EUROPEAN CENTRE FOR MEDIUM RANGE WEATHER FORECASTS

TECHNICAL REPORT NO. 2

****

DECEMBER

1976

****

***

*

THE EFFECT OF ARITHMETIC PRECISION

ON

SOME METEOROLOGICAL INTEGRATIONS

BY

A.P.M. Baede, D. Dent, A. Hollingsworth

EUROPEAN CENTRE FOR MEDIUM RANGE WEATHER FORECASTS

Fitzwilliam House

Skimped Hill

BRACKNELL, BERKS.

UNITED KINGDOM

C O N T E N T S

A B S T R A C T

In order to evaluate the effect of arithmetic precision on long range integrations of atmospheric prediction models, experiments with two different models were carried out both on a CDC 6600 (48 bit mantissa) and an IBM 360/195(24 bit mantissa).

In the first experiment an adiabatic baroclinic spectral model was integrated over 8 days on both computers. A comparison of both runs was made on the basis of their mass and energy conserving properties. The non-conservation of mass turned out to be completely independent of arithmetic precision and can be shown to be due almost entirely to time truncation. The non-conservation of energy however is determined by round off errors, but is very small anyhow. The observed changes of the total energy can be shown to be caused by a single bit change. The conclusion is that the arithmetic precision has no significant influence on the results of integrations with this model.

In the second experiment the N24 version of the GFDL general circulation model was integrated over 10 days on both computers. The space and time development of the discrepancies between both runs was studied in some detail. A further comparison was made on the basis of RMS errors and correlation coefficients. It was found that the initial rapid discrepancy growth, in particular in the tropics, was caused by the Moist Convective Adjustment. We have not been able to further isolate the cause of this problem. It was concluded that this part of the code requires high precision arithmetic and very careful coding.

1. Introduction

Different studies have been devoted to the evaluation of the relative influence of physical and numerical factors on the practical limits of predictability of long range forecasting models. A factor to which relatively little attention has been given, is the computer system (software and hardware). In this report we describe two experiments that were undertaken to shed some further light on this area; one with an adiabatic hemispheric spectral model (HOSKINS and SIMMONS, 1975) and one with the N24 GFDL general circulation grid point model (MIYAKODA, 1973).

Both models were run on a 24 bit mantissa IBM 360/195 and on a 48 bit mantissa CDC 6600. Moreover, the spectral model although designed in single precision was run in both single and double precision on the 195. In the GFDL model all of the moist processes are run in single precision on the 6600 and double precision on the 195.

In section 2 we discuss previous work in this area, mainly the work of WILLIAMSON and WASHINGTON (1973) and KURIHARA and TULEYA (1974). After a very brief description of the differences between both computers in section 3, we discuss the adiabatic spectral experiment in section 4, and conclude that a 24 bit mantissa is probably adequate. Section 5 describes the experiment with the general circulation model. Here large differences in the tropical wind and temperature forecasts are found between the two integrations. These differences are attributed to the moist convective adjustment, which is apparently sensitive to differences between both systems. It is concluded that this part of the code requires high precision arithmetic and that careful coding is necessary.

2. Previous Work

WILLIAMSON and WASHINGTON (1973) presented the results of experiments designed to test the importance of arithmetic precision in short term forecasts and long term simulations with the NCAR global circulation model. The same model was run using a 48-bit mantissa arithmetic and 24-bit mantissa arithmetic. They concluded that, with the NCAR general circulation model, the shorter word length was satisfactory for short range forecasts, since typical observational errors are greater than the round-off errors and since the model exhibited a fast error growth which was attributed to the moist processes. For 80-day integrations they found that there was no tendency for the round-off error to dominate.

KURIHARA and TULEYA (1974), in a response to the work of WILLIAMSON and WASHINGTON, reported that they had had difficulty in preserving mass and in maintaining consistent energetics in a hurricane simulation model when converting from a 27-bit mantissa to a 24-bit mantissa. They made a number of recommendations regarding techniques to be used to retain as much significance as possible in numerical computations. These recommendations have been followed in the GFDL code, made for the IBM 360/195. In particular the moist processes are calculated using double precision. Moreover, double precision is used with the special "rounding up and down" procedure recommended by KURIHARA and TULEYA whenever a small number is added to a much larger number as in the time extrapolation. On the other hand WILLIAMSON and WASHINGTON maintain that round off error accumulation in low precision arithmetic would not be an important cause of error in two-day forecasts. The reason is that random errors (introduced by errors in data for example) grow faster than the accumulation of round off error.

SEARLE and DAVIES (1975) compare the eddy kinetic energy computed from their model by three runs on two computers with different arithmetic precision. They find significant differences after about 1500 time steps without, however, giving a quantitative evaluation of their results.

It is clear that the conclusions from these experiments might be highly model dependent and cannot be accepted as general conclusions. It is therefore essential to repeat these experiments with other well established models.

3. Differences between both computers

The present experiments were performed on a CDC 6600 and on an IBM 360/195. It is impossible to evaluate the relative importance of system dependent factors such as the compiler and the number representation. It is noteworthy that on both the IBM and CDC all arithmetic computations are truncated rather than rounded.

The differences in number representation are summarised in Table I. However, it should be noted that on the IBM, after renormalization, the mantissa may have lost up to three bits due to its hexadecimal representation.

	IBM 360/195		CDC 6600	
	single	double	single	double
word	32	64	60	120
mantissa	24	56	48	108
representation	hexadecimal		binary	

TABLE I. Word and mantissa length (bits) in single and double precision

4. An 8-day integration with an adiabatic baroclinic spectral model

4.1 The model and the experiment

The model used for this experiment is the adiabatic hemispheric 5-layer spectral model developed by the U.K. Universities Atmospheric Modelling Group in Reading and described extensively by HOSKINS and SIMMONS (1975). All calculations in this model are done in single precision and no special measures have been taken to minimise round off errors. The code was written for a CDC 7600. From this model a version for the IBM 360/195 was prepared in which all calculations are done in double precision. In all runs a triangular truncation at total wave number 21 (T21) was used. A semi-implicit time extrapolation was used with timesteps of 30 and 90 minutes respectively.

The initial data are the same as those of HOSKINS and SIMMONS (1975) in their study on a CDC 7600 of a growing baroclinic wave: a differential solid body rotation with a small initial perturbation in the (8,9) vorticity coefficient.

In total six 8-day integrations were made. On the CDC 6600 two single precision integrations were made with timesteps of 30 and 90 minutes respectively, to check if the results were consistent with earlier integrations on a CDC 7600 (HOSKINS and SIMMONS,1975). Four integrations were made on the IBM 360/195: single and double precision for both timestep lengths.

4.2 Results and discussion

We examined the amplitudes and phases of the dependent variables which were printed to four significant figures. After eight days there was no difference to this accuracy in any of these

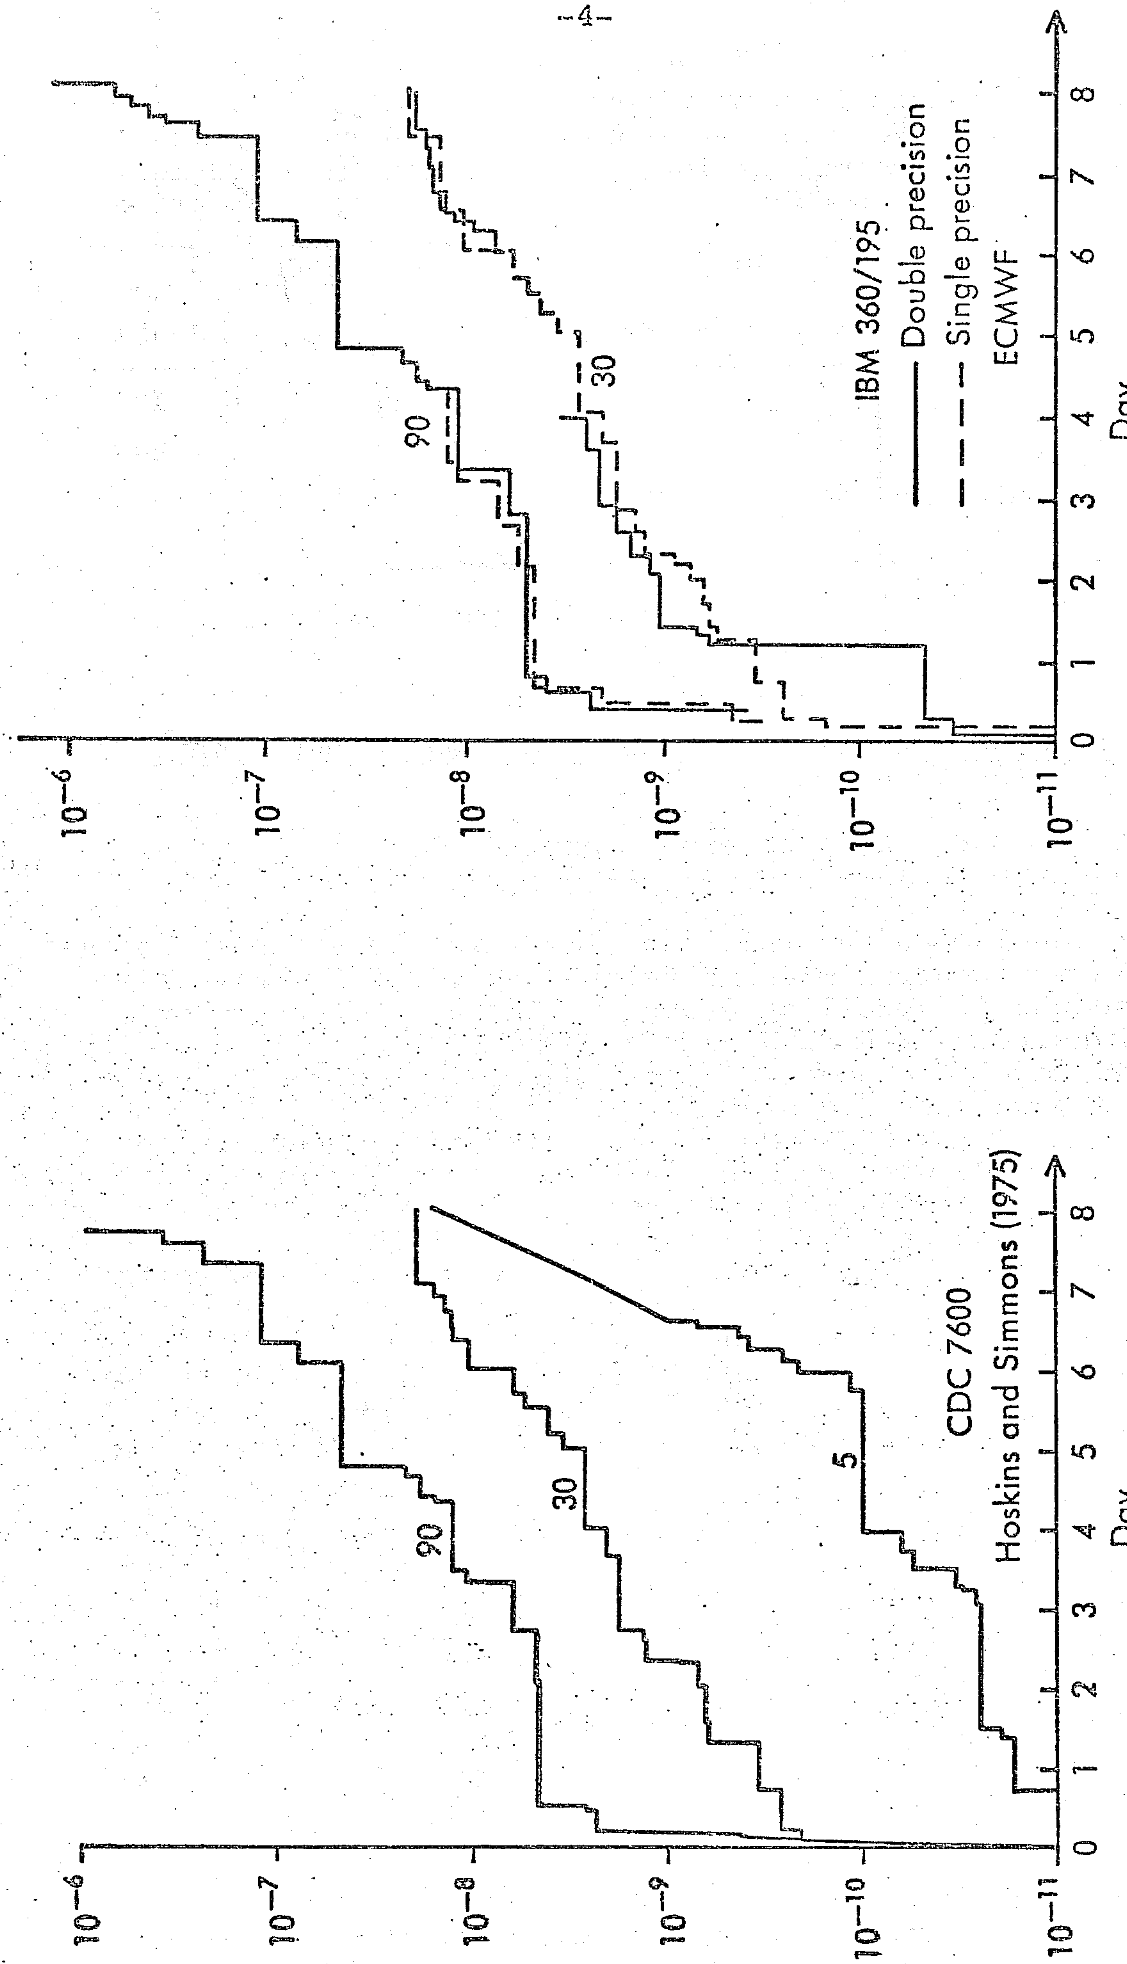

Fig. 1 The spurious change of total mass normalised by the mass itself. The value plotted at a time is the maximum up to that time. The labels indicate the time step length in minutes.

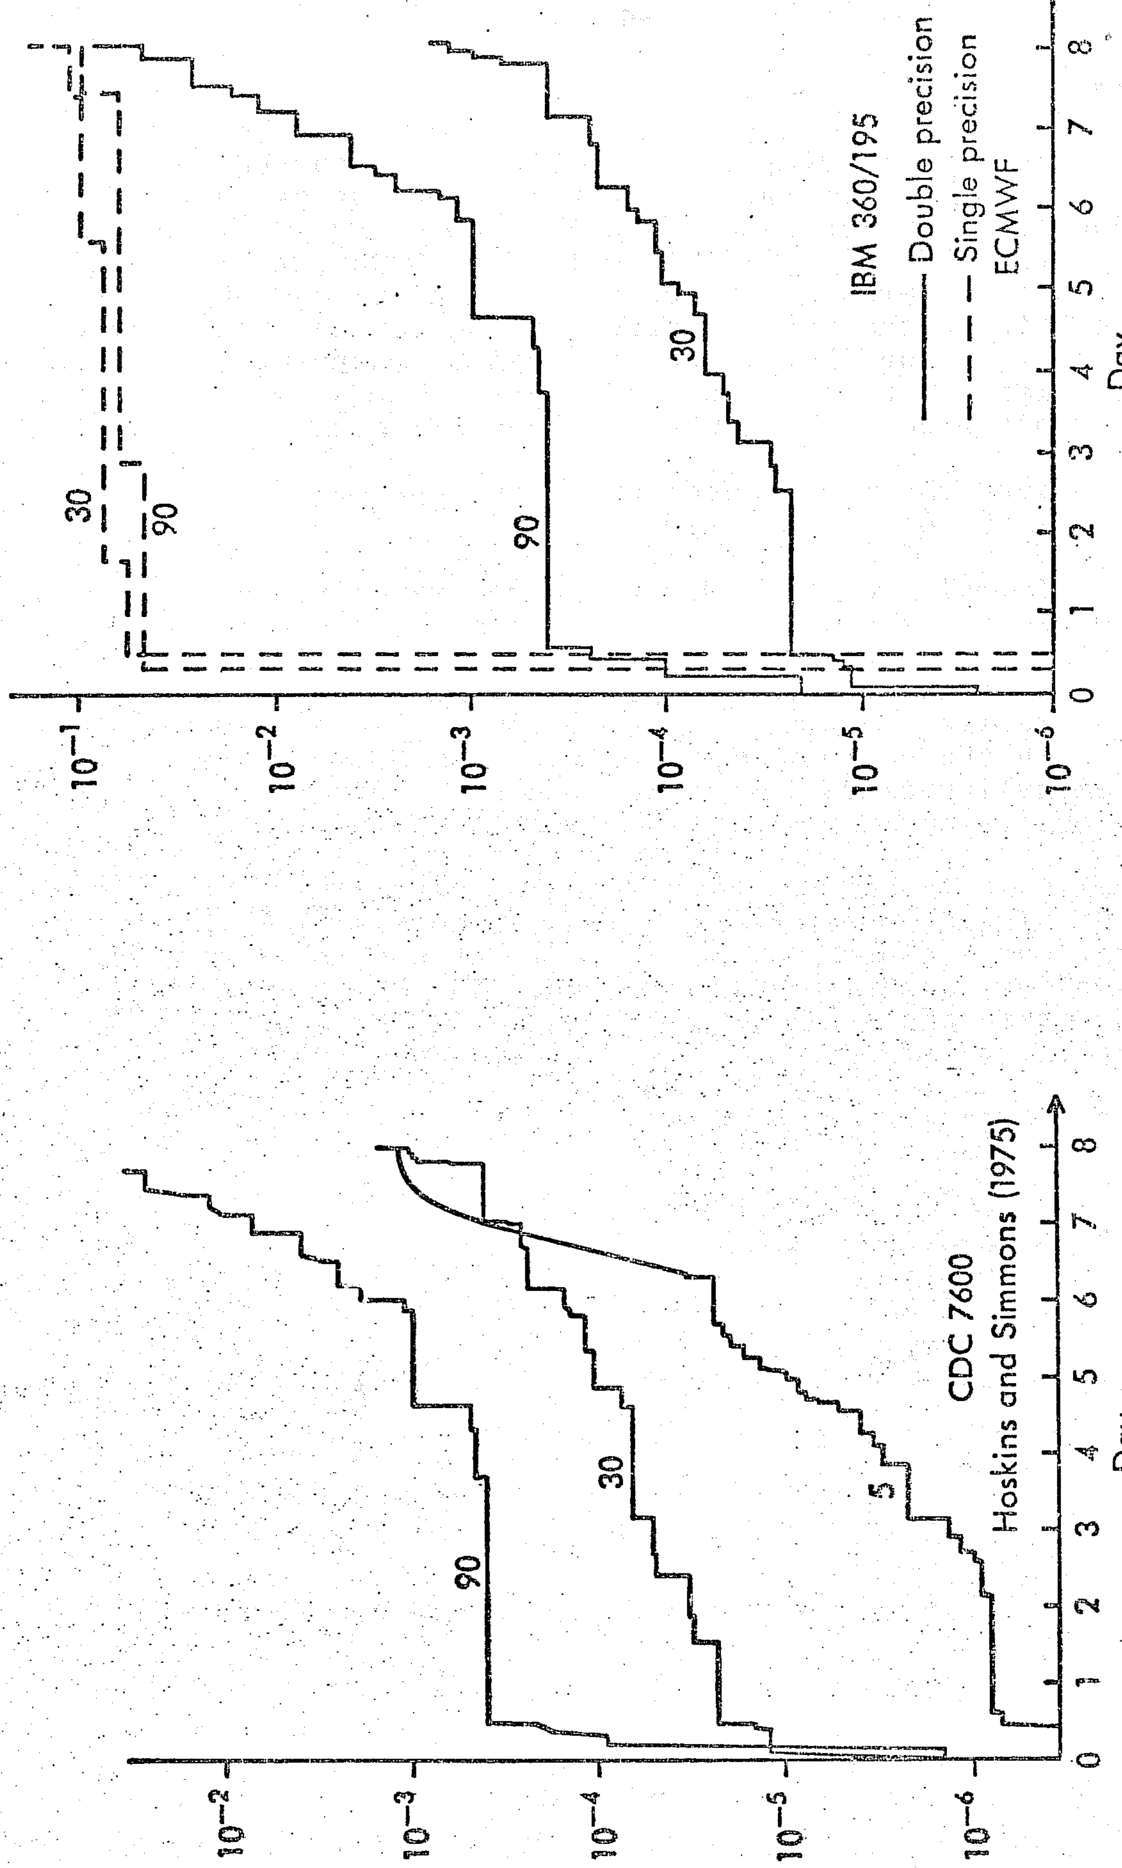

Fig. 2 The spurious change of total energy normalised by the change of the kinetic energy at day 6. Otherwise as fig. 1.

components between the single and double precision runs on the IBM machine. As between the single precision run on the 6600 and the double precision run on the 195 there were differences in the amplitude of the dominant vorticity component of one or two units in the fourth significant figure after one day. The difference between the runs remained in the fourth figure for this component during the succeeding days of exponential growth with an e-folding period of about 2 days. The differences between the other components were of the same order.

We also compared the mass conservation and energy conservation properties of the different integrations.

The mass is formally conserved, apart from time truncation. Figure 1 shows the spurious change of total mass normalised by the mass itself for runs in single and double precision using a 30 minute and 90 minute time step. The value plotted at a given time is the maximum change up to that time. The errors are almost identical with the corresponding curves from HOSKINS and SIMMONS. This agreement indicates that time truncation errors rather than round-off errors dominate the non-conservation of mass.

It may seem surprising that accuracies of order 10^{-9} are claimed for the IBM single precision run. The reason is that the variable used is not surface pressure but rather the natural logarithm of surface pressure normalised by 1000 mb. The latter variable is a small quantity and its gradients can be represented more accurately than gradients of the total pressure.

The mass conservation equation is

$$\frac{\partial}{\partial t} \ln p_* = -\int \underline{V} \cdot \nabla (\ln p_*) \, d\sigma - \int \nabla \underline{V} \, d\sigma$$

Ignoring time truncation we have

$$p_* \frac{\partial}{\partial t} \ln p_* = -\int \underline{V} \cdot \nabla p_* \, d\sigma - \int p_* \cdot \nabla \underline{V} \, d\sigma = -\int \nabla (p_* \underline{V}) \, d\sigma$$

This last transformation will be exact in the spectral model apart from round-off errors. Finally, apart from round-off errors we have

$$\int_{\text{area}} \nabla \underline{A} \, dS \equiv 0$$

Consequently $\int p_* \frac{\partial}{\partial t} \ln p_* \, dS = 0$ to the same accuracy and so

$\frac{\partial}{\partial t} \int p_* \, dS = 0$ to within round-off errors and time truncation errors.

The mean value of the normalised pressure in these runs is, initially, 1.0 + .306230E-6. The model predicts changes in the small number .3E-6. Precision errors would initially change the last digit of this small number, giving an error of 10^{-12} in the pressure. Time truncation errors change the third or fourth place within a day and change the first place by the eigth day. Thus the non conservation of mass in these runs is due almost entirely to time truncation. This feature accounts for the remarkable agreement between the runs for this quantity.

We turn now to the question of energy conservation. The double precision curves for the IBM runs are barely distinguishable from the corresponding curves for the CDC run in Fig. 2. It appears that the single precision runs, particularly with the longer time step, are quite inaccurate. However, the situation is not as bad as it appears. If we take the total energy as 1.0 then the total kinetic energy is 10^{-3} and the change in the kinetic energy over the first six days is 2.5%, i.e., 2.5×10^{-5} of the total energy. Due to the hexadecimal representation and renormalisation on the IBM machine a single bit change may cause a change of order $1:2^{20} \simeq 1:10^{6}$. This would appear on the curve as an error of order 4x10-2. The observed change is of this order of magnitude and may therefore well be explained by a single bit change.

The lcrge difference between the kinetic energy and the total energy in this run is fairly typical for the atmosphere. This example points out the difficulties inherent in studies of atmospheric energetics due to the large imbalance between potential and kinetic energy. However, if we are interested in the dynamics then it is clear from this run that a low resolution (T21), adiabatic model gives at least 4 figure accuracy in velocities in an 8 day integration.

5. A 10-day forecast with the N24 GFDL general circulation model

5.1 The model and the experiment

In this second experiment the 9-level global general circulation model, developed at GFDL, was used. The horizontal resolution is N=24. The time differencing was the same as that used for the experiment described in ARPE et al (1976) . The reader is referred to this report for an extensive description of the model.

Most of the dynamics and physics, apart from the moist processes, are carried out in single precision, although at many points in the GFDL code double precision arithmetic is used. In this respect the recommendations of KURIHARA and TULEYA (1974) have been followed, as mentioned above. In particular the double precision "rounding up and down" procedure was applied in the time extrapolation.

In the IBM run the original model was used, with the moist processes in double precision, but on the CDC 6600 a model version was used in which these processes were computed in single precision. A comparison of a one day integration on the CDC 6600 with double and single precision code for the moist processes

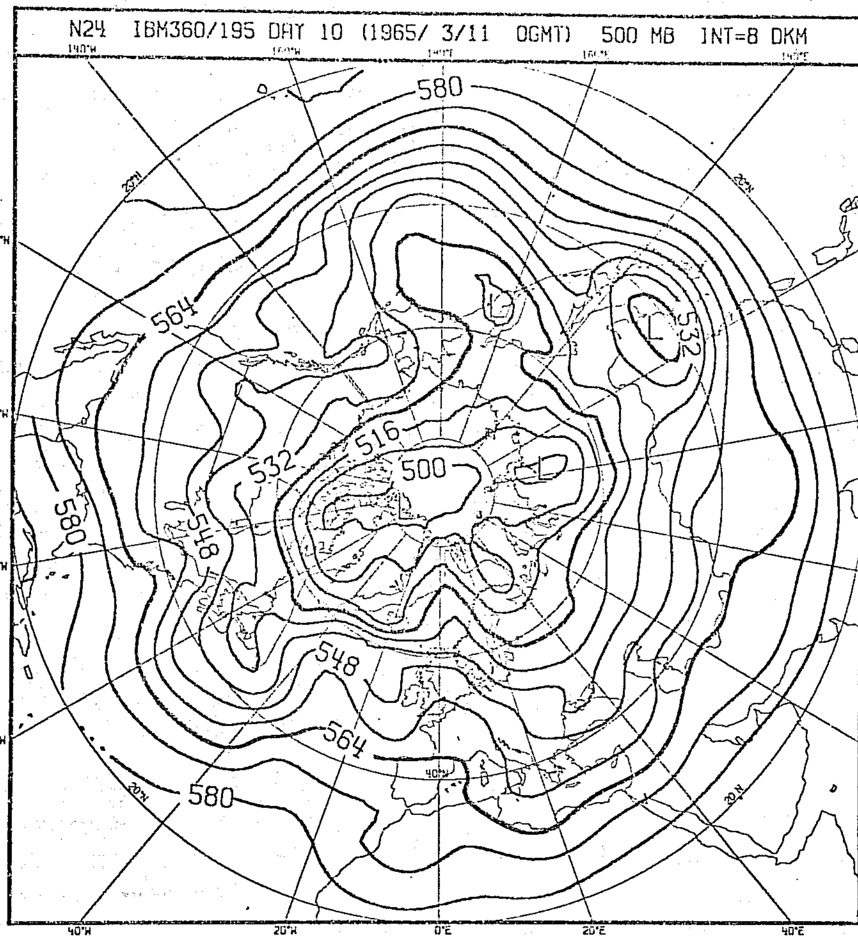

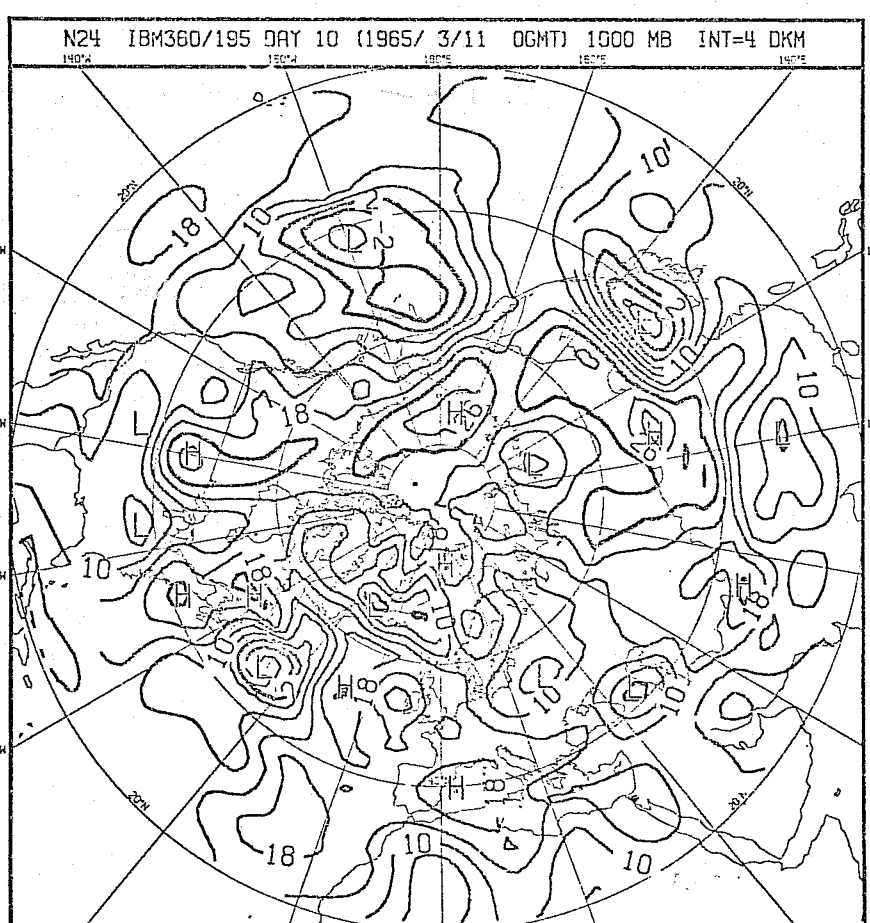

Fig. 3 Result of a 10-day forecast of the height of the 500 mbar (upper) and 1000 mbar (lower) level on the IBM 360/195

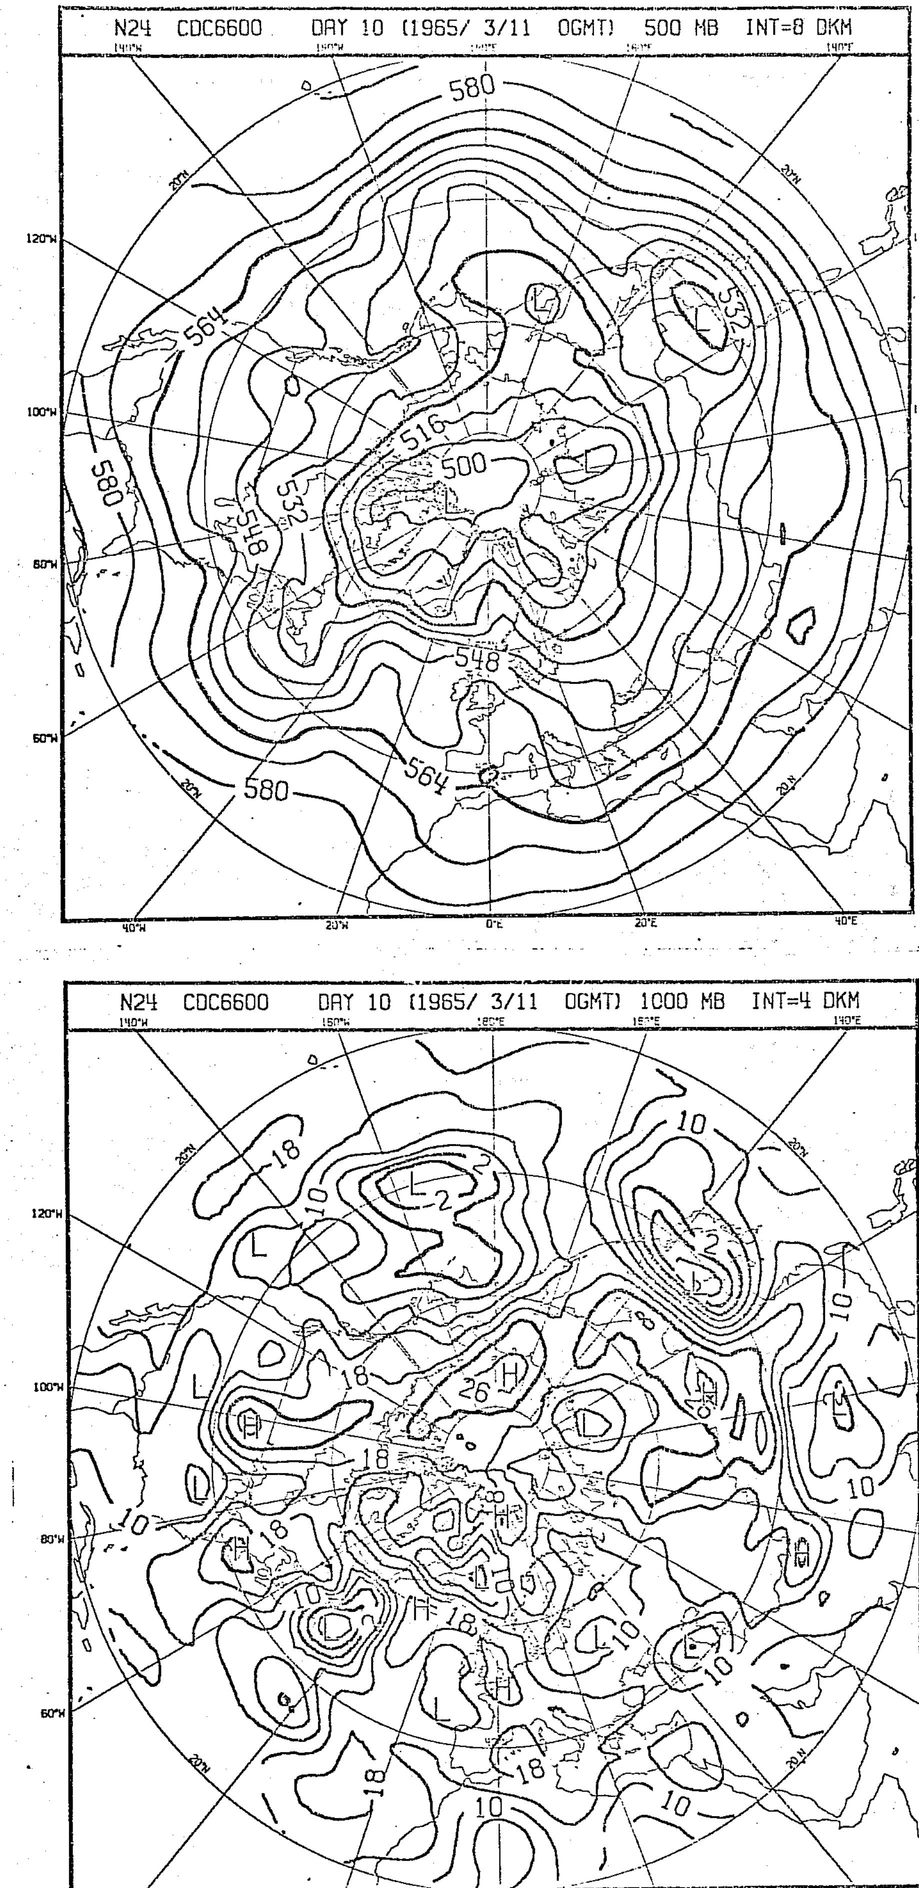

Fig. 4 Result of a 10-day forecast of the height of the 500 mbar (upper) and 1000 mbar(lower) level on the CDC 6600

showed differences in the temperature field of only $1:10^9$.

A 10-day forecast was run on both computers using the initial data of 1 March 1965, the same data used in the experiments described in ARPE et al (1976). The analysis and intercomparison of the results were carried out on the CDC using the standard diagnostic programs provided by GFDL.

5.2 Definition of measures of discrepancy growth

Two types of measures of discrepancy growth are used in this report: rms discrepancies and correlation coefficients over horizontal domains, which are defined as follows:

G : Global
NH : Northern Hemisphere : north of 25^oN
T : Tropical : $25^oS - 25^oN$
E : Equatorial : $10^oS - 10^oN$

Moreover, the rms errors may or may not be vertically averaged over the pressure. The rms discrepancy of a quantity F over the different domains is now defined as

$$E_F(t) = \sqrt{\overline{(\Delta F)^2}^{\lambda,\vartheta,p}}$$

A second measure is the correlation coefficient p for the deviation of a quantity F from the "climatological" mean F_o. For convenience the climatological mean is defined to be the 10 day average of the CDC-run. Defining $\delta F=F-F_o$, the correlation coefficient averaged over the different domains is now defined as

$$\rho_F(t) = \frac{\overline{(\delta F)_{CDC}\,(\delta F)_{IBM}}^{\lambda,\vartheta,p}}{\sqrt{\overline{(\delta F)^2_{CDC}}^{\lambda,\vartheta,p}\;\overline{(\delta F)^2_{IBM}}^{\lambda,\vartheta,p}}}$$

5.3 Synoptic comparisons

A synoptic comparison between the day 10 1000- and 500 mb height charts of both N24 runs reveals no significant differences (compare Fig. 3 and Fig. 4). The position and strength of all synoptic features in the 500 mb chart are almost identical. The differences between the 1000 mb charts are larger, but again not substantial. Fig. 5 shows that the differences are mainly confined to mid latitudinal belts on both the Pacific and the Atlantic Ocean, in particular those areas with high baroclinic activity.

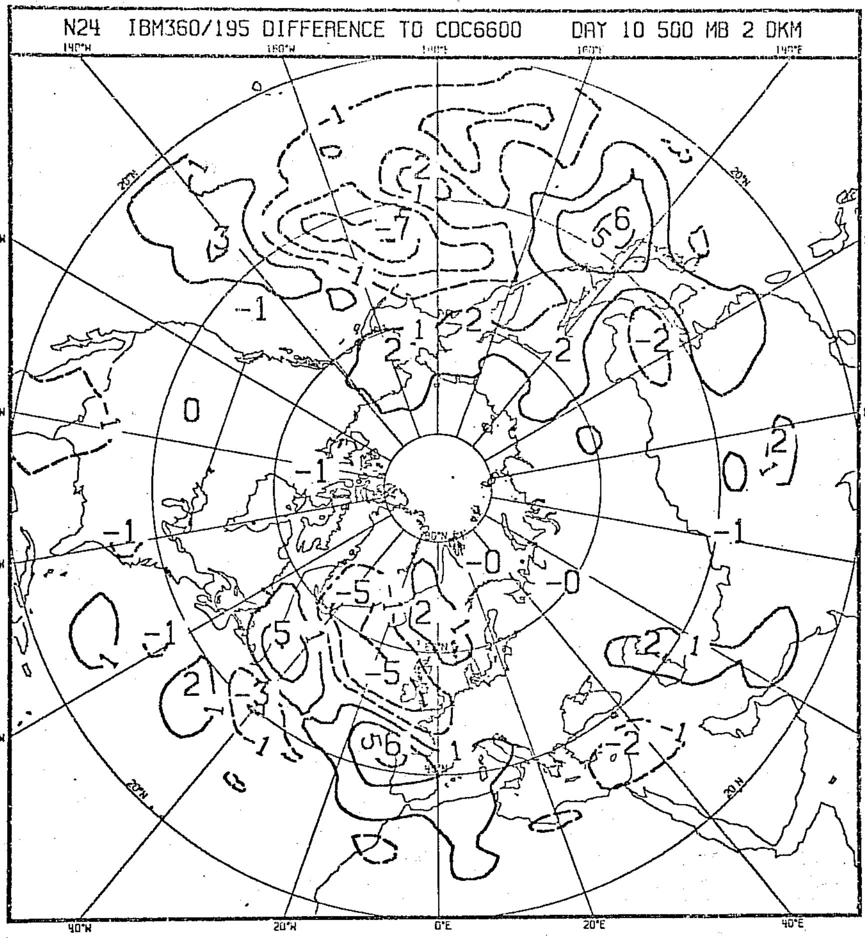

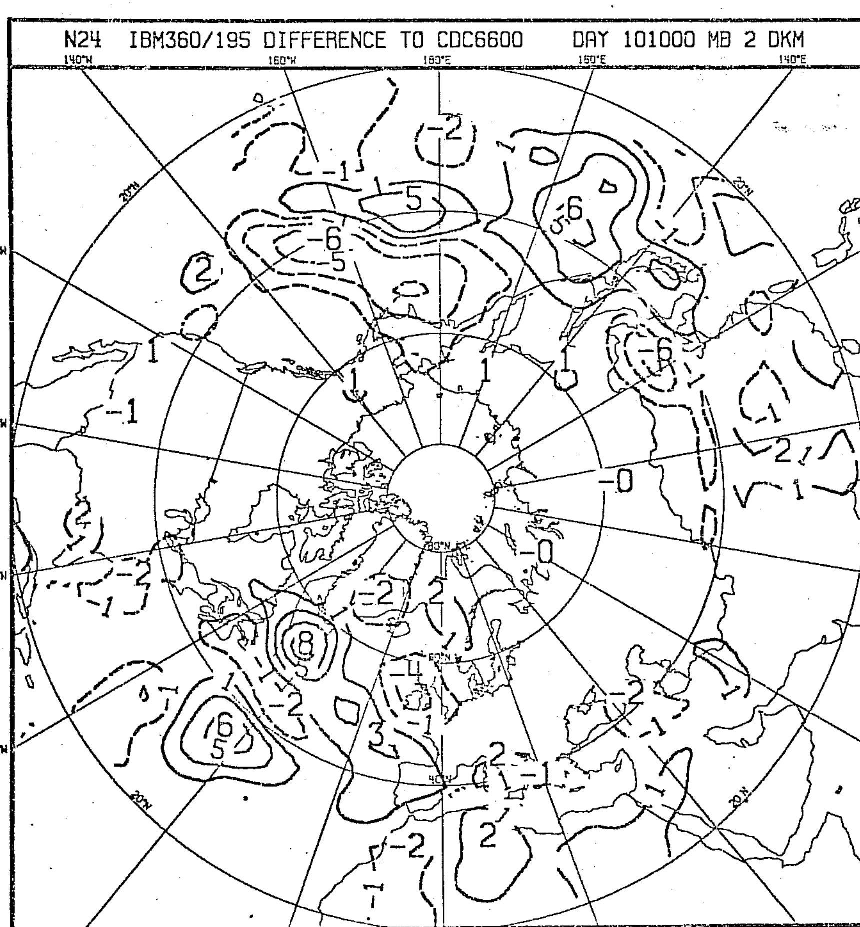

Fig. 5 Differences between the day 10 500 mb (upper) and 1000 mb (lower) heights on both computers.

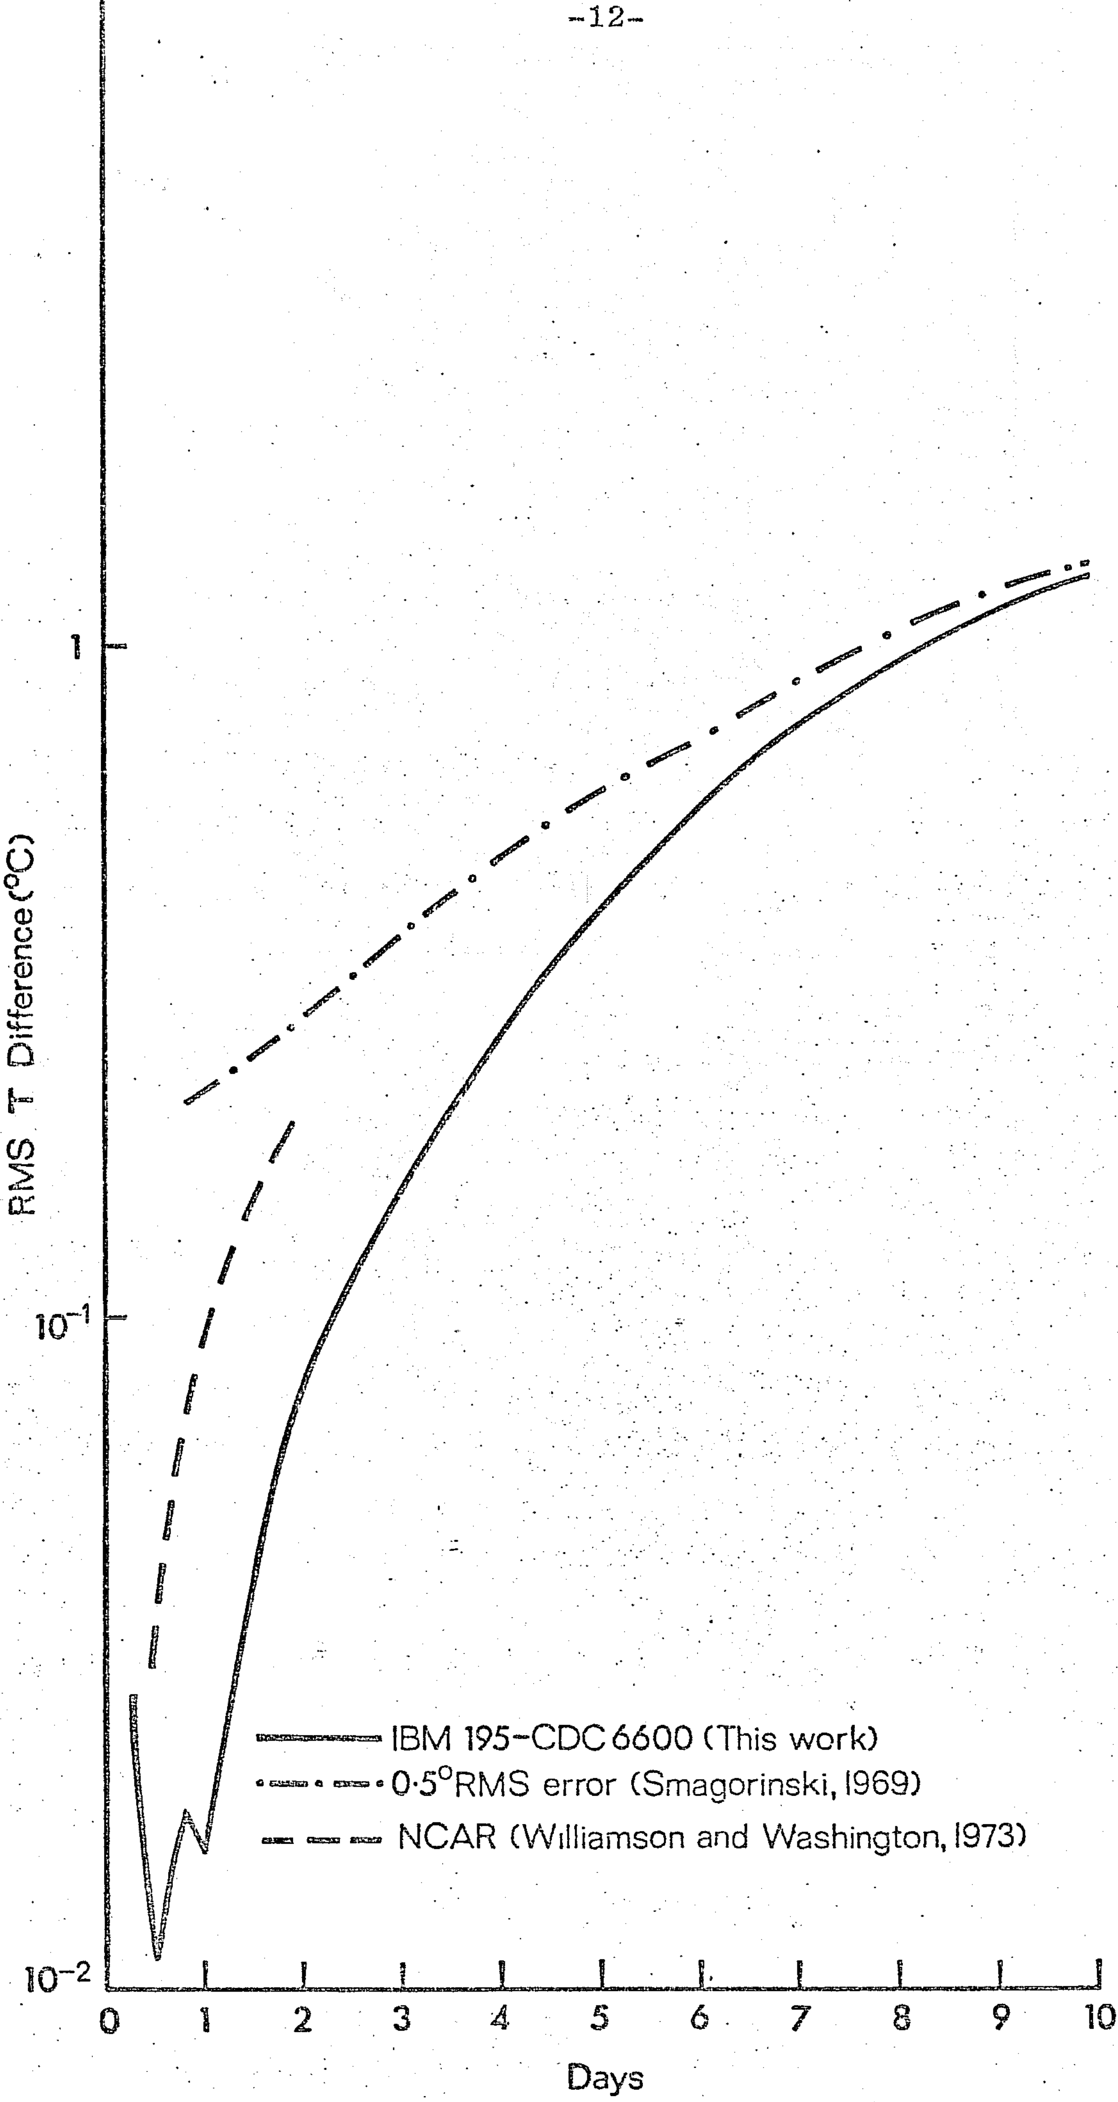

Fig. 6 Global rms temperature difference between the runs on

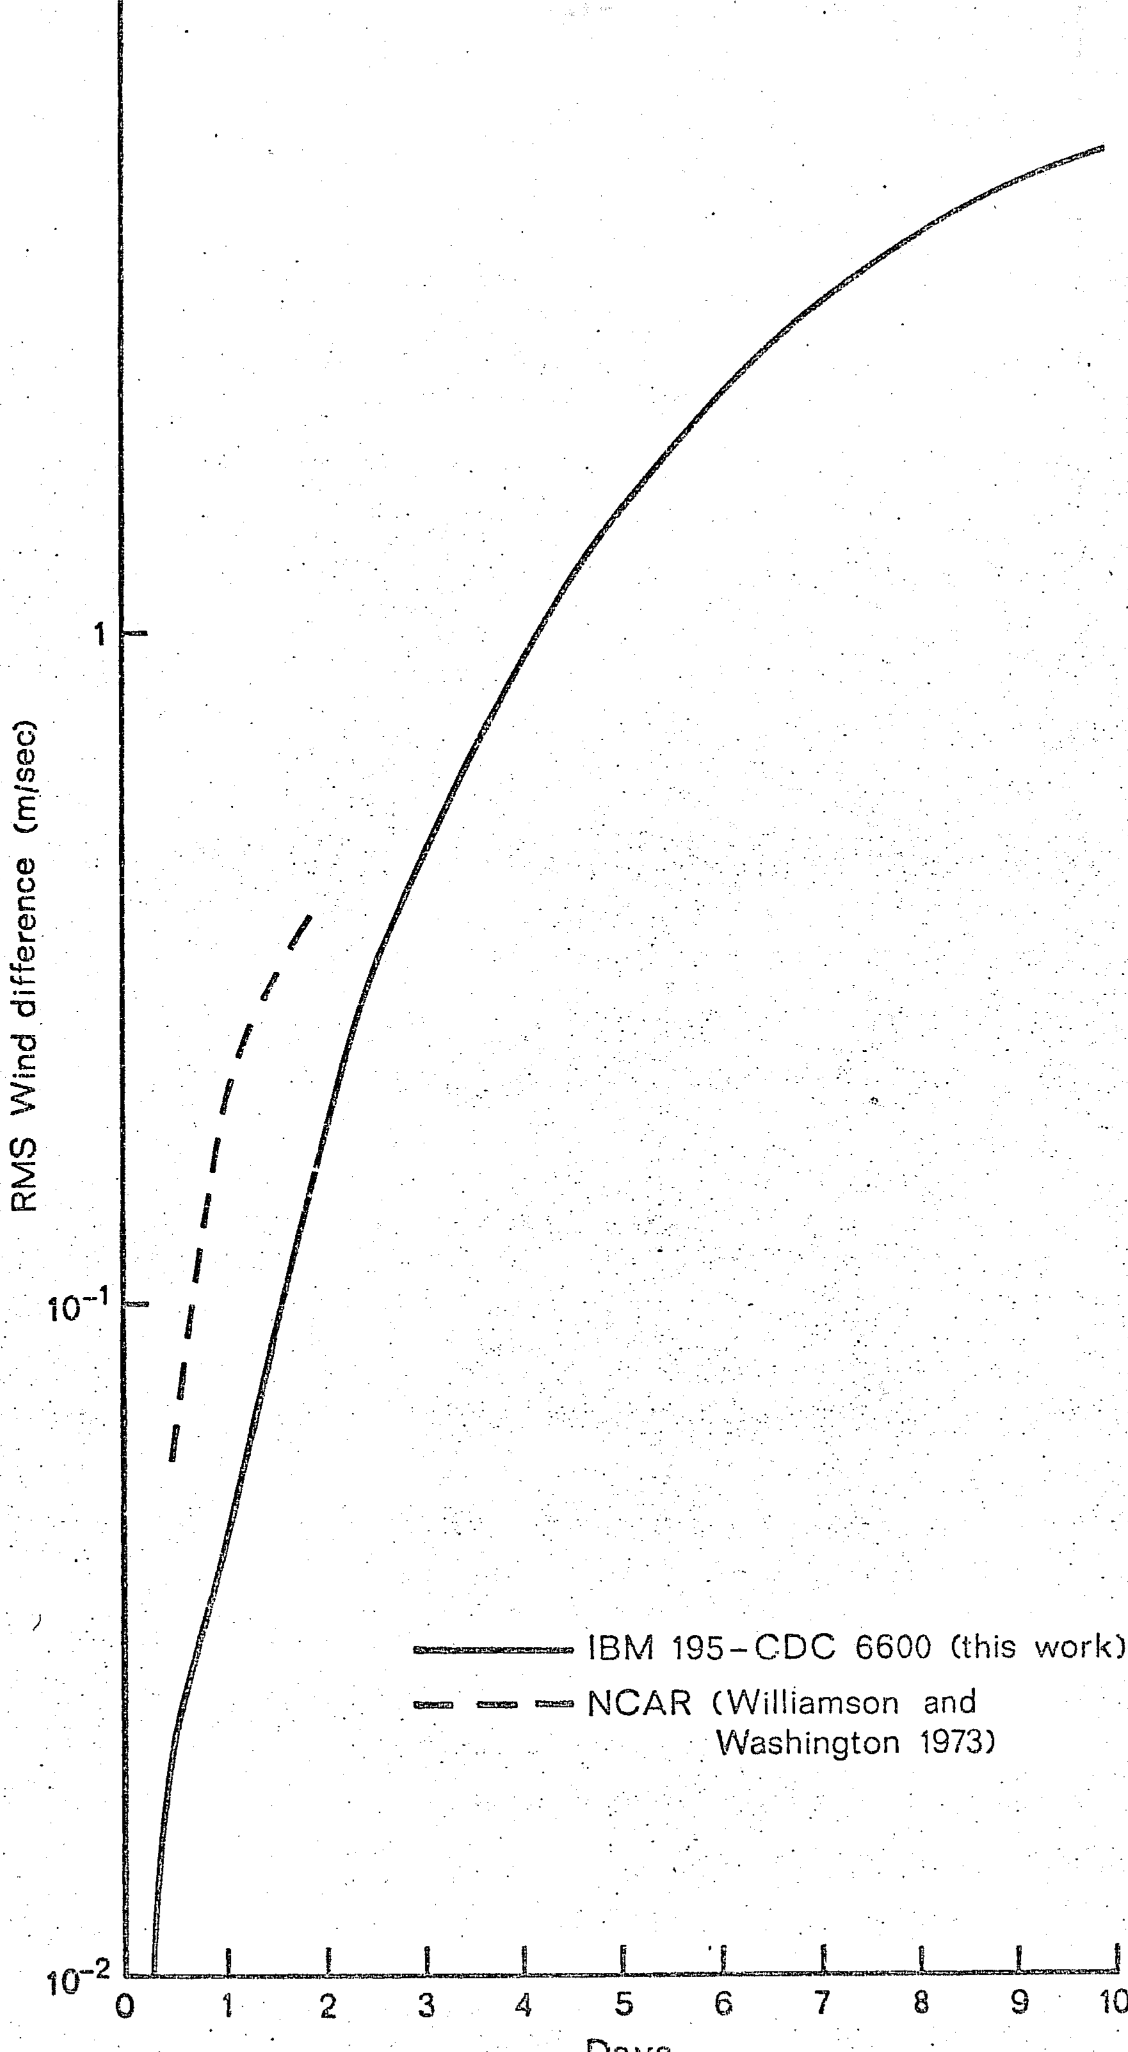

Fig. 7 Global rms wind difference between the runs on both computers as a function of forecast time. Results of

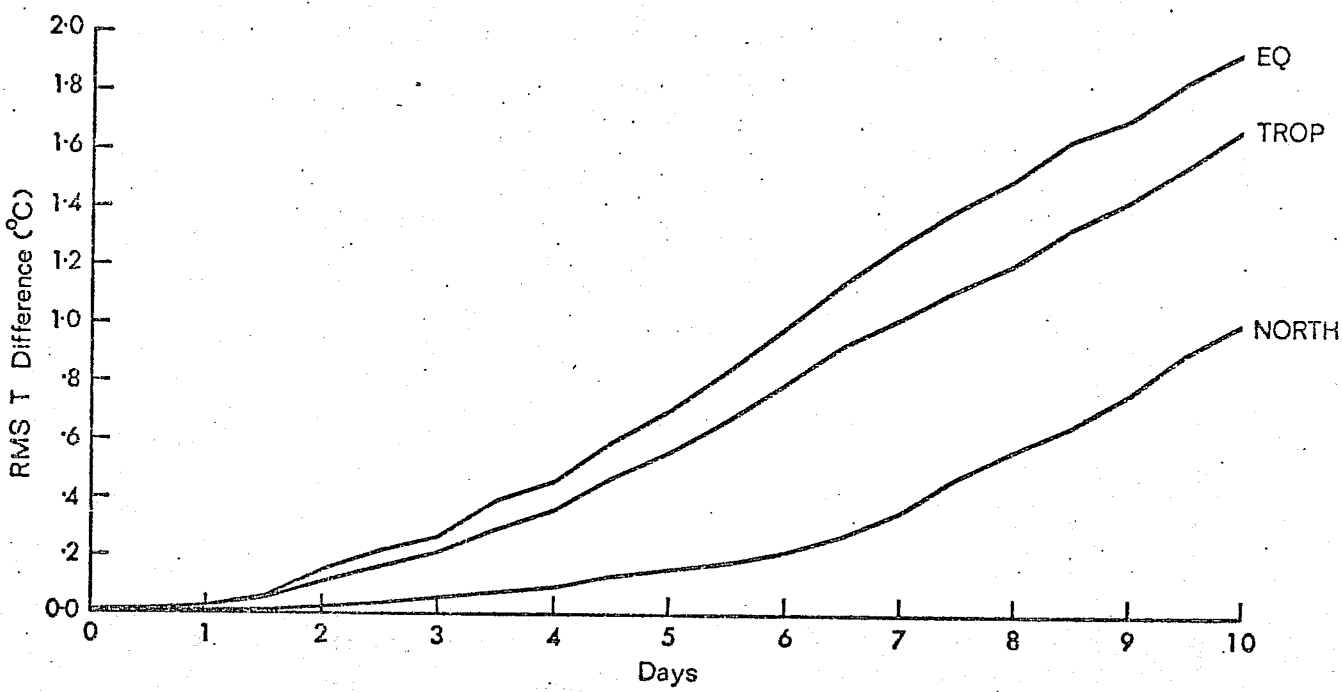

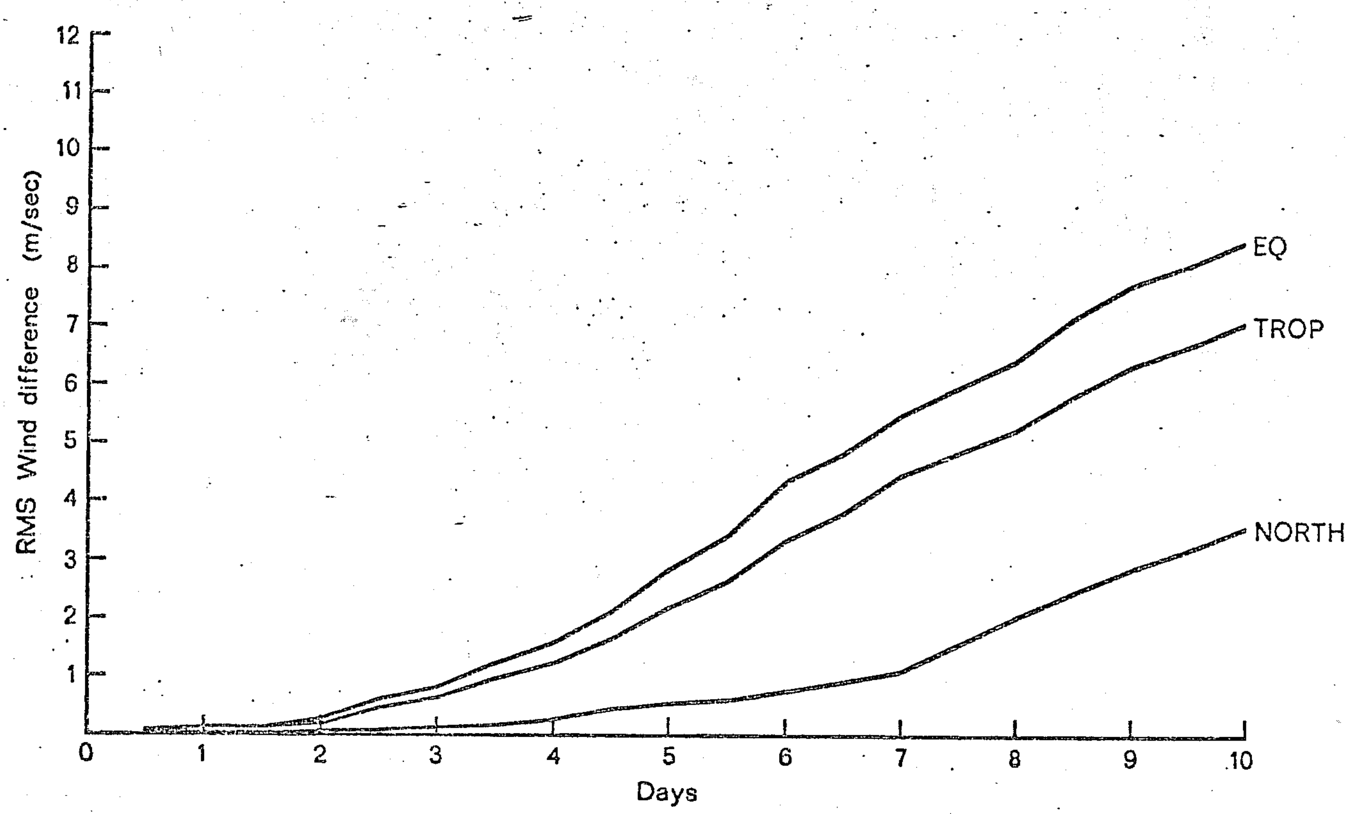

Fig. 8 Global rms temperature (upper) and wind (lower) differences for three different horizontal domains, defined in the text.

To appreciate the differences between both runs, it is apparently not enough to compare only the height charts. Instead a closer look at the temperature and wind forecasts in different areas is necessary.

5.4 Description of the results

Overall growth of errors

The overall growth of the global vertically integrated rms errors of wind and temperature is shown in fig. 6 and fig. 7. Both figures show a very fast growth rate during the first few days followed by a much slower rate thereafter. The temperature rms error moreover is very noisy during the first day. In fig. 6 the results of the hemispheric predictability experiment by SMAGORINSKY (1969) are also shown. It is seen that after 8 days both magnitude and growth rate are comparable. In this figure the global rms discrepancy found by WILLIAMSON and WASHINGTON between rms with different word length on the same computer is also shown. Although their discrepancies are larger, the growth rate shows the same fast initial growth. Figures 8 and 9 show the rms errors and the correlation coefficient over the different horizontal domains. From these figures it is clear that large discrepancies occur in tropical and equatorial regions. Indeed the correlation coefficient after 10 days shows that the temperature forecasts of both runs in the equatorial region bear only a minor semblance. A general feature is the smaller error growth in the southern hemisphere than in the northern hemisphere, which is typical for a northern hemisphere winter situation.

Time - height development of rms errors

Figures 10 and 11 depict the time height evolution of the rms errors of the wind $\sqrt{\overline{(\Delta u)^2 + (\Delta v)^2}^{\lambda,\theta}}$ and temperature $\sqrt{\overline{(\Delta T)^2}^{\lambda,\theta}}$ both in the equatorial region and in the northern hemisphere. The initial error growth of temperature and wind is much faster in the equatorial region in agreement with the features described above. In the equatorial region the initial temperature error primarily grows near the ground and in mid-troposphere. After about 5 days two strong maxima develop, one at low levels, the other in the stratosphere. In the northern hemisphere the initial temperature error growth primarily occurs near the ground and in the stratosphere. Eventually a low level tropospheric maximum develops with only a minor maximum near the tropopause. Two wind error maxima appear both in the equatorial and the northern hemisphere regions.

The development of a low level maximum in the northern hemisphere temperature error growth is also found in SMAGORINSKY's (1969) predictability experiment. In contrast with his results, however, we did not find a stratospheric maximum.

Meridional propagation and distribution of discrepancies

From the results it is clear that the wind and temperature discrepancies originate in the equatorial region and propagate into mid-latitudes. The wind discrepancies start growing near the equatorial tropopause whereas the development of temperature discrepancies starts near the ground and 500 mb. The meridional distribution of rms discrepancies after 10 days is shown in figure 12. This figure shows clearly the strong maxima in the equatorial region. Substantial maxima can also be observed in mid-latitudes, in particular in the northern hemisphere.

5.5 Discussion

A clue to the explanation of the initially rapid equatorial growth of discrepancies is found in the observation that already after 6 hours local temperature discrepancies of a few degrees are found in the equatorial and tropical regions. At first these discrepancies are observed near the ground, but later on also in mid troposphere. This suggests that the moist convective adjustment (MCA) is responsible for the discrepancies. A similar explanation was suggested by WILLIAMSON and WASHINGTON (1973) who showed that the rapid initial growth does not occur when the release of latent heat was set equal to zero.

In order to see whether the actual calculation of the MCA is strongly computer dependent, this part of the code was isolated and allowed to operate, both in single and double precision, on a prescribed supersaturated temperature and humidity profile. The adjusted temperatures all agreed to within $1:10^7$ and the humidities to within $1:10^6$. This result suggests that it is not the actual calculation of the MCA that causes the discrepancies, but that the procedure that determines whether or not MCA will be applied is critically computer dependent.
This latter decision depends on weighted vertical differences due to precision. One should remember, however, that the decision making is done in double precision on the IBM 195 system and in single precision on the CDC system. Granted that the systems reach different conclusions about a case of near neutral stability it is remarkable that the convective scheme can bring about changes near neutral conditions that grow to the order of 1 or 2 degrees in a few hours. This is not well understood and will be the subject of further studies.

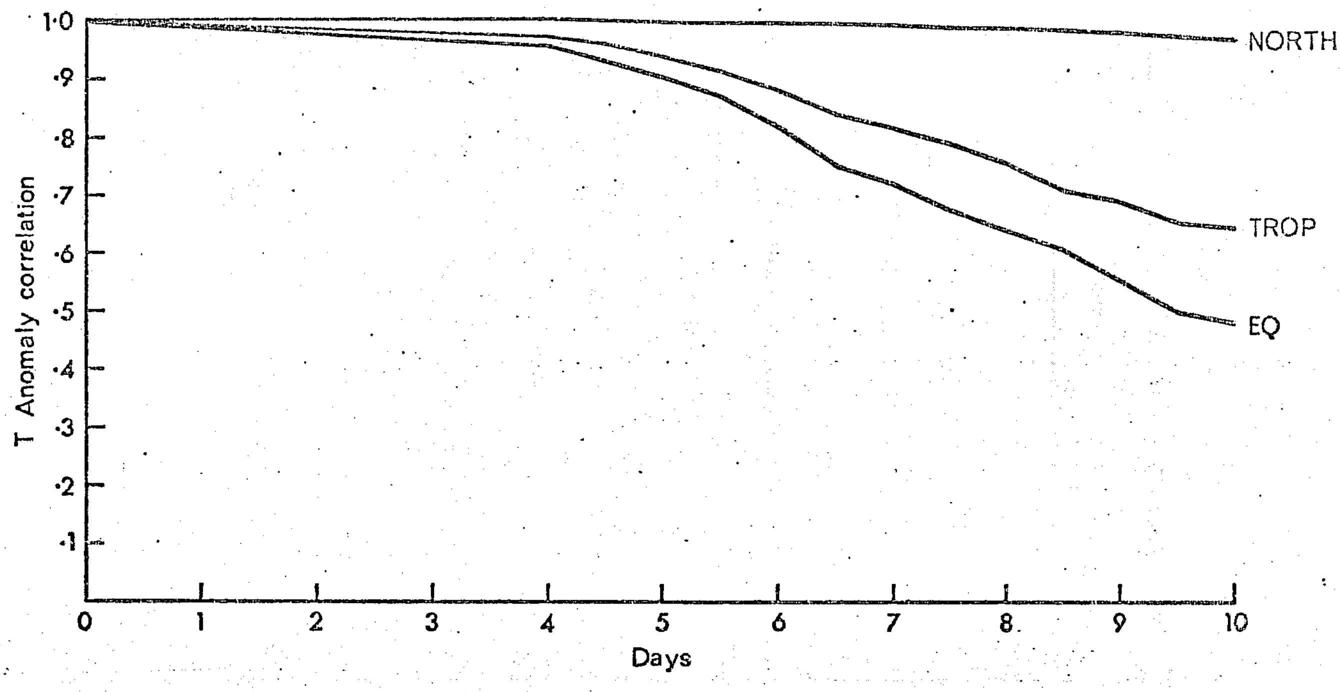

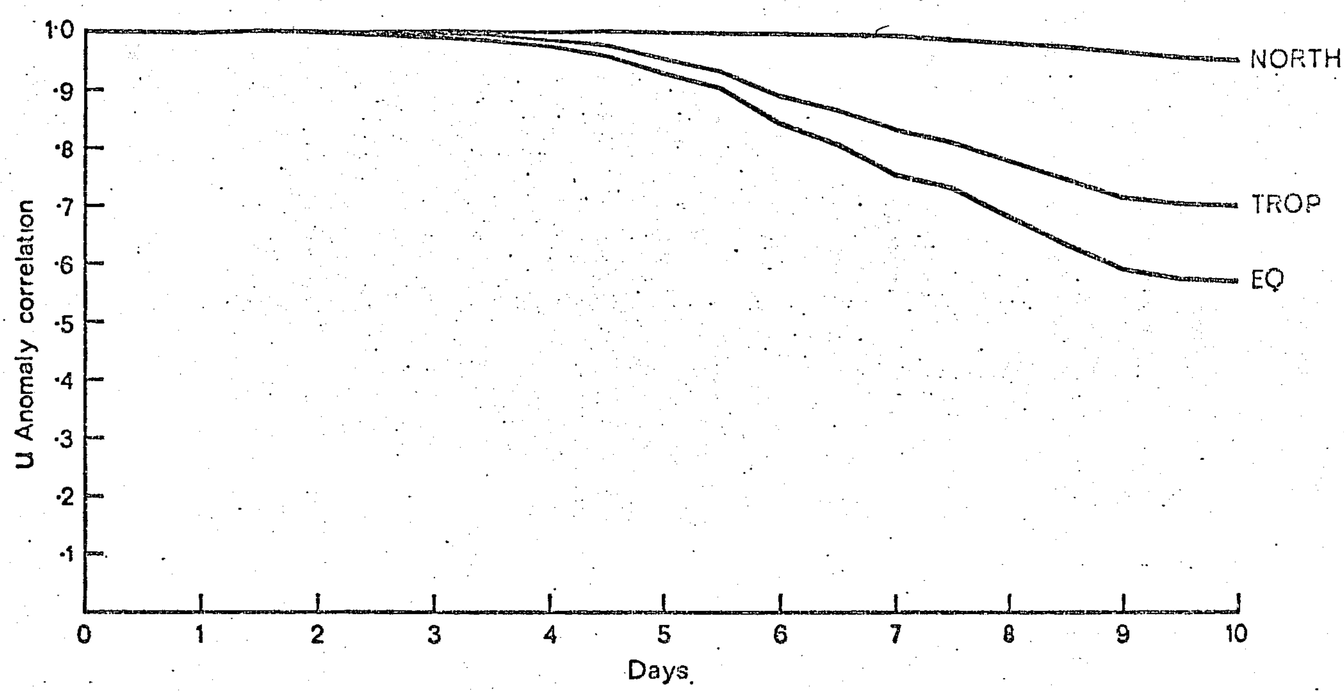

Fig. 9 Global anomaly correlation of temperature (upper) and u-component of the wind (lower) for three different horizontal domains, defined in the text.

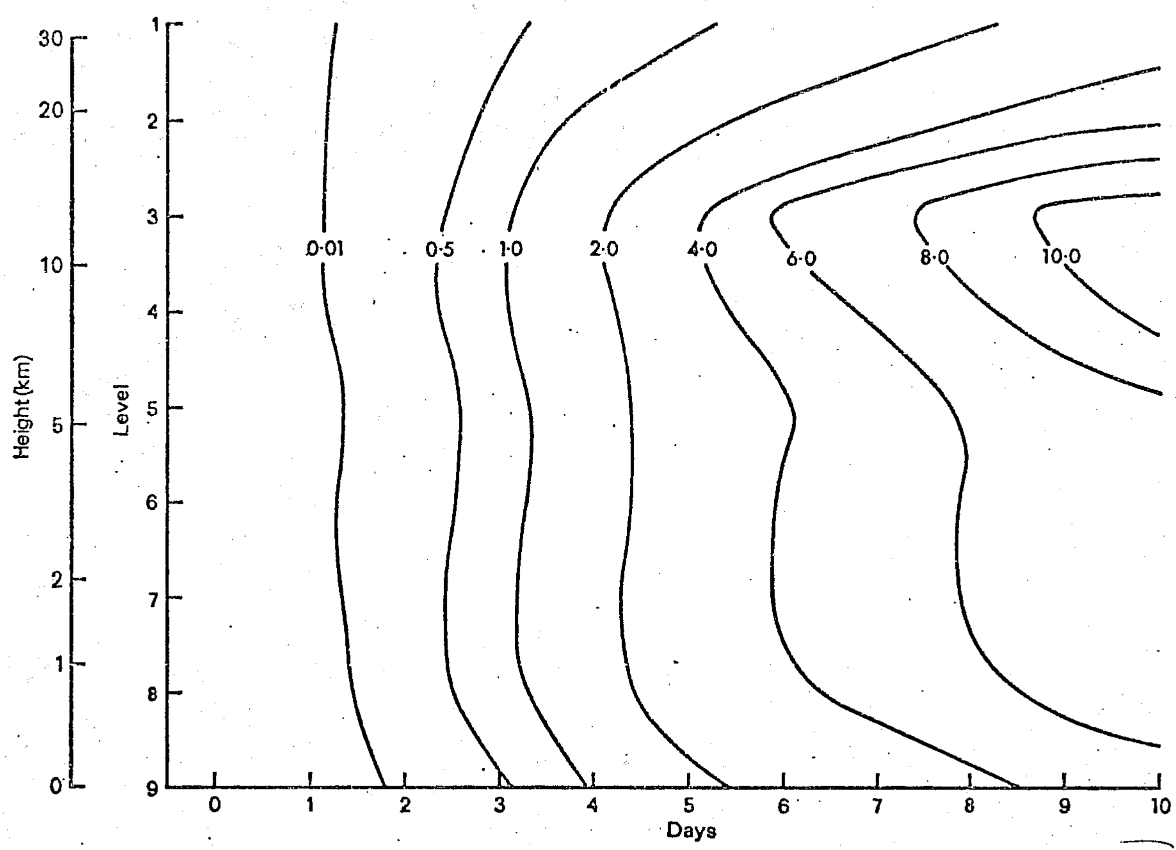

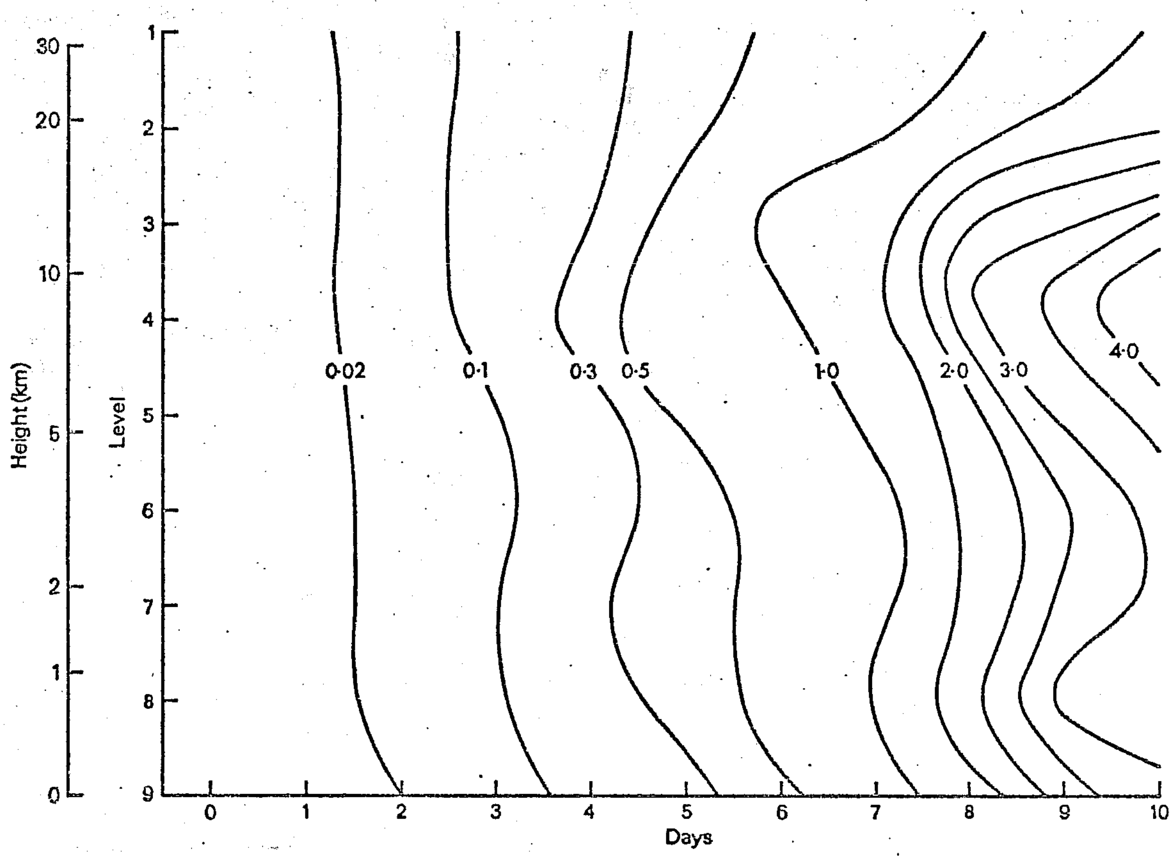

Fig. 10 Vertical structure of the rms wind difference as a function of forecast time for the equatorial region (upper) and for the northern hemisphere (lower)

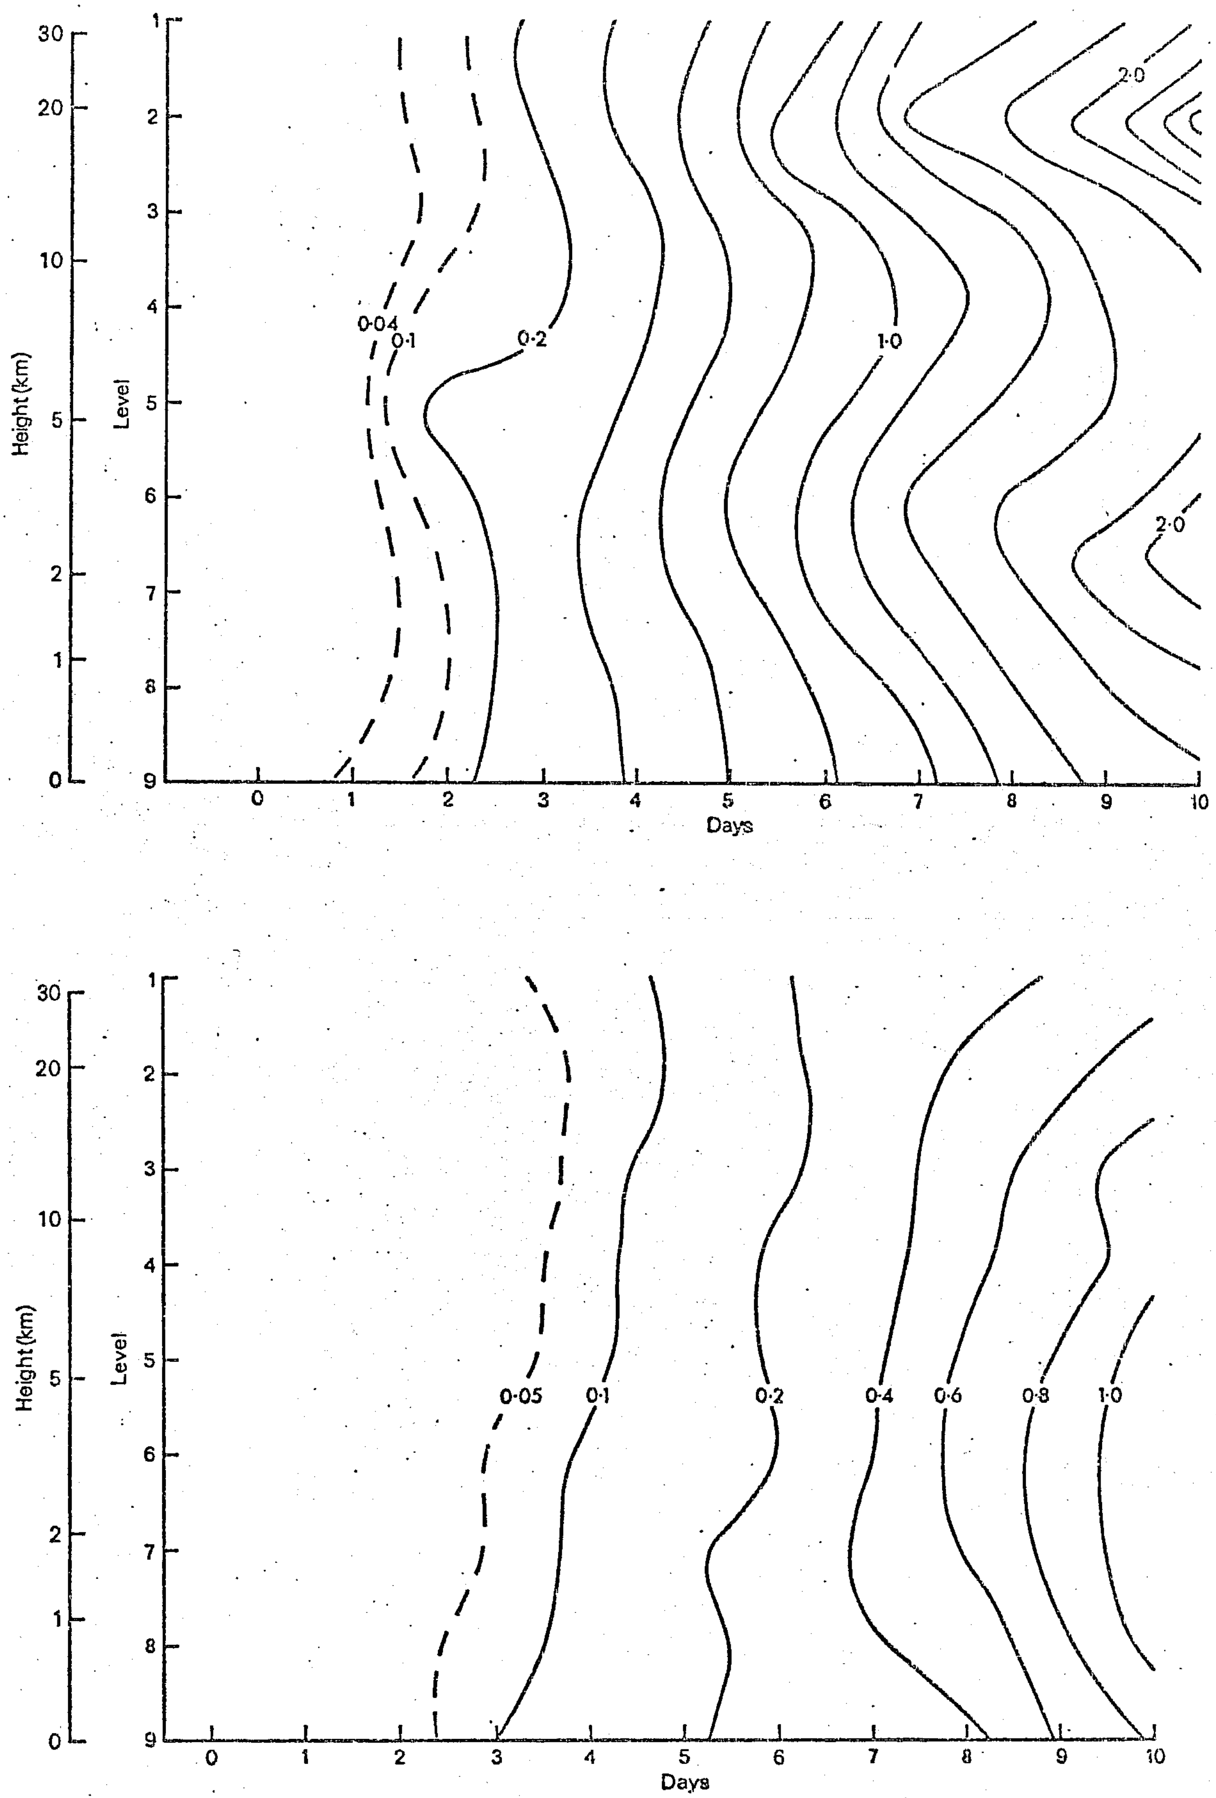

Fig. 11 Vertical structure of the rms temperature difference as a function of forecast time for the equatorial region (upper) and for the northern hemisphere (lower).

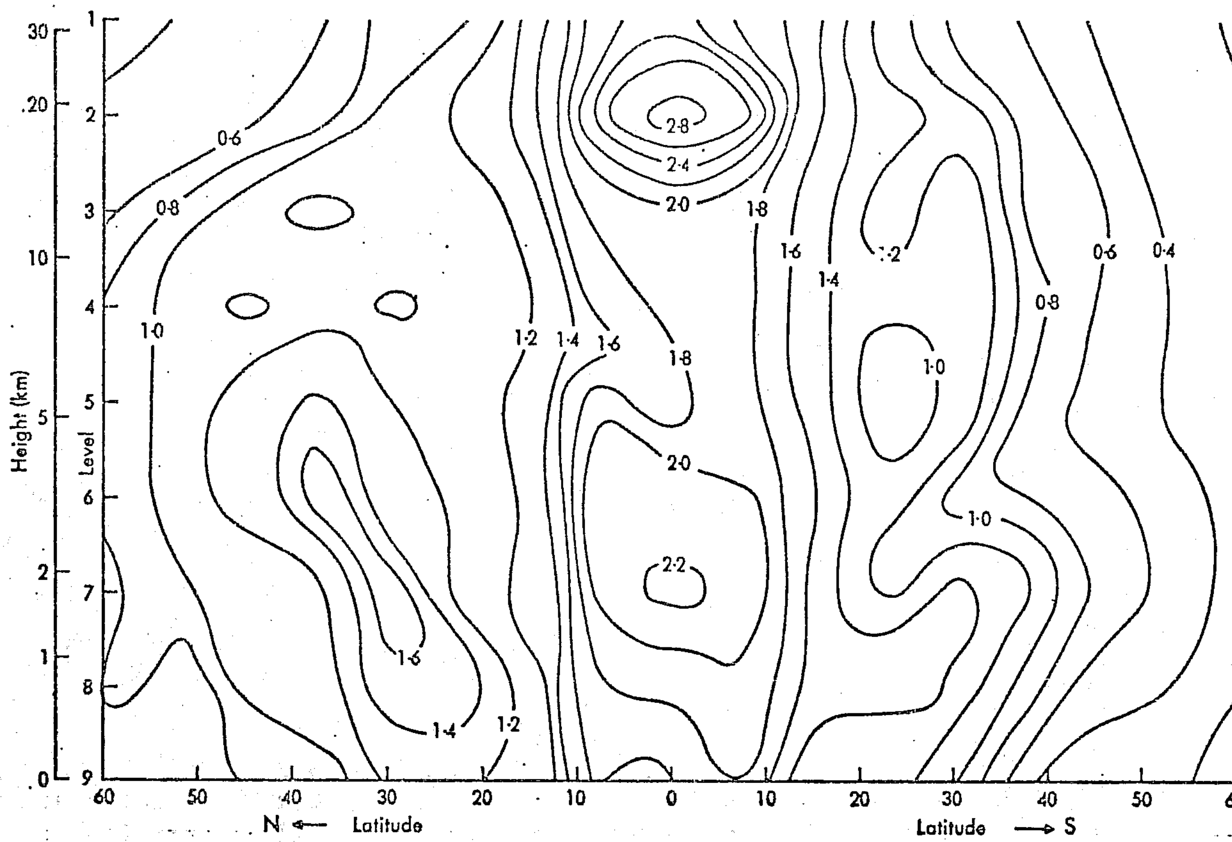

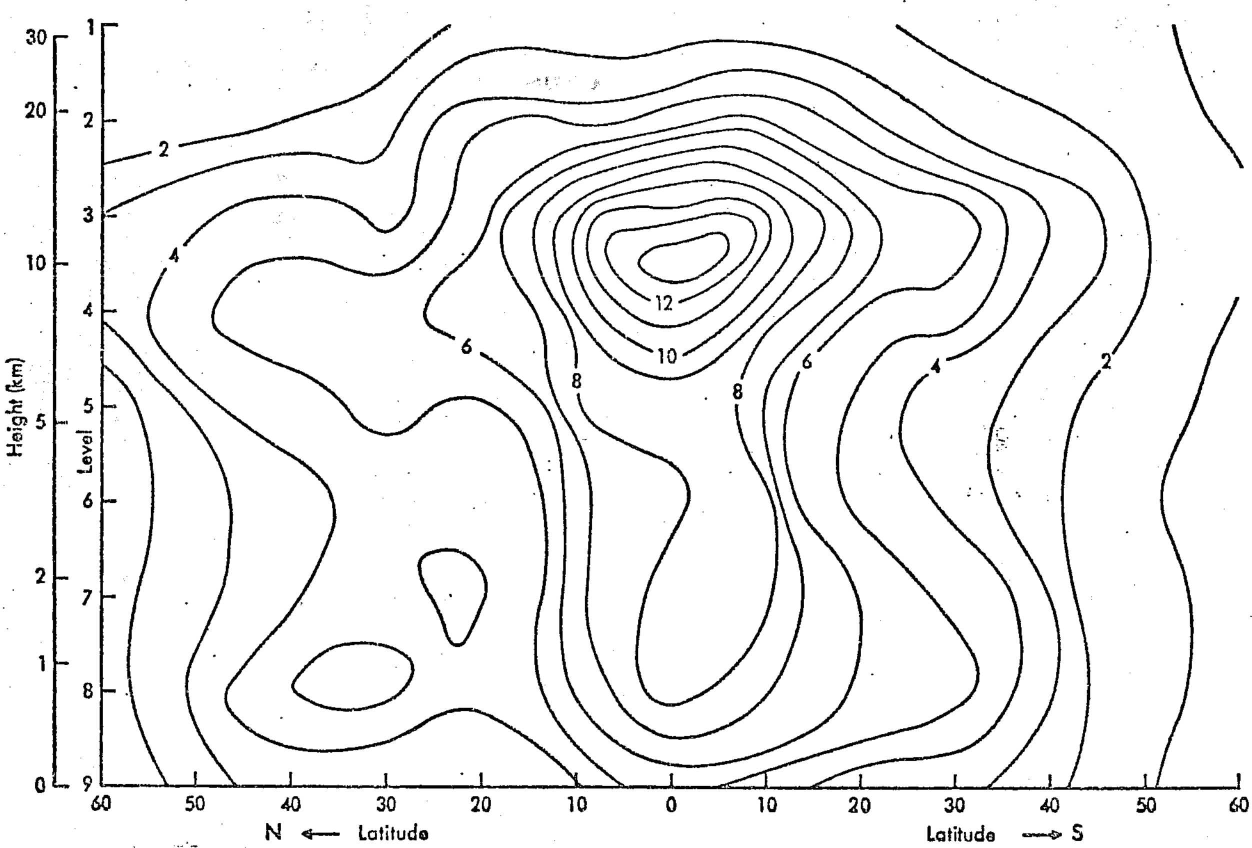

Fig. 12 Meridional section at day 10 of the rms temperature difference (upper) and the rms wind difference (lower).

6. Conclusion

From the adiabatic spectral experiment it can be concluded that for the purposes, for which such models are used, 24-bit mantissa arithmetic provides sufficient accuracy. The non-conservation of mass is almost completely dominated by time truncation errors rather than round-off errors. The observed non-conservation of energy may be shown to be caused by changes in the lowest two bits in the mantissa of the total energy.

The results obtained in the present comparative study with the GFDL general circulation model are in agreement with the findings of WILLIAMSON and WASHINGTON (1973). In particular the large differences between both runs should be attributed to the Moist Convective Adjustment (MCA) and, in our case, to the arithmetic in the process of deciding whether a particular ascent is stable or not. In their experiment the difference in word length was clearly responsible for the effect, because this was the only difference between their runs. In our case the situation is much more complicated. It should be emphasised again that the MCA part of the code is run in single precision on the CDC and in double precision on the IBM, which implies that the length of the mantissa was comparable in both runs (see table 1). As a consequence compiler differences play an important role, a factor which is not easily evaluated. Therefore we have not yet been able to isolate the cause of the problem.

Our experiments, together with the results of WILLIAMSON and WASHINGTON, justify the conclusion that this part of the code requires high precision arithmetic and that very careful coding is necessary.

Acknowledgement

The authors express their gratitude to the Geophysics Fluid Dynamics Laboratory/NOAA for supplying their model and the initial data. Similarly they thank Drs. Hoskins and Simmons of the U.K. Universities Atmospheric Modelling Group in Reading for making their model available. Thanks are due to Dr. W. Blumen and Dr. L. Bengtsson for their critical reading of the manuscript and to Mrs. A. Dinshawe for typing the manuscript.

References

Arpe, K., Bengtsson, L., Hollingsworth, A., and Janjić, Z., 1976; A case study of a 10 day prediction. ECMWF Technical Report No. 1.

Hoskins, B.J. and Simmons, A.J., 1975; A multi-layer spectral model and the semi-implicit method. Quart. J.R. Met. Soc. 101, 637.

Kurihara, Y. and Tuleya, R.E., 1974; Comments: " On the Importance of Precision for Short-Range Forecasting and Climate Simulation". J. Appl. Met. 13, 601.

Miyakoda, K., 1973; Cumulative results of testing a meteorological-mathematical model. The description of the model. Proc. Roy. Irish Academy 73A, 99.

Searle, J.W. and Davies, D.R., 1975; A note on the effect of grid-point numerical errors on the large scale flow characteristics of a four level model of the atmosphere. Tellus, XXVIII, 2.

Smagorinsky, J., 1969; Problems and promises of deterministic extended range forecasting. Bull. Am. Met. Soc. 50, 286.

Williamson, D.L. and Washington, W.M., 1973; On the Importance of Precision for Short-Range Forecasting and Climate Simulation. J. Appl. Met. 12, 1254.

ECMWF LIBRARY