Forecasting Development

Forecasting Research Technical Report No. 361

The Feasibility of an Ensemble Low Cloud Prediction System

By C. Bryden

August 2001

Met Office , Forecasting Development , Room R321, London Road , Bracknell , Berkshire ,RG12 2SZ,
United Kingdom

00/0486

Forecasting Research
Technical Report No.361

The Feasibility of an Ensemble Low Cloud Prediction System

by

C. Bryden

August 2001

Local Forecasting R&D
Forecasting Development
Room R316
The Met Office
London Road
Bracknell
RG12 2SZ
United Kingdom

Local Forecasting R&D, Forecasting Development, Room R316, The Met Office,
London Road, Bracknell, RG12 2SZ, United Kingdom

The Feasibility of an Ensemble Low Cloud Prediction System

C. Bryden

Abstract

The simulation of the formation and dissipation of boundary layer stratiform cloud is known to be critically dependent on two main factors: the description of physical processes occurring between the cloud and its environment, and the initial and forcing conditions imposed upon the simulation. To investigate the feasibility of addressing these critical uncertainties, the site-specific model system (SSFM) was extended with the latest non-local physical process schemes and an ensemble approach.

It was found that unless changes are made to the forcing methodology, it is not feasible at present to ensemble on the physical processes. However that ensembling on initial and forcing data is feasible, and could add skill to the forecast , deriving from the improvement in understanding of model biases during periods of low cloud cover.

Contents

Page

1. Introduction

The simulation of the formation and dissipation of boundary layer stratiform cloud is known to be critically dependent on two main factors: the description of physical processes occurring between the cloud and its environment, and the initial and forcing conditions imposed upon the simulation. The aim of the Low Cloud Prediction Project (CI056 under the 2000/01 Corporate Investment Programme) was to use an extension to the site-specific model system (SSFM), with the latest non-local physical process schemes and an ensemble approach, to investigate the feasibility of addressing the critical uncertainties in low cloud formation and dissipation. In the longer term, the technique could be used as a corrective input to automated forecasting systems such as OpenRoad and in fog and stratus prediction.

This short report outlines the work that has been undertaken; briefly present the results that have been obtained; and makes suggestions for further work. (More detail is available at http://fr1400/ssfm/LFRD/LCP/LCP.html.)

2. Implementation of MOSES II in the SSFM

2.1. Modelling

At the start of the project, the operational version of the SSFM (SSFM 4.5) incorporated the 5A boundary layer scheme and 7-tile surface exchange scheme (MOSES I), and did not include the Edwards-Slingo radiation scheme.

In the UM, development of the 5A scheme had taken two paths: the 6A scheme was developed by Adrian Lock based on non-local mixing and six boundary layer morphologies; the 7A scheme was based on 5A, but incorporated the 9-tile surface exchange scheme (MOSES II). These two new schemes were brought together in 8A, which exists as a set of modifications on the 7A scheme.

The first task was therefore to incorporate the 8A boundary layer scheme and MOSES II into a test version of the SSFM (SSFM 8A). Parallel work on incorporating the Edwards-Slingo radiation scheme (ES) into the operational version of the SSFM meant that it should also be incorporated into SSFM 8A.

The incorporation of ES was smooth, but there were a number of problems with 8A. The MOSES team had sent a new version of the modifications, which had to be dovetailed with the existing SSFM modsets. Problems in firstly getting the model to run and then producing reasonable results could be broadly categorised as:

- accounting for the different higher-level structures of the SSFM and UM
- mismatches between arguments to the call and declaration of subroutines
- discrepancies in variable declarations between routines
- calculation of tile albedos and net radiation on tiles
- inconsistencies in calculation and use of roughness lengths and blending heights

The implementation of MP Section's version of MOSES raised concerns over the way in which the scheme does the calculations over tiles. In order to economise on processor time, MOSES II ignores tiles that are not in the fetch as the fetch is the grid-box aggregate and constant with respect to wind direction. This becomes a problem in a site-specific model as the fetch is allowed to change. Some variables (eg surface temperature by tile) depend on their lagged value. This problem was solved in the short term by cutting out the economy and performing calculations over all tiles.

Finally, some of the results for MOSES diagnostics are poor. The results for leaf area index and canopy height depend largely on initialisation data, but there may be a problem with canopy capacity and other diagnostics. This was given a low priority but should nevertheless be addressed.

A satisfactory version was eventually produced, giving results reasonably close to those from the operational SSFM. Results are briefly discussed below in Section 2.2.

2.2. Verification and Case Studies

SSFM 8A was first verified subjectively by comparison of selected diagnostics with results from the operational SSFM (pre Edwards-Slingo), and by comparison of cloud fractions with reflectivity observations from the Chilbolton 94 GHz radar.

SSFM 8A was then verified objectively against observations using the LFRD standard Winter 97/98 SSFM trial procedure. This is reported on separately.

For the more subjective verification, the dates of the test runs were chosen by looking through the Chilbolton images for cases of stratocumulus. Two periods were picked out: 7-9 September 2000 and 2-4 October 2000.

Figure 1 shows a comparison of SSFM 8A with SSFM 4.5 for the 00Z run on 2 October 2000. SSFM 8A tends to be warmer, except at around 18Z. This bodes well as the current operational version shows a cold bias at 00Z, 06Z and 12Z, and a warm bias at 18Z during October. SSFM 8A also shows increased humidity between 12Z-24Z. A number of changes were made to the roughness length and blending height code, but it is possible to account for amplified roughness length by the increase in energy in the system, and increased stability during the period of lower temperatures.

Figure 2 shows the plot of cloud fraction against the corresponding image of reflectivity from the Chilbolton 94GHz radar. So subjectively, the SSFM 8A simulation gives quite good results. For a more objective test, the Chilbolton reflectivity data (see Acknowledgment) was converted to ice water content, and this was compared with cloud ice results from SSFM 8A. Owing to the low ice water content of stratocumulus cloud, this analysis proved inconclusive. Nevertheless, the use of Chilbolton radar and lidar imagery in SSFM verification may bear further investigation.

3. Experimenting with Ensembles

The two main factors affecting the simulation of low cloud are: the description of physical processes occurring between the cloud and its environment; and the initial and forcing conditions imposed upon the simulation. The feasibility of an ensemble prediction system was tested on both these factors.

The ensemble system was set up using Fortran namelists. To ensemble on an independent parameter, the parameter in the Fortran was replaced with a namelist variable. To ensemble on a dependent variable, a multiplicative/additive factor was set up as a new variable and again passed via namelists. These namelist variables were adjusted in the calling scripts.

The model was again run for 36 hours beginning at 00Z on 2 October 2000.

3.1. Options in the Non-Local Boundary Layer Mixing Scheme

The science behind the 6A boundary layer scheme is described in Lock et al (2000). It includes a representation of non-local mixing (driven by both surface fluxes and cloud-top processes) in unstable layers, either coupled to or decoupled from the surface, and an explicit entrainment parametrization. The scheme is formulated in moist conserved variables so that it can treat both dry and cloudy layers.

The scheme was introduced into version 4.5 of the UM. Lock et al (2000) presents examples of its performance in single-column model tests, but although the SSFM is based on the SCM, the scheme has not yet been incorporated into the operational version of the SSFM.

Evaluation of the 6A boundary layer scheme seemed to indicate that, once it was incorporated into the test version of the SSFM, there would be the possibility of ensembling on several parameters.

The main criteria for choosing the parameters on which to test the ensemble system were:

- feasible in the model code
- physically meaningful

The latter could just mean that there is uncertainty over the value of the parameter or that there is a likelihood of showing skill, but the main emphasis was on the parametrizations that affect the evolution of low cloud as directly as possible, such as entrainment velocity.

The parametrization of the entrainment velocity is quite complex, and there were a number of possible parameters that could be tested for ensemble skill. Quite substantial variations in these parameters seemed to have little effect on the standard 36-hour simulation, so a blunter instrument was used in the form of a multiplicative factor directly applied to entrainment velocity. Although the factor was varied between 0.5 and 2.0, it had little impact on the simulation, with no obvious variation in cloud clearance or persistence. The model was then run out to 4 October 2000, by initialising each run on the 06Z diagnostic output of the previous run. However, the effects did not intensify appreciably over time.

Further ensembling on parameters used in the diagnosis of boundary layer type also showed little sensitivity in the model.

This insensitivity seems to be the result of the way in which the SSFM is forced using Mesoscale model output (see Dunlop and Clark, 1997). Above the top of the boundary layer, the SSFM is identical to the MES. Just below the boundary layer, SSFM fields relax to MES fields at a timescale of 5 minutes (10 model timesteps). The timescale increases to 1 hour at the surface, so that the MES has little or no effect. Now entrainment occurs at cloud top, which is also defined as the top of the boundary layer. That is, even though changing the entrainment velocity will affect the moisture content and temperature of the air near the cloud top, within 5 minutes the effect will be washed out by the forcing (ie the advection in of air with 'environmental' properties). In order to ensemble on cloud properties, therefore, the relaxation profile would have to be modified such that it extends further into the atmosphere and there is a longer timescale of relaxation at cloud top. Alternatively, the way in which the SSFM is forced could be modified, either by turning off the moisture and temperature forcing, or by using geostrophic forcing.

3.2. Ensembling on Initial and Forcing Data

The SSFM is based upon the concept of driving a 1D model (primarily concerned with the surface and boundary layer) using output from the 3D NWP models, the most appropriate of which is the mesoscale model (MES).

Errors in the MES (ie SSFM forcing) can be resolved into:

- subgrid physics
- error covariance from the previous forecast and derivation of gridded observations

The next section describes the use of random walks to try to capture the step-to-step errors in the subgrid physics, then the combination of random walks with variation in initial conditions to try to capture the error covariance.

The first step, however, was to take a more deterministic approach to gauging the sensitivity to the initial conditions and forcing. In the 2 October 2000 run, the initial and forcing data for the moisture profile in the boundary layer were multiplied by a factor, constant throughout the boundary layer and in time. This was intended to mimic the effects of mixing in air from outside the boundary layer with different moisture content. The simulation was run nine times, including the control, with the factor set at 0.05 intervals between 0.8 and 1.2.

The results for cloud fraction are shown in Figures 2 (control) and 3, and the runs can be grouped in threes – 0.8-0.9, 0.95-1.05 and 1.1-1.2. The main points to note are:

- the first and second groups show increasing cloud persistence with moisture
- the third group shows increasing depth with moisture, but persistence does not increase substantially

The implications of variations in cloud depth/amount and cloud clearance time have implications for short-wave and long-wave fluxes, and hence temperatures. Figure 4 shows that:

- LW and SW fluxes show increased sensitivity to drying and cloud removal as the factor decreases from 1.2 to a threshold of 0.9
- detail in the LW flux is washed out by increasing moisture, ie cloud becomes homogeneous stratiform
- eyeballing the plot of screen temperature, for which observations are available and shown in the Figure, the run with factor 1.1 looks as though it gives the results closest to observations
- when the front comes in at T+20, the variation in surface temperature is more than 1K, which at temperatures closer to zero would have important implications for, for example, road surface forecasts

This last point suggests that looking at more marginal cases would be a useful line of enquiry. It would be more worthwhile pursuing this using the random walk methodology outlined in the next section.

3.3. Random Walks

For the random walk analysis, a multiplicative factor was again applied to the moisture profile, but this factor was allowed to take a random walk into ensemble space. Two experiments were conducted. In each case, the control was taken to be the case with factor held constant at 1.

Experiment 1

The random walk was run for 100 simulations, in each case with an initial factor of 1.

A step was taken in the random walk every model stepcount (30s) over a forecast length of 36 hours, to try to capture the effect of errors in the physics. This timescale approximately corresponds to the middle of the inertial subrange of the turbulent velocity spectra (Kaimal et al, 1973). An interval between steps of 10 minutes was subsequently tested. This is closer to the timescale of buoyancy overturning of the whole boundary layer, so could be more realistic given that the factor was applied to moisture across the whole of the layer. However, this experiment has not yet been pursued.

The step was taken to have a standard normal distribution (mean 0, SD 1), reflecting the Gaussian turbulence, multiplied by an amplitude (set at 0.001). This step was then added to the multiplicative factor. The factor was constrained by an arbitrary lower limit of 0.2, though in practice this was unnecessary as the factor varies between about 0.5 and 1.5. No upper limit was set.

Figure 5 shows the spread of the runs, overlaid with the control case. Observations were available for screen temperature and low cloud fraction (converted from oktas). Climatology was available for screen temperature only.

Considering low cloud fraction, in general the model predicts less cloud than was observed, and the ensembles add little to the forecast. For screen temperature, the runs are fairly close to observations. They show a cold bias during the first 12 hours of the forecast period and a warm bias during the last 12 hours, though the former could be accounted for by the difference in initial conditions.

The climatology plots are of minimum, average and maximum temperatures over 1984-2000 at each hour from midnight on 2 October to midday on 3 October. For the most part, the walks lie within the climatological range, although there are slight deviations at T+11 and T+36. The ensemble spread is too narrow to construct a probabilistic temperature forecast. In this chosen case, uncertainty is relatively small (see also, for example, the agreement among ensembles on zero low cloud fraction between T+18 and T+24). However, other cases may well repay study.

For each variable, snapshots of results were taken at 6-hour intervals to form ensemble spaces. An FFT was applied to map this raw data from a 'spatial' to a frequency domain, and this indicated no evidence of dominance among ensembles.

The ensemble space was then binned into 10 bins of equal size between the minimum and maximum values. The modal bin was picked, and the median/mean of this bin compared with the control. In most cases, the control was in the modal bin. Figure 6 shows the bin distributions, with the blue crosses marking the control on the graph base.

The distributions are roughly normal, with the exception of cloud fraction, which tends to be heavily skewed towards zero or one. The spread in specific humidity tends to increase with forecast time, but the spread in temperatures is more affected by cloud fraction.

Experiment 2

The random walk was run 401 times with initial conditions equally spaced between 0.8 and 1.2, thus trying to capture the effects of the error covariance as well as errors in the sub-grid physics.

The above analysis was repeated with a couple of modifications.

Firstly, the snapshots were adjusted to account for the initial condition. The methodology was fairly crude, based on the assumption that the linearity in the initial condition would hold through the forecast. So a linear regression was calculated for the raw data, and the data rebased to their average by applying the distance from the regression. Figure 7 shows how this works for screen temperature at the T+24 snapshot. The top-left chart shows screen temperature at T+24 for each of the 401 ensemble members (along the x-axis). The variation is accounted for by both the initial condition and the subsequent random walk. Superposed on this are the control temperature at T+24, and the regression. The average temperature is about 282.2K. The top-right chart shows the ensemble members after they have been rebased. This was not appropriate for cloud fraction, for which a strong skew in the distribution makes the regression meaningless.

Secondly, greater ensemble numbers allowed 20 bins to be used.

Figures 8 and 9 show the timeseries of diagnostics and the distribution of the snapshot ensemble spaces. As expected, the spread is wider, with cloud fraction especially sensitive for much of the forecast period. The bin to which the control is assigned varies little when moisture is considered (though note that the control is always above the ensemble average), but the spread of bins for surface and screen temperature is quite wide. Moreover, the distributions of surface and screen temperature are not normal, especially at T+12. The higher temperatures occur in the ensemble members with initial factors between 0.8 and 0.9 - the air becomes unsaturated, cloud amount decreases (see Figure 3) and surface SW radiation increases.

4. Conclusions and Way Forward

4.1. Evaluation of Completed Work

The overall conclusion is that ensembling on initial and forcing data is feasible, and could add skill to the forecast, but that unless changes are made to the forcing methodology, it is not feasible at present to ensemble on cloud parameters.

The improvement in forecast skill derives from the improvement in understanding of model biases during periods of low cloud cover.

Comparison of the ensemble with observations (see Figures 5 and 8) shows a cold bias in screen temperature from T+0 to T+12. (The warm bias from T+20 is of less interest, as there is no Sc.)

As regards specific humidity, Figures 6 and 9 show that the distribution in bins of the 100-member ensemble is broadly Gaussian, whereas the 401-member ensemble is skewed towards the higher bin numbers. The control is higher than ensemble mean at all snapshot times in the 401-member ensemble, and at all-bar-one in the 100-member ensemble, implying that the model is too moist.

That is to say, during periods of Sc cover, the SSFM is too cool and wet. Hence, because it is looking at more specific conditions, this analysis gives added value over the LFRD standard Winter 97/98 SSFM trial.

4.2. Future Priorities

- To investigate the impact of changing the relaxation profile in SSFM forcing (see Section 3.1)
- To investigate how to use this work operationally

4.3. Sensitivity of OpenRoad

Ensembles could also have an important role in road surface temperature forecasts. At present, MORST produces deterministic forecasts of temperature, indicating whether or not frost is expected. However, an indication of the *probability* of frost may be more helpful in deciding whether or not to grit.

Work has therefore started on the production of ensembles for a case where surface temperatures are close to zero. The criteria for choosing the case are:

- stratocumulus cloud at Chilbolton
- Chilbolton forcing data available from the MES
- observations available for an OpenRoad site close to Chilbolton, and showing temperatures within half a degree of zero

The case of 14-23 January 2000 has been chosen. Analysis is ongoing.

4.4. Further Options for Ensembling

As stated in the introduction, the simulation of the formation and dissipation of boundary layer stratiform cloud is known to be critically dependent on two main factors: the description of physical processes occurring between the cloud and its environment; and the initial and forcing conditions imposed upon the simulation.

Ensembling on the initial and forcing data was broadly successful, but developments could be made in both the ensembling methodology and in the analysis of results. There is substantial expertise in the Ensemble Forecasting Research Group, which could be tapped.

However, it was found that ensembling on cloud parameters was infeasible given the present methodology of forcing the SSFM. The further option here, therefore, is to change the forcing methodology and retest the ensembles. Adrian Lock (personal comm.) has been investigating the impact of changing the entrainment rate in the SCM, with forcing specified as 'large-scale advection' warming/moistening increments that don't change if the model column warms or dries. His preliminary results seem to indicate that the SCM is insensitive, because increasing the entrainment rate, and drying out the cloud layer, decreases the drizzle rate and hence gives a similar amount of cloud.

Finally, the ensembles described in Section 3 are on the model *state*, whereas for the OpenRoad work it might also be fruitful to look at the model *timing* (eg of a cold front relative to the diurnal radiation cycle). This would be provided by a full 3-D model approach. However, it might be feasible to doctor the forcing data to achieve this in the SSFM.

5. References

Dunlop, CC, and PA Clark: 1997, 'Forcing the Single Column UM from the Mesoscale Model', Forecasting Research Technical Report No.215, June 1997.

Kaimal, J C: 1973, 'Turbulence spectra, length scales and structure parameters in the stable surface layer', Boundary-Layer Meteorology, 4, 289-309.

Lock, AP, et al: 'A New Boundary Layer Mixing Scheme. Part I: Scheme Description and Single-Column Model Tests', Monthly Weather Review, 3187-99, September 2000.

6. Acknowledgment

We thank the Radiocommunications Research Unit at the Rutherford Appleton Laboratory for providing the 94 GHz Galileo radar data. The Galileo radar was developed for the European Space Agency by Officine Galileo, the Rutherford Appleton Laboratory and the University of Reading, under ESTEC Contract No.10568/NL/NB.

Figure 1: Comparison of SSFM 8A with operational version 4.5

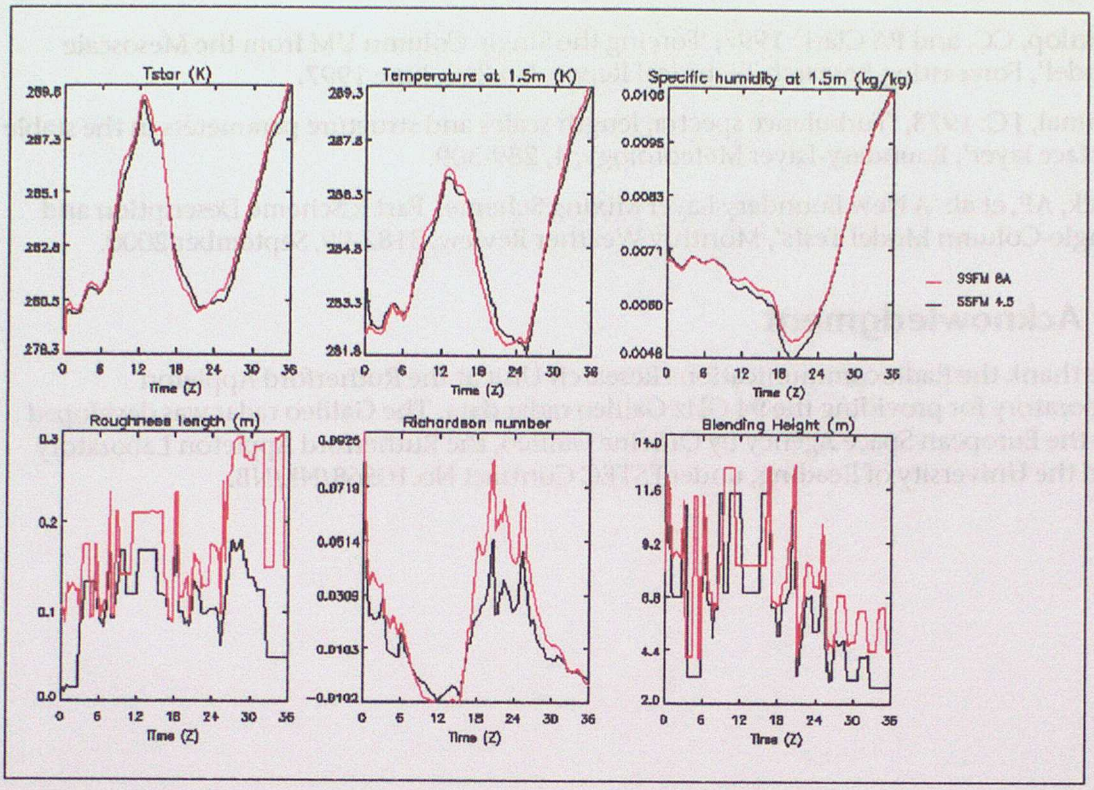

Figure 2: Comparison of SSFM 8A cloud fraction plot with Chilbolton 94 GHz radar image

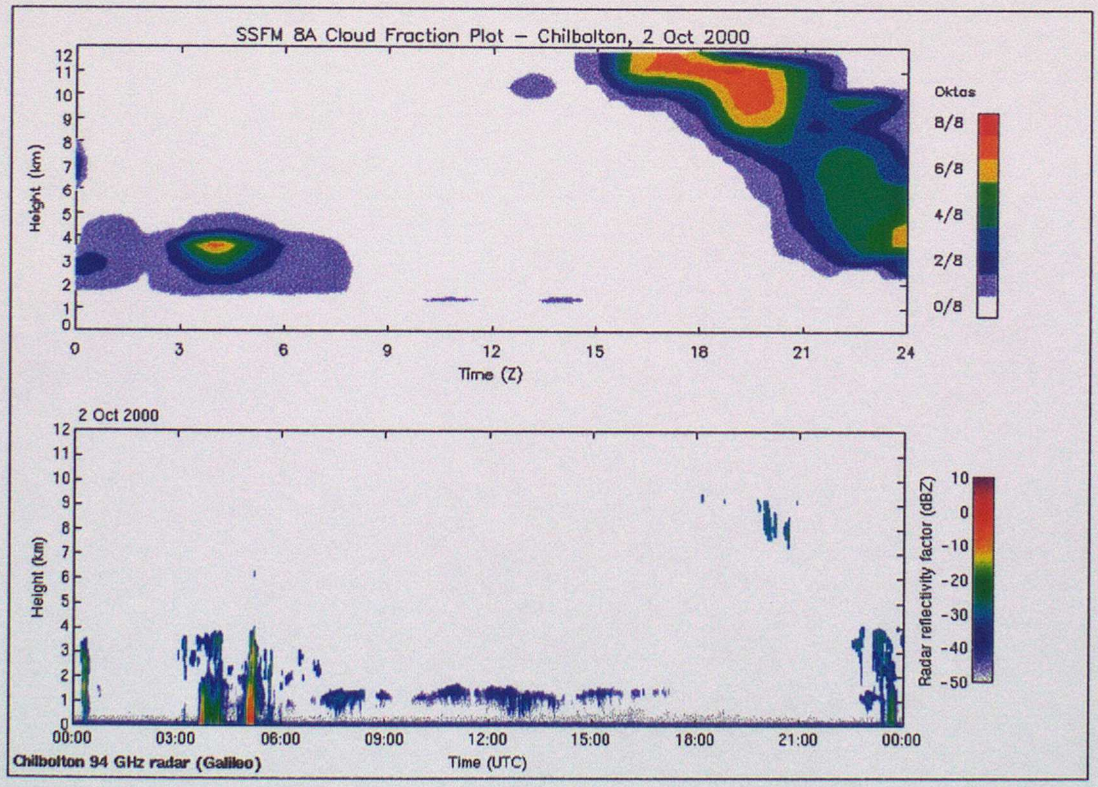

Figure 3: 9-member ensemble, multiplying initial and forcing moisture profile by factors between 80% and 120% – plot of cloud fraction

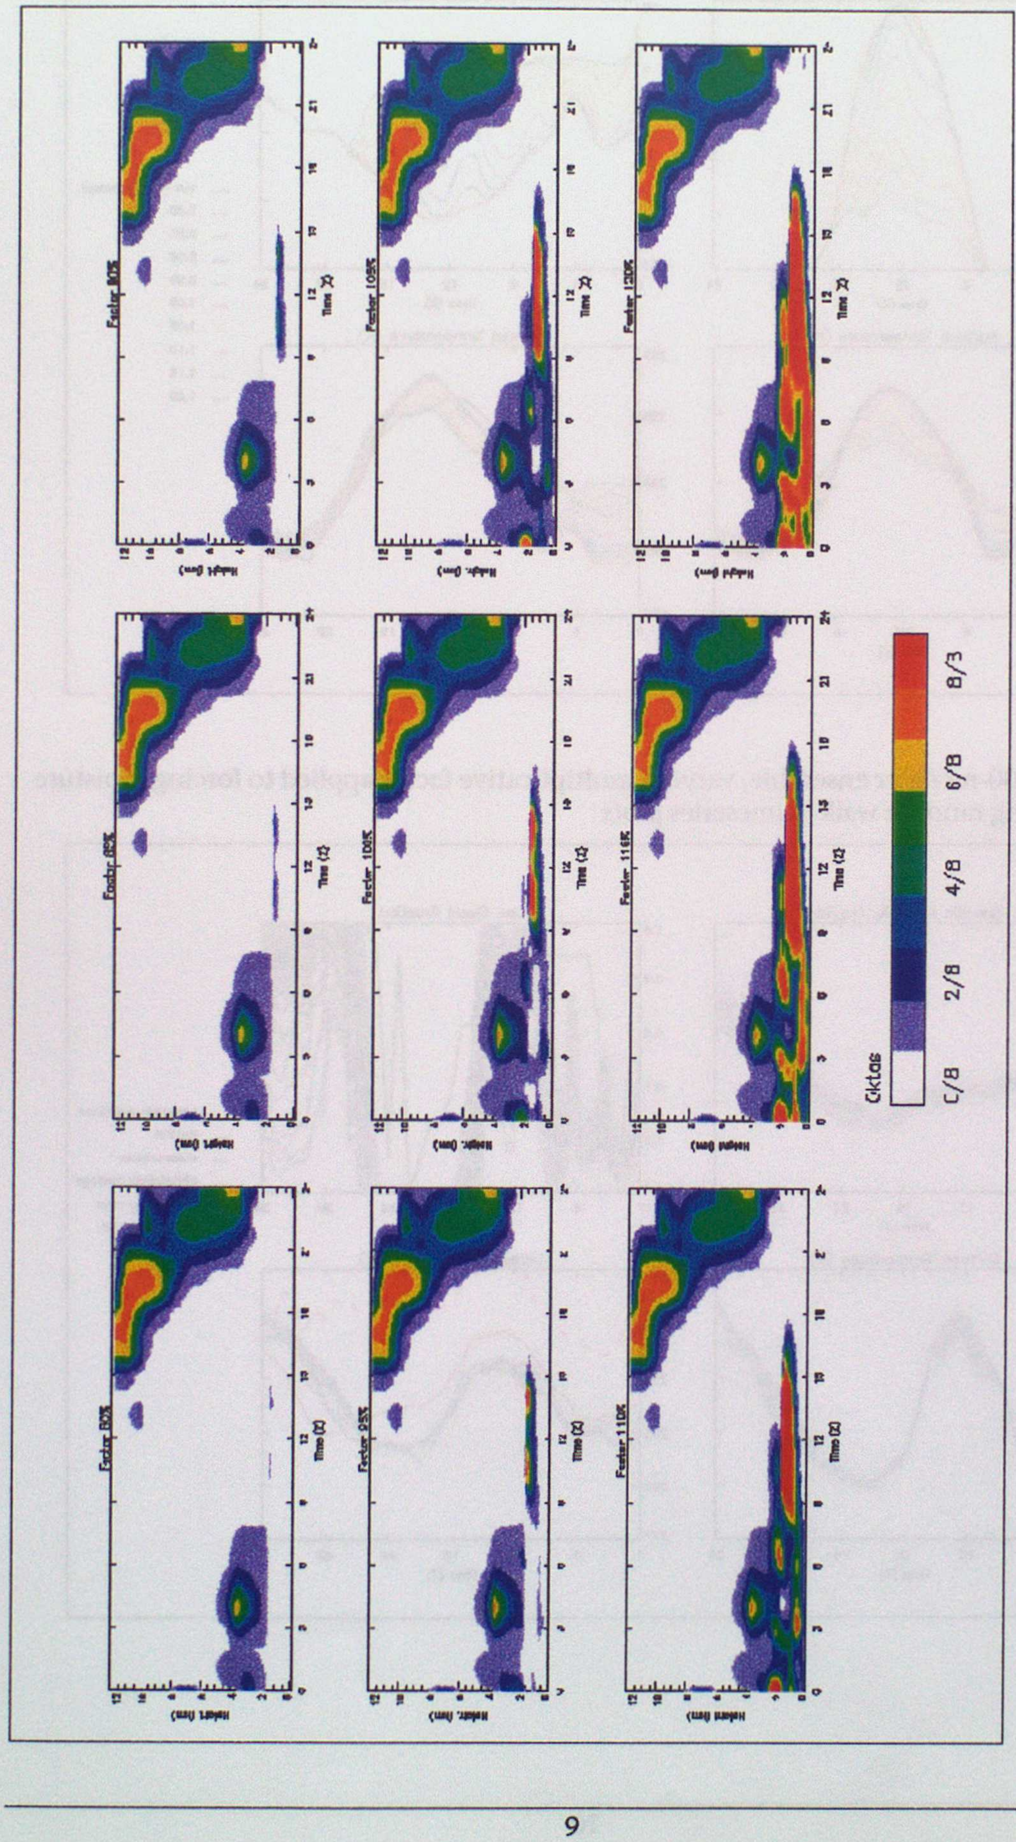

Figure 4: 9-member ensemble on moisture profile – plots of other diagnostics

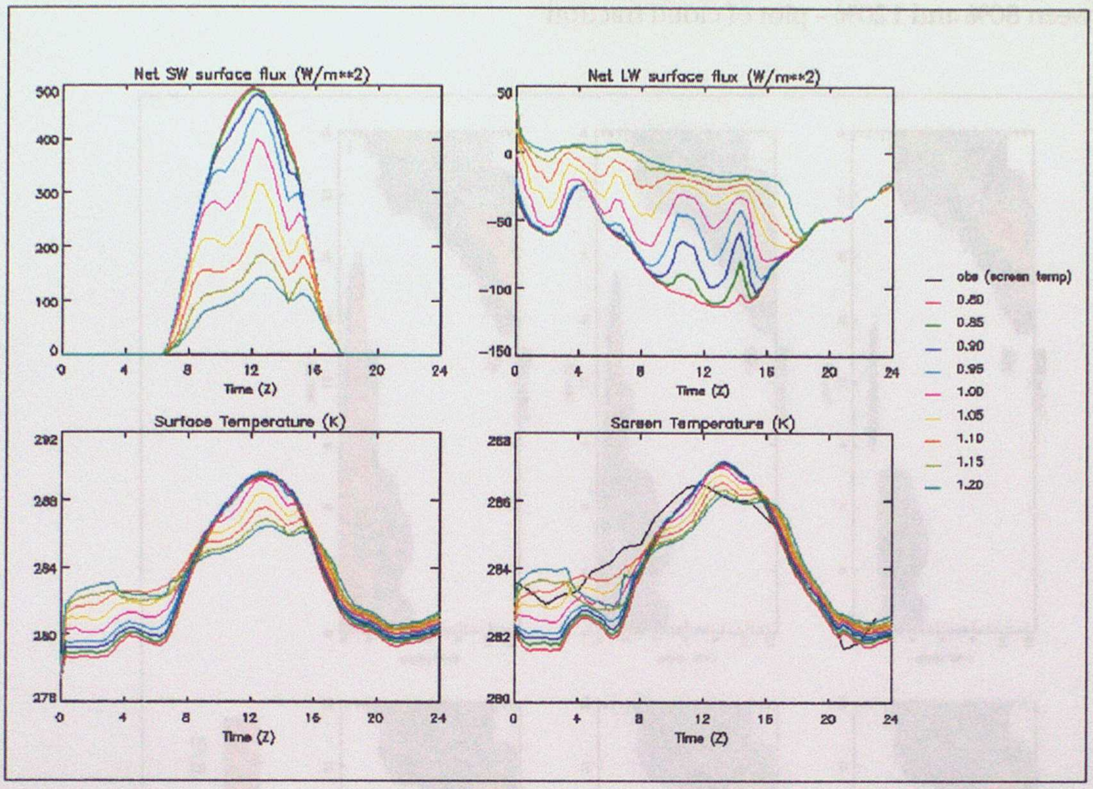

Figure 5: 100-member ensemble, varying multiplicative factor applied to forcing moisture profile using random walk – timeseries plots

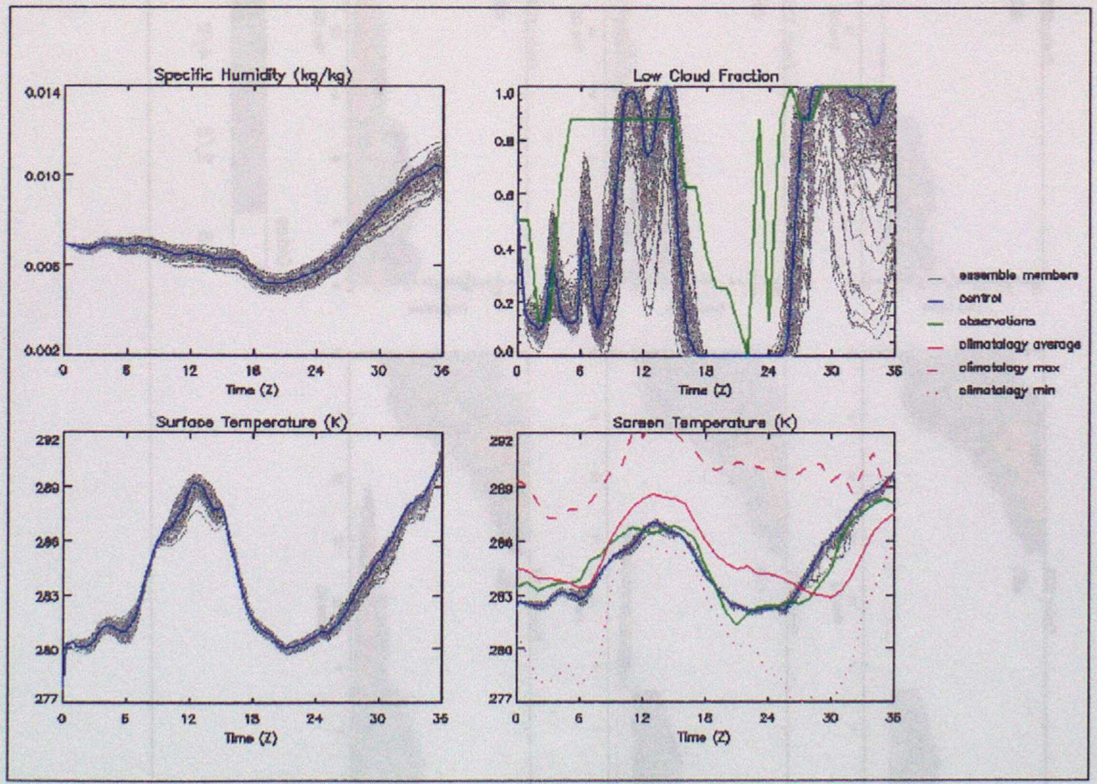

Figure 6: 100-member ensemble on moisture profile, showing 10 probability bins of ensemble space at 6-hour intervals

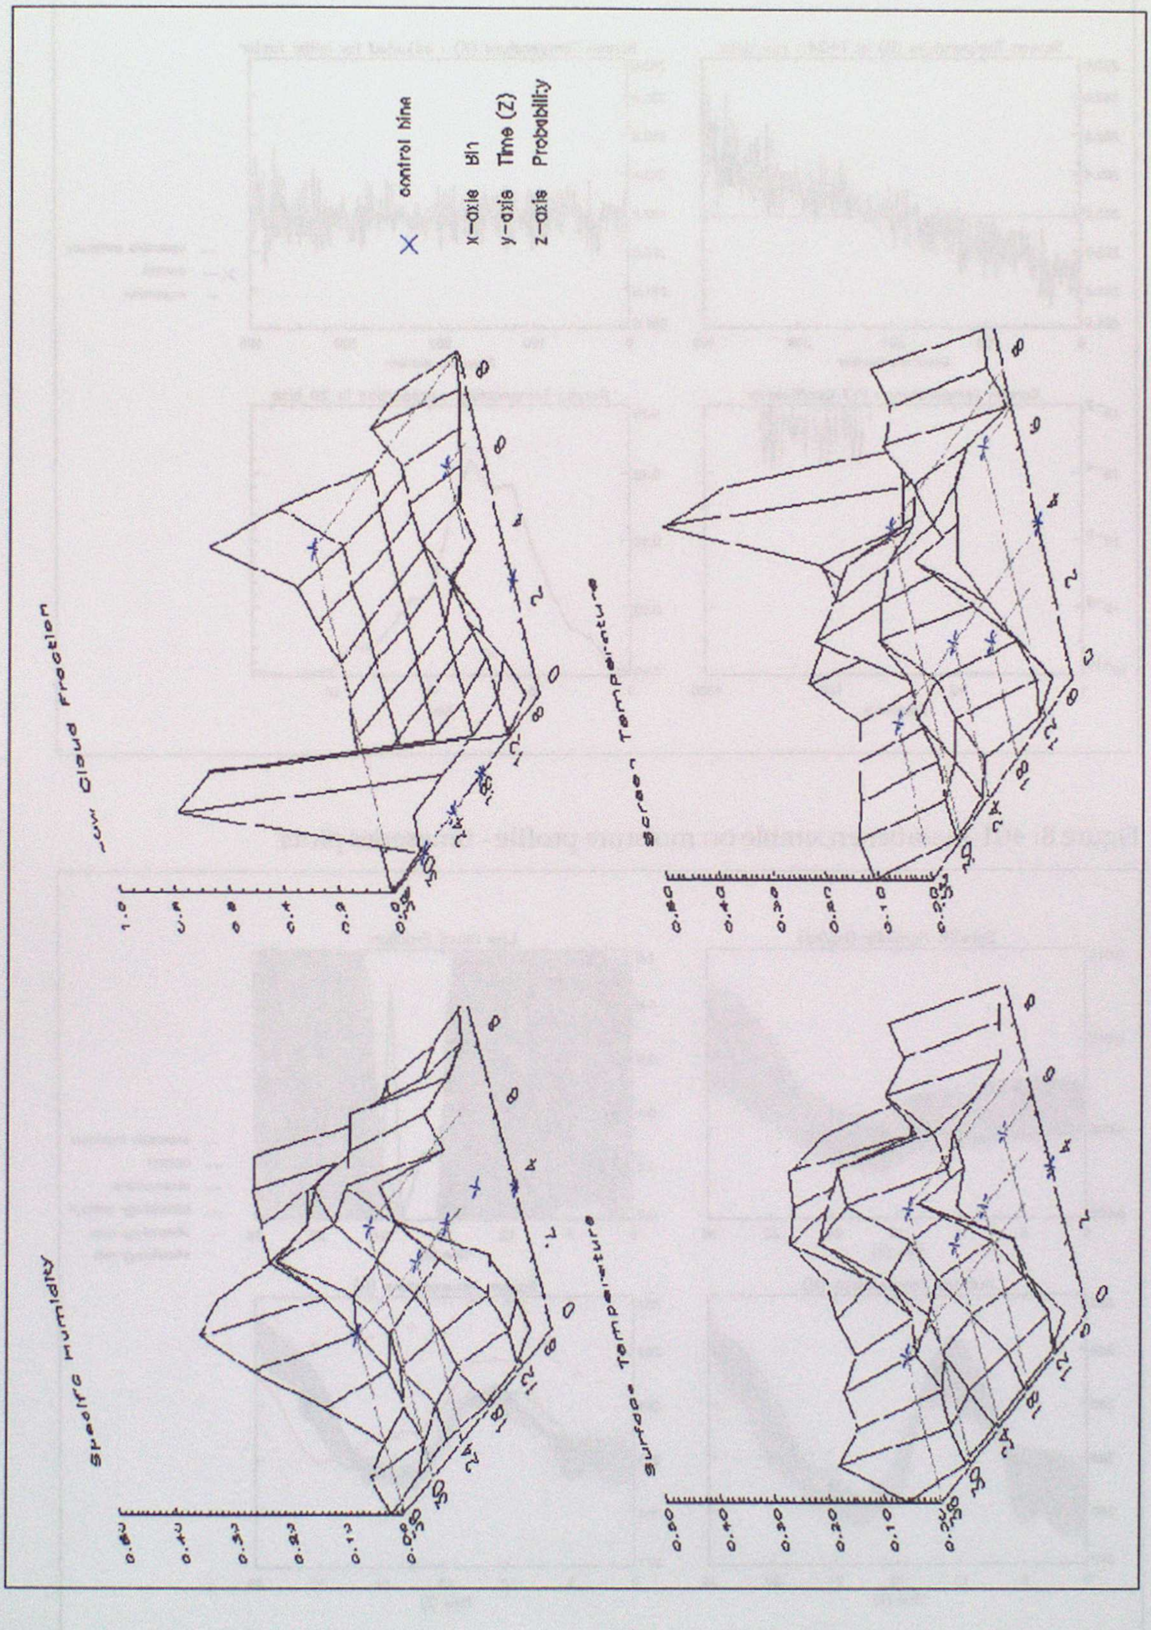

Figure 7: 401-member ensemble on multiplicative factor applied to initial moisture profile, then factor applied to forcing varied using random walk – example of analysis, applied to screen temperature at snapshot time T+24

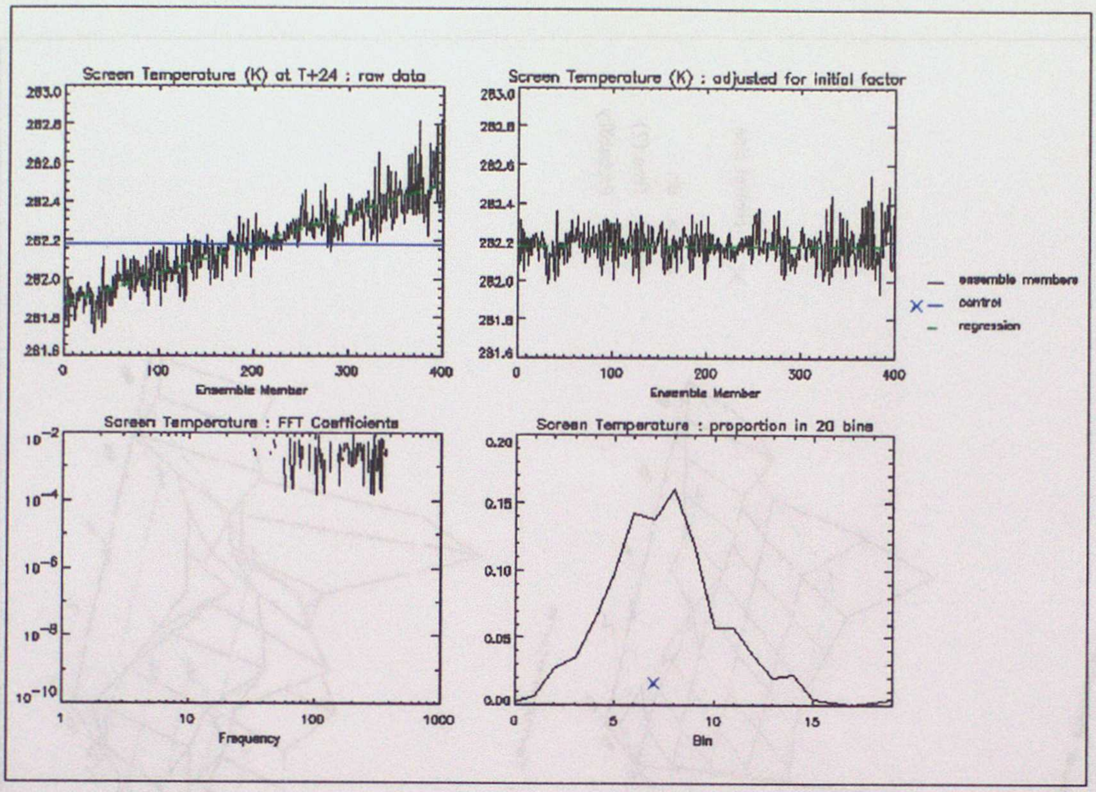

Figure 8: 401-member ensemble on moisture profile - timeseries plots

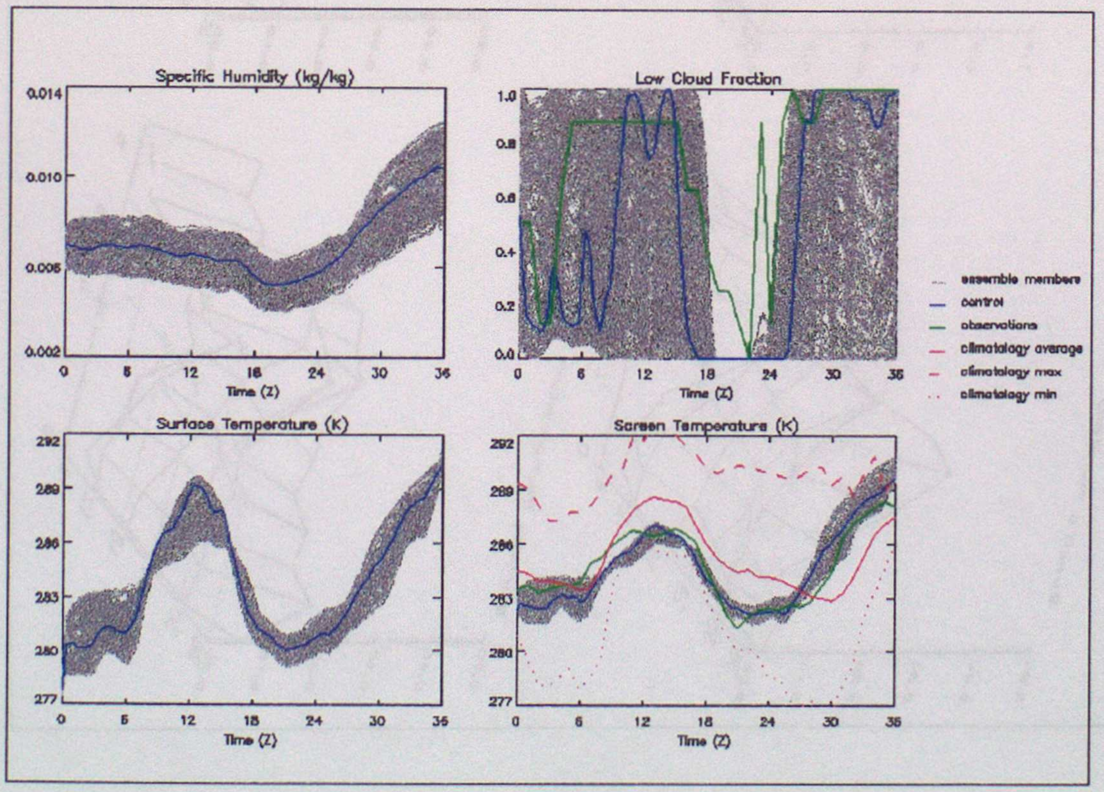

Figure 9: 401-member ensemble on moisture profile, showing 20 probability bins of ensemble space at 6-hour intervals