TECHNICAL MEMORANDUM

33

Analysis code: Operational implementation

G.J. Cats

Research Department

June 1981

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

ABSTRACT

Tests of a time reduced analysis scheme are reported. CPU time savings of about 17% are achieved, but some meteorological changes required benchmarking of the scheme. The benchmark results suggest that the scheme is acceptable for operational purposes.

1. INTRODUCTION

Although the analysis code has been written to be very efficient, there is still room for optimisation. Particularly, the double inversion of each covariance matrix can be replaced by one inversion and one solution of a linear system of equations.

As written in Technical Memorandum No. 26 (Cats, 1981, hereafter referred to as (I)), this optimisation, which in itself is unmeteorological, must go together with some other changes. These are of an unphysical nature; they do, however, slightly influence the estimated analysis error, and, consequently, the data checking. It must be anticipated, therefore, that some observations might be treated in a different way by the optimised version. This will in particular happen with data that are on the borderline of being rejected. Because these data have in general a large deviation from the preliminary analysis without the use of the checked datum, different treatment of these data will change the analysis substantially.

Section 2 describes the effect of the optimisation and associated changes in a one cycle test analysis for each change separately.

Because of the changes in the data selection, it was considered a suitable opportunity to test another change there as well, namely a reduction in the number of data selected to check data (this follows a suggestion by O.Talagrand). This change can neither be defended nor fought from physical principles, but it will have a meteorological effect. It reduces the time spent in data checking considerably (an I/O time saving is expected as well). The effect of this change in the test analysis is described in section 3.

In section 4 the total effect of all changes is assessed in another test case and section 5 gives the results of a benchmark against a (pseudo-) operational analysis and forecast.

2. OPTIMISATION CODE

2.1 Introduction

The introduction of SYMSOL, the subroutine to solve a linear system of equations, to replace SYMINV, the matrix inversion subroutine, required many changes (I). Their effect will be described separately for each change. The effects have been checked using a test analysis (7 September 1980, 1200 GMT). In the following the acronym gp-run will be used for the calculation of mass and wind grid point values. I assume it to follow a data checking (dc) run. The present operational code will be called "the present code" and the code under testing the test code.

2.2 Use of unchecked data

If in an analysis volume (a slab in a box) there are more than 191 data, not all data will be checked in the dc run. If some of the checked data are rejected, the present code then also selects unchecked data in the gp run, to be used for the final analysis. This is, fortunately, unlikely to occur, due to the formation of super observations etc. in data dense areas. Yet, it must be considered a coding error.

Although it is not directly related to the introduction of SYMSOL, it was decided to change the code, so that in the gp-run only checked data are selected. The test analysis did not show any change at all, indicating that the present code in this case did not use unchecked data in the gp-run.

2.3 Observation error correlation

In the present code two formulae are used for the observation error correlation c_{ij}. The first, reading

$$c_{ij} = c_o \mu_{ij} h_{ij} \tag{1}$$

(where c_o is a constant (0.8), μ_{ij} the vertical and h_{ij} the horizontal correlation matrix) is used if the data are in different reports (satem-thicknesses). The second:

$$c_{ij} = \mu_{ij} \tag{2}$$

is used for data within one report (temp heights and satem-thicknesses). Clearly (2) is not the limit of (1) for small horizontal distances. From a physical point of view this is not correct ((2) implies that, if in a temp the height of a pressure level is reported twice, these two heights are the same, which is not necessarily true in practice; furthermore, this case would be treated differently to a case where a temp report comes in twice).

The introduction of SYMSOL requires that the uncorrelated part of the observation error have a certain minimum value (I). Therefore, (2) was replaced by

$$c_{ij} = c_o \mu_{ij} \tag{3}$$

The uncorrelated observation error is then at least $(1-c_o)$ times the observation error variance. If this was smaller than the minimum required by SYMSOL (which, in practice, is seldom, if ever, the case) it was replaced by that minimum in the test gp run.

The effect of the change from (2) to (3) was large. The maximum difference in height analyses occurred at 140°E, 19°N, 1000 mb. The present code accepts a synop of 992 mb MSLP and rejects a temp of -17m 1000 mb height. The new code prefers the temp report, with a resulting increase in the 1000 mb height of 34m. Surprisingly enough, the two reports had the same call sign EREBO, the same observation time, but different (by 100 km) positions. These data were found to be borderline cases before (Shaw et al., 1980). In many analyses for other days this ship is recognised as peculiar as well.

The large change in surface pressure went, of course, together with a large change in surface wind (6 m/s). Other changes were less important (in height less than 20m, in wind less than 2 m/s, apart from some points in the stratosphere).

The number of ill-conditioned matrices reduced from 3 to 1. This is in accord with the theory (Cats, 1980).

2.4 SYMSOL

In the next test analysis SYMSOL was used to solve a linear system rather than invert the matrix in the gp run. The analysis error calculation was shifted from the gp to the dc run. The changes are all as expected.

Data checking and rejection did not change at all. The mass and wind analysis changed by less than 10^{-9}m and 10^{-9}m/s, apart from one analysis volume, where the difference was 0.5m in height and 0.1 m/s in wind speed. This analysis volume was the one with the ill-conditioned matrix (section 2.3). The present code cures ill-conditioning by an unmeteorological increase of the observation error. SYMSOL cannot detect the ill-condition (I) and therefore uses the matrix in unchanged form. There is no reason for worry here, because the criterion used in the present code is far more severe than actually needed (compare (I) with Cats and Robertson, 1980). (This is the first opportunity to estimate the effect of the matrix fix-up in the present code; up to now, no baseline existed because it would have required the inversion of a matrix too ill-conditioned to be inverted). The normalised analysis error changed by less than 0.1. This change originates from the use of unchecked data (I). From the one cycle test it cannot be inferred whether this change has a meteorological effect. The analysis error is added to an estimated forecast model error to obtain the first guess error for the next cycle. Because the forecast model error estimate is very rough, the analysis error change is of unphysical importance.

Time savings are estimated to be 40s per analysis run, assuming that the CAL version of SYMSOL is four times as fast as the Fortran version (in analogy with SYMINV).

The analysis error estimate cost 40s (with about maximal optimisation I was able to reduce this by 5s). This time seems unreasonably large as compared to, e.g. inversion of the matric (70s) or calculation of N48 15 level grid point values (70s).

2.5 Conclusion

Using SYMSOL in the analysis scheme will save an estimated 40s. per cycle. The introduction of SYMSOL requires some changes that are acceptable or even desirable from a physical point of view, but they may have an unpredictable effect via the data checking. The changes are a different specification of the analysis error and a slightly different set of observations used to estimate the analysis error. The analysis changes sometimes considerably (by up to 4mb MSLP), but this is brought about by rejecting or accepting data that are on the borderline of being acceptable. This implies that it cannot be decided uniquely whether these data should be accepted or not, and the analysis scheme (both present and test versions) must be expected to accept or reject these data essentially at random. It might be argued that often from other information sources (e.g. time continuity, structure and likelihood of hurricanes), the reliability of borderline data can be inferred. However, because the analysis scheme does not know this information, it will sometimes take the right and sometimes the wrong decision, dependent on the value of an essentially random number, (see section 3), both in present and test versions.

3. REDUCED DATA SELECTION FOR DATA CHECKING

Because the introduction of SYMSOL required changes that are of consequence to the data selection, it was considered a good opportunity to test some other changes in the data checking as well. In the analysis dc run, a datum is flagged with flag k if k is the highest number i, $i \leq 3$ for which

$$\sigma > a_i \, (\varepsilon^2 + \delta^2)^{\frac{1}{2}} \tag{4}$$

where $\sigma = |O - A|$

O = observed value

A = analysis without the datum being checked

ε = modelled observation error

δ = analysis error estimate (without the datum being checked)

$a_1 = 3$, $a_2 = 4$, $a_3 = 5$

From the values of a_k (k=1,2,3) it is clear that these numbers have not (yet) been established with any accuracy. Further, δ and ε are rough estimates and O contains a random observation error. It is therefore nonsense to spend much time on calculating A, as done presently.

A considerable time saving can be achieved by using less data to check a datum. Code has been tested in which the horizontal search radius from the edge of the box for data checking was reduced to b, the horizontal correlation length. This saved almost 60s (out of 220s for the dc run). As expected, some data received a different final flag. The subjective impression is that one datum was wrongly rejected, and one rightly by this test run, reflecting the random treatment of borderline data.

I have doubts on the rejection of a MSLP datum of 978 mb on the Antarctic Peninsula. Fig. 1 demonstrates that the analysis was hardly changed by this rejection, because of data abundancy. The rightly rejected datum was on the equator, 20°W. Its reported MSLP of 1022 mb was more than 4 mb in excess of any other datum or the first guess. Its rejection reduced the analysis by 4 mb to produce a more reasonable tropical pressure field.

Two more data were treated differently, but without substantial effect on the analysis.

A slight reduction in time (2s) was further obtained by not using data that were more than two levels away from the nearest level in the analysis volume. The gross rejection did not change, but two data had their flags reduced from 3 to 2 (i.e. from wrong to probably wrong - both classes are not used in the final analysis).

4. AN OVERALL TEST

All changes described in sections 2 and 3 have been used in a rerun of the operational midnight analysis on 11 January 1981. This case was chosen because of the erroneous acceptance of the stratospheric wind data (apart from the duly rejected 20 mb wind) in a temp. at 53°N, 174°E. The test run accepted them as well, and even the 20 mb wind. Because of the use of the 30 mb wind in both analyses, the change was small (Fig. 2).(The strong cross isobar flow originates mainly from ageostrophic flow in the first guess). The case as a whole shows remarkable similarity to a case I reported on before (Cats, 1979).

The test analysis accepted many more data than the operational, especially over East Asia. Because of data abundancy, however, the changes in the analysis in that area were very small.

5. OPERATIONAL BENCHMARK

The operational forecast from 8 April, 1981, 1200Z, was repeated from initialised fields that were obtained after five cycles of analyses with the test scheme. A recently proposed change in the stratospheric pressure to sigma interpolation (van Maanen et al., 1981) was included; due to an administrative error the soil moisture field was specified in a slightly different way. A clear comparison is, however, available between the test run and another experiment.

The five test analyses deviated from the operational analyses mainly in that they accepted slightly more data. The largest difference, induced by this, occurred in the last cycle, in the southern hemispheric, 1000 mb, analysis (Fig. 3). The first guess field at the position of ship GWAN (52^{o}S, 96^{o}W) was 997 mb and at first sight the operational analysis seems to produce a better analysis by rejecting the ship. There are hardly any other surface pressure observations nearby, but the one or two PAOBS around do not confirm the ship report. The time sequence of the ship's reports is: 1010 mb at 0804, 0000Z, 1011 mb at 1200Z and 1014 mb at 0904, 1200Z.

The operational analyses, following 8 April, 1200Z, started developing the ridge that is present already in both the test and the verification analyses at 1200Z. Both the test forecast and the comparison one do not develop the ridge well, but the test forecast has clearly a better pressure distribution around Cape Horn in the first four days of the forecast (Fig. 4). In view of the unphysical character of the differences between the two analysis schemes, this superiority of the test forecast is not an argument in favour of the test scheme, but merely not an argument against it.

Objective verification scores for the test and the verification forecast are shown in Fig. 5. They show remarkable similarity over both hemispheres.

Over four consecutive analysis cycles the following CPU time savings were achieved with the test scheme:

During data checking: 2 minutes out of 9.

During grid point evaluation: 4 minutes out of 11.

For the analyses (including humidity analysis, interpolation, etc.): 6 minutes out of 32.
(This includes the effect of shifting the analysis error evaluation from the grid point evaluation run to the data checking run).

6. CONCLUSION

A number of changes have been tested to decrease the CPU time taken by the analysis, without decreasing its quality. Savings over four cycles were of the order of 6 minutes CPU (17% of the total analysis CPU time). The scheme has, via unphysical arbitrariness in the decision whether to accept or reject data, meteorological consequences and has therefore been benchmarked. The benchmark results do not indicate that there is any danger in introducing the test code operationally.

The sensitivity of the analysis to unphysical quantities via data rejection is a worrying aspect of our scheme (and many other schemes). It is therefore of the utmost importance to develop and use a better data checking scheme, even if it turns out to be far more expensive computationally than the present scheme.

REFERENCES:

Cats, G.J., 1979: A Small-scale Cyclone in the Stratosphere. Memorandum, ECMWF, dated 12.12.80.

Cats, G.J. and D.A. Robertson, 1980: Condition of the Covariance Matrix and Stability of its Inversion Algorithm. ECMWF Working Paper No.1/6/GJC/071/1980.

Cats, G.J., 1980: Construction of Correlation Matrices for the ECMWF Analysis Scheme. ECMWF Working Paper No.1/10/E/GJC/073/1980.

(I): Cats, G.J., 1981: A Method for Solving a System of Linear Equations efficiently in order to optimise the Analysis Code. ECMWF Technical Memorandum No. 26.

Shaw, D., O. Talagrand, G.J. Cats and P. Lönnberg, 1980: Planned Changes to the Operational Suite. Memorandum, ECMWF, File: R2314, 7.11.80.

van Maanen, J., J. Andersen, D. Shaw and A. Simmons, 1981: Change in the Vertical Interpolation in the Stratosphere. Memorandum, ECMWF, File: R2318/No. 286, 13.5.81.

Figure Captions

Fig. 1 Control (upper) and reduced data checking (lower) analysis for 1200 GMT, 7 September 1980.

Fig. 2 Operational analysis (upper panels) and rerun (lower panels) of height and wind at 0000 GMT, 11 January 1981, at 20mb (large panels), 30mb (upper of each pair) and 70mb (lower of each pair of small panels). The dot gives the position of the erroneous wind report.

Fig. 3 Comparison (left panel) and test (right) analysis from benchmark experiment. The figure shows MSL pressure over the area with the largest difference between the two analyses. Analyses are 5th cycle, validity time: April 8th, 1981 1200Z

Fig. 4 Test (lowest panels) and comparison (middle panels) forecasts for 2 (left panels) and 4 days (right panels) from 08 04 1981, 1200Z. Shown are MSL pressure (full) and 1000mb temperature (dashed) over an area similar to that of Fig. 3. The top panels give the verifying operational analyses, on a slightly different scale, (here dashed lines show 850mb temperature).

Fig. 5 Objective verification scores for the test forecast (F06) and the verifying forecast E86). Ignore ECF (operational forecast) and F14 (a deep soil parameter experiment, described elsewhere).

Fig. 5a : Northern hemisphere

Fig. 5b : Southern hemisphere.

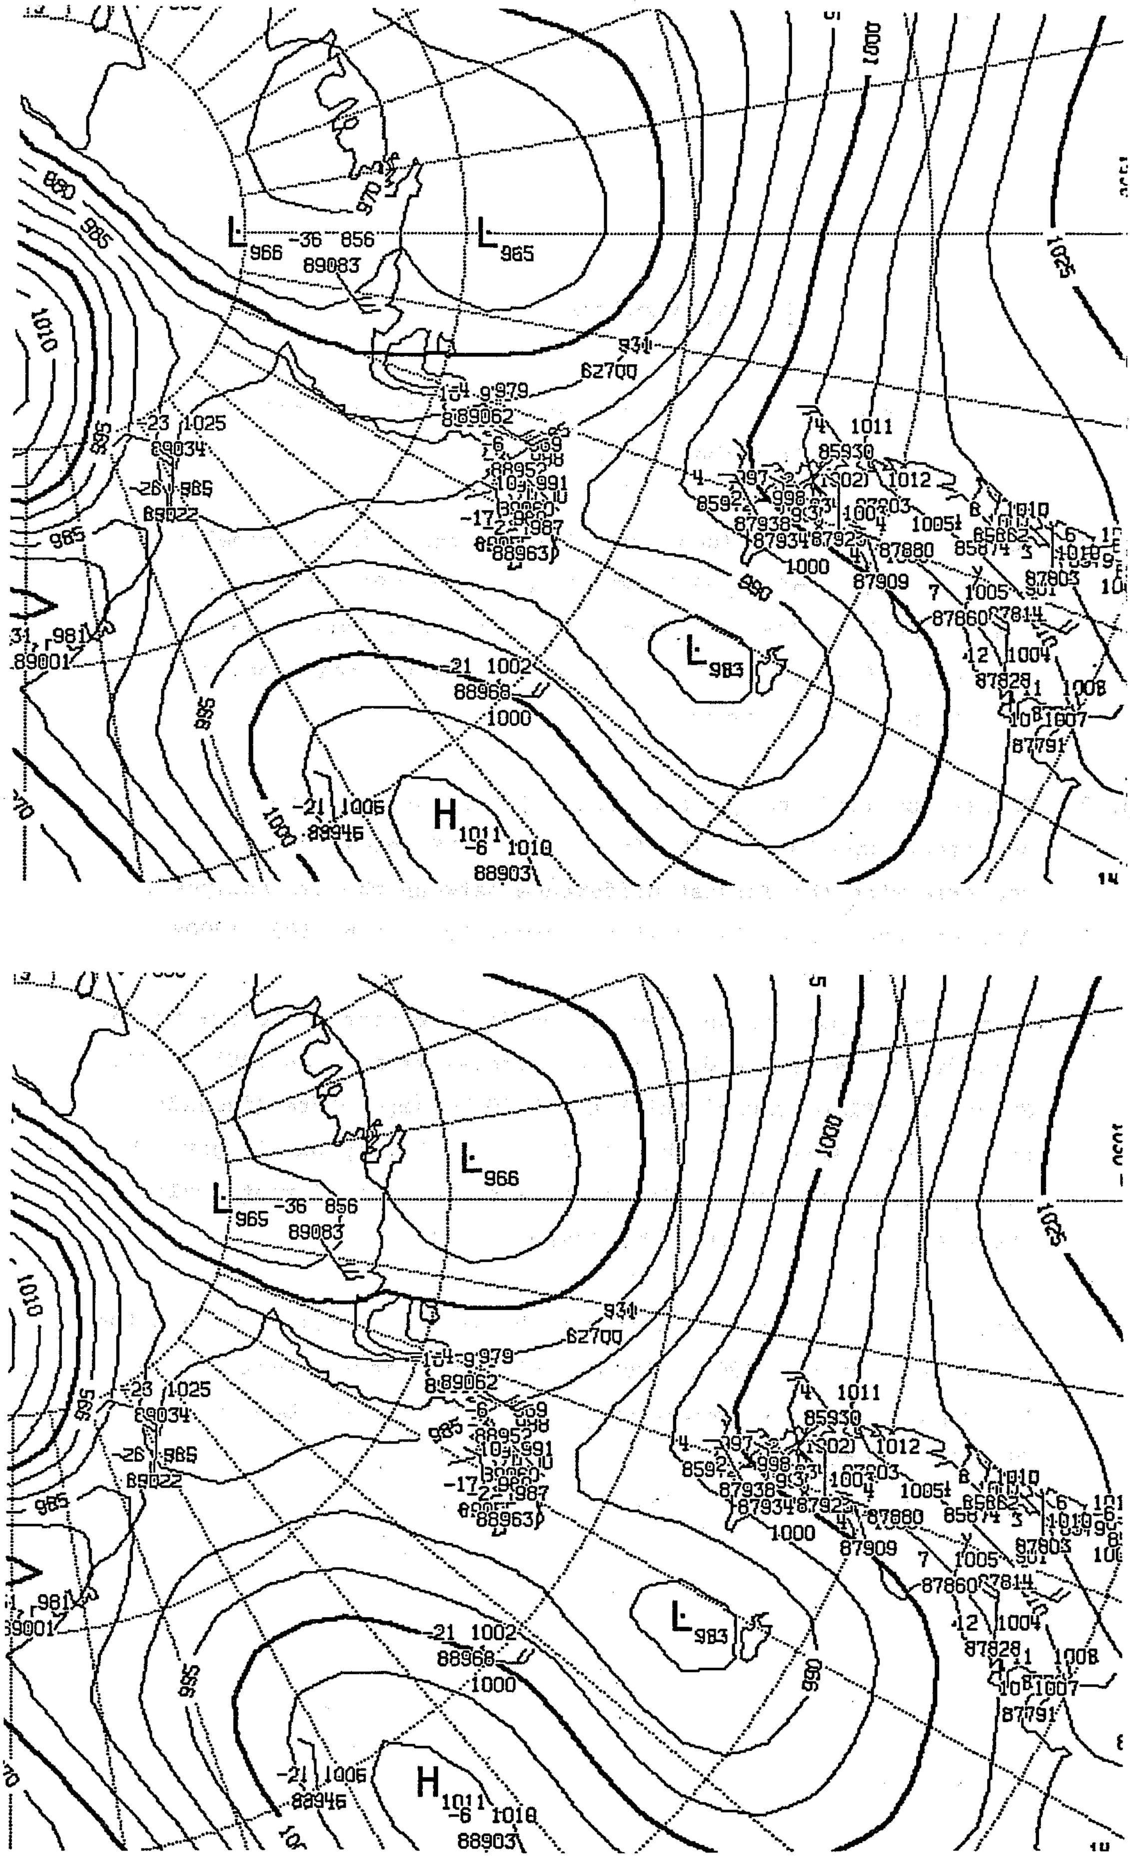
L 966
L 965
-36 856
89083
931
62700
979
89062
-23 1025
89034
-26 985
89022
981
89001
-21 1002
88968
1000
L 983
H 1011
-6 1010
88903
-21 1006
88946
1011
85930
1012
87909
1005
87860
1004
87828
1008
1007
87791
960
970
985
995
1010
1025
990

Fig. 1

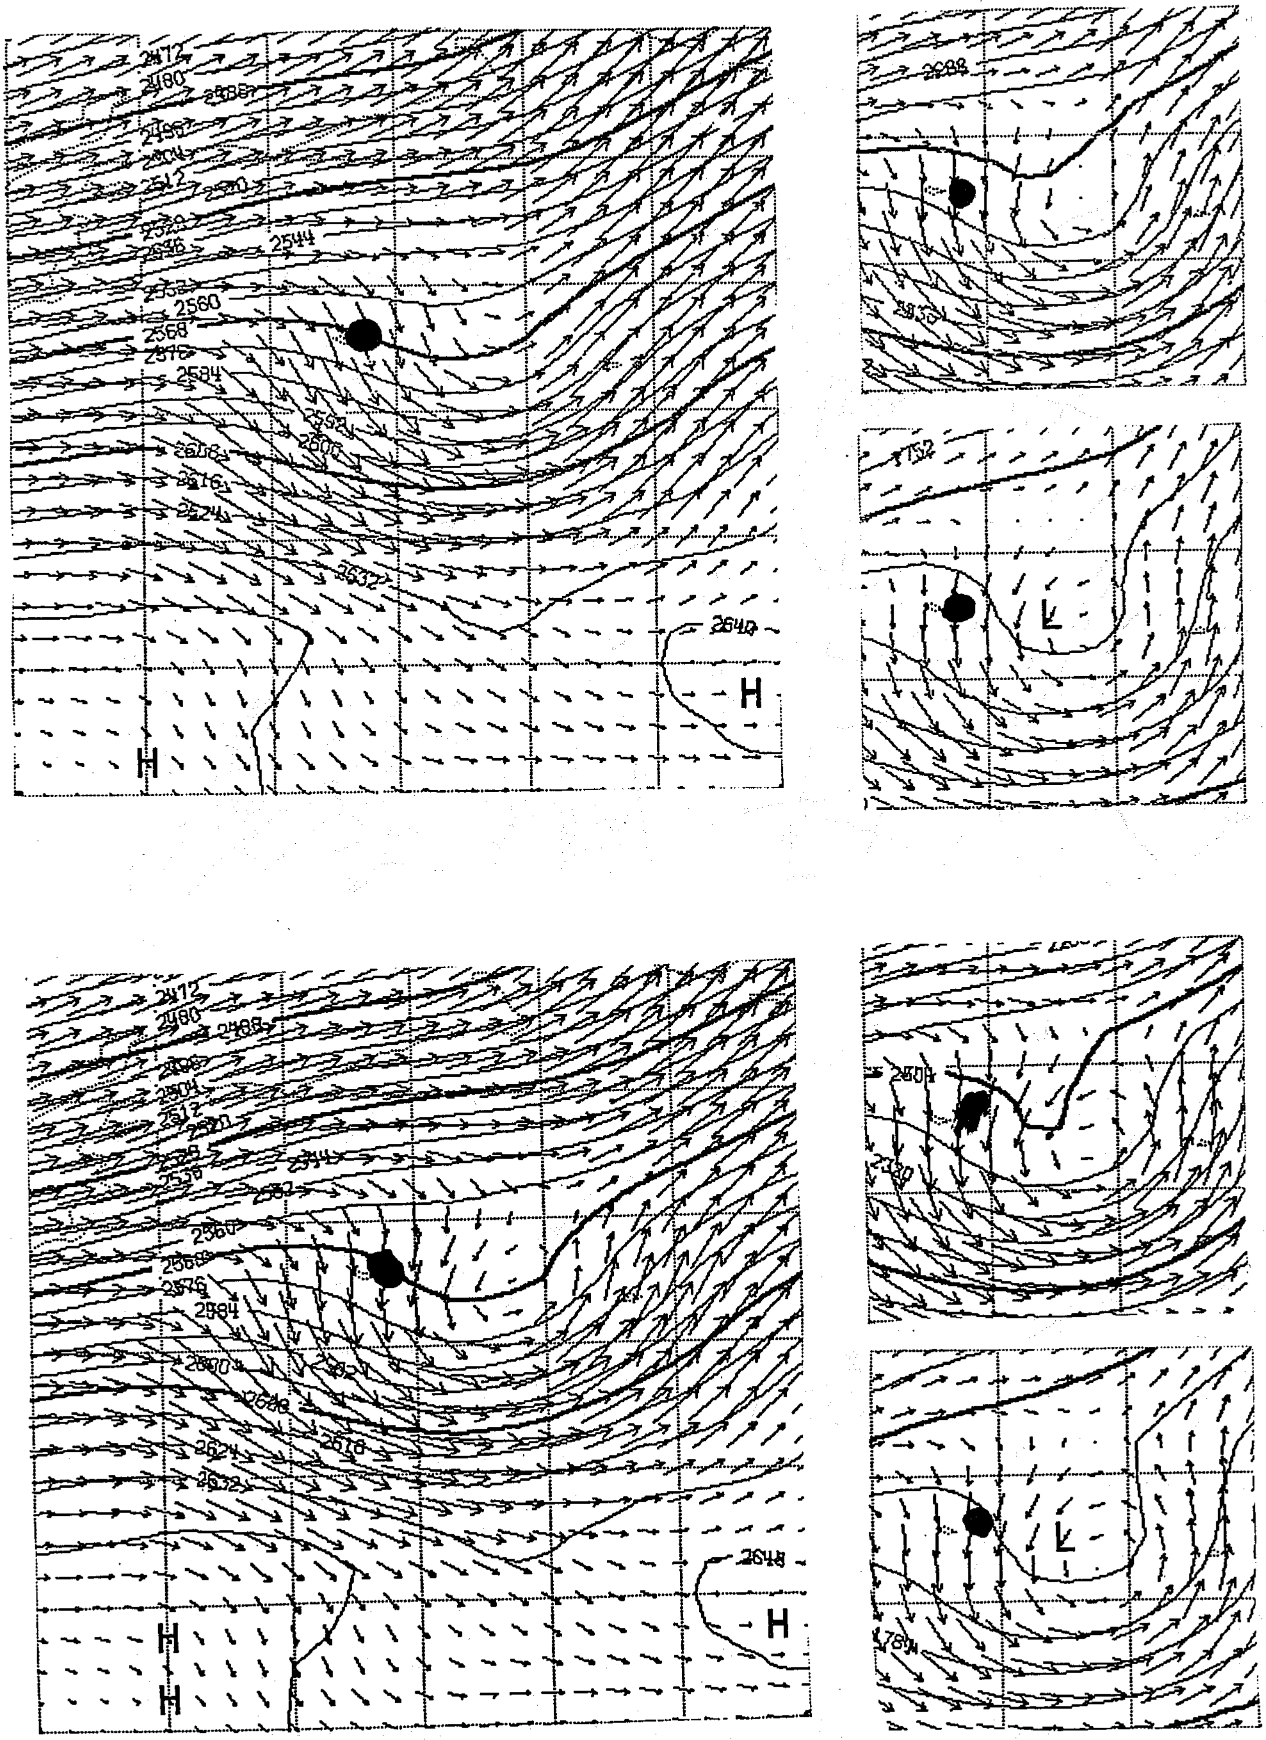
2472
2480
2488
2496
2504
2512
2520
2528
2536
2544
2552
2560
2568
2576
2584
2592
2600
2608
2616
2624
2632
2640
H
H
2688
2752
L
2472
2480
2488
2496
2504
2512
2520
2528
2536
2544
2552
2560
2568
2576
2584
2592
2600
2608
2616
2624
2632
2648
H
H
H
L

Fig. 2

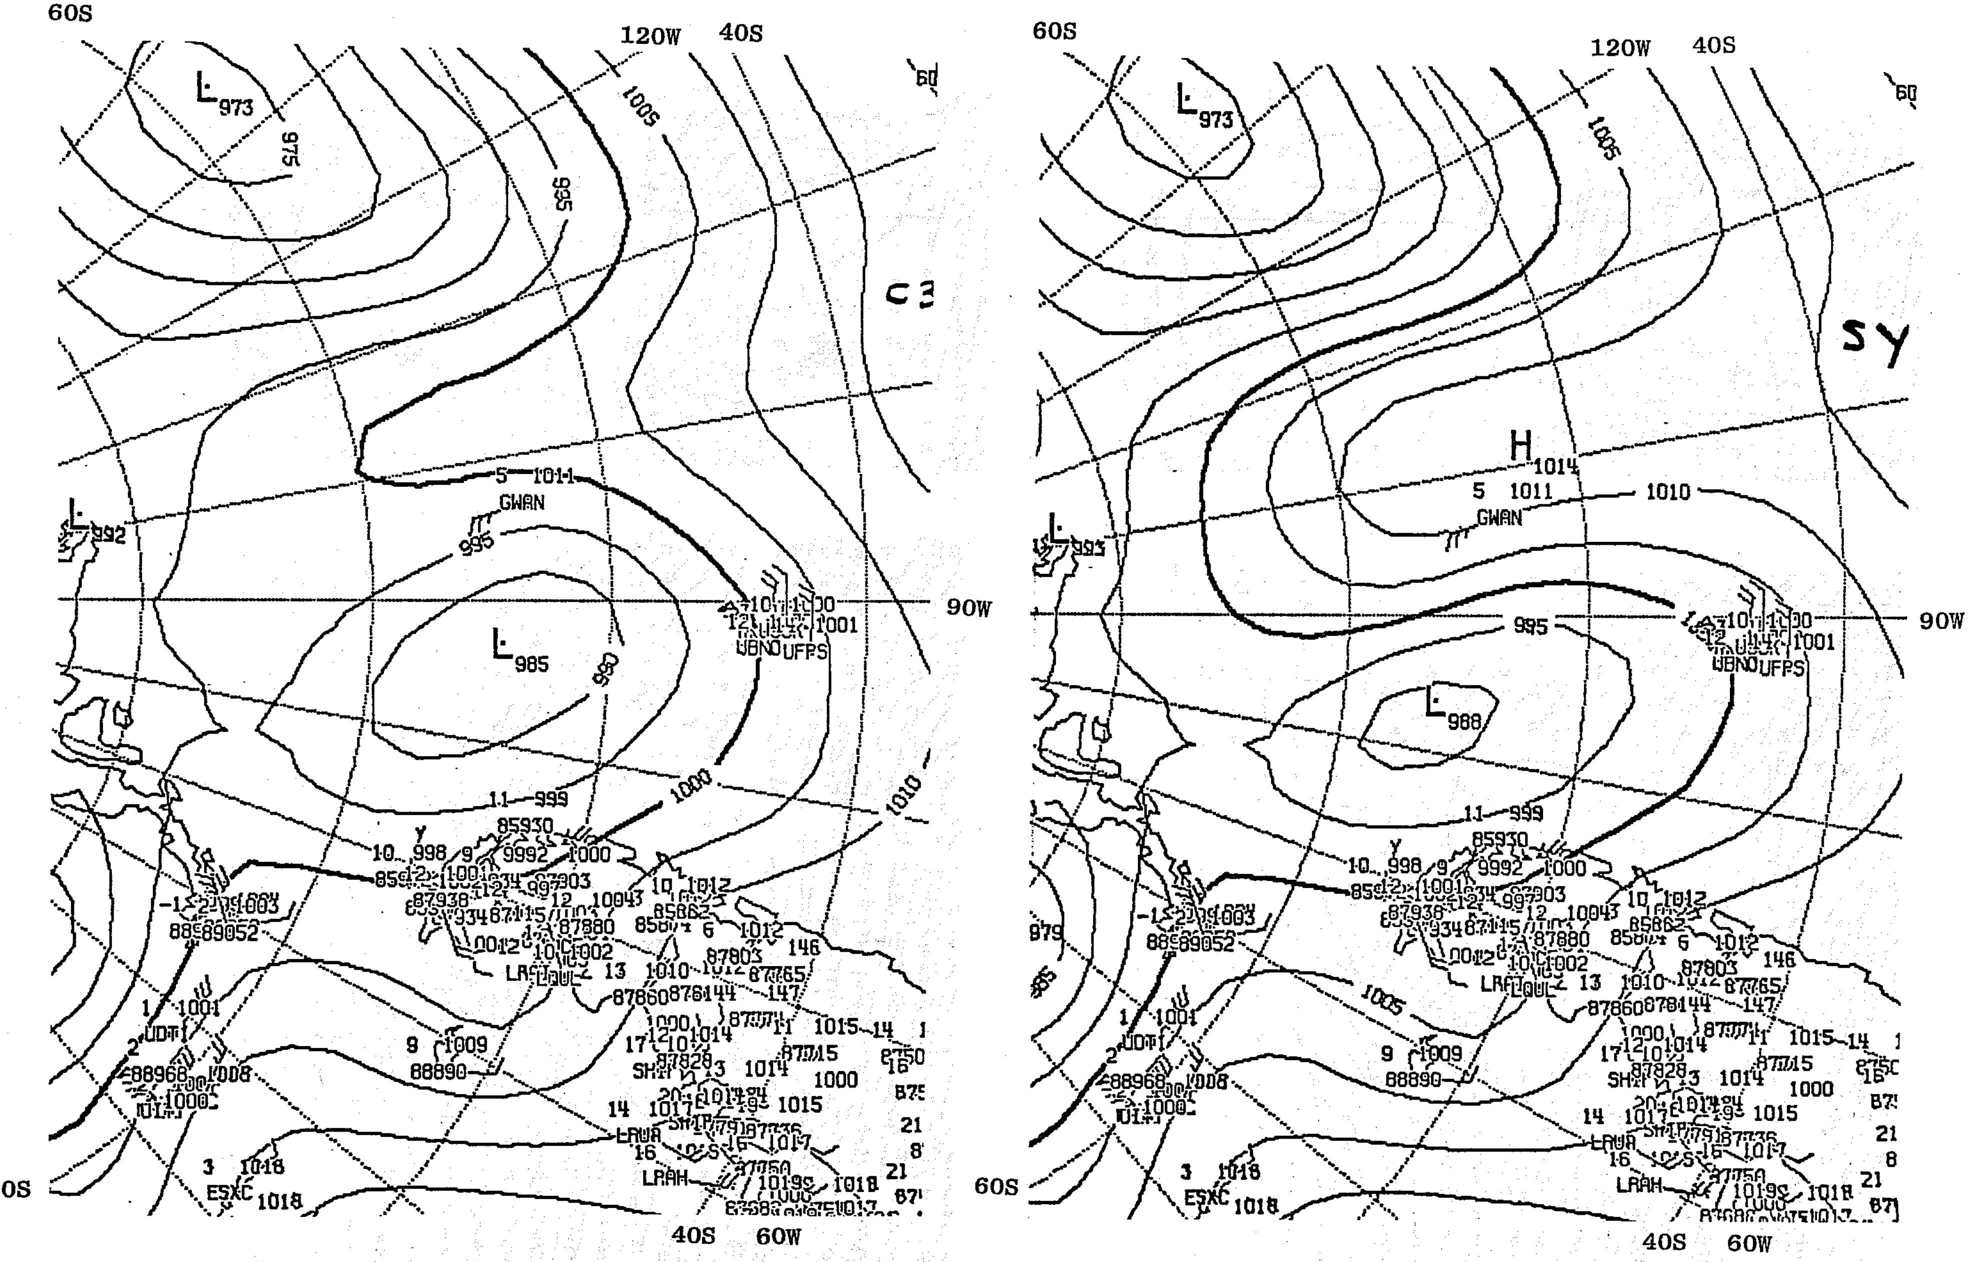
60S
120W
40S
60
C3
L 973
975
995
1005
5 1011
GWAN
L 992
995
L 985
990
90W
1000
1010
11 999
60S
40S
60W
SY
L 973
1005
H 1014
5 1011
1010
GWAN
L 993
995
L 988
1005
ESXC 1018
UBNO UFPS

(a)

(b)

Fig. 3

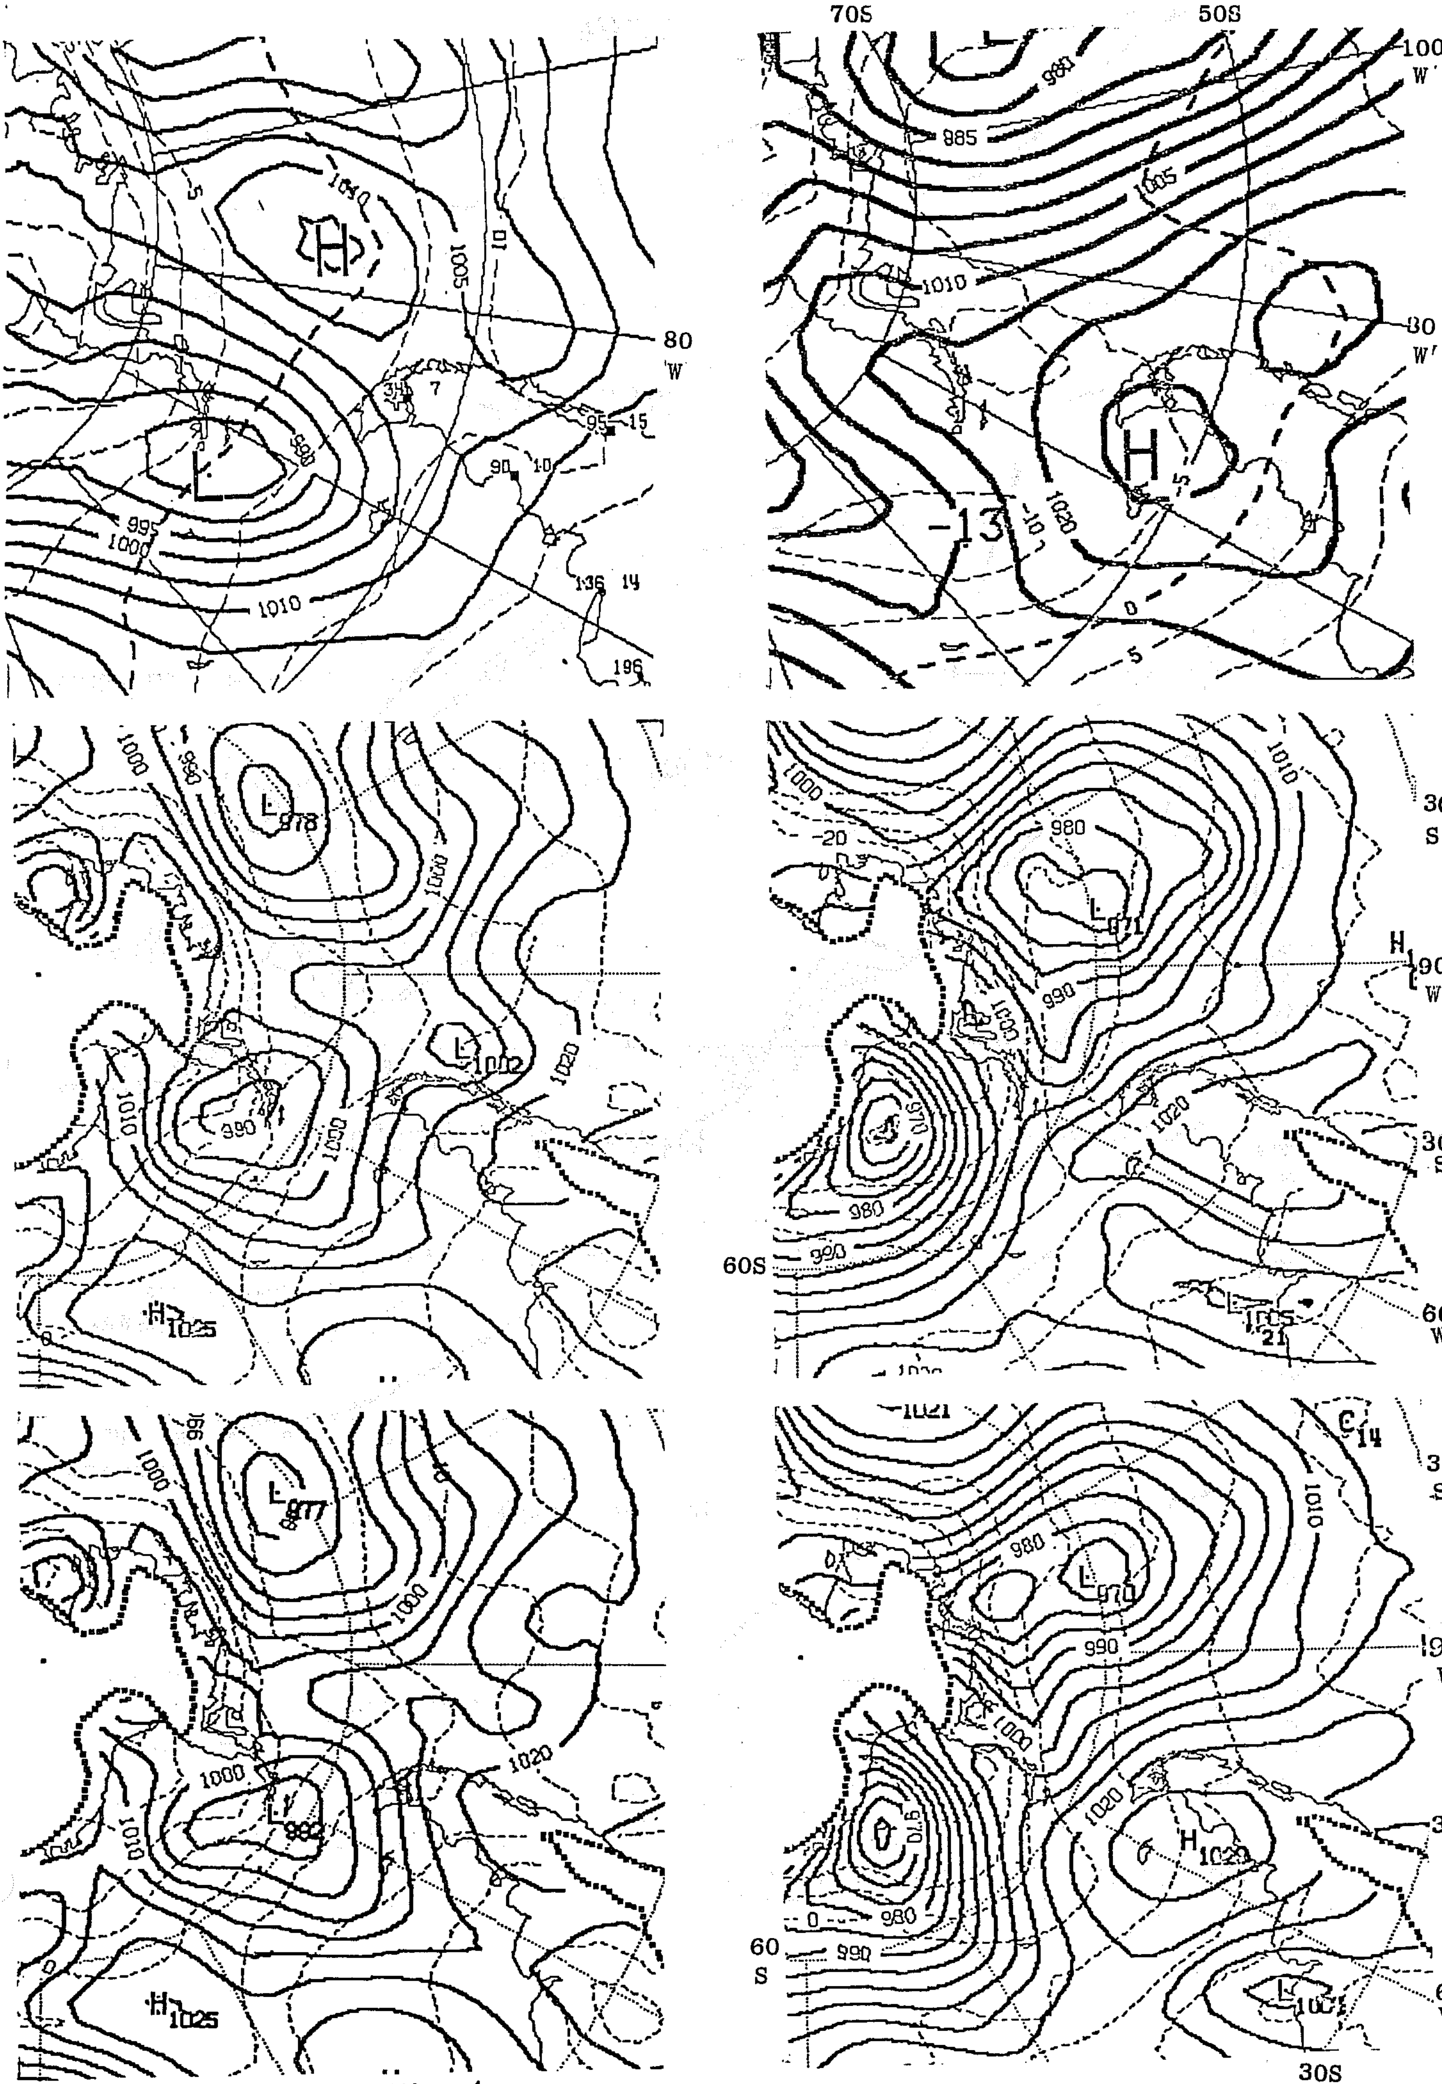
70S
50S
100 W
80 W
80 W
H
L
1040
1005
990
995
1000
1010
985
1010
1020
-13
L 978
L 1012
990
1000
1010
1020
1030
H 1025
L 971
980
990
970
1000
1020
-20
60S
30 S
90 W
60 W
L 1025
L 977
982
1000
1010
1020
H 1025
1021
L 970
980
990
1000
1010
1020
H 1020
60 S
30S

Fig. 4

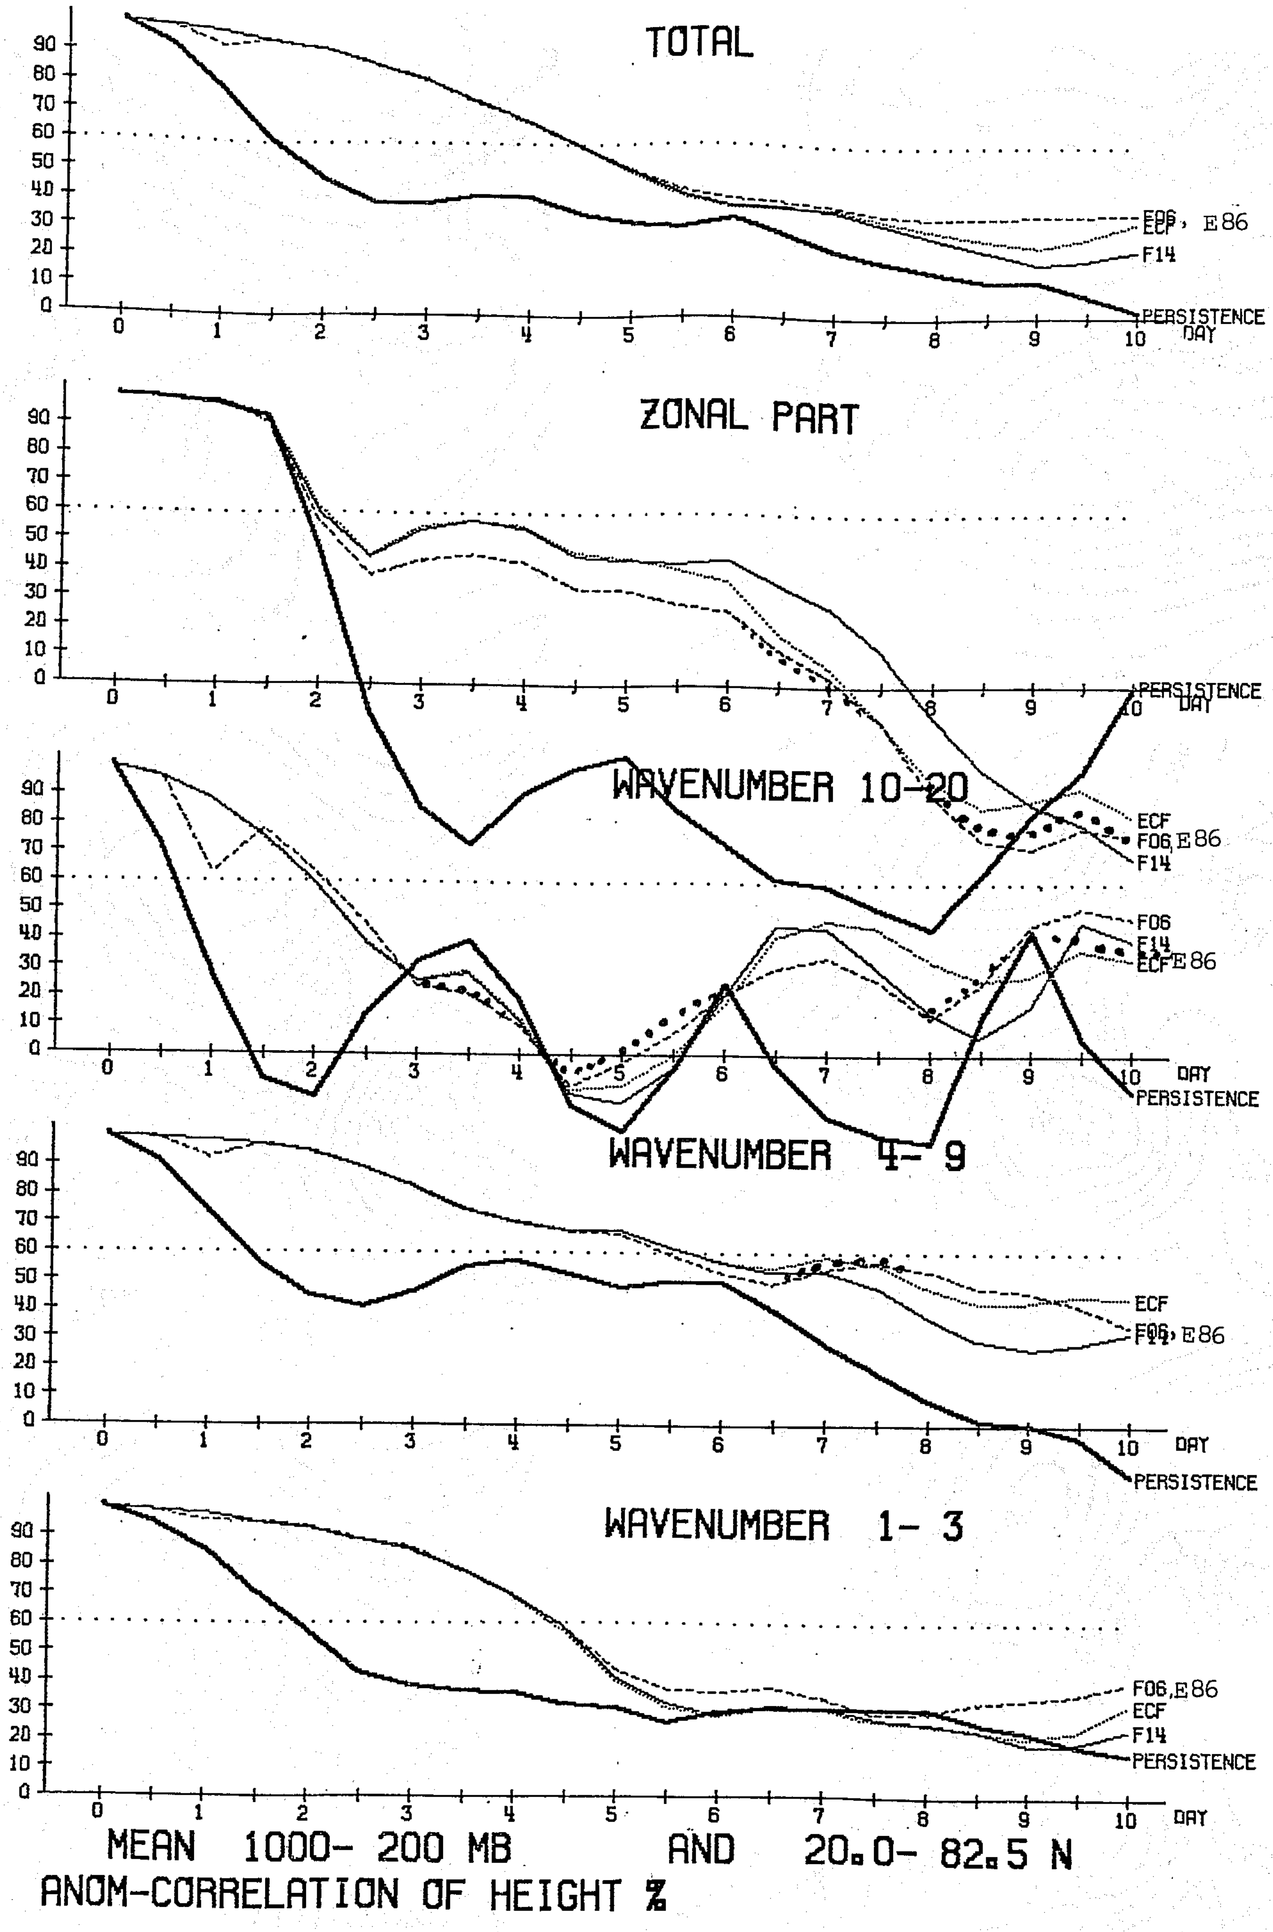
TOTAL
ZONAL PART
WAVENUMBER 10-20
WAVENUMBER 4- 9
WAVENUMBER 1- 3
0
10
20
30
40
50
60
70
80
90
0 1 2 3 4 5 6 7 8 9 10
DAY
F06, ECF, E86
F14
ECF
F06, E86
PERSISTENCE
F14
F06
ECF E86
F06, F14, E86
MEAN 1000- 200 MB AND 20.0- 82.5 N
ANOM-CORRELATION OF HEIGHT %

Fig. 5a

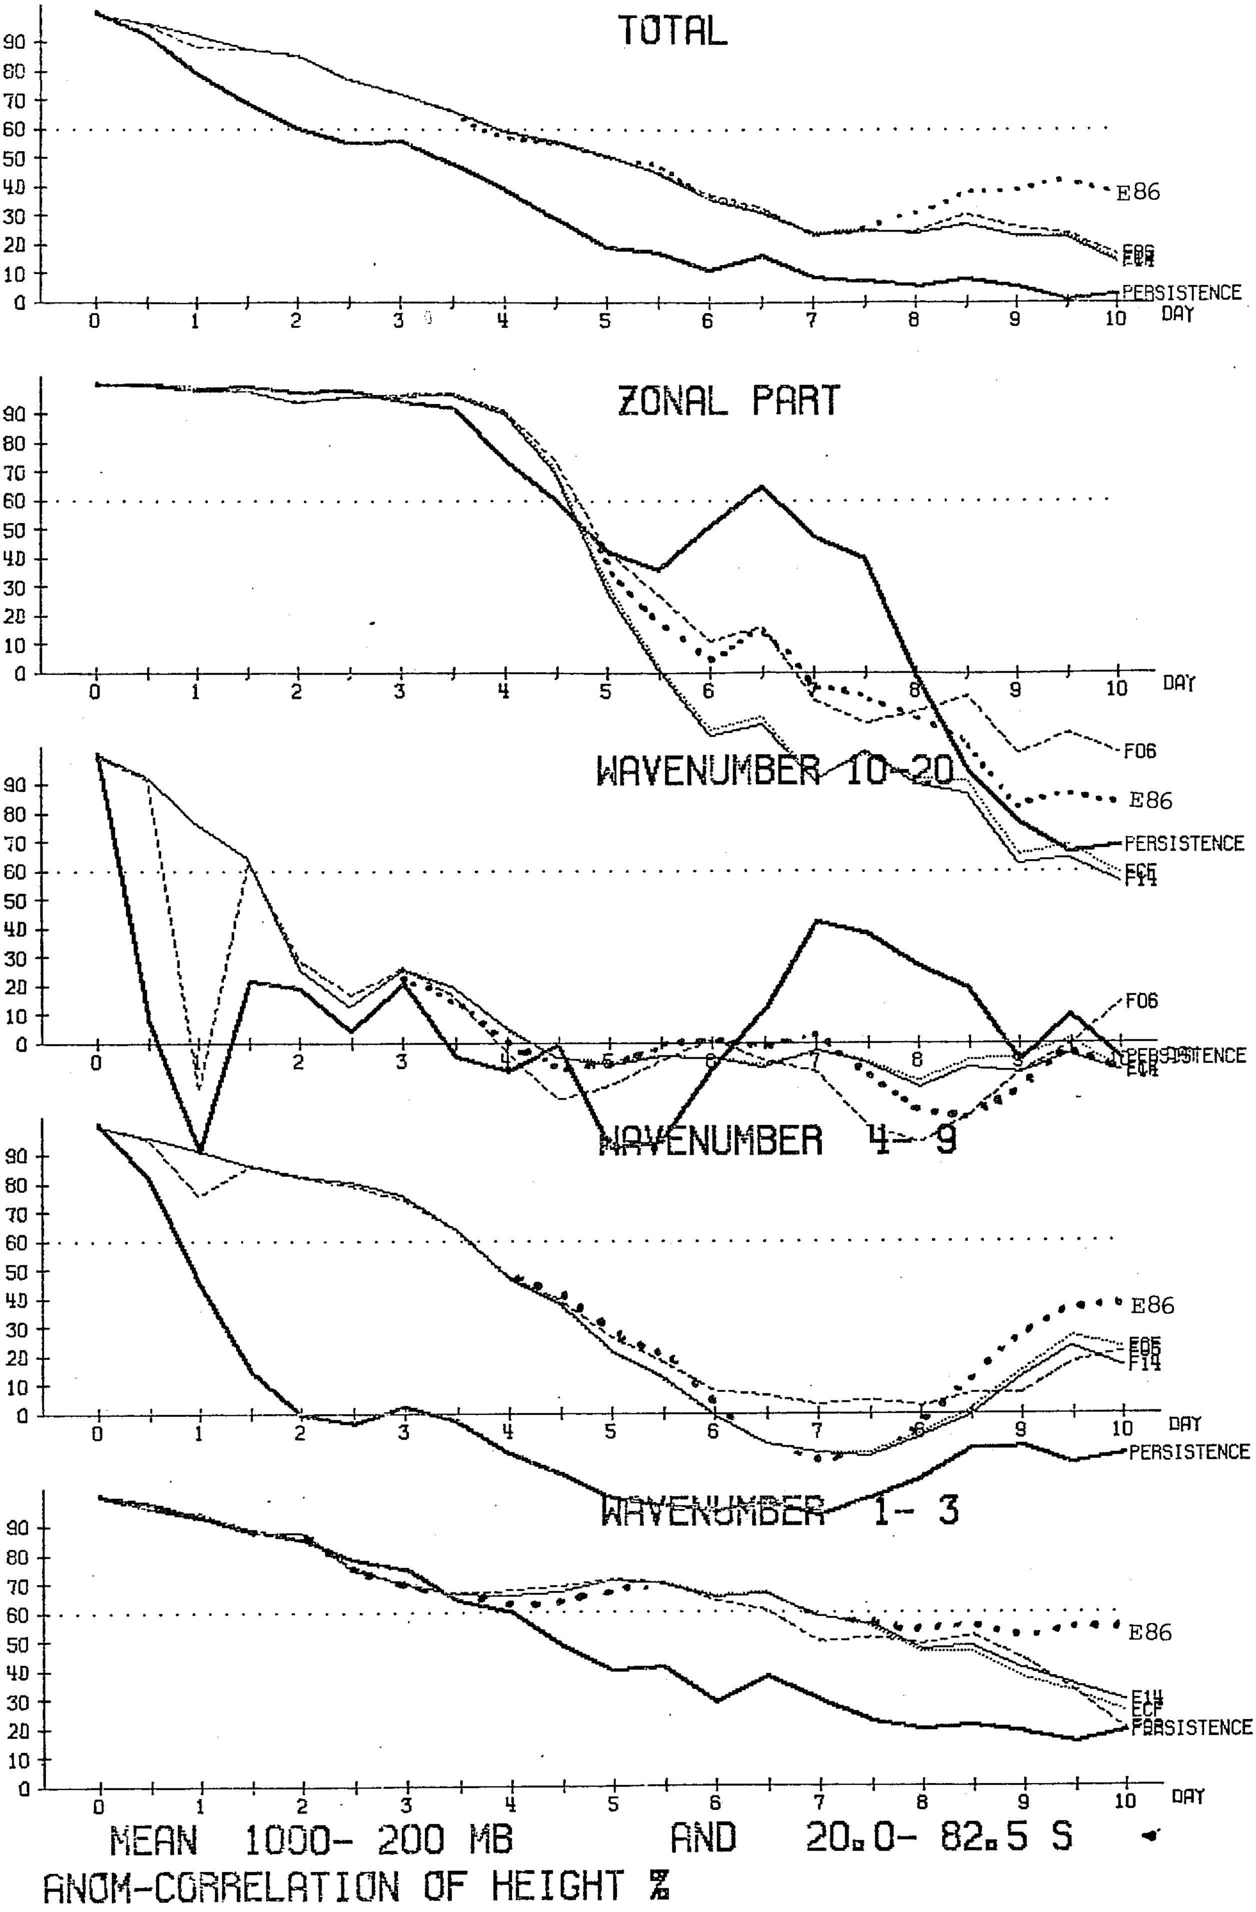
TOTAL
ZONAL PART
WAVENUMBER 10-20
WAVENUMBER 4- 9
WAVENUMBER 1- 3
0
10
20
30
40
50
60
70
80
90
0
1
2
3
4
5
6
7
8
9
10
DAY
E86
F06
F14
ECF
PERSISTENCE
MEAN 1000- 200 MB AND 20.0- 82.5 S
ANOM-CORRELATION OF HEIGHT %

Fig.5b