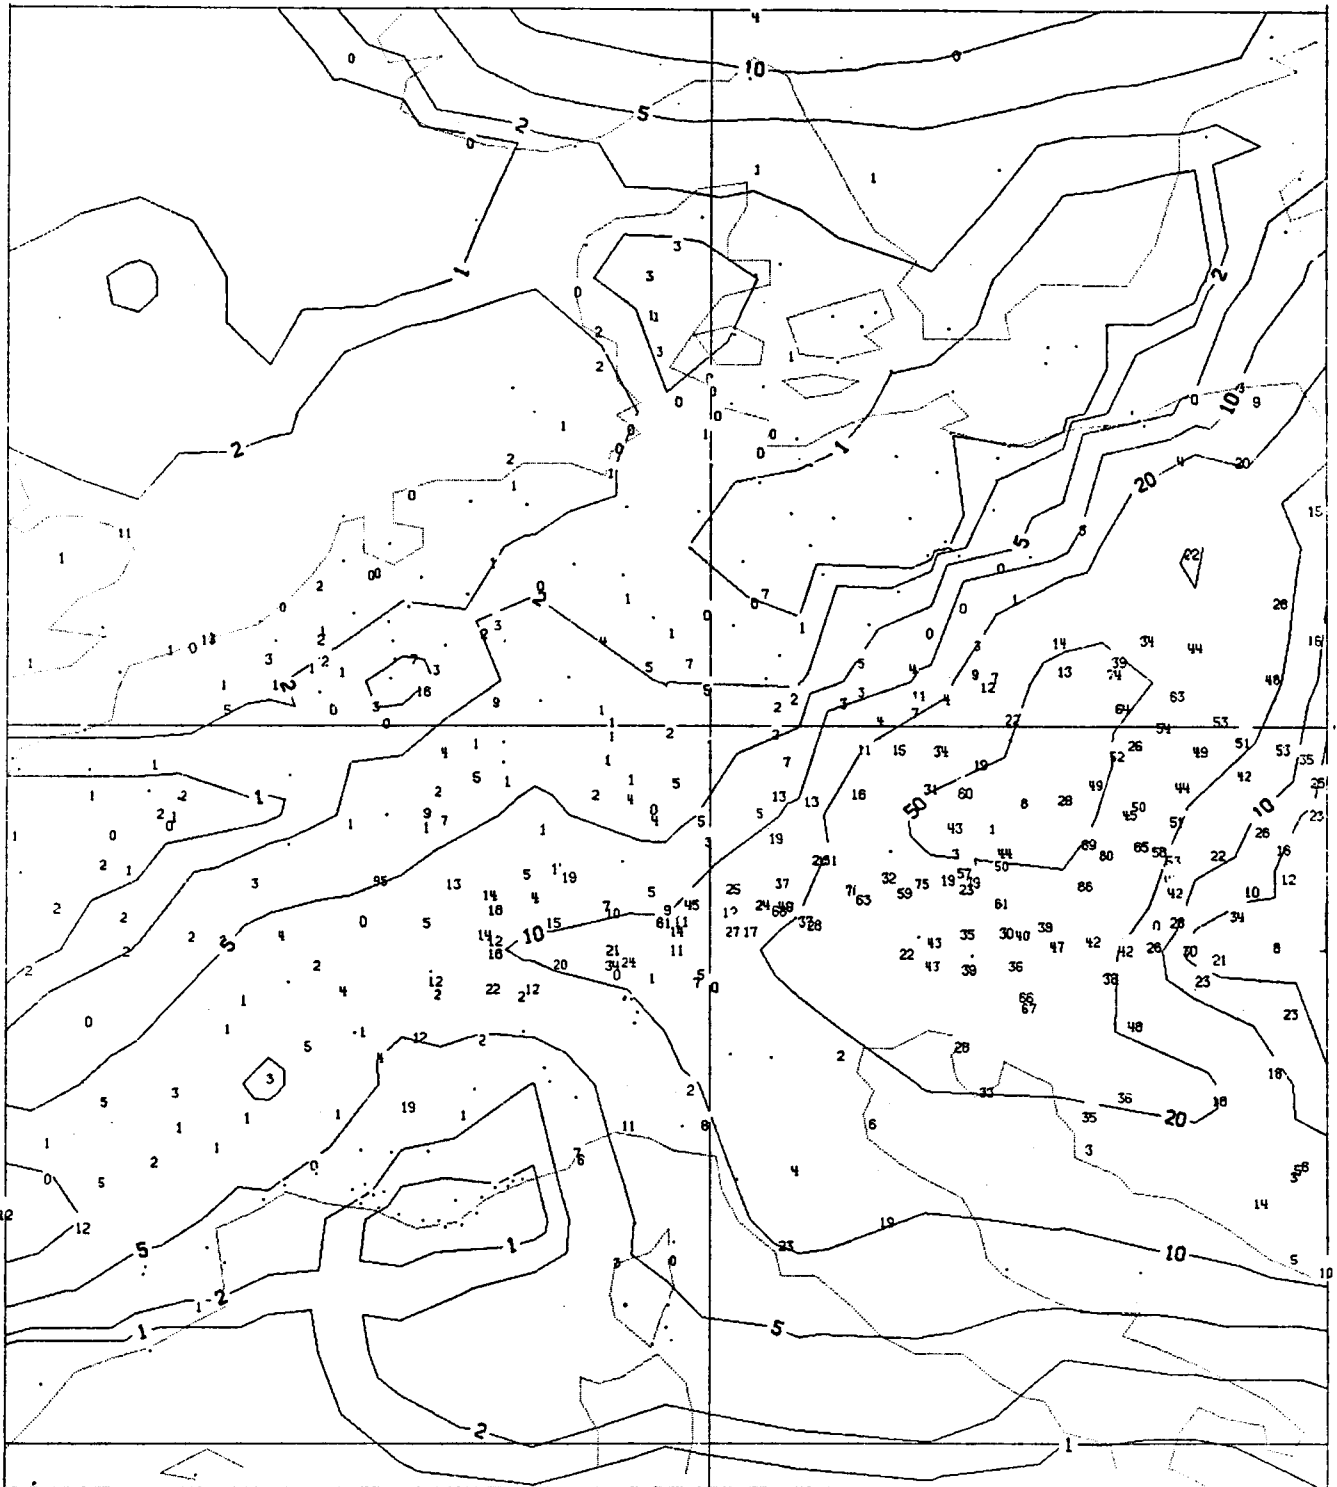


Fig. 2: Observed and predicted rainfall over Central Europe, when serious flooding occurred in the Alps





Fig. 3a: Difference between experimental and operational forecast for 1000 mb height, D+8 from 23 July 1985

Fig. 3b: Mean error of D+10 forecasts at 1000 mb, for July 1985



Conclusions

After elimination of some typical "teething problems" the forecasts produced by the new model are being found to give an improved performance and provide increased usefulness. Further research will be carried out on the treatment of surface exchanges and convective processes, and the remaining problems with rainfall predictions should be solved as a result of this work.




Fig. 4: Daily verification scores 17 July-16 August 1985

- Herbert Pümpel

INVESTIGATION OF SYSTEMATIC ERRORS BY RELAXATION EXPERIMENTS

Although diagnostic investigations of model forecasts have led to comprehensive documentation of systematic errors, little is known for certain about the sources of the errors. In order to improve our understanding, experiments have been carried out in which the forecasts were relaxed towards objective analyses in regions which are assumed to be the locations of major error sources. At each timestep the tendencies of temperature and/or momentum received an extra contribution which was proportional to the current error. For example, the equation for temperature becomes

$$\frac{\partial T}{\partial t} = \dots + \lambda (T_a - T)$$

where T_a is the analysed temperature and λ is a relaxation co-efficient which determines the magnitude of the remaining forecast error after relaxation.

The experiments were performed with the T42 version of the spectral model - increasing the horizontal resolution beyond T42 has little effect on the character of the systematic error.

For the investigation of error propagation from a specific region three types of experiments were carried out. One forecast with global relaxation (A), one forecast with regional relaxation (B) and one control forecast without relaxation (C). By looking at the appropriate differences between these experiments, the total error (C-A) can be subdivided into errors originating from the relaxed area (C-B) or from the non-relaxed area (B-A).

Patterns of the mean wind error had suggested that a good deal of forecast errors in mid-latitudes are forced by errors in the tropics. To test this hypothesis the forecast fields of mass and wind were relaxed towards analyses in a tropical belt from 20°S to 20°N.

The large scale mean height errors for a case with large errors are shown in Figs. 1 to 3. The total error (Fig. 1) represents quite well the mean seasonal error distribution of heights which is too low over NW Alaska and the Norwegian Sea. The Pacific error pattern which indicates a northward shift of the subtropical jet is especially suspected of being forced by tropical errors. However, after separating the total error into the extratropical (Fig. 2) and tropical part (Fig. 3) we must reject the hypothesis of tropical errors forcing large errors in the extratropics. The tropical errors in the Pacific and in the Atlantic represent only a small contribution to the total error compared to the extratropical errors.

For several years orographic forcing has been suspected of being an important source of the systematic error of the model. A first step in testing the importance of orographic forcing for the mid-latitude forecast errors was




Fig. 1: Total forecast errors of height at 500 mb, 5 day mean, spectrally filtered (T10). Units: dkm




Fig. 2: Extratropical forecast errors of height at 500 mb. (Forecast with tropical relaxation minus analysis.) 5 day mean, spectrally filtered (T10). Units: dkm

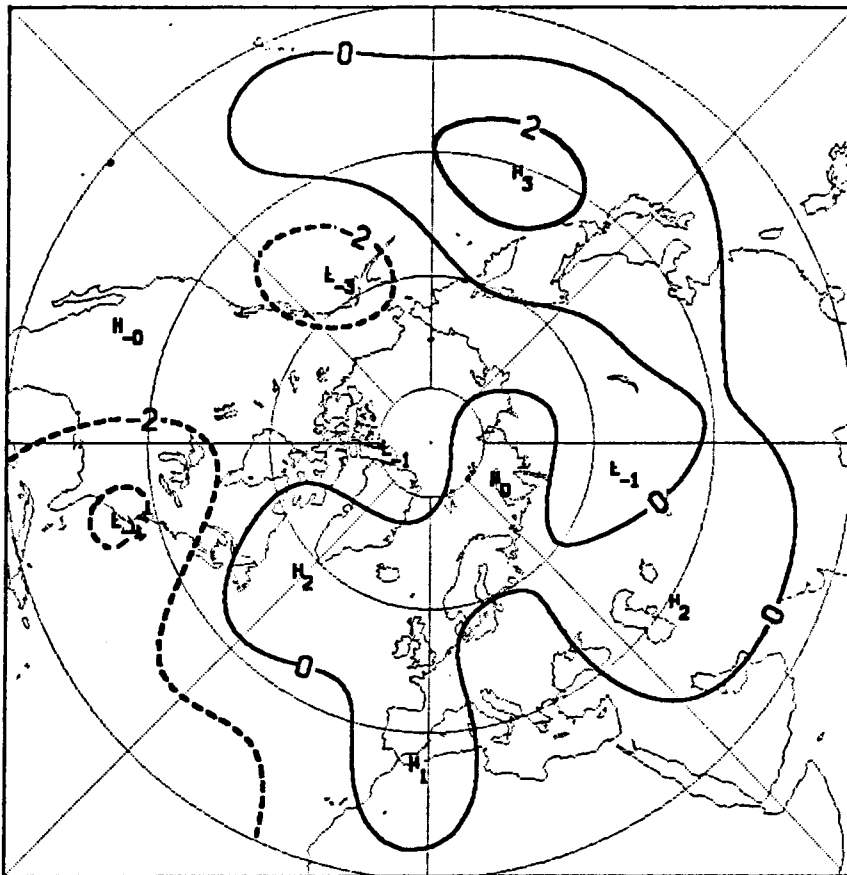


Fig. 3: Tropical forecast errors of height at 500 mb. (Non-relaxed forecast minus tropical relaxed forecast.) 5 day mean, spectrally filtered (T10). Units: dkm

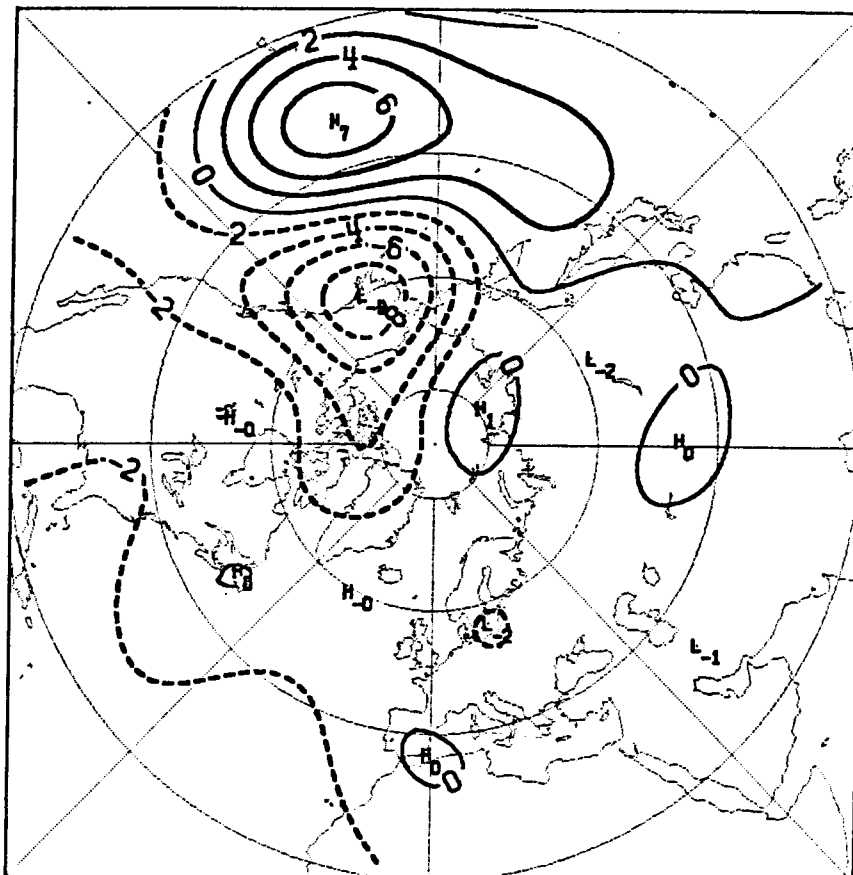


Fig. 4: Sea forecast errors of height at 500 mb. (Forecast relaxed over land minus analysis.) 5 day mean, spectrally filtered (T10). Units: dkm

made by relaxing the forecast towards the analysis over land only. The map for the large scale mean height error for this experiment (Fig. 4) shows a different response from the error structure in the Pacific and in the Atlantic. Whereas little impact was found over the Pacific - and this was independent of the inclusion or exclusion of the Himalayas and the Tibetan Highlands - a large reduction in the errors in the Atlantic region was found by relaxing the forecast towards analyses over North America. Inclusion or exclusion of the Rocky Mountains revealed a great sensitivity to mountain forcing over North America in the Atlantic errors. These experiments suggest that the errors connected with the Pacific jet are forced locally, whereas the errors in the Atlantic jet region are strongly influenced by erroneous mountain forcing over North America.

Further sensitivity experiments with global relaxation were carried out to investigate the impact of errors in different atmospheric layers. They show that the upper troposphere is most important for the systematic error; relaxing the wind field towards analyses at these levels leads to clear improvements in tropospheric forecasts.

The results presented here are similar to 3 other cases of the winter season 1984/85. They all show the great dependence of the Atlantic error on errors of mountain forcing. It is hoped that this relaxation technique can be used to determine the correct forcing in the region of mountains like the Rocky Mountains.

- Ernst Klinker

* * * * *

THE NEW ECMWF TELECOMMUNICATIONS SYSTEM

In autumn 1983, the Centre's proposal to replace the present telecommunications system as a matter of priority was endorsed by the Technical Advisory Committee and approved by the Council. With the co-operation of a working subgroup of the Technical Advisory Committee, a thorough investigation of Member States' requirements and expectations of a future telecommunications system was undertaken. An Invitation To Tender, based on the results of those enquiries, was issued on 19 July 1984, with tenders to be submitted by 12 October 1984.

In January 1985, after a thorough evaluation of the tenders received, the solution tendered by Software Sciences Ltd. (SSL) was selected and detailed negotiations, including the drawing up of a purchase contract, commenced. Having been examined by the Technical Advisory Committee Subgroup on the New Telecommunications System, the draft purchase contract was scrutinised by the Technical Advisory Committee who also took advantage of the advice of its subgroup. Having received the endorsement of the Technical Advisory Committee, the draft contract was approved by Council at its 22nd session (May 1985) and signed on 11 July 1985. The majority of the hardware was delivered and installed on 29 July and development is now progressing well.

Description of the new telecommunications system

The system is based upon four VAX 11-750 computers manufactured by Digital Equipment Corp. (DEC). They share disk storage and have access to Member State communications lines by small front-end processors. In addition, the systems are connected by an Ethernet local area network to gateway systems supporting the International Packet Switching Service, local terminals for software support and access to Telex and dial-up facilities. There will be a centralised operating facility. Fig. 1 shows the full configuration.

A total of four VAX 11-750 processors, each equipped with 4 million bytes of main memory forms the basis of the system. The VAX processors will be configured into a VAX cluster to provide access to shared disk space via a 70 million bits per second proprietary local area network called Computer Interconnect. Six disk drives, each holding 456 million bytes of data will be shared, configured into three mirrored pairs of disk space providing 1,368 million bytes of data. Each unit of information will be written to two separate disk units to provide resilience against the failure of a single disk drive. The cluster controllers which provide access to the shared disks will also be duplicated to eliminate single points of failure.

One pair of disks will hold the system files and the queues necessary to run the system. The other two pairs will hold Member State files and the files for Member State data dissemination, acquisition and remote job entry.




Fig. 1: Configuration of the new ECMWF Telecommunications System

In addition, there are two 205 million byte removable disks that will be used for transfer of software, the development and support team and for back-up of essential system files. One VAX in the cluster will be connected to the LCN to provide access to the other mainframes at the Centre, using hardware and software purchased from Control Data.

The private medium speed lines connecting the Centre to the Member State lines will be connected to the four VAX systems via KMS-11 front-end communications microprocessors. Each VAX will have its own KMS-11 front-end communications microprocessor and will be able to handle up to 8 medium speed lines, though under normal circumstances it is likely that 4 to 6 lines will be handled by each VAX. The front-end microprocessors will run software to handle the X25 level 2 protocol. Line switching equipment external to the VAXs will be installed to allow Member State lines to be moved between the VAX systems as necessary.

The 4 VAX computers will be interconnected by an Ethernet local area network, operating at 10 million bits per second. A terminal server gateway will also be connected to the Ethernet to provide access to the VAX computers for asynchronous terminals for ECMWF support staff. The terminal server gateway is based on a PDP 11-24 minicomputer. Low speed lines will be connected, via the switching equipment, to the normal terminal line on each VAX.

The new system will also provide various facilities in addition to those available on the old system. An X25 network server gateway is connected to the VAX systems via the Ethernet. This will provide connection between the Public Packet Switched Network and the VAXs for terminals using X29 protocols and providing file transfer facilities. The X25 gateway is based on a PDP/11-24 minicomputer. There will be access to telex and autodial modems via normal terminal lines available on each VAX.

The hardware and software necessary to connect the Centre's existing graphics VAX computer to the VAX cluster and the Ethernet used for the new telecommunications system will be purchased to enable the graphics VAX to be used as the fourth VAX during acceptance testing, to ensure that performance requirements are met.

It will also provide other benefits, since the graphics VAX can initially be used as a back-up communications VAX, prior to the installation of the fourth telecommunications VAX. It will also provide an alternative path to LCN in the event that the new telecommunications connection to LCN becomes unavailable. Finally, the magnetic tape deck on the graphics VAX can be used as a system load device, thus saving the cost of a magnetic tape deck for the new system.

The system is based on the standard systems software provided by DEC. The operating system will be VMS in the case of the VAX computers and Micro VMS in

the case of the Micro VAX II (used to enable centralised operator control). DECNET will be used to provide communication between the VAX systems and with the gateway systems connected on the Ethernet. DEC's X25 level 2 package will be used as the basis for connecting to the Member States using dedicated lines, and the X25 level 3 package will be used to connect to the Public Packet Switched Network.

Additional purpose-made software will be produced to provide various facilities, including the ECMWF end to end protocol and file transfer protocol, software for network management and monitoring, error processing and software to allow the normal facilities of VMS to be made available to Member States via the interactive liaison capability.

- Pam Prior

* * * * *

IDEAS FOR REDUCING THE NUMBER OF MAGNETIC TAPES USED

The Centre is facing a potential problem in the number of magnetic tapes in use. On the one hand the general workload is increasing substantially, leading to a demand for more and more magnetic tapes, on the other hand the amount of free space in the tape library is decreasing rapidly. The latest estimates show that we can only sustain 2 more years of growth at the current rate before we fill the tape library completely.

One way to help contain the problem is for users to ensure that they make the most efficient usage of the potential storage capacity on each tape. This should help slow the present growth in tapes used, until a time when it is hoped that a more efficient storage medium can be provided.

Initial planning

The first step is to study your own data storage needs from a very basic viewpoint - must all the data be stored, or by taking some simple steps can the volume be decreased considerably? Consider the following points at the initial planning stage:

- will they ever be used? Do not store data that will never be used.
- Can some of the data be easily derived from a basic subset? If so, store the subset only.
- Can the data be represented adequately by reduced accuracy? If so, pack 2 (or more) data words into one computer word. The GBYTE, SBYTE series of routines in ECLIB provide simple packing/unpacking capabilities.
- Remember the simple technique of using a base value, all other values being a scale multiple of that base value. Often the scale values require less bits to store than the full data values.
- Vary the frequency of data points to suit the rate of data variation. For example, a slowly varying field can be stored using far fewer data points than one which is rapidly varying. Do not store more data than is necessary to convey the maximum amount of information required.
- The Centre is using the GRIB code (a WMO bit-oriented code) as its internal data storage method in the archives. For meteorological data this is an efficient means of storage and is well worth considering. For a description of the code, and ECMWF's extensions to it, see ECMWF Technical Memorandum No. 102, Extensions to AF82 GRIB (Gridded Binary) by J. K. Gibson. Appendix 1 of that memorandum details the code itself.

- Finally, you never get something for nothing. Remember that to obtain a reduction in data storage using the above ideas you will certainly have an increase in computing complexity, and probably an increase in computer time as well. Users must use their judgement to find the right balance between reduced storage needs and increased complexity.

Technical tricks

The next step is to look at the storage methods available. Some, or all, of those given below could make a substantial difference to the number of tapes a project will use. It should be noted that some of these methods can also be used in existing projects. Often a significant saving can be made as a result of a little effort on the part of users.

- There are data compression techniques available to go through the data looking for repeated patterns of bits, which can be replaced by simple code groups. This is especially valuable on text data, where space savings of up to 50% can be made. However, experiments have shown that the savings on normal meteorological data are small, unless there are some special circumstances. From knowledge of his data a user should be able quickly to ascertain whether it is worth pursuing this idea or not.
- Normal Cyber storage methods (SI tapes) use a fixed block size of 512 words on tape. Such a block is relatively short compared to the gap between blocks. By making each block larger, up to 30% more data can be contained on a given tape for very little penalty. The changes required are:
 - (i) declare the tape to be an L tape on the REQUEST card;
 - (ii) if the tape is being used by a Cyber job it will require a FILE card;
 - (iii) write/read the tape using BUFFER OUT/IN.

Accessing L tapes via Fortran programs with BUFFER IN/OUT is the normal method. There are, however, some restrictions when trying to copy such tapes with standard utilities. The basic advice is to use COPYBF, COPYBR (or COPYCF, COPYCR). Known not to work is COPY itself, plus the PRINTBR series of utilities.

- Multifile tapes are a very useful technique when there are lots of relatively small files to be stored. They are to be strongly recommended. However, the CDC manuals on how to use multifile tapes are not easy to follow. A simple introduction has therefore been written to follow this article.
- On the other hand, multi-reel tape files, i.e. a file which spans more than one tape, are to be avoided. Such tapes seem to cause great problems and are definitely not worth the effort entailed in trying to get them to function correctly.

Conclusion

The storage problem will not be solved quickly. It is to everyone's advantage to use tape storage as efficiently as possible. The ideas presented here are intended as food for thought. When a new project commences, effort should be made from the outset to reduce the number of tapes required to the absolute minimum. For existing projects, consider whether perhaps L tapes or multifile tapes can reduce the existing usage. The resulting savings can make all the difference as to whether this problem can be contained or it grows out of control, literally.

- Andrew Lea

* * * * *

HOW TO USE MULTIFILE TAPES

As is stated in the previous article using multifile tapes is a very useful technique and is to be strongly recommended. By this technique many relatively small files can be stored on one tape, so considerably reducing the number of tapes used.

Many users at the Centre already use multifile tapes successfully. However, since it is almost impossible to find out from CDC manuals how to use them, this article provides some simple examples, together with some hints which have been discovered the hard way.

The basic ideas to remember are:

- specify the MF parameter on the REQUEST card;
- when starting a new multifile tape (or overwriting an old one) specify the N (New) parameter; subsequently use the E (Existing) parameter whenever using that tape; using N on an existing tape will destroy the previous contents, so caution must be exercised;
- use LABEL cards to identify each specific file on the tape;
- files are identified by their position on the tape and by no other means, so it is quite possible to have the same name for several files on the same tape.

Some examples follow. The first shows how to create a new multifile tape putting two files on it. The MF parameter on the REQUEST statement informs the system that it is a multifile tape, N says that this is a new tape. The M and L parameters on the LABEL statement inform the system that it applies to a member of a multifile tape (M=mftape), and that the file is to have the name given by the L parameter. Notice the recommended use of LISTMF at the end, this is to check that the content of the tape is as you expect it to be.

The second example shows how to extend this tape by putting two more files onto it. These will be added at the end and become the third and fourth files on the tape. Note the N parameter on the REQUEST statement has been replaced by E. Failure to do this results in the first two files being overwritten.

The final example shows how to read files 1 and 4 on this tape. Notice that the first file has been located by specifying the label name only, while the fourth has been found by specifying its position. To avoid possible confusion because of identical label names, it is always advisable to identify a file precisely by its position. Use of LISTMF will always show you what file is in what position.

Finally some closing comments:

- the way in which multifile tape usage is described in manuals leads to the assumption that a file is uniquely defined by the label field. This is wrong, as you can see from the examples.
- The use of LISTMF to check the contents of a multifile tape is advisable. This check must definitely be performed, if a job breaks down in the middle of writing to a multifile tape.
- Do not use multifile tapes which need more than one physical reel (i.e. avoid continuation tapes).
- A description of how to use a multifile tape in a Cray job is contained in Computer Bulletin B2.7/1 chapter 3.3.3. This facility is widely used at the Centre. Here a check of the tape contents by running LISTMF is even more important, since in the case of transmission failure, the CRAY-CYBER link will issue a rerun of that transmission job. The file will be copied to tape eventually but the structure of the multifile tape becomes unpredictable and the CRAY logfile will not necessarily tell you about the incident.

EXAMPLE 1: Creating a multifile tape

```
uid,T7,GE1.  
ACCOUNT,ECXYZ.
```

```
VSN,mftape=nnnnnn.  
REQUEST,mftape,GE,RING,MF,N.
```

Generating a new multifile tape or overwriting an existing multifile tape requires the 'N' parameter on the REQUEST statement. All subsequent jobs adding files to the tape must not have the 'N' but an 'E' instead.

ATTACH,cydata,ID=uida.
ATTACH,crdata,ID=uida.

Attach some files to copy
---ID=uida---

LABEL,tf,W,M=mftape,L=cydatacopy.

Open new file on multifile
tape.

COPY,cydata,tf
RETURN,cydata

Copy file to multifile tape

LABEL,tf,W,M=mftape,L=crdatacopy.

Open second file on multi-
file tape

COPY,crdata,tf

Copy file to multifile tape

LISTMF,M=mftape.

LISTMF to check labels and
position numbers

EXAMPLE 2 Extending a multifile tape

uid,T7,GE1.
ACCOUNT,ECXYZ.

VSN,mftape=nnnnnn.
REQUEST,mftape,GE,RING,MF,E.

The tape must be requested
with an 'E' on the request
card. An 'N' would initiate
the tape as a new multi-
file tape and therefore
destroy the previous
contents.

ATTACH,cydata,ID=uidb.
ATTACH,crdata,ID=uidb.

Attach some other files to copy
---ID=uidb---

LABEL,tf,W,M=mftape,L=cydatacopy.

Open new file on multifile tape.
This will become the third file on
the tape. Identical labels are
possible. A file on tape is
uniquely defined by its position
number, not the label field.

COPY,cydata,tf.
RETURN,cydata.

Copy file to multifile tape.

LABEL,tf,W,M=mftape,L=crdatacopy.	Open another file on multifile tape.
COPY,crdata,tf.	Copy file to multifile tape.
LISTMF,M=mftape.	LISTMF to check labels and position numbers.

EXAMPLE 3 Reading a multifile tape

uid,T7,GE1.
ACCOUNT,ECXYZ.

VSN,mftape=nnnnnn.
REQUEST,mftape,GE,NORING,MF,E.

The tape must be requested with an 'E' on the request card. An 'N' in conjunction with the parameter 'RING' would initiate the tape as a new multifile tape and therefore destroy the previous contents.

LABEL,tf,R,M=mftape,L=cydatacopy.

Read first file having the specified label field.

REQUEST,cydata,PF.
COPY,tf,cydata.

Copy file to disk

CATALOG,cydata,ID=uida.
RETURN,cydata.

Restore file as permanent file
---ID=uida---

LABEL,tf,R,M=mftape,P=4.

Open file to be found at position no. 4. This is the second file with the label 'crdatacopy'.

REQUEST,crdata,PF.
COPY,tf,crdata.

Copy file to disk

CATALOG,crdata,ID=uidb.

Restore file as permanent file
---ID=uidb---

where:

uid,ECXYZ.	are your user identifier and account number
mftape	is the logical (multi)file name you assign to the tape
nnnnnn	is the tape VSN

tf is the lfn for a single tape file

cydatacopy are examples of label names given to each file on the
crdatacopy tape

cydata,crdata are the pfns of the two files to be copied, owned by user
"uida" in the first example, by "uidb" in the second.

- Norbert Kreitz

* * * * *

STILL VALID NEWS SHEETS

Below is a list of News Sheets that still contain some valid information which has not been incorporated into the Bulletin set (up to News Sheet 179). All other News Sheets are redundant and can be thrown away.

<u>No.</u>	<u>Still Valid Article</u>
16	Checkpointing and program termination
19	CRAY UPDATE (temporary datasets used)
56	DISP
67	Attention Cyber BUFFER IN users
73	Minimum Cyber field length
89	Minimum field length for Cray jobs
93	Stranger tapes
118	Terminal timeout
120	Non-permanent ACQUIRE to the Cray
121	Cyber job class structure
122	Mixing FTN4 and FTN5 compiled routines
127	(25.1.82) IMSL Library
130	Contouring package: addition of highs and lows
135	Local print file size limitations
136	Care of terminals in offices
140	PURGE policy change
141	AUTOLOGOUT - time limit increases
144	DISSPLA FTN5 version
152	Job information card
158	Change of behaviour of EDIT features SAVE, SAVEX. Reduction in maximum print size for AB and AC
164	CFT New Calling Sequence on the Cray X-MP
165	Maximum memory size for Cray jobs
166	Corrections to the Contouring Package
167	CFT 1.13 improvements
170	NOS/BE level 604
171	" " "
172	" " "
	Change to CFT Compiler default parameter (ON=A)
174	Warning against mixing FTN4 and FTN5 compiled routines.
176	Archival of Cyber permanent files onto IBM mass storage
177	RETURNX, REWINDX
178	TIDs on Cray include 2 chara. TID plus 3 chara source computer ID. Caution with ACQUIRE on RERUN jobs
179	Job submission from Cray job to Cray

* * * * *

ECMWF PUBLICATIONS

- TECHNICAL MEMORANDUM NO. 99: A valley-filled orographic representation in numerical weather forecast models

- TECHNICAL MEMORANDUM NO. 100: Systematic errors in surface and upper lows in ECMWF's operational spectral model for the 1983-1984 winter season and comparing these to the relevant errors in the grid point model

- TECHNICAL MEMORANDUM NO. 101: Predictability studies with the ECMWF model for extended range: the impact of horizontal resolution and boundary layer forcing

- TECHNICAL MEMORANDUM NO. 102: Extensions to AF82 GRIB (gridded binary)

- TECHNICAL MEMORANDUM NO. 103: Practical considerations concerning the use of the planetary boundary layer top as a co-ordinate surface

- ANNUAL REPORT 1984: English Version

* * * * *

CALENDAR OF EVENTS AT ECMWF

- 9-13 September 1985 Annual ECMWF seminar: "Physical Parameterisation for numerical models of the atmosphere"

- 16-18 September 1985 14th session of Scientific Advisory Committee

- 18-20 September 1985 10th session of Technical Advisory Committee

- 25-27 September 1985 35th session of Finance Committee

- 20-21 November 1985 22nd session of Council

- 22 November 1985 ECMWF 10th anniversary celebration

* * * * *

INDEX OF STILL VALID NEWSLETTER ARTICLES

This is an index of the major articles published in the ECMWF Newsletter plus those in the original ECMWF Technical Newsletter series. As one goes back in time, some points in these articles may have been superseded. When in doubt, contact the author or User Support.

	<u>No.*</u>	<u>Newsletter Date</u>	<u>Page</u>
<u>CRAY</u>			
Bi-directional memory	25	Mar. 84	11
Buffer sizes for jobs doing much sequential I/O	14	Apr. 82	12
CFT 1.11 Subroutine/function calling sequence change	19	Feb. 83	13
COS 1.12 and products	23	Oct. 83	17
Cray X-MP/48 - description of	30	June 85	15
- hints on using it	26	June 84	10
Dataset storage	13	Feb. 82	11
Multifile tapes - disposing of	17	Oct. 82	12
Multitasking ECMWF spectral model	29	Mar. 85	21
Public Libraries	T5	Oct. 79	6
<u>CYBER</u>			
Arithmetic instructions - comparative speeds of execution on the Cyber front ends	14	Apr. 82	17
Cyber front ends - execution time differences	15	June 82	9
Buffering or non-buffering on Cyber?	15	June 82	10
CMM-Fortran interface	10	Aug. 81	11
Cyber 855 - description of	21	June 83	18
Dynamic file buffers for standard formatted/ unformatted data	3	June 80	17
Formatted I/O - some efficiency hints	4	Aug. 80	9
FTN4 to FTN5 conversion	6	Dec. 80	15
FTN5 - effective programming	9	June 81	13
	& 10	Aug. 81	13
- optimisation techniques	14	Apr. 82	13
	& 15	June 82	10
Graphics - hints on memory and time saving	T6	Dec. 79	20
- a summary of planned services	17	Oct. 82	10
Magnetic tapes - hints on use	T2	Apr. 79	17
- making back-up copies	1	Feb. 80	9
Public libraries	T5	Oct. 79	6