

Final report on workshop on mathematical techniques for improving forecasting of hazardous weather

University of Reading, 16 to 20 June 2003

(Report updated 13 November 2003)

1. Introduction

The workshop addressed the question of the feasibility of storm-scale NWP, with particular attention to prediction over the UK. The key questions which the workshop addressed were the following:

- i) What has to be achieved to meet the user requirement?
- ii) What level of predictability is scientifically possible, both deterministically and stochastically?
- iii) What can be realistically attempted within the next 10 years using affordable computing facilities, existing data assimilation and modelling techniques, and available observing technology?
- iv) If there are problems that cannot be solved using existing techniques, are there techniques available from other application areas which could be used?
- v) If there are major scientific questions to be answered before techniques can be developed, what sort of research should be undertaken to address them?

The answers to these questions were developed in three rounds of discussion groups. Reports from these, and the ensuing plenary sessions follow. These are lightly edited records of the actual discussions, so they are not very homogeneous or necessarily well-focussed. After these, a summary is given of the main strategic conclusions and the answers to the key questions above.

2. What can be achieved with planned evolutionary development of existing systems?

2.1 Are NWP models, with planned evolutionary developments, up to it?

'It' is model capability given perfect initial conditions for forecasts up to 12 hours ahead of convection and other small-scale phenomena. The formulation of current cloud-resolving models provides an adequate basis for meeting this requirement, since they are non-hydrostatic, use terrain-following coordinates, and can treat a no-slip boundary condition properly. Areas that require improvement include coupling between resolved dynamics and turbulence models. Ideally a cloud-resolving model needs a 20m grid so that convective motions are well-resolved and turbulent processes can be separated from the main convective motions. However, such resolutions will not be attainable for many years. The practical problem is to identify on what scales the predictions are reliable if affordable resolutions are used. Another problem is that the convective motions are intimately linked with boundary-layer motions, which also have to be sufficiently resolved. Also important, but less critical, were representations of cloud microphysics, radiation, aerosols, and surface fluxes.

In plenary discussion, the key problem raised was that of physical processes and coupling on the grid scale (more of an issue than, for example, cloud microphysics) – since there is no scale separation. Paul Mason said that cloud microphysics is important in cloud resolving models (and in climate models) but is it so important at the convective scale?

2.2 Can all expected observation types be used effectively within existing/planned assimilation systems?

Doppler winds could be used. Radars have the best potential. Doppler winds in clear air are needed. Information about boundary layer depth could also be obtained from Doppler radar.

Satellite image sequences (e.g. MSG) could also be used effectively. Microwave radiances needed work before they could be used over land. Infra-red radiances could only be used for observing conditions favourable for initiation (i.e. in clear air). Infra-red sounders would be tested in the boundary layer.

GPS soundings from space did not have enough spatial resolution. From the ground the temperature/humidity ambiguity was a problem, the TPW time series could be used. Tomography was possibly feasible but expensive.

Lidar could give boundary layer winds, but not under clouds. However, it would not be available in the near future.

Surface observations were crucial in measuring low-level humidity. Acquisition and development of more automated stations was needed. In order to use the data good surface physics, initialisation and physiography was required.

In plenary discussion the view was that the observational coverage was probably OK in radar-dense areas. Keith Browning said that we should be observing/diagnosing the depth of the boundary layer (integrated measure of convergence). Alan Gadian emphasised the need to observe surface fluxes and exchanges. There was a need for research on conceptual models of boundary layer growth etc. (how well is the boundary layer represented in DA?)

2.3 Are existing/planned assimilation systems up to it?

The ‘task’ was considered to be forecasts up to 12 hours ahead of convective weather over the UK. Most current experience in data assimilation is on synoptic scales. Existing systems have analysis resolutions of 40km (Hirlam), 9-10km (Aladin) and 12km (UKMO). Most current systems have broad structure functions, though Aladin uses special structure functions.

Existing assimilation systems should be able to provide the information on the large-scale forcing. There is room for improvement in current performance on the 100km scale, and the need to exploit more data (e.g. MSG).

On convective scales there are no simple balance relationships. Therefore schemes that use the model such as 4dVAR and EnKF may be better because the model can

take the place of a balance relationship. Experiments with cloud-resolving models may help to derive suitable balance relations if they exist. Structure functions probably need to be flow-dependent and orographically varying (e.g. wavelets, recursive filters, 4dVAR and the EnKF). Choices of control variable should incorporate knowledge of boundary layer structure and surface information. They should also reflect any balance relations that can be derived from storm structures.

Spin-up problems are clearly a major issue.

The cost of incremental 4dVAR is only 3 times that of 3dVAR, so it should not be ruled out.

Convection is clearly not Gaussian, so an assimilation strategy based on this may not be appropriate. The outer iteration loop in 4dVAR may become more important. It is difficult to initiate convection, when none is present in the first guess. Ideally a 'nearness to convection' term should be included in the cost function. Nonlinear observation operators are required. It is not clear whether the incremental 4dVAR approach will be adequate or whether 'full' 4dVAR will be needed.

If explicit probability forecasts are needed, the EnKF approach will automatically give an ensemble. This could be combined with 3dVAR. The ensemble size and strategy was an issue.

In plenary discussion the points made were: Is nudging dead? Is vertical velocity important? Can we use ellipticity conditions to diagnose the breakdown of balance?

2.4 Are ensemble generation techniques available?

Present: George Craig, Martin Ehrendorfer (chair), Brian Golding, Gordon Inverarity, Sebastian Reich, Bill Skamarock, Edriss Titi.

Regarding an ensemble as being a sample from an unknown probability density function (pdf) of the weather, the need to capture important characteristics of this evolving pdf was emphasised. At its most general, the problem consists of identifying which regions of the initial probability phase space map into regions of the final space corresponding to specific outcomes. More specifically, can ensembles determine where convection develops and where it is absent? Is convection limited to particular locations and times? Can possible arrangements of convection cells be classified in different clusters? Can the maximum intensity or severity of convection be identified? How many members are required to reliably estimate the mean and covariance of the pdf?

It was considered essential that ensemble members should incorporate perturbations to the initial conditions, model equations (e.g., by adding stochastic noise to the model physics), lower boundary conditions and, in the case of limited-area models, lateral boundary conditions.

The following approaches for generating ensembles were considered: singular vectors, bred modes, Monte Carlo (white noise superposition), lagged forecasts and lagged initialisation (take several analyses valid at slightly different times then

consider them all to be valid at a different common time). Particular difficulties were then discussed. For example, singular vectors optimised over several hours might be inappropriate in situations where rapid growth subsequently saturates due to nonlinear effects. Even so, singular vectors were deemed worthy of further study on small scales and over short timescales when growth is exponential. Next, if the initial ensemble is not sufficiently close to the truth, the evolved ensemble may not include important real events. Finally, it is important to examine the sensitivity of the ensemble to model errors.

In the group's opinion, ensembles should contain on the order of 100 members and, at the very least, boundary conditions should be perturbed when generating an ensemble. It was also recommended that singular vectors and bred modes be studied further, particularly for spatial scales of around 1km.

In plenary discussion, Paul Mason pointed out the need to focus on the role of stochastic sub-grid processes in ensemble generation

Eddriss Titi asked where are we going to put the stochasticity? – Paul Mason said it would be in the stresses and fluxes (based on physical considerations).

3. What new research should be undertaken to address issues which cannot be addressed by the evolution of current systems

3.1 Developments in high resolution modelling techniques.

The group considered new research directions for non-hydrostatic models where balance constraints were not important. So resolutions were 50m to 1km.

Areas of interest for the resolved dynamics were the representation of coherent structures, the use of Hamiltonian and symplectic methods, and total derivative techniques. The choice of variables and the method of averaging the equations were considered important. Adaptive grids were not considered a priority for UK forecasting.

Reference nonlinear solutions were required for model validation. These could use simplified forcing and physics. Transient solutions were needed. The assessment could be using gross measures like spectra and rates of convergence to the exact solution.

In sub-grid modelling, key areas were modelling of microphysics- how complicated does it have to be? Should they be Lagrangian particle based methods? Are prognostic schemes required or are diagnostic schemes sufficient? Can the idea of 'super-parametrisation' (i.e. embedded models of much higher resolution) be used? It was necessary to quantify uncertainty. It was not clear how knowledge of storm structures, or of convective boundary structure suitable for storm development should be used. Surface and sub-surface modelling was also important, especially identifying the aspects important for triggering convection.

Turbulence modelling was another key area, particularly the coupling back to larger scales and the representation of diffusive processes. In this context the methods of

operator splitting needed to combine the resolved and parameterised terms were important.

In plenary discussion, Paul Mason raised the question of diffusion along terrain following coordinate surfaces, which is unrealistic. A good test was a statically stable atmosphere at rest over orography in a viscous fluid. Both coordinates and numerical methods needed to be appropriately chosen. Others stated that the use of constant height surfaces gave unsatisfactory results.

There was discussion about the importance of microphysics. Predicting initiation correctly does not require microphysics, predicting subsequent development does.

3.2 Dynamics and numerical simulation of multiphase flows.

Multiphase methods are more properly called multifluid methods, i.e. different parts of the fluid may have separate equations, whether phase changes are important or not. Two cases were discussed. The first was a prognostic multiphase model designed for use at resolution around 5km, where convection was only partly resolved and a prognostic multifluid formulation could replace the mass flux parametrisations used at lower resolutions. This leads to a system of conservation laws with source terms that could be analysed and solved using well-known techniques for hyperbolic conservation laws.

Mitch Moncrief mentioned previous work on intermediate models. George Craig was concerned that instabilities would be fudged. Brian Golding pointed out that closure was the main issue, which still had to be solved independently. Paul Mason thought that the method had potential advantages over normal schemes which scrambled all the fluids together.

The second case was for use in cloud-resolving models with resolution $<1\text{km}$. The problem then is representing turbulent entrainment into/out of updraughts/downdraughts. A special LES with about 20m resolution should be used to calibrate schemes for entrainment/detrainment into convective updraughts and downdraughts.

Mitch Moncrief quoted slugging in pipes as a successful application of multiphase techniques. George Craig noted that anisotropy was required in entrainment schemes for a 1km model.

3.3 New inverse problem techniques.

Present: Dan Cornford (chair), Amos Lawless, Sue Ballard, Sarah Dance, Per Uden, Fathalla Rihan..

The group had no inverse problems experts, so discussed new data assimilation techniques for storm-scale prediction. The main areas of research noted under observation errors and characteristics were in non-Gaussian and correlated errors, intelligent pre-processing, deal with scale disparities, and characterising the errors of radiances and radar reflectivities. The group also discussed how to represent model error, particularly model bias.

Under algorithms for assimilation, convergence issues were discussed, though it was noted that existing methods are never converged. This might depend on the type of observation, and whether the observation operator was linear or nonlinear. The suitability of the incremental method at high resolution was discussed. The optimum time-windows had to be established for different resolutions. It was not clear whether 4dVAR could capture all the necessary timescales with a very short assimilation window. Conversely, if a long assimilation window is used, 4dVAR may not be able to capture short timescale motions. Longer timescales would have to be linked to the treatment of the boundary conditions. All the algorithms would have to be parallelised. The issue of balance was discussed, is it appropriate and is the definition state-dependent? What control variables should be used? Flow dependent structure functions are needed at the top of the boundary layer and the tropopause?

More strategically, there was an issue of whether data assimilation schemes could 'change regimes', e.g. create convection where none existed. In principle, 4dVAR with the full equations can do this. There was also an issue whether displacement errors were more useful than the standard forms of error.

In plenary discussion, Browning raised the issue of failing to detect key features as more important than misidentification. How is tropospheric height and lightning centre activity used by ECMWF? There was a discussion about geostrophy as a form of filter in 4dVAR. What sort of filters are appropriate at storm scales? Experience shows that these filters should be used in correlation models, not applied to increments. However, on convective scales we only have subjective feelings, not quantitative relationships. It was suggested that there might be relationships involving the history of the system, rather than the instantaneous values. Can 4dVAR exploit these through the model constraint?

3.4 Stochastic modelling techniques, are there alternatives to ensembles?

Present: George Craig, Martin Ehrendorfer, Brian Golding, Gordon Inverarity, Terry Lyons (chair), John Norbury, Bill Skamarock.

Some fundamental theoretical results were clarified during this discussion.

The importance of having an appropriate probability measure for the problem of interest was emphasised. It was noted that an ensemble, which itself constitutes a probability measure, is to be preferred over a probability density in systems with many dimensions. For example, on encountering a dubious word, a spell-checker presents an ensemble of alternative words rather than giving the user an entire dictionary and set of associated probabilities for each word contained therein. Furthermore, the inadequacies of using only the mean and covariance to represent the underlying distribution were noted.

This discussion combined with another on the nature of Ito stochastic calculus, in which the addition of noise can be thought of as an infinitesimal version of the process:

The system evolves according to a deterministic differential equation (=vector field), but is subject to random perturbations. These noisy perturbations are also modelled by the differential equation. In each unit of time, a random amount of a perturbative vector field is superimposed on the primary field modifying the dynamics. The perturbation is different every unit of time for a (random and possibly negative) unit of time.

Combining the two discussions, it was then noted that the order in which the noise is added is important in such systems and that covariances alone cannot describe this effect. Mathematically, the new vector fields associated to the noisy perturbation can be used to constrain the probabilistic evolution in phase space, and their choice should be determined from physical, rather than mathematical reasoning. For example, the noise could represent the small-scale convective instabilities generated on the lower boundary or uncertainties in the microphysics parametrizations. For limited-area models, it is also important that the noise on the lateral boundaries is appropriate. The perturbed vector fields probably need to conserve basic quantities.

The different ways in which stochasticity is viewed in physics and mathematical finance were also explored, noting that the random steps are regarded as occurring concurrently in the former (leading to a stratonovich integral). In the financial setting one might well approximate by discrete steps with the integrand (investment strategy) being determined using knowledge up to but not including the current stock movements. This leads to the Ito integral.

Having established the theoretical preliminaries, the view taken was that ensembles do indeed represent an effective stochastic technique for representing uncertainty. Recognising that the perfect model assumption is inadequate and noting that the addition of noise causes solution trajectories to drift away from their "perfect" counterparts, the group's recommendation is that the ways of introducing this noise should be investigated. In particular, it might be hoped that, by combining these ideas with filtering ideas, one could better inform deterministic scenario methods by better choosing ensembles of initial conditions for deterministic models conditional on the models.

In plenary discussion, the points were made that there seems no reason to separate random perturbations that do not change the Hamiltonian from those that do, so that energy producing noise should be added along the path. Terry Lyons responded that the important effect of the noise was the resulting drift, even though the magnitude of the drift was much less than that of the noise. Paul Mason asked whether the noise was supposed to represent model errors or missing physics. The theory of backscatter was only well known for 3d isotropic turbulence. Research was required to characterise it in other cases, and give it a physical basis. It should be noted that ensemble-based forecasting does not preclude approaches such as those based on mixture distributions.

4. What are the key strategies that should be followed to deliver storm-scale NWP, and when will it be practicable.

4.1 What are the predictability limitations?

The group felt that in some sense, this covered the whole subject matter of the workshop, so they de-scoped it to "what are the limits of predictability?"

There appear to be two options for the output of a NWP model: to predict there will be a thunderstorm at point X at time T (deterministic), or to predict that there is likely to be thunderstorms at X,T (probabilistic).

Looking at the relationship between the time taken to produce a forecast and the predictability of the features of interest, we note that a storm cell (predictability <1hr) is not forecastable, assuming that a forecast takes one hour to produce. A large thunderstorm (predictability ~2hr) is marginally forecastable, while a MCC/MCS (predictability ~4hrs) is forecastable.

Major issues identified were: (1) Observations - this was not pursued, (2) Analysis/Initial conditions (3) Forecast. An important research topic was to identify the relative contributions of (2),(3) for these scales - which would presumably be phenomenon dependent.

Three types of systems were identified: (1) Synoptically forced, (2a) Strongly surface forced (2b) Weakly surface forced. Keith Browning suggested that the latter did not produce hazardous weather, but in response the example of the Great Plains of the USA was offered as an example of weak surface forcing that often produced very damaging weather. However, Andrew Crook said (post-meeting comment) that most examples of damaging convective storms over the USA are associated with some kind of upper level or lower level forcing. He had yet to see an example of a severe storm which was demonstrably not associated with some kind of forcing.

A key issue is verification - in order to define when a forecast is better than its predecessor. Should this be based on "what sells best"? Is there an objective measure? Recall that if a feature is displaced by half a wavelength from its correct position, but is otherwise perfectly forecast, the value may be substantial, while the threat score and many other scores are the worst possible.

In plenary discussion, George Craig suggested tuning the forecast to specific types of forcing - e.g. using a coarser scale ensemble for synoptic forced situations, and a high resolution model for internal forcing.

It was noted that the ECMWF system assumed doubling times of 1-2 days, giving climate variance by 2 weeks. Does this sort of calculation make sense at small scale - a research topic is to identify the average doubling times at this scale, and the contributions to this of different systems. To do this would require separating large and small scales. Another contributor pointed out that we know the doubling time for convection is short - ~10-15mins, but that it also saturates very quickly. It was suggested that this could be investigated using a CRM, and it noted that Andrew

Crook's forecasts showed a q perturbation giving yes/no results for the existence of convection after 2 hours integration.

4.2 What is the best way of producing a probability forecast?

Present: Peter Clark, George Craig (chair), Martin Ehrendorfer, Gordon Inverarity, Terry Lyons, John Norbury, Glenn Shutts.

It was agreed that all forecasts are probabilistic even if they are not presented as such. The nature of a probability forecast itself depends on the statistical question being asked (i.e., are we interested in marginal or conditional distributions etc.), which in turn depends on the users and their risk profiles. Furthermore, many users require thresholds at which action will be taken, in which case simply providing mean and covariance may not be appropriate. Ultimately these considerations can all be addressed from the underlying probability distribution, whose estimation was considered to be the goal of the discussion.

The two elements of producing a probability forecast focussed on were those of generating an ensemble and model output statistics (MOS) postprocessing (essentially producing the climatology conditioned on the observations). In the former, it is important to characterise the sources of uncertainty in the whole system. A large ensemble is not necessarily required to address simple questions not involving spatial and temporal correlation (e.g., the temperature at a particular time and place), assuming that it is generated from a truly random set of perturbations. However, the so-called poor man's ensemble, in which the ensemble is generated from different forecast models, was not considered to be well-founded. Given the degree of cross-fertilisation of ideas between models, the resulting perturbations cannot be considered to be sufficiently random. How to choose the optimum combination of ensemble size, physics and resolution was also considered to be an important question. Although computationally cheaper than using an ensemble, the MOS approach was not considered to be well suited for forecasting extreme events.

Although not strictly concerned with producing an individual probability forecast, verification is an important element in the development of a forecasting framework. This introduces the additional requirement to model the observation process so that predicted model states can be compared with actual observations but adds another element of uncertainty to the process. It is also important to determine the skill (the improvement over forecasting no change in the weather) and accuracy of a forecast.

Overall, ensembles emerged as the preferred approach. Further work is required to determine how best to ensure suitably random perturbations when generating an ensemble and also in finding the optimum combination of ensemble size, physics and resolution. Additional studies of forecast verification would also be beneficial.

In plenary discussion, the group stated that mode and variance were not useful outputs - rather each user has a threshold at which they will act. However, it was noted that for a Gaussian distribution, the mode and variance were sufficient to define other probability points. If the distribution is not Gaussian, then not only are mode/variance

insufficient, but also a large set of probability points is required to define the required information.

The group stated that it was important to establish how to tune an ensemble system - size of ensemble, physics, resolution etc... - it was commented that the ensemble does not have to be large if we are only interested in a low-dimensional output subspace, and the random inputs are carefully selected - defining the means of selecting these perturbations at storm scale is an important research topic.

MOS was offered as an alternative to ensembles, but it was noted that dynamical feedbacks were not included.

It was suggested that ensembles are only really needed if the output distribution is non-Gaussian. For Gaussian situations, there may be simpler ways of modelling the output distribution.

The fact that 4Dvar works, suggests that up to its window length, the linear/Gaussian statistics assumption is reasonably accurate.

4.3 Is it possible to do storm-scale NWP fast enough to be useful?

The group made the assumptions that to achieve adequate speed, the process would be fully automatic, including delivery to end users. They also assumed that the output would focus on heavy rain (e.g. >15mm/3hr).

A key issue was predictability. If prediction through 3 system lifecycles was achievable, that would give an upper limit of 3 hours for an isolated thunderstorm and ~24hr for an organised convective system.

In order to resolve the thunderstorm scale; 1km scale up/down-draughts needed to be resolved, which would require a 250m grid. However, initiation of such an integration could probably be achieved with data at 1km scale. Given the lack of analytical balance at these scales, 3dVAR would probably not be usable as an initialisation tool - a DA technique involving use of the full model would be required, leading to the need to use either 4dVAR, EnKF, or a hybrid. Given the fast error growth, a time window of 15min was suggested as being appropriate for either technique.

The current 12km integration of the UM takes roughly 1 minute per hour of forecast. This integration speed has to be maintained to keep up with the data flow. The cost of a 1km model is estimated as x2 vertical resolution, x12³ horizontal resolution, x1/3 area = x1150. 4dVAR takes x3 the cost of 3dVAR per cycle, x20 iterations = x60. Alternatively, EnKF takes ~50 members. Therefore either 4dVAR or EnKF requires x60000, which will not be achievable for a very long time. However, the real essential requirement is that the data assimilation can keep up with real time. This means that 1 hour of assimilation has to be completed in 1 hour of real time (not 1 minute). Allowing some time for the forecast step, this will require a speedup of about x1200.

For the MCS scale we are looking at 100km structures. It is not obvious what initialisation resolution is required, but say it is 3km. Thus the solution given above,

with incremental DA at 3km applied to a 1km model might enable us to carry out a useful 1km forecast. The cost would be reduced by a factor of 27, given requirements of x2000 and x50 respectively. The latter speedup is probably feasible within the next 5 years.

The key issue is that the data assimilation process should keep up with the real time data acquisition.

In plenary discussion, Keith Browning asked whether balance between the 4D fields of precipitation and velocity could help. The workshop was not confident that one could use correlations at this scale to reduce the dimensionality of the problem, e.g. by defining a new control variable. The feeling was that it would be necessary to incorporate this information via the full model. Dan Cornford suggested that the filtering approach would be more efficient, as it involved continuous assimilation, rather than in windows, and retained the covariance structure at the end of each assimilation window, leading to a preference for EnKF or a similar approach. Andrew Crook suggested that 3km data assimilation and a 1km forecast should provide valuable information on individual large thunderstorms not just organised convective systems.

4.4 What new types of observations would be of most benefit, beyond those we expect to have?

Those discussed by the group are listed, with comments.

Boundary layer depth could be inferred from either radar refractive index observations, or satellite observations of the cloud top height of small cumulus. Such data would be difficult to use in DA.

The existence of, but not the magnitude of boundary layer convergence could be inferred from Doppler radar returns from insects, or from satellite observations of cloud lines. Such data would also be difficult to use in DA.

Near surface humidity could be obtained using the Fabry technique for refractive index. However, the information will be at a variable distance from the ground. This could be assimilated in existing DA schemes.

Upper tropospheric PV structures could be inferred from water vapour imagery. Assimilation would be via 4dVAR, or using a conceptual model to deduce dynamical information before assimilation.

Precipitation type and microphysical variables could be inferred from polarisation properties measured by polarisation radar.

Path integrated rain rate could be inferred from polarisation radar, but would be difficult to assimilate.

3D airflow could be inferred from radar reflectivity and assimilated using 4dVAR.

Convective Inhibition can be deduced from temperature and humidity profiles that could be obtained from targetted regions using unmanned aircraft, either in situ, or using dropsondes. Data assimilation could be straightforward, but maximum benefit would be obtained only if the forcing of the boundary layer fluxes were also changed.

Wind velocity could be obtained from Doppler radar in the form of either radial winds or VAD wind profiles. This is technically achievable but needs to be pursued. Ideally greater sensitivity in the radars is required.

Humidity can be obtained using the Raman lidar technique. This would be expensive and provide only sparse data.

Precipitation can be obtained from TRMM, and is planned for GPM, which will deliver a global distribution every 1.5 hours. It will be of limited use over land, where better quality information is available from ground-based radar.

In plenary discussion, George Craig enquired whether there was value in using bistatic receivers to obtain dual Doppler winds. This was felt to be too expensive. The lack of vertical temperature and humidity profiles was felt to be a key issue, and hence the potential contribution of unmanned aircraft was an important area of investigation.

5. Key strategic points arising from meeting

- i) Assimilation techniques not using the model need a balance constraint or equivalent to work successfully. In storm-scale prediction it looks unlikely that there is such a constraint. Therefore the more expensive forms of data assimilation (4dVAR, EnKF) will be essential. This contradicts previous thinking that a simpler method (e.g. 3dVAR) would have to be used.
- ii) Ensemble techniques still look the best bet, though there is much research to be done on the best way of using them on this scale.
- iii) The data requirements for the assimilation can be met if enough money is spent. If the level of the investment made in the USA was made in the UK, prediction would be possible in principle.
- iv) The main constraint on implementation is that the assimilation can be done fast enough to keep up with the real-time data flow. If this is achieved, the forecast/dissemination will also be possible on a useful timescale.
- v) It was not possible to determine what level of skill is actually achievable based on current evidence. There appears to be room for considerable improvement in the accuracy of the larger scale (50km) predictions, which will have to be achieved before storm-scale prediction has any chance.

The following is the main response to the questions asked at the beginning.

vi) What has to be achieved to meet the user requirement?

See Dr Golding's introductory talk.

vii) What level of predictability is scientifically possible, both deterministically and stochastically?

It was assumed that predictions of about 3 lifecycles was possible, giving time ranges of ½ hour for individual cells to 36hrs for organised mesoscale convective systems. However, there is no practical/theoretical evidence to back this up.

viii) What can be realistically attempted within the next 10 years using affordable computing facilities, existing data assimilation and modelling techniques, and available observing technology?

Predictions of organised convection, supported by a data assimilation on a 3km scale over the UK. This could also give useful information on severe thunderstorms.

ix) If there are problems that cannot be solved using existing techniques, are there techniques available from other application areas which could be used?

There were a number of techniques suggested, see the reports in section 3.

x) If there are major scientific questions to be answered before techniques can be developed, what sort of research should be undertaken to address them?

See the reports in section 3.