Task Order Report

Data Assimilation Techniques for Application to the RAMS and WRF models

Contract #FA2521-04-D-0003, Delivery Order #5

Prepared for:

AFTAC – Air Force Technical Applications Center
Patrick Air Force Base, FL

and

Yocke and Company
Novato, CA

Prepared by:

ATMET, LLC
PO Box 19195
Boulder, Colorado 80308-2195

Craig J. Tremback
John S. Snook
Robert L. Walko

August 2005

REPORT DOCUMENTATION PAGE	Form Approved OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
August 2005	Task Order Report	

4. TITLE AND SUBTITLE

Data Assimilation Techniques for Application to the RAMS and WRF Models

5a. CONTRACT NUMBER

FA2521-04-D-0003

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Craig J. Tremback
John S. Snook
Robert L. Walko

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

ATMET,LLC
PO Box 19195
Boulder, CO 80309-2195

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Technical Applications Center
Patrick AFB, FL

Yocke and Comp.
Novato, CA

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

A - Distribution is Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT

15. SUBJECT TERMS

(1) Data Assimilation (2) RAMS (3) WRF

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Craig J. Tremback
U		U		70	19b. TELEPHONE NUMBER *(Include area code)*

Table of Contents

List of Acronyms

3DVAR	3-Dimensional Variational Assimilation Scheme
4DDA	4-Dimensional Data Assimilation
4DVAR	4-Dimensional Variational Assimilation Scheme
AFTAC	US Air Force Technical Applications Center
AFWA	US Air Force Weather Agency
CIRA	Cooperative Institute for Research in the Atmosphere
CALPUFF	Air quality dispersion model from EarthTech
CAMx	Comprhensive Air Quality Model with Extensions
CMAQ	Community Multiscale Air Quality model
CSU	Colorado State University
EDAS	ETA Data Assimilation System
ETA	Forecasting model from NCEP
ECMWF	European Center for Medium-range Weather Forcasting
EnKF	Ensemble Kalman Filter
FFT	Fast Fourier Transform
FSL	NOAA Forecast Systems Laboratory
GFS	Global Forecasting System
GRIB	Gridded Binary data format
HIRLAM	HIgh Resolution Limited Area Model
HYPACT	Hybrid Particle and Concentration Transport model
JMA	Japan Meteorological Agency
KF	Kalman Filter
LAPS	Local Analysis and Prediction System
MLEF	Maximum Likelihood Ensemble Filter
MM5	Mesoscale Model version 5
MPI	Message Passing Interface
NCAR	National Center for Atmospheric Research
NCEP	National Centers for Environmental Prediction
netCDF	network Common Data Format
NOAA	National Oceanographic and Atmospheric Administration
NOAH LSM	Land Surface Model for MM5 and WRF

OI	Optimal Interpolation Assimilation Scheme
OLAM	Ocean Land Atmosphere Model
OMEGA	Operational Multiscale Environment model with Grid Adaptivity
RAMDAS	Regional Atmospheric Model and Data Assimilation System
RAMS	Regional Atmospheric Modeling System
RAMS/ISAN	RAMS Isentropic Analysis
REVU	RAMS Evaluation and Visualization Utilities
RRTM	Rapid Radiative Transfer Model
RUC	Rapid Update Cycle
SSMI	Special Sensor Microwave Imager
SST	Sea Surface Temperature
TEB	Town Energy Balance model
UKMO	United Kingdom Meteorological Office
WRF	Weather Research and Forecasting model
WRF-ARW	WRF Advanced Research WRF core
WRF-EM	WRF Eulerian Mass core
WRF-NMM	WRF Non-hydrostatic Mesoscale Model core
WRF-SI	WRF Standard Initialization

1 Introduction

Data assimilation, which can be loosely defined as the techniques to best utilize observational data in conjunction with a numerical simulation, has continued to be a challenging problem for the atmospheric numerical modeling field. Past techniques of objective analysis (Cressman, Barnes, etc.) were simple attempts to interpolate randomly-located observations of basic state variables to a regular grid structure.

Within the past decade, schemes based on variational numerical methods have become popular. Variational schemes use the concept of the minimization of a cost function. A simple example of a cost function could be the domain average of the squared differences between an analyzed field and the observations. By adjusting the analyzed field, the size of the cost function is affected. The resultant solution is the analyzed field when the cost function is at a minimum.

Variational data assimilation has the main advantage that it is possible to use non-linear observation operators and the resulting freedom to use any observed quantities as input to the assimilation instead of a derived product. This makes it possible to use remote sensing data, for example satellite radiances in the variational data assimilation schemes, in addition to the basic state variables. It is not required to first convert the remote data to the state variables as in traditional data analysis, although a relationship must be established between the observed and model variables. It is also possible to include physical constraint terms in the cost function (e.g., non-divergence or gradient wind balance). Weights (error coefficients) are applied to each term to determine the relative importance in the final analysis.

The variational data assimilation schemes can be divided into two groups, called 3DVAR or 4DVAR. A 3DVAR scheme will only analyze the fields at a given point in time, while the 4DVAR schemes will also take into account the time variation.

A 4DVAR scheme uses a cost function as in a 3DVAR scheme, but the terms in the cost function encompass multiple times. Minimization of the 4DVAR cost function is accomplished through a series of iterative steps in which the full non-linear model is run forward in time, then an *adjoint* of a simplified linear version of the model is run backward in time. This cycle is typically repeated on the order of 100 times for adequate convergence of the solution.

It has been pointed out (Zupanski 2005; Kalnay et al. 2000) that both 3DVAR and 4DVAR, as well as the commonly used data assimilation method of Optimal Interpolation (Daley 1991) are approximations to Kalman filtering theory, and that with certain assumptions, the variational schemes will converge to the basic Kalman Filter. There is a relatively new technique, also an approximation to the Kalman Filter, that is starting to gain favor from many researchers. It has been termed an *Ensemble Kalman Filter (EnKF)* (or *Ensemble Kalman Smoother*) and provides similar results to 4DVAR

schemes without the need of adjoint or linear models. EnKF started being investigated about 10 years ago in meteorology and oceanography.

The EnKF technique has the advantages that only forward runs are used (no adjoints are required) and the convergence properties have a similar (or less) level of computational effort to 4DVAR. Also, since only forward runs are used, all aspects of full model (non-linearity, all physics scheme, etc.) can be utilized in creating the ensembles.

This project was intended to serve as a preliminary investigation into the variational schemes, with the goal of determining a reasonable research path to attain the best possible data assimilation scheme(s) for both the Weather Research and Forecasting model (WRF) and the Regional Atmospheric Modeling System (RAMS), with a focus on AFTAC's applications.

We executed and tested the WRF 3DVAR scheme in the context of the WRF model. We also developed a converter to transform the WRF output files into the RAMS input ("varfile") format, allowing us to use the 3DVAR fields in RAMS simulations. Simulations were performed over two cases with the two models in several configurations. Literature reviews of the current state of the 4DVAR and EnKF schemes were performed. An important result of this project are the recommendations on how best to proceed in the development and usage of 3DVAR data assimilation schemes in conjunction with RAMS and WRF and the feasibility of implementation of the 4DVAR /EnKF schemes in the models.

It should be noted that this project used WRF v2.0.3.1. A new WRF version 2.1 was released in early August, which included a new version of the 3DVAR scheme. Due to time constraints, we were unable to test the new version for this report. However, we have reviewed the code and release notes and have commented in the sections below.

2 The WRF 3DVAR scheme

The following section is partially summarized from Barker et. al. 2003 (hereinafter, Barker03). Figure 1 shows the main components and a flow diagram of the 3DVAR scheme.

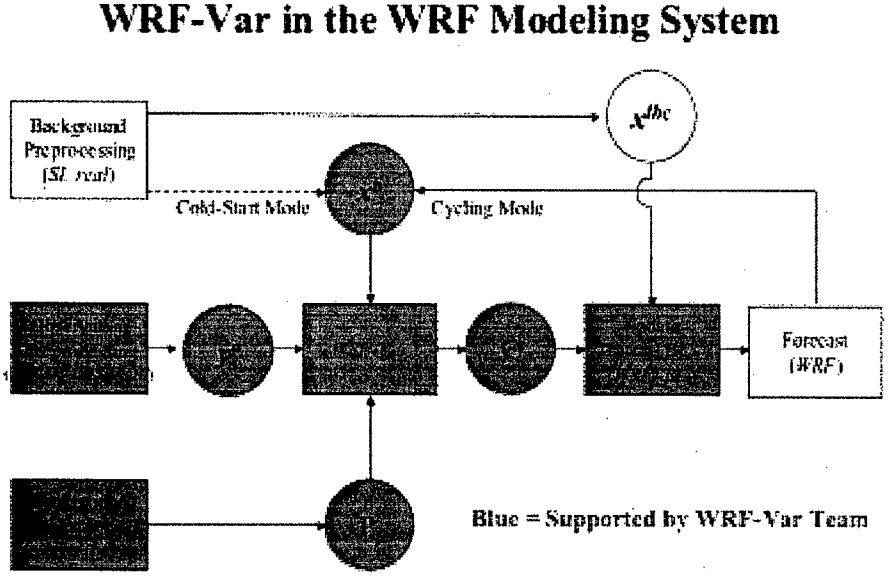

Figure 1: Components of the WRF 3DVAR system (from http://www.mmm.ucar.edu/wrf/WG4/wrfvar.htm)

The WRF/MM5 3DVAR code was developed at NCAR, but follows closely the basic design as implemented operationally at the UKMO (Lorenc et al. 2000), on which most 3DVAR schemes implemented in forecast centers are based. The main features of the 3DVAR system include (from Barker03):

- Quasi-Newton minimization algorithm (Liu and Nocedal 1989) (for minimization of the cost function)
- Computations are performed on an unstaggered grid (Arakawa A grid). The results are then interpolated to the target MM5 (B grid) or WRF (C grid), and also for the results below, the RAMS grid (C grid). Computations are performed on the sigma-height levels of MM5, or mass-coordinate levels of WRF.
- For efficiency, preconditioning is accomplished with a "control variable transform": U defined as $B=UU^{T}$. Preconditioned control variables are chosen as streamfunction, velocity potential, unbalanced pressure and a choice between specific or relative humidity. Preconditioning in this manner removes most of the error correlations between the analysis variables.

- Linearized balance terms relating the mass and wind fields (geostrophic and cyclostrophic terms) are used to define a balanced pressure.
- Climatological background error covariances (the ***B*** matrix) are estimated with the "NMC-method" of averaged forecast differences. The details of this computation can be found in Barker03, but primarily it is based on average forecast differences from one forecast time to another (e.g., differences in the model fields taken between the 12-hour and 24-hour forecast).

According to Barker03, 3D variational techniques have the following advantages over earlier techniques (e.g., Cressman, Barnes, Optimum Interpolation (OI), etc.):

- Observations can be assimilated directly without the need for retrieval algorithms (i.e., conversion to state variables). This implies a more consistent treatment of all observations and, since the observation errors may be less correlated (with each other and the background errors), some simplifications to the analysis algorithm are attained. (However, the only non-state variable able to be handled by the current scheme is precipitable water.)
- The solution is found using all observations simultaneously, unlike the OI technique for which a data selection into artificial sub-domains is required.
- Asynoptic data can be assimilated near its valid time through the use of frequent updates of the analysis.
- Balance constraints (e.g. geostrophy, hydrostatic) can be built into the scheme.

A 3-D variational scheme can be viewed as the solution of the vector $\boldsymbol{x}$ which minimizes a *cost function.* In the WRF 3DVAR scheme, the cost function takes the form (in matrix notation):

$$J(\boldsymbol{x}) = \frac{1}{2}(\boldsymbol{x}-\boldsymbol{x}^b)^T \boldsymbol{B}^{-1}(\boldsymbol{x}-\boldsymbol{x}^b) + \frac{1}{2}(\boldsymbol{y}-\boldsymbol{y}^o)^T (E+\boldsymbol{F})^{-1}(\boldsymbol{y}-\boldsymbol{y}^o) \qquad (1)$$

Alternatively, this can be written in vector notation as:

$$J(\boldsymbol{x}) = \frac{1}{2}\sum_b w_B (\boldsymbol{x}-\boldsymbol{x}_b)^2 + \frac{1}{2}\sum_O (w_E + w_F)(\boldsymbol{y}-\boldsymbol{y}_O)^2$$

where

x – model variables defined at grid points

x_b - "first guess" field (e.g., from larger-scale forecast or analysis)

y - model variables transformed to observation type and location - $y = Hx$ where H is defined as the observation operator, a function or scheme to perform the transformation.

y_O - observed variable

w_B - weighting factor for each grid point and variable dependent on expected errors between model and first guess field

w_E - weighting factor dependent on observation instrument error

w_F - weighting factor dependent on "representivity" error (inaccuracies in H)

In the matrix form, $\boldsymbol{B}$, $\boldsymbol{E}$, and $\boldsymbol{F}$ are more commonly described as the background, observation (instrumental), and representivity *error covariance matrices.*

The quadratic cost function defined in (1) assumes that the statistics of the observation and background error covariances are described using Gaussian probability density functions with zero mean error. Alternative cost functions may be used which relax these assumptions (e.g. Dharssi et al. 1992), but these are generally more complex and time-consuming. Also, (1) additionally neglects correlations between observation and background errors.

Under these assumptions, the minimization of the cost function which defines the field $\boldsymbol{x}$ is then accomplished by an iterative solution using standard numerical methods. Convergence of the algorithm as implemented takes on the order of about 50 iterations.

Barker03 identifies the following drawbacks of the 3DVAR algorithm:

- The quality of the output analysis depends crucially on the accuracy of prescribed errors, especially B.
- The method allows for the inclusion of linearized dynamical/physical processes, but actual errors in the models may be highly nonlinear. This limits the usefulness of variational data assimilation in highly nonlinear regimes, for example, on the convective scale or in the tropics.

Expanding on a few of the points above:

- While the 3DVAR cost function would most commonly incorporate standard meteorological observations (i.e., rawinsondes and surface observations), inclusion of nonstandard observations is straightforward. For example, data collected by satellites, aircraft, Doppler radar, and wind profilers, which occur at irregular locations, are easily included in the cost function provided that a means for relating these observations to the model state vector is provided. The easiest way to do this is to interpolate the state variables to the location of the observation and formulate the corresponding difference term of the cost function in observation space rather than in model grid-point space. Observed quantities such as satellite radiances that are not elements of the model state vector can be obtained from the state vector using supplementary computation. Missing or incomplete observational data does not pose any particular numerical problem for 3DVAR; it simply results in fewer terms in the cost function.

- The 3DVAR implementation for MM5 and WRF is performed on an Arakawa A grid and subsequently interpolated to the B grid of MM5 or the C grid of WRF. This interpolation introduces additional error that is not subject to the minimization procedure of the variational cost function. Interpolation of the optimized initial state from the A grid to the RAMS grid (which is Arakawa C as in WRF) introduces the same type of error, but this procedure is no less valid than it is for MM5 and WRF.

- Hydrostatic, geostrophic, and cyclostrophic balance terms have been included in the MM5/WRF 3DVAR system. Imposition of these balances in numerical forecast initial conditions reduces unrealistic oscillations during the initial model integration and is likely to provide better forecast results. However, the relative success of this procedure is highly dependent on the degree to which these balances exist in the flow. Successful test results are described in Barker et al. (2004) for a hurricane, which is one of the most highly balanced systems that exists in the atmosphere. Similar improvement to the solution is not to be expected for less balanced systems, such as tropical convection and many mesoscale circulations.

- As with all 3DVAR (and 4DVAR) schemes, the cost function may have more than one local minimum, and the numerical procedure for minimizing the cost function results in convergence to only one of those minima (usually the one directly down-slope from the initial conditions). This minimum may not be as low as another, and hence may not produce the most optimal of all possible state vectors.

- There is no straightforward means of determining optimal values for coefficients of the covariances in the cost function. These coefficients determine the relative weight of each term and hence have a strong influence on the solution. In practice, one must determine empirically which coefficient values result in the best numerical forecasts. It is helpful to have a large database of forecast results, which is usually available only in an operational setting. Matters are further complicated by the fact that optimal coefficients depend on model grid resolution and the geographic size and location of the model domain, so any changes to these require a new search for best values. Moreover, the best results are obtained with coefficients that are allowed to vary both spatially and temporally, at least in some broad sense such as by latitude, altitude, and season. Because of these factors, the entire procedure of developing a highly optimized 3DVAR analysis system is far more complicated than the relatively straightforward tasks of setting up and minimizing the cost function. The search for optimal coefficients remains an active area of research.

- The current WRF (and MM5) 3DVAR scheme only can input state variables (u, v, T, P, vapor), except for column precipitable water. Reflectivities and radiances have not been implemented. ***Update***: the newly-released WRF v2.1 3DVAR is supposed to process radar reflectivity. There is a mention of "pre-operational

trials in Korea" of this capability. However, no documentation, presentations, or papers could be found as of this writing. The change logs and code imply that an observation operator has been added to convert the reflectivities to condensate species and vertical motion through the use of a 1-D cloud model. The cloud model is based on the model in the Anthes-Kuo parameterization from the 1970's.

3 WRF 3DVAR Test Cases

We selected two disparate cases to begin our look at the performance of the WRF 3DVAR scheme: 1) a large-scale, strongly forced winter storm case over the eastern US that moved from the Midwest into the Northeast, and 2) a smaller-scale, weakly-forced sea breeze case over Florida. For each case, both WRF and RAMS were run in different configurations, as noted below. Of course, with an issue as complex as variational data assimilation, two cases will not provide a definitive answer as to the performance of the scheme. However, the two cases allowed us to have a starting point to gain familiarity with the scheme and to be able to compare our results and impressions with other tests found in the literature.

3.1 The WRF-to-RAMS Converter

In order to initialize RAMS from the WRF 3DVAR fields, we developed a converter to transform the WRF fields written in netCDF format to the RAMS input format ("varfiles"). The converter was developed rather generally. It will work on output from the WRF-SI files (Standard Initialization, which is a simple interpolation from GRIB-format gridded datasets, such as GFS or Reanalysis fields), the WRF-3DVAR files, or output files from a WRF model run. In order to properly prepare the RAMS files, several steps and considerations needed to be taken into account:

- Any WRF projection can be used (only the Lambert-Conformal has been tested) as long as the RAMS domain fits within the WRF coarse grid.
- The RAMS grid points will be defined from the highest resolution WRF grid present at the grid point locations.
- The vertical levels do not need to match between the models. Care is taken to conserve momentum, energy and water mass in the vertical interpolations.
- Several hydrostatic balance calculations are performed to ensure an appropriate balance for the RAMS input fields.

The output of the converter is a set of RAMS "variable initialization" files that can be directly used to start a model run. RAMS/ISAN is not needed in this case.

3.2 *Winter snowstorm case*

3.2.1 Meteorological summary

A significant winter storm developed over the upper Midwest and traversed through the Ohio Valley into the mid-Atlantic states from 22-24 January 2005. The origins of the system were associated with a 500 mb short-wave embedded in northwest flow coming out of central Canada. Figure 2 shows the evolution of the system at 500 and 1000 mb. At 0000 UTC on 22 January, the 500 mb short wave is entering the northern US with an associated, small 35 ms-1 jet maxima. Near the surface, the system is reflected as an inverted trough from Minnesota southward and a cold front is evident across Nebraska. Significant development is noted over the next 12 hours as the short-wave tightens and the jet maxima increases to 45 ms-1. A closed off surface low pressure system is centered over northern Indiana and a large high pressure circulation, moving south out of Canada, is reinforcing the cold front.

The jet maxima begins to round the base of the trough over the lower Ohio Valley at 0000 UTC on 23 January that causes the surface low, now centered over Washington DC, to turn eastward. The jet maximum continues to round the base of the 500 mb trough over the next 24 hours as the surface low intensifies off the Eastern Seacoast. Cold air behind the cold front has now reached well into the Gulf of Mexico. By the end of the time period (1200 UTC 24 January), the system has closed off at 500 mb over the Canadian Maritime Provinces and the surface has moved out to sea east of Maine.

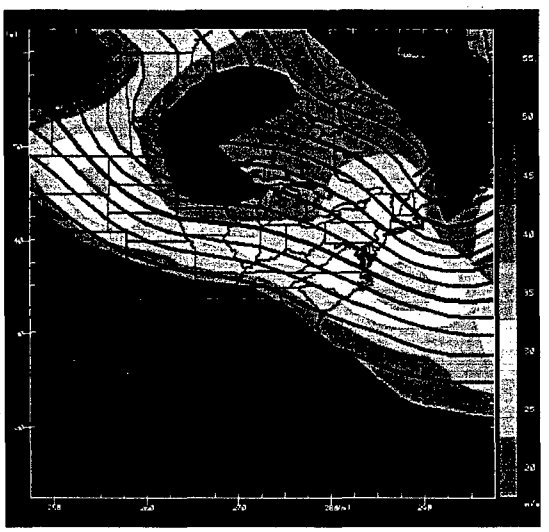

0000 UTC 22 January 2005

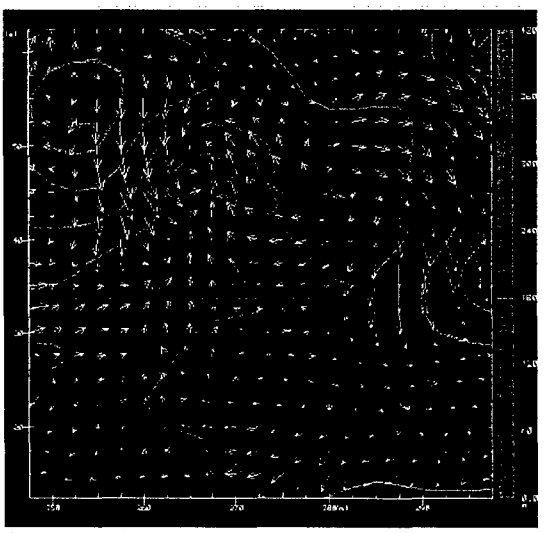

0000 UTC 22 January 2005

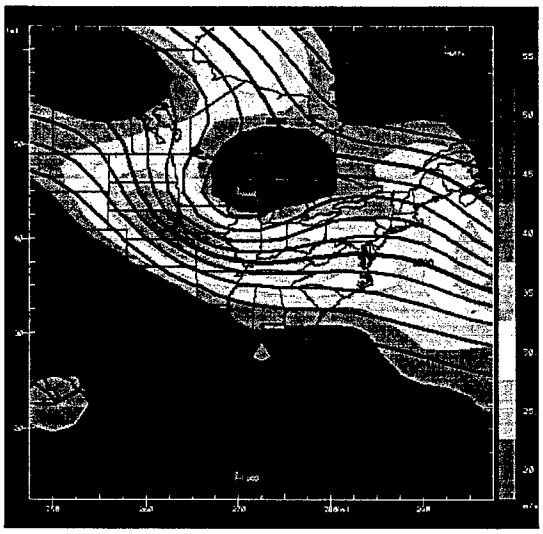

1200 UTC 22 January 2005

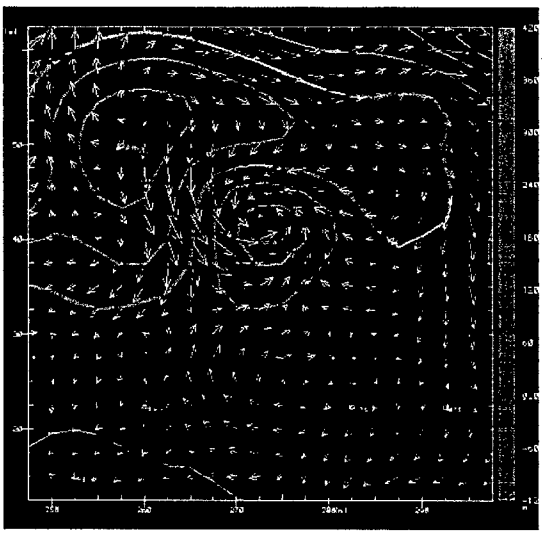

1200 UTC 22 January 2005

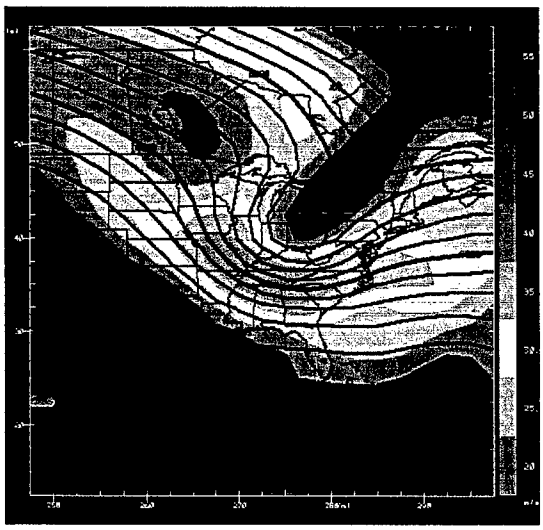

0000 UTC 23 January 2005

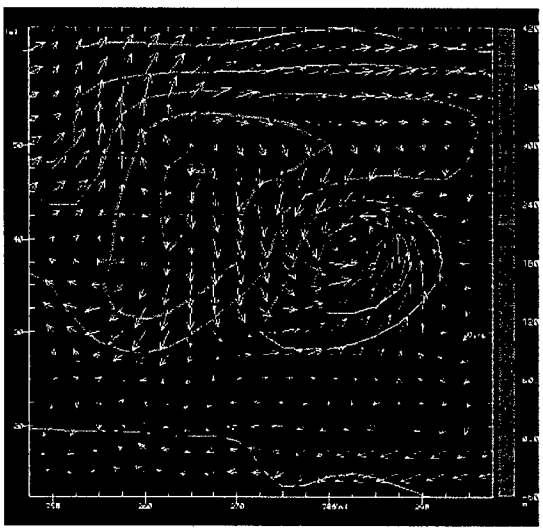

0000 UTC 23 January 2005

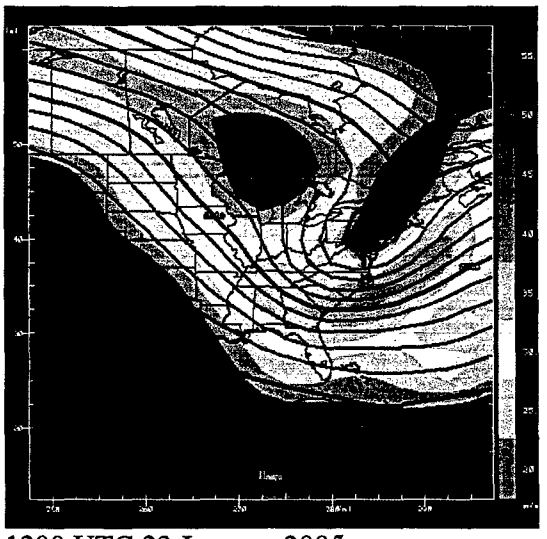

1200 UTC 23 January 2005

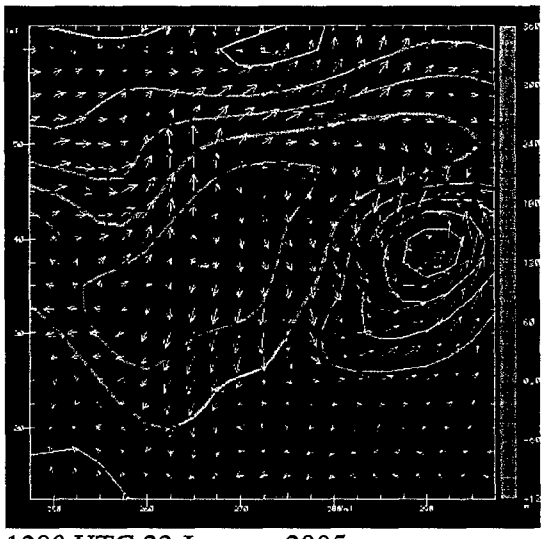

1200 UTC 23 January 2005

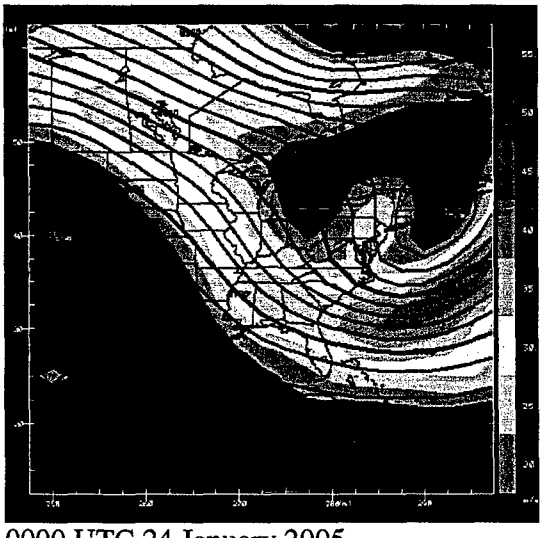

0000 UTC 24 January 2005

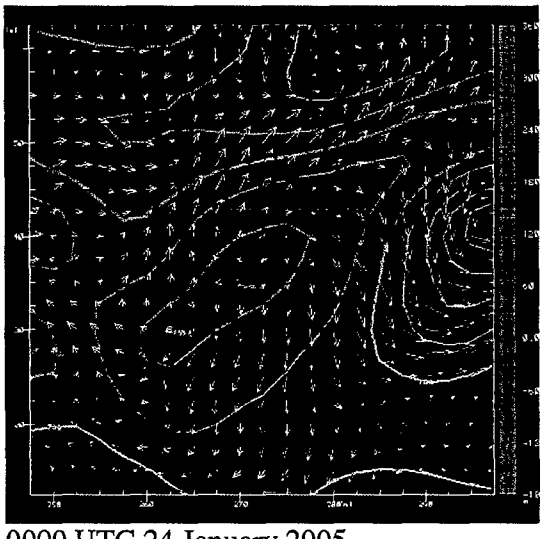

0000 UTC 24 January 2005

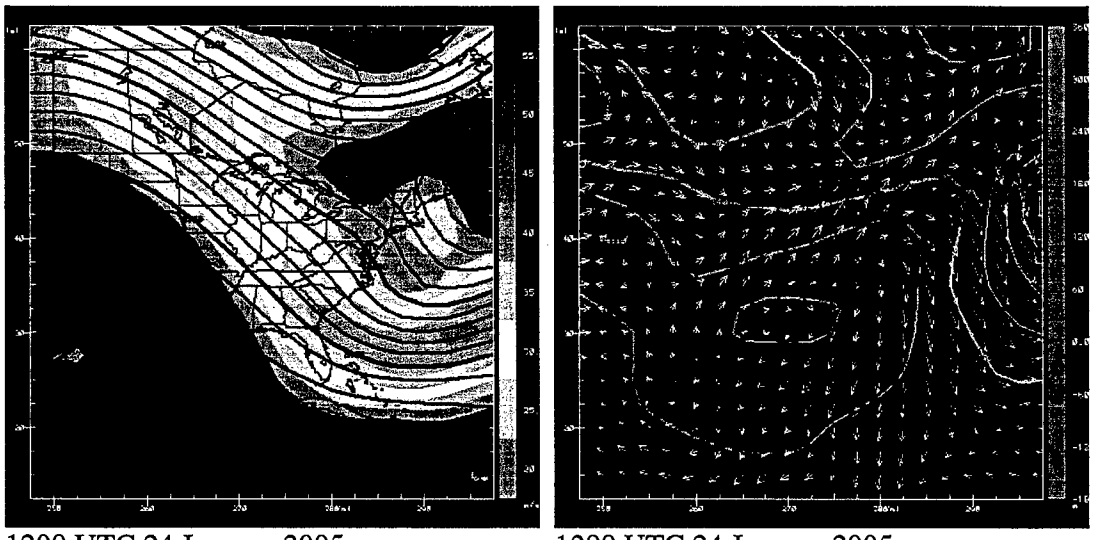

1200 UTC 24 January 2005 1200 UTC 24 January 2005

Figure 2: Left column shows 500 mb height (black lines, contour interval 60 m) and wind speed (> 20 ms^{-1}, color image, contour interval 5 ms^{-1}). The right column shows 1000 mb height (cyan lines, contour interval 30 m) and wind vectors at every model grid point for 0000 UTC 22 through 1200 UTC 24 January 2005.

3.2.2 Model configurations

3.2.2.1 WRF Configuration

A relatively large domain (Figure 3) was selected to cover the entire region affected by this winter storm. To keep compute time within reason, a rather coarse, 30-km grid spacing was employed with a horizontal grid size of 126x126 and 31 vertical levels. The WRF Standard Initialization (WRF-SI) package was used to generate static fields, such as topography, for the model domain. Initial condition and forecast lateral boundary condition data were derived from the NCEP Global Forecast System (GFS) model analysis data that were obtained from archives available at NCAR. The WRF-SI package was used to interpolate the GFS data to the WRF model grid.

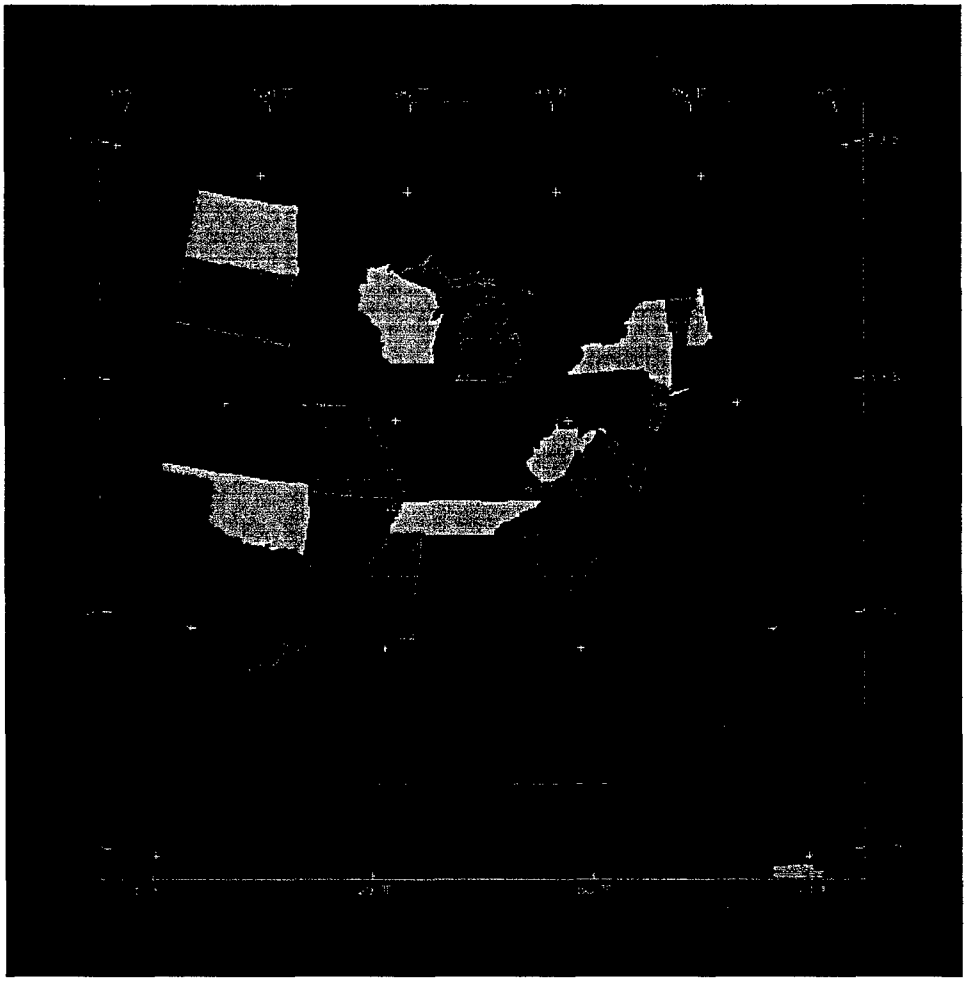

Figure 3: Winter case study model domain.

Two sets of model simulations were completed – control and 3DVAR. The control runs were initialized with the GFS data grids as output by the WRF-SI package. Initialization of the 3DVAR runs was accomplished using grids derived from the WRF-3DVAR package. As discussed in the previous section, the package starts with an initial background field and additional observational datasets are assimilated in to form a new set of three-dimensional atmospheric analyses. The package was configured to use the WRF-SI GFS grids generated for the control runs as the background field. AFTAC provided two observational datasets for January 2005 that were already formatted for direct use by the WRF-3DVAR package. The first set included point observations that contained METAR, ship, aircraft report, rawinsonde, and satellite cloud drift wind data. The second set contained SSM/I data. These datasets were assimilated directly into the GFS grids to form the 3DVAR initial conditions. WRF uses tendencies from the initial condition as its forecast lateral boundary condition, and these tendencies were adjusted based on the new 3DVAR initial condition fields.

Initially, two WRF simulations, control and 3DVAR, covering the entire event were completed. Both simulations were initialized at 1200 UTC on 22 January and results were generated at 1-hour increments out to 72 hours. Five additional sets (control and 3DVAR) of forecasts were generated to provide more comparison data. These simulations were initialized at 6-hour increments (1200, 1800 UTC 22 January, 0000, 0600, and 1200 UTC 23 January) and results were output at 2-minute increments out to 24 hours for the 1200 UTC 22 January runs and out to 6 hours for the other simulations.

The WRF model has a variety of physics packages available. The comparison of physics packages was not an objective of this study and hence, the physics selected were based primarily on experiences with the MM5 model that contains many physics packages similar to WRF. Table 1 contains a summary of the model configuration. The WRF single-moment 6-class microphysics scheme with ice, snow and graupel processes is most suitable for working with a wintertime precipitation event. It should be noted that using the NOAH land surface model is preferable based on results from previous ATMET studies. The NOAH LSM, however, requires a skin temperature/SST initialization and these fields are not available in the NCAR GFS dataset.

Table 1: WRF model physics options used for winter case study simulations.

Physics	Option
Microphysics	WRF single-moment 6-class
Land Surface	5-layer thermal diffusion: Soil temperature only
Longwave Radiation	RRTM
Shortwave Radiation	Dudhia scheme
Surface Layer	MM5 similarity
Planetary Boundary Layer	Yonsei University scheme (modified MRF scheme)
Cumulus Parameterization	Kain-Fritsch scheme

3.2.2.2 RAMS Configuration

RAMS version 6.0 was used for the simulations for this project. The model was configured in a similar grid configuration as WRF, using 125x119x38 grid points with a horizontal grid spacing of 30 km. The RAMS' horizontal domain needed to be slightly smaller than the WRF domain so that the 3DVAR fields could be interpolated to all grid points due to the slight mismatch between the RAMS polar-stereographic projection and the WRF Lambert-Conformal projection. The various physics options are listed in Table 2.

Table 2: RAMS model physics options used for winter case study simulations.

Physics	Option
Microphysics	Full 7 species, prognostic ICC
Land Surface	LEAF 3 vegetation and soil; 8 soil layers; 2 land patches
Longwave Radiation	Chen-Cotton
Shortwave Radiation	Harrington
Surface Layer	Louis
Planetary Boundary Layer	Mellor-Yamada TKE
Cumulus Parameterization	Modified Kuo scheme

Three RAMS runs were made for this case, varying the details of the definition of the initial and boundary conditions:

1) RCTL ("Control") run – straight interpolation of the GFS files
2) ISAN run – RAMS/ISAN used to blend observations with GFS first-guess fields
3) R3DV run – WRF 3DVAR files interpolated to RAMS grid

3.2.3 WRF simulation analyses

There are obviously many figures we can show of the model fields for these runs, but the main indications of model performance can be summarized in the statistical evaluation in the following section. Here, we will present the initial surface temperature analysis (Figure 4) from the WRF control and 3DVAR runs. Note there are some differences apparent in the amplitude of the features, especially the ridges near the southern Tennessee border and over Wisconsin.

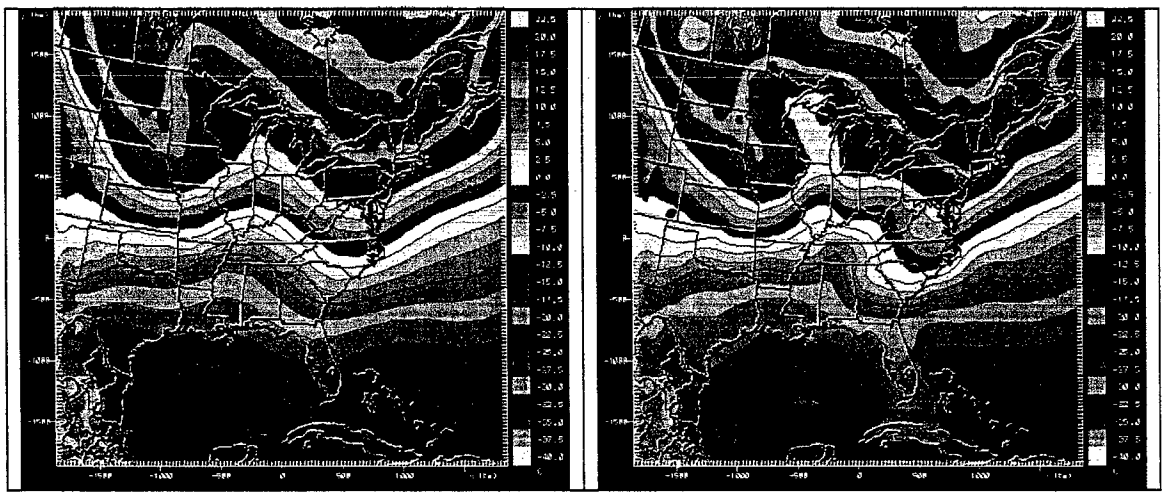

Figure 4: Initial surface temperature from WRF control run (left) and 3DVAR run (right). Contour interval is 2.5° C.

To provide an overall summary of the differences between the control and 3DVAR runs, several methods were employed. These included: 1) traditional statistical evaluation of point observations with model results interpolated to observation locations, 2) subjective evaluation of simulated time series of pressure and vertical motion, and 3) discrete Fourier transform analysis to objectively evaluate power spectrum and noise characteristics of the simulations.

3.2.3.1 Statistical evaluation

The archived National Weather Service METAR surface data were obtained from NCAR. These data were decoded into a format compatible for use in the RAMS statistical analysis package. WRF gridded forecasts were then interpolated to all available METAR locations within the model domain for direct comparison with the observations. This required development of the RAMS Evaluation and Visualization Utilities (REVU) package to process the WRF output files.

Figure 5 illustrates domain-wide average forecast temperature, dew point, and wind speed from the control and 3DVAR simulations along with observations (approximately 1,800 surface stations were used). Mean absolute error statistics are also displayed. In all cases, the 3DVAR run initial condition showed significant improvement over the comparable control simulation initial condition. This is not surprising, since the control run initial conditions were a straight interpolation from the coarse resolution GFS fields, while the 3DVAR assimilated the observations.

However, the improvement dropped off very rapidly with time. Overall quality was virtually identical after 8 simulation hours for temperature and 12 hours for dew point. And after 36 hours, the control runs showed slightly better forecasts than the 3DVAR runs. It should be noted that both the control and 3DVAR runs showed a significant warm temperature and dew point bias through the entire forecast period, most likely partly due to the lack of snow cover in the model. The 3DVAR wind speed forecast did, however, show small improvements over the control runs through the forecast cycle. But again, the improvement was small when compared to the large high wind speed bias compared to the observations.

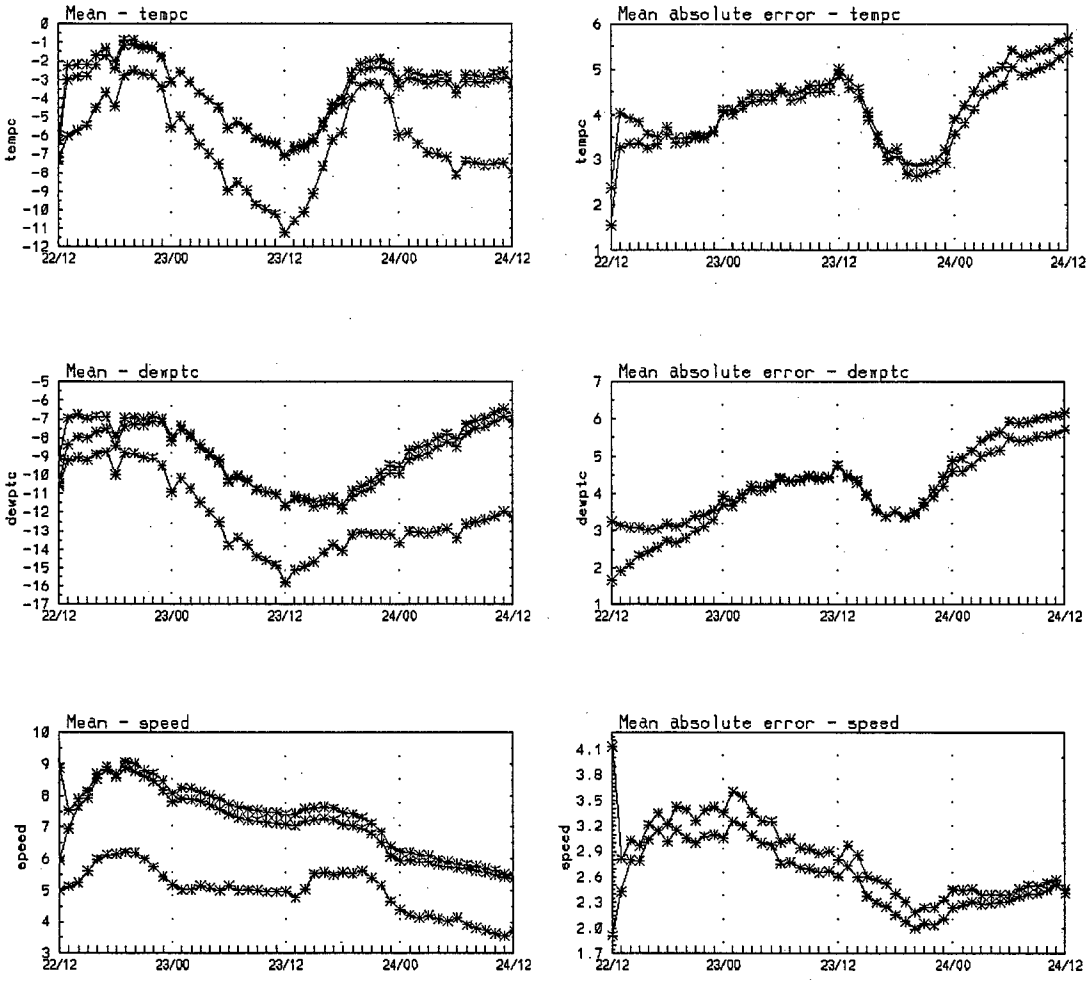

Figure 5: Domain-averaged statistics for temperature (C), dewpoint temperature (C), and wind speed (m/s). The left column contains the domain-averaged quantities; the right column is the mean absolute error. Red lines – average of the observations; blue lines – control run results; brown lines – 3DVAR results.

3.2.3.2 Time series analysis

Imbalances between mass and wind in the model initial condition can manifest into spurious gravity wave features that propagate at high speeds compared to other meteorological phenomena. These imbalances will affect the initial conditions of the model run, but also, if the analysis fields are used in an "analysis" nudging 4DDA scheme, the imbalances can affect the entire simulation. Since one of the main advantages of 3DVAR schemes is to be able to include balance terms, we tested for the degree of balance which was attained with the current WRF scheme. The imbalances will

typically be exhibited as high frequency oscillations in simulated time series of pressure and vertical velocity.

Time series of pressure and vertical velocity were examined from the simulations with the 2-minute interval output to subjectively evaluate and compare noise characteristics from the control and 3DVAR simulations during the early portions of the forecasts. Individual grid point locations were first reviewed. Figure 6 shows a representative sample of pressure at the lowest model level and vertical velocity at model level 5 (approximately 500 m above ground level). The four points are located over different terrain characteristics: 1) central Iowa, 2) mountainous terrain along the Virginia-West Virginia border, 3) the North Carolina coast, and 4) the Atlantic Ocean.

A wide variety of oscillations are evident at different temporal scales. Careful examination of each time section does suggest that the 3DVAR simulations exhibited somewhat less of the highest frequency perturbations compared to the control runs during the first several hours of simulation, however, the amplitude of the somewhat longer frequencies are larger. It should be noted that high frequency oscillations are not necessarily spurious, especially in areas of mountainous terrain and convection. For example, the control simulation vertical velocity time series over the Appalachian Mountains indicated a 6-hour period of high frequency oscillations which are likely influenced by the terrain. The 3DVAR simulation also showed a period of high frequency oscillations, but of shorter duration. Without detailed observational data, which are not obtainable from the conventional observational network, it is difficult to validate how much noise is real versus the spurious noise resulting from the initial condition analyses. This type of analysis, however, does provide information as to how well-balanced the initial fields produced by the 3DVAR scheme were.

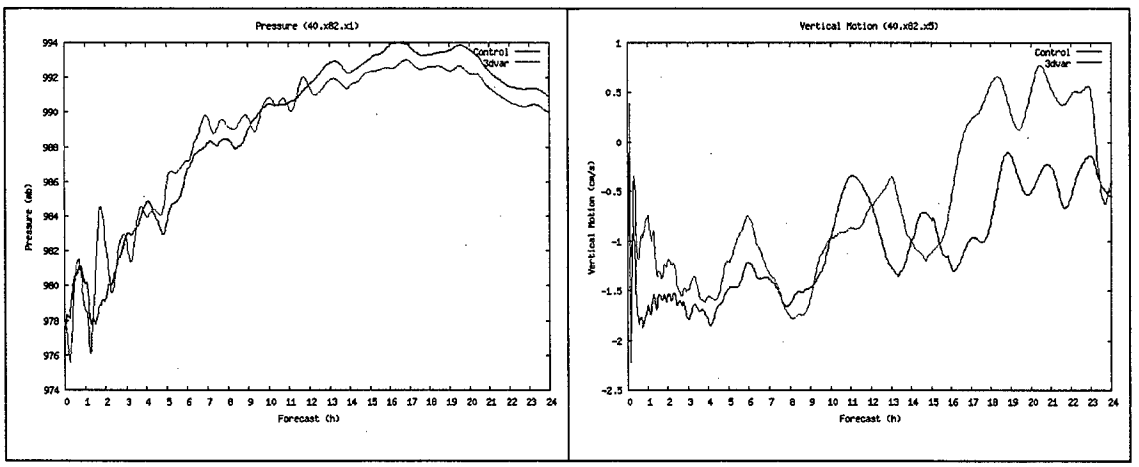

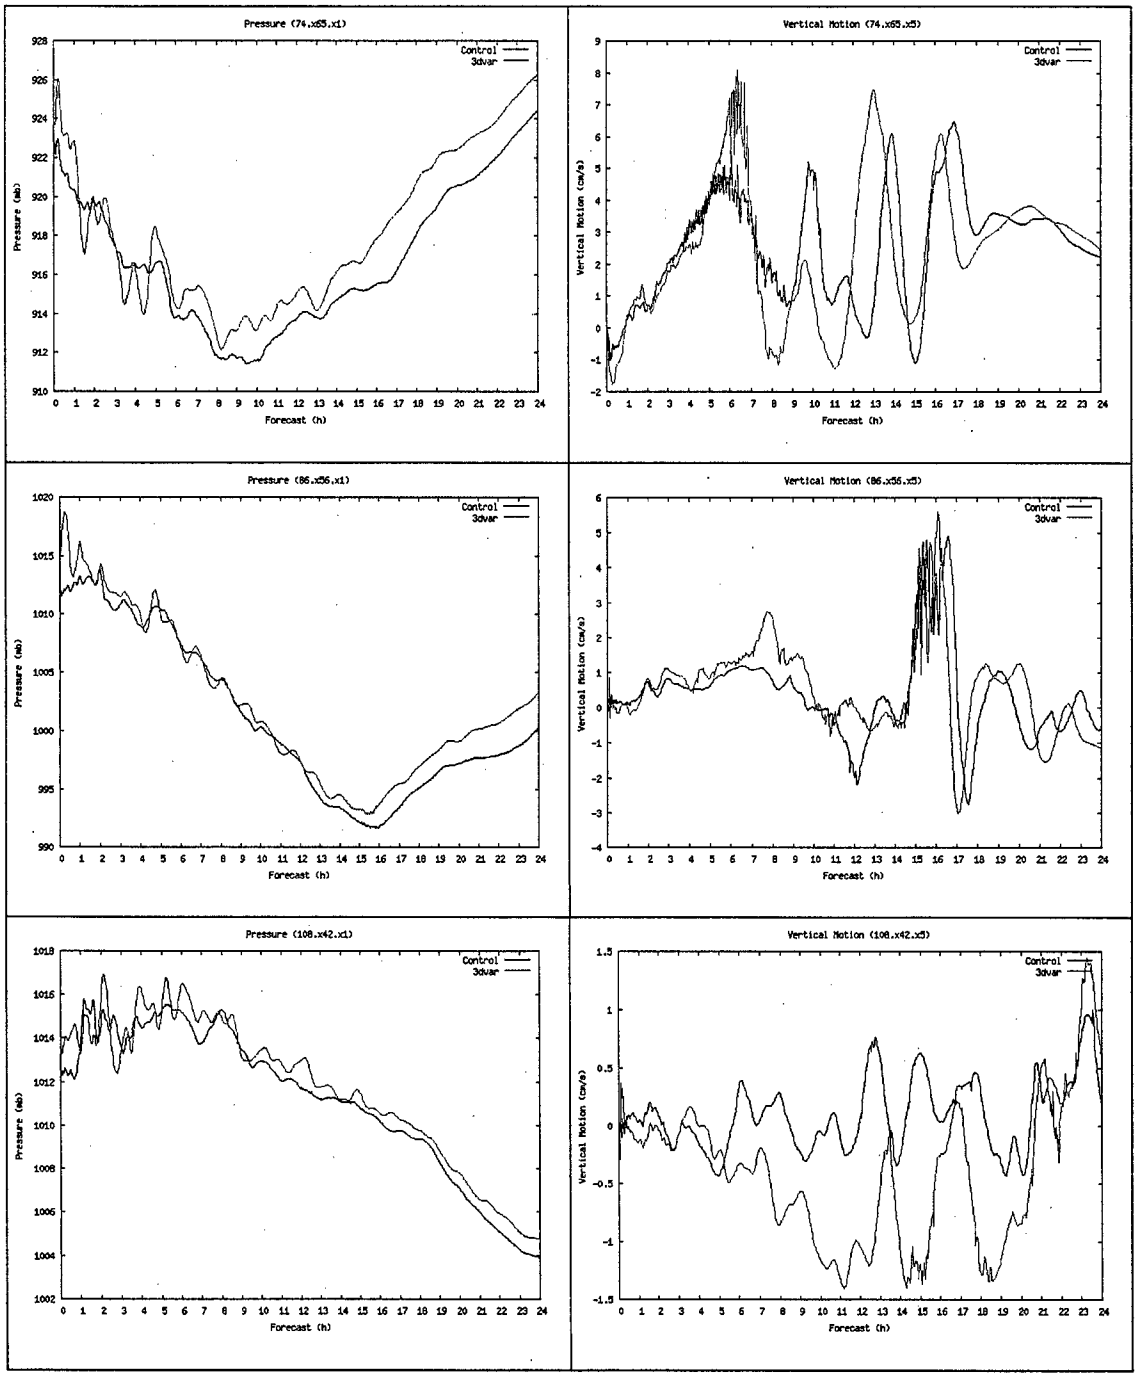

Figure 6: Surface pressure (mb) and 500m vertical motion (cm/s) time series from the WRF control run (red) and 3DVAR run (green) at four representative grid points: 1) central Iowa, 2) mountainous terrain along the Virginia-West Virginia border, 3) the North Carolina coast, and 4) the Atlantic Ocean.

A broader subjective view is presented by averaging the forecast time series across all domain grid points (Figure 7). Domain averaging will reduce the evidence of gravity wave oscillations since nearby grid points will offset each other as some locations are in

the pressure trough while other points are in the pressure ridge. As expected, the averaged results did indicate much smoother fields than the individual time series, although high frequency oscillations were evident during the first 12 minutes of prediction, and the magnitude was about equal for both control and 3DVAR simulations.

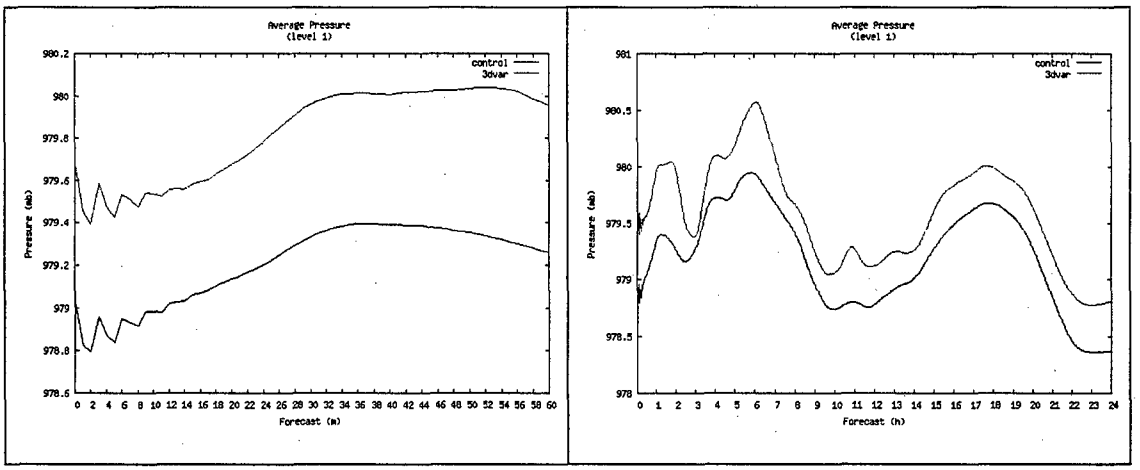

Figure 7: Domain-averaged surface pressure (mb) for the control (red) and 3DVAR (green) runs. Left panel is a close-up of the first 60 minutes of the simulations; the right panel is for the first 24 hours.

3.2.3.3 Fourier time series analysis

Discrete Fourier transform analysis is a method that can be used to objectively evaluate and quantify periodicities in a data series and the relative amplitudes of the wavelengths. By applying Fourier transform techniques to the simulated time series discussed in the previous section, we analyzed the temporal frequency characteristics of the model simulations and objectively evaluated whether WRF-3DVAR reduced the influence of high frequency noise early in the model simulations. This technique will also remove the influence of temporal offsets from point to point, as was the case of the domain-averaged analysis in the previous section.

The implemented technique uses a Fast Fourier Transform (FFT) on the desired time series. The result is a sequence of complex numbers with the same length as the original time series. The complex series is symmetrical around the mid-point and represents the positive and negative frequency components of each wavenumber. Finally, the power spectrum is derived by taking the complex modulus, defined as the magnitude of the real and complex components, of the transformed series and scaling the result so that the sum of the components equals one. Amplitude versus wavenumber plots of the power spectrum reveals the relative importance of each wavenumber that comprises the original time series. Wavenumbers will range from 1 to (n/2)-1 where n is the number of points in the time series.

The FFT technique was applied to the five simulation sets (control and 3DVAR) with different initialization times and 6-hours of output at 2-minute increments. Fourier transforms were applied to each individual grid point time series using two time spans: 3 hours in length (91 time points) and 6 hours in length (181 time points), providing 3-hour and 6-hour power spectrum results at every domain grid point for each simulation. A comparison of 3-hour and 6-hour FFT applications showed similar findings, thus only the 6-hour results are presented.

Domain-averaged power spectrum results were then compiled for each simulation. Further averaging was accomplished by averaging the results of all five runs together for the control and 3dvar simulations respectively. Figure 8 shows these results that are averaged over all five simulations for the 6-hour time span. It should be noted that these averaged results had similar characteristics of results from the individual simulations. The plots do not include wavenumbers 1 and 2 which removes the longer-term trends and allows closer examination of the scaled results for the remainder of the wavelengths.

The power spectrum curves indicate consistent decreasing influence from low to high wavenumbers in both the control and 3DVAR simulations. The 3DVAR results indicate a greater influence of low wavenumbers when compared to the control runs, and a crossover occurs at wavenumber 10 after which the higher wavenumbers show less influence in the 3DVAR forecasts. This suggests that the 3DVAR initialization has acted to remove some of the higher frequency variations and shift them to longer wavelengths.

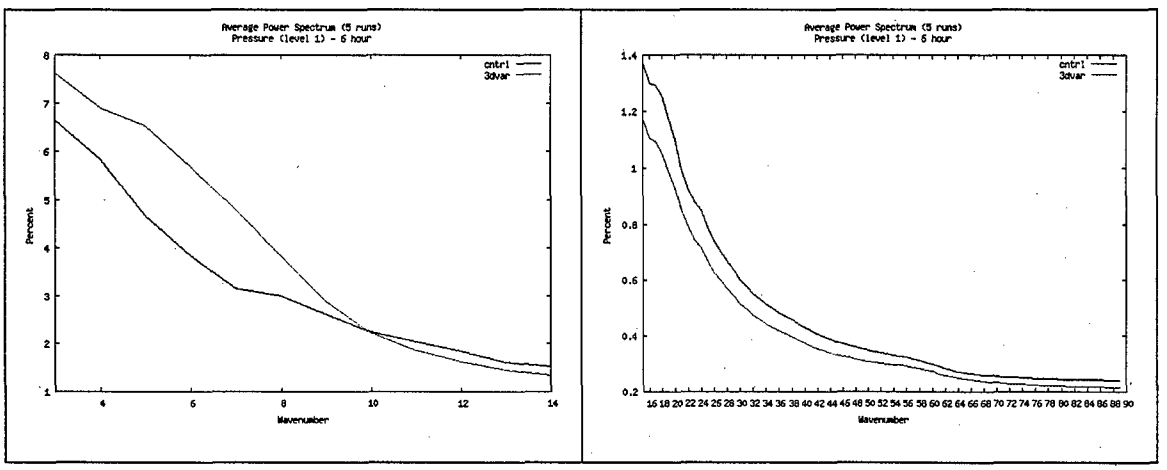

Figure 8: Average power spectrum over a combination of 5 runs of the control (red) and 3DVAR (green) runs (see text for details). Left panel shows wavenumbers 3-14, right panel shows 15-90.

3.2.4 RAMS verifications

Figure 9 shows the same initial surface temperature analysis as presented above for the WRF runs, comparing the RAMS/ISAN analysis with the WRF 3DVAR field interpolated to the RAMS grid. Overall, they are very similar. However, there is a notable

difference in a couple areas. In the RAMS/ISAN analysis, note the orientation of the contours off the coast of the Carolinas and to the east of New Jersey north to Massachusetts. The contour lines tend to be parallel to the coast. This feature is an artifact of the Barnes objective analysis scheme, which manifests itself when moving from a data rich region (over land) to a data poor region (over water). Values from the land areas are wrongfully extended over the water. Although the simulated fields tend to adjust rather quickly as the simulation progresses, the effects of the initial analysis over the water can linger for several hours. The 3DVAR analysis, because of the use of the observations *and* the first guess field in the cost function, can produce a smoother, more consistent analysis. This is not implying that the 3DVAR field is "correct" (since there are no observations to compare it to), only that it is more consistent with the first-guess fields.

Figure 9: Initial surface temperature from RAMS ISAN run (left) and R3DV run (right). Contour interval is 2.5° C.

The same three validation methods: statistical evaluation, subjective evaluation of forecast time series, and Fourier analysis, were applied to the three RAMS simulations.

3.2.4.1 Statistical evaluation

RAMS gridded forecasts were interpolated to all available METAR observation locations within the model domain for direct comparison. Domain-wide average observations and simulations of temperature, dew point, and wind speed are shown in Figure 10. Mean absolute error statistics are also displayed.

As with the WRF runs, the ISAN and R3DV-initialized simulations showed improvement over comparable control simulations at the initial time. The length of improvement varies by field. For temperature, the ISAN simulation retained an small improvement over the

control run through the entire forecast period, while the R3DV improvement lasted for about 24 hours and then showed a degradation of quality after 30 hours. ISAN and R3DV dew point results were nearly identical and showed improvement through the first 12 forecast hours. The ISAN run showed a small improvement which continued through the entire 48 hour forecast period, while the R3DV results faded and lost advantage over the control run after about 22 hours. Similar to the WRF simulations, a large wind speed improvement was initially noted in both the ISAN and R3DV simulations. The initial improvement was, however, lost quickly as the control and ISAN simulations were nearly identical and R3DV showed a small degradation after 28 hours.

Overall, the RAMS results were better than the WRF findings with much smaller overall biases noted in all three fields. RAMS improvements by the 3DVAR initialization were only found early in the simulation period and actually showed a degradation during the later periods for all fields. Meanwhile, the ISAN initialization showed improvement throughout for temperature and dew point, and about equal quality compared to the control run for wind speed. Note that the large error in initial wind speed in the control run was mostly due to the vertical interpolation from the GFS fields. The same feature can be seen in the WRF statistics.

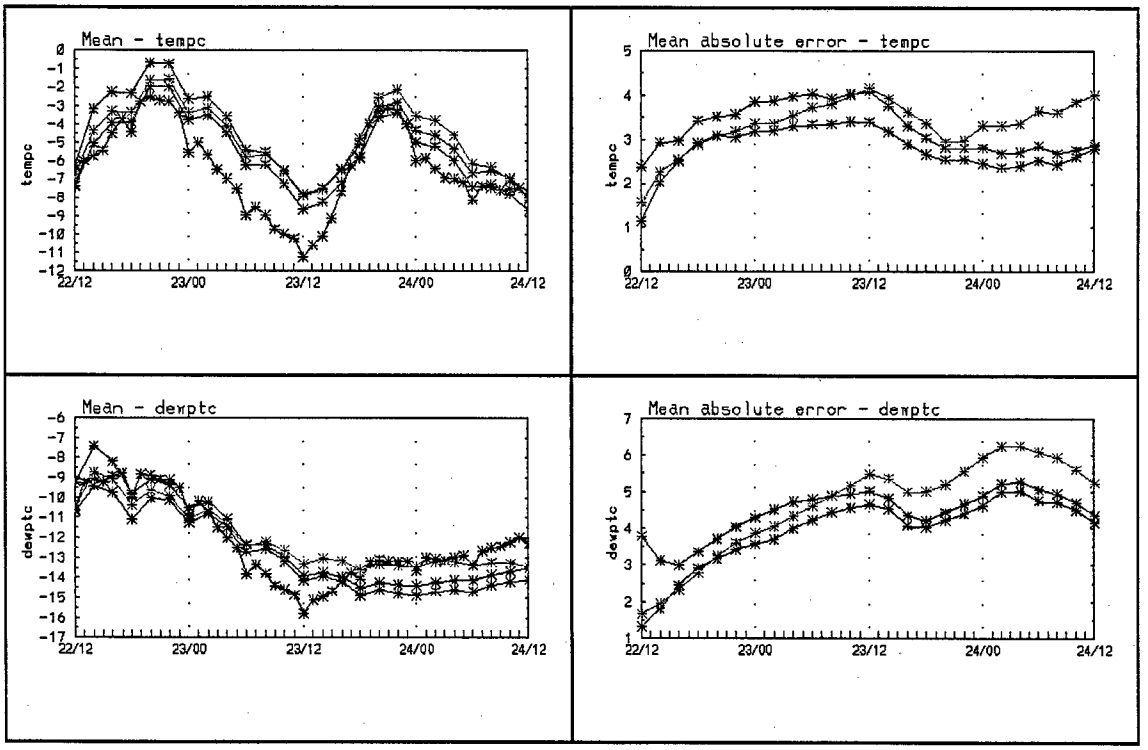

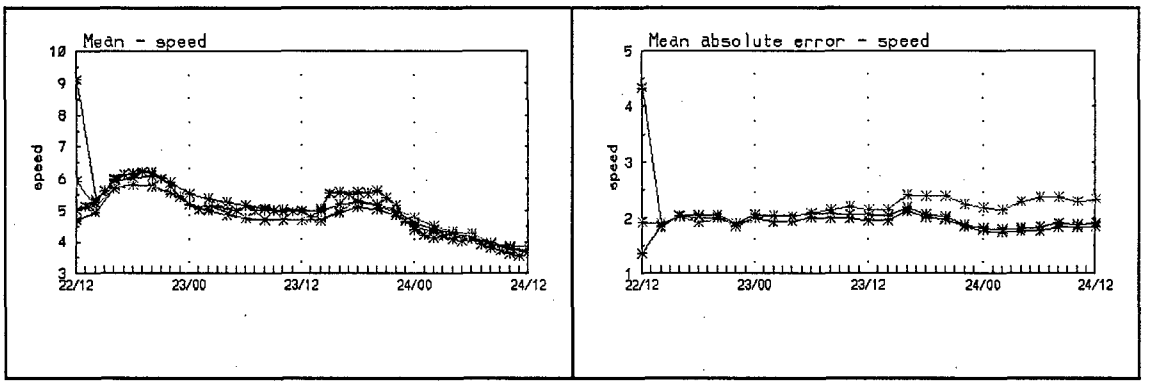

Figure 10: Domain-averaged statistics for temperature (C), dewpoint temperature (C), and wind speed (m/s) from the 3 RAMS runs. The left column contains the domain-averaged quantities; the right column is the mean absolute error. Red line – average of the observations; blue lines – RCTL run results; brown lines – ISAN run results; green lines – R3DV results.

3.2.4.2 Time series analysis

Time series analysis of RAMS pressure at individual locations were examined and compared to results from the WRF simulations. Figure 11 shows forecast surface pressure time series at the same representative locations examined in the WRF discussion (Figure 6). Two-minute output increments were only produced out to six hours with RAMS, which is less than the 24 hours completed with WRF. Nonetheless, the very high frequency oscillations noted in the WRF simulations were conspicuously absent in the RAMS runs. There was a suggestion of other middle frequency oscillations that were apparent in the ISAN and R3DV-initialized runs, but not the control runs, indicating imbalances caused by the inclusion of observations in the analysis. The R3DV time series also showed a significant drop in pressure compared to the control and ISAN simulations. Domain-averaged results (Figure 12) also show this pressure drop. The middle frequency oscillations are, however, not apparent in the domain-averaged time series, in contrast to results from WRF (Figure 7).

These pressure differences could have been caused by a difference in the way RAMS and WRF were run. RAMS used the WRF 3DVAR output to generate "varfiles" at 12 hour intervals for the 72 hours of simulations. Therefore, the RAMS' boundary conditions included the effects of the observations and the 3DVAR analysis. However, the WRF code has no capability to use the 3DVAR analysis for the lateral boundary conditions. Only the first-guess fields (GFS in this case) are able to be used.

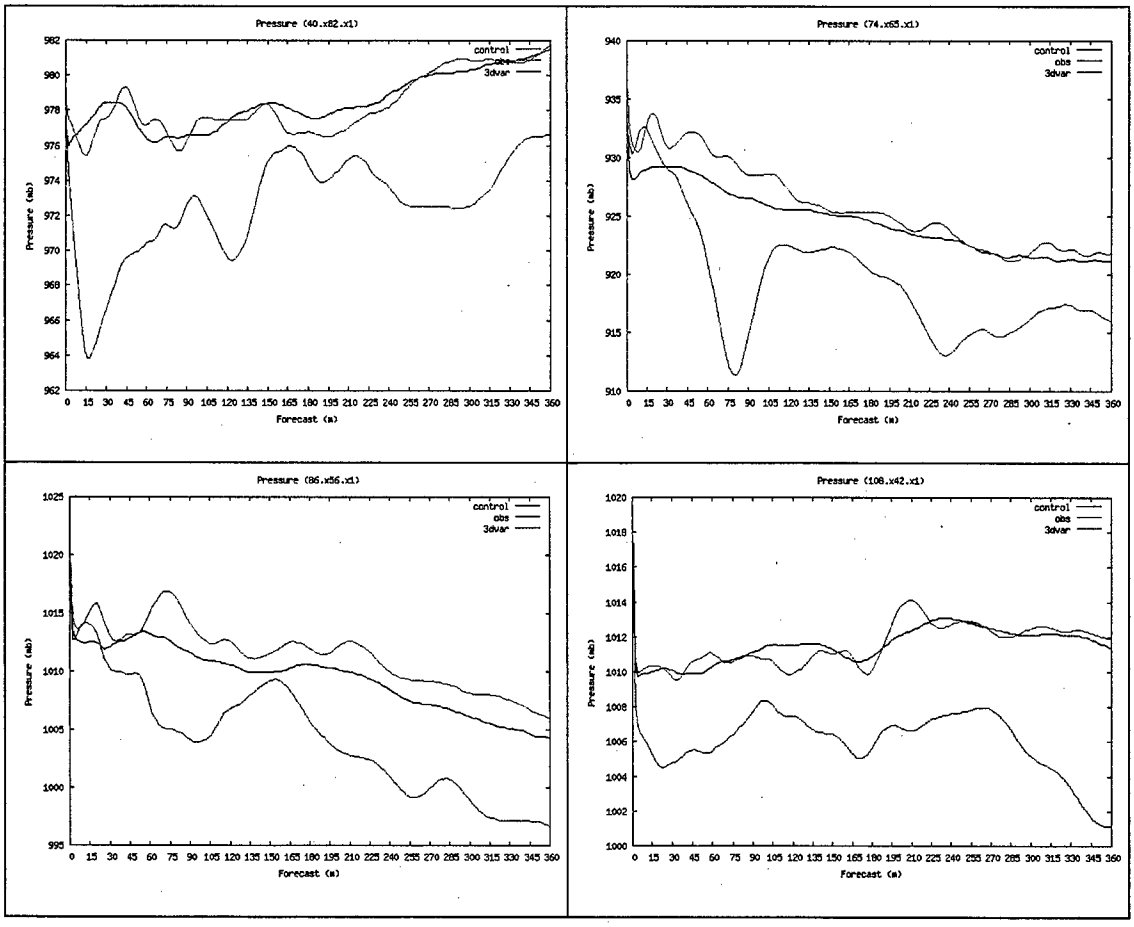

Figure 11: Surface pressure (mb) time series from the RAMS control run (red), ISAN run (green), and R3DV run (blue) at four representative grid points: 1) central Iowa, 2) mountainous terrain along the Virginia-West Virginia border, 3) the North Carolina coast, and 4) the Atlantic Ocean.

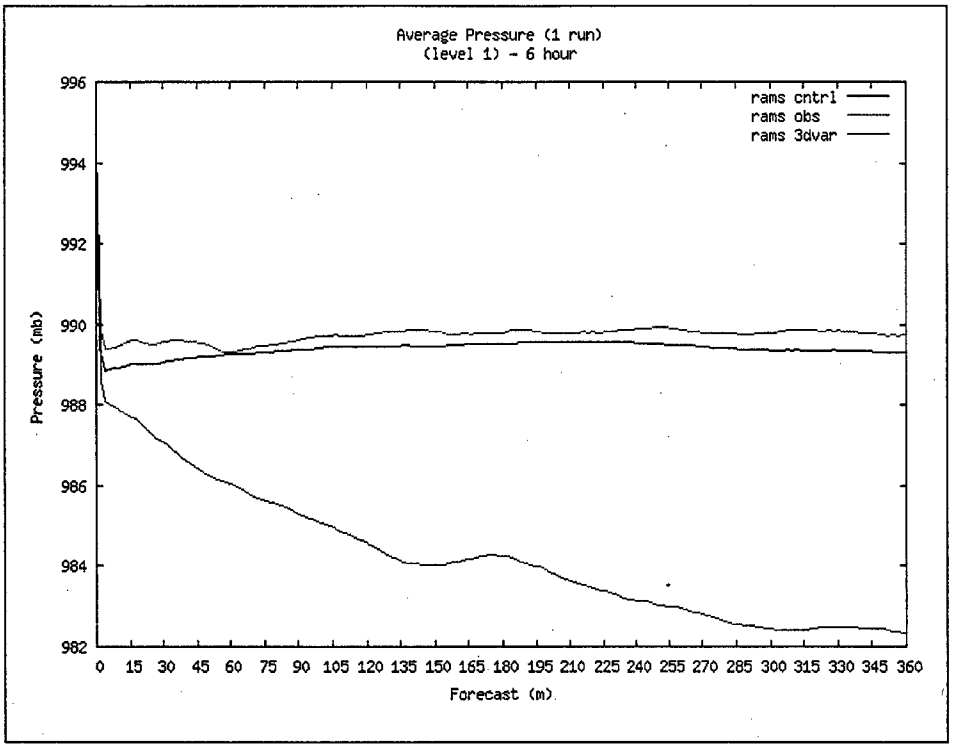

Figure 12: Domain-averaged surface pressure (mb) from the RAMS control run (red), ISAN run (green), and R3DV run (blue).

3.2.4.3 Fourier time series analysis

Domain-averaged, 6-hour power spectrum results were generated using the same methodology as described for the WRF runs, with the exception that a smaller subset of equally spaced grid points were used to calculate the domain average. Results (Figure 13) were similar to the WRF findings with consistent decreasing influence from low to high wavenumbers in both the control and R3DV simulations. Also similar to WRF, the R3DV results indicated a greater influence of low wavenumbers when compared to the control runs, and a crossover occurred at wavenumber 9 (compared to wavenumber 10 in WRF) after which the higher wavenumbers show less influence in the R3DV forecasts. As suggested previously, the 3DVAR initialization acted to remove some of the higher frequency variations, which may be a desirable effect for this resolution of simulation. The ISAN results indicate power spectrum characteristics that were nearly identical to R3DV at higher frequencies above wavenumber 23. Some differences were, however, observed at lower wavenumbers where ISAN shows lower, nearly identical to control, power at and below wavenumber 4, then higher power through the middle wavenumbers, and a higher wavenumber crossover point (wavenumber 17) to lower power at the high wavenumbers.

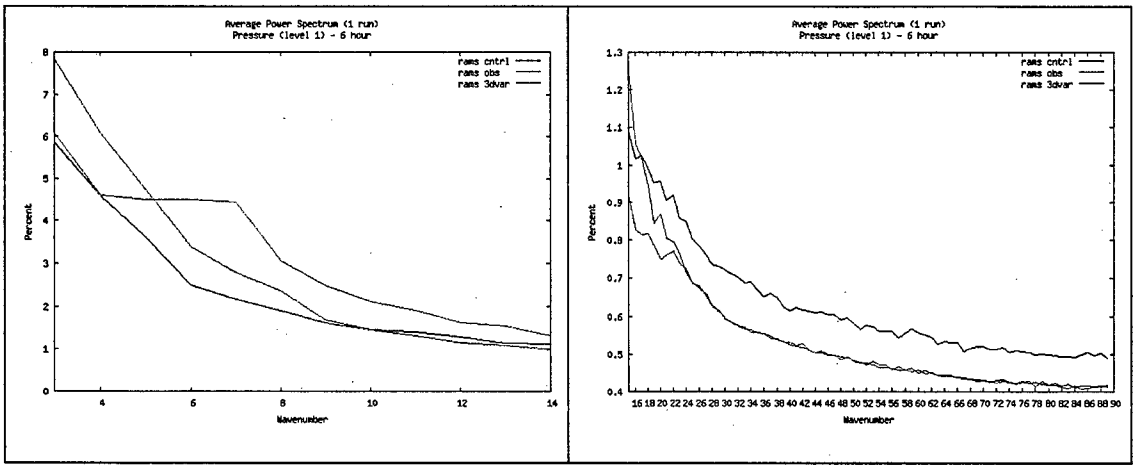

Figure 13: Average power spectrum from the RAMS control run (red), ISAN run (green), and R3DV run (blue) (see text for details). Left panel shows wavenumbers 3-14, right panel shows 15-90.

3.3 Sea breeze case

3.3.1 Meteorological summary

A second summer case study was selected to contrast with the first scenario. This was a sea-breeze scenario over the Florida Peninsula where synoptic-scale forcing was minimal. Figure 14 shows 500 and 1000 mb GFS analyses from 0000 UTC 18 June through 0000 UTC 19 June 2004. A weak subtropical high was evident at 500 mb that was centered over the Florida Panhandle through the period. Light easterly flow south of the high covered the entire Peninsula. Near the surface, anticyclonic circulation was evident over the Atlantic Basin with weak on-shore flow along the east coast of Florida. The spatial resolution of the GFS analyses was insufficient to resolve any sea-breeze features, but a well-defined sea-breeze circulation surely occurred during this weak synoptic forcing period.

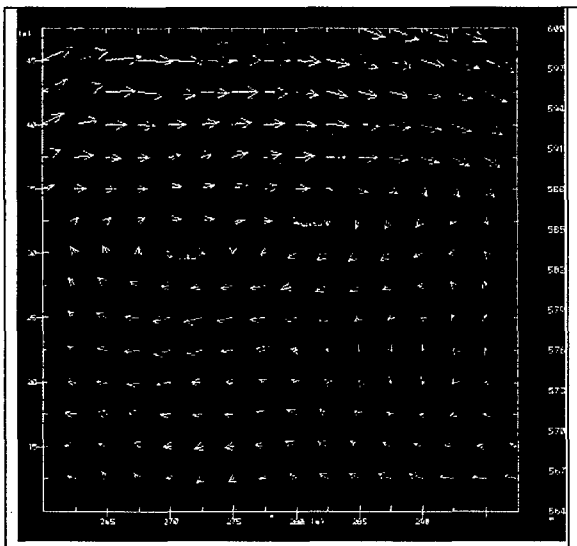

0000 UTC 18 June 2004

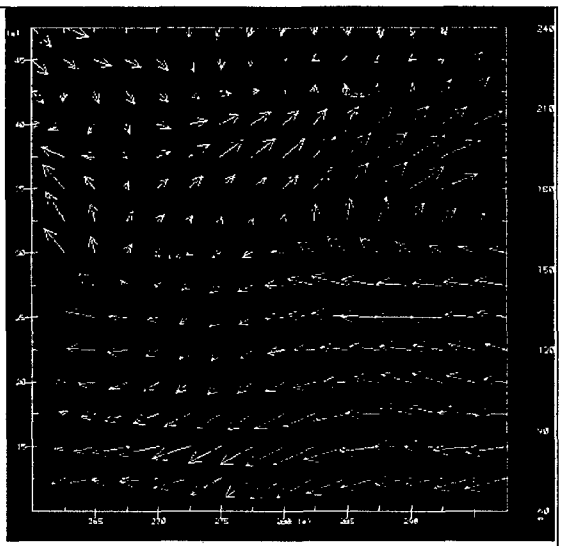

0000 UTC 18 June 2004

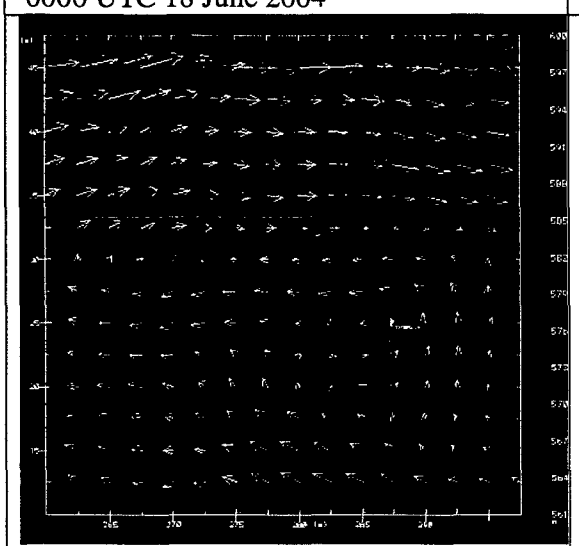

1200 UTC 18 June 2004

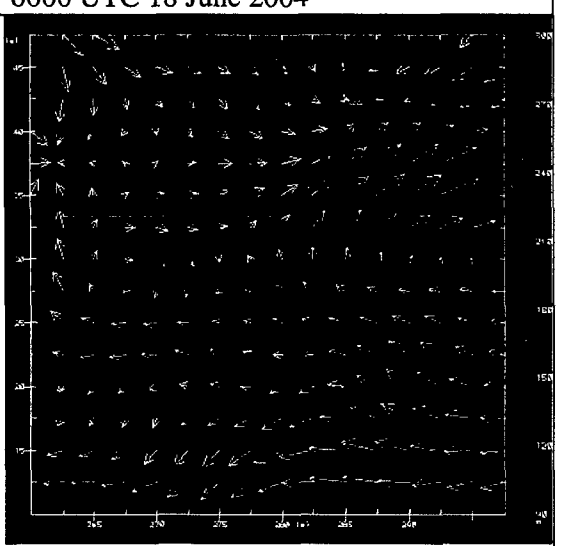

1200 UTC 18 June 2004

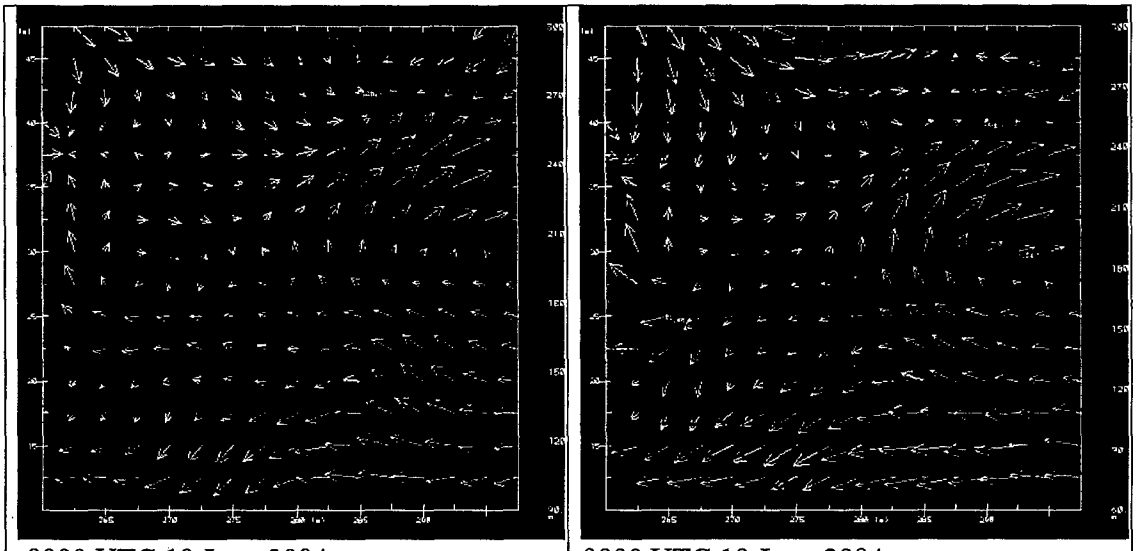

Figure 14: Left column shows 500 mb height (black line, contour interval 30 m) and wind vectors and right column shows 1000 mb height (black line, contour interval 30 m) and wind vectors at every model grid point for 0000 UTC 18 through 0000 UTC 19 June 2004.

3.3.2 Model configuration

3.3.2.1 WRF Configuration

A smaller domain with much higher resolution was selected to cover the entire Florida Peninsula (Figure 15). Initially, a double-nested grid configuration was initially used with a 2-km inner grid centered over Cape Canaveral. WRF-3DVAR simulations, however, consistently failed with this configuration despite using a variety of model configurations. The failure was determined to be in the radiation scheme, but a fix was not able to be found in a reasonable time. Hence, a new single grid system was implemented with a 3 km horizontal grid spacing. The grid size is 201 x 201 with 40 vertical levels. The WRF-SI package used 30 second datasets to generate static fields such as land use type for the model domain. Since sea-surface temperature is very important to correctly simulating a sea-breeze circulation, archived NCEP ETA model data was obtained from the NOAA Forecast Systems Laboratory. This dataset includes sea surface temperature and soil temperature and moisture, none of which is available in the NCAR GFS dataset. The WRF-SI package was used to interpolate the ETA data to the WRF model grid.

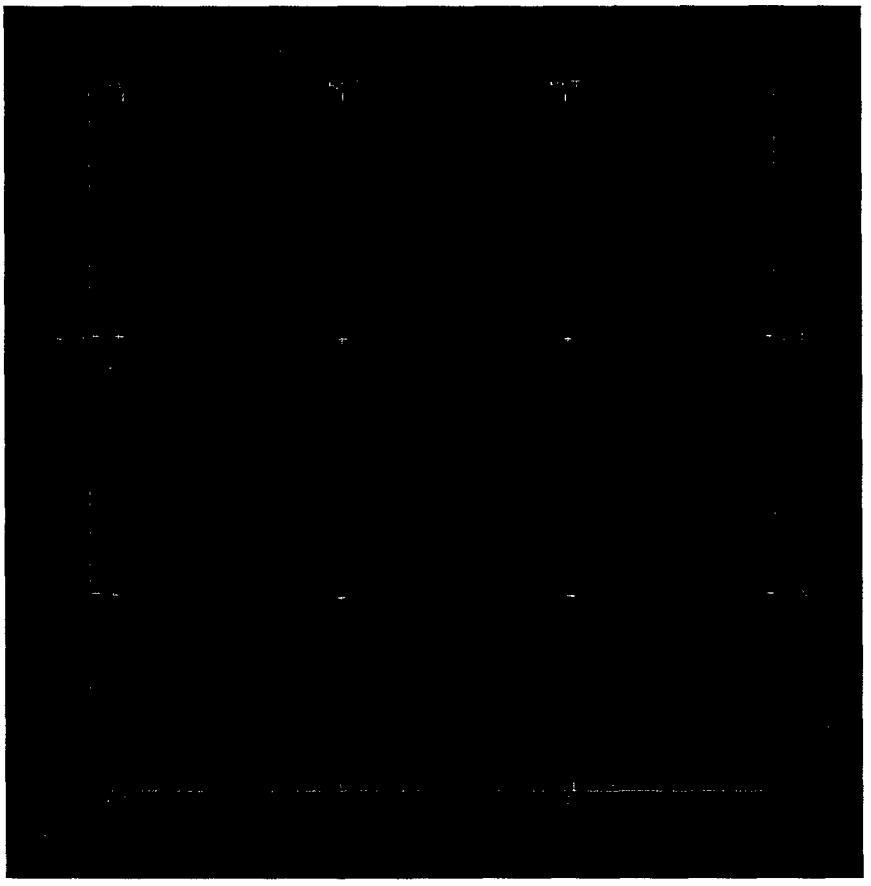

Figure 15: Summer case study model domain.

The experiment design was similar to the first case study with two sets of model simulations – control and 3DVAR. The control runs were initialized with interpolated ETA data as output by the WRF-SI package. Initialization of the 3DVAR runs was accomplished using grids derived from the WRF-3DVAR package. The package was configured to use the WRF-SI ETA grids generated for the control runs as the background field. AFTAC point observation data were not available for this time period, hence archived METAR data was obtained from NCAR and code was developed to convert this data into a format useable by the WRF-3DVAR package. This dataset was then assimilated directly into the ETA background fields to form the 3DVAR initial conditions. The WRF model forecast boundary conditions were adjusted accordingly.

Two WRF simulations, control and 3DVAR, were initially completed that covered a full diurnal cycle of 30 hours at one hour increments from 0000 UTC 18 June through 0600 UTC 19 June 2004. Four additional runs (control and 3DVAR) were generated to provide further comparison data. These simulations were initialized at 6-hour increments (0000, 0600, 1200, and 1800 UTC 18 June) and output was generated at 2-minute increments out to 6 hours. WRF model physics remained the same as in the first case study (Table 1), except the convective parameterization was not used.

3.3.2.2 RAMS Configuration

RAMS version 6.0 was again configured very similarly to WRF for these runs. The model was configured with 199x199x38 grid points with a horizontal grid spacing of 3 km. The various physics options used were the same as for the winter case (Table 2), except the convective parameterization was not run.

Again, three RAMS runs were made for this case, varying the details of the definition of the initial and boundary conditions:

4) RCTL ("Control") run – straight interpolation of the ETA files
5) ISAN run – RAMS/ISAN used to blend observations with ETA first-guess fields
6) R3DV run – WRF 3DVAR files interpolated to RAMS grid

3.3.3 WRF verifications

Figure 16 shows the initial surface temperature analysis from the two WRF runs. Very little difference can be seen between the two analyses, as the ETA first-guess fields also would have assimilated the same observations, albeit at a coarser resolution.

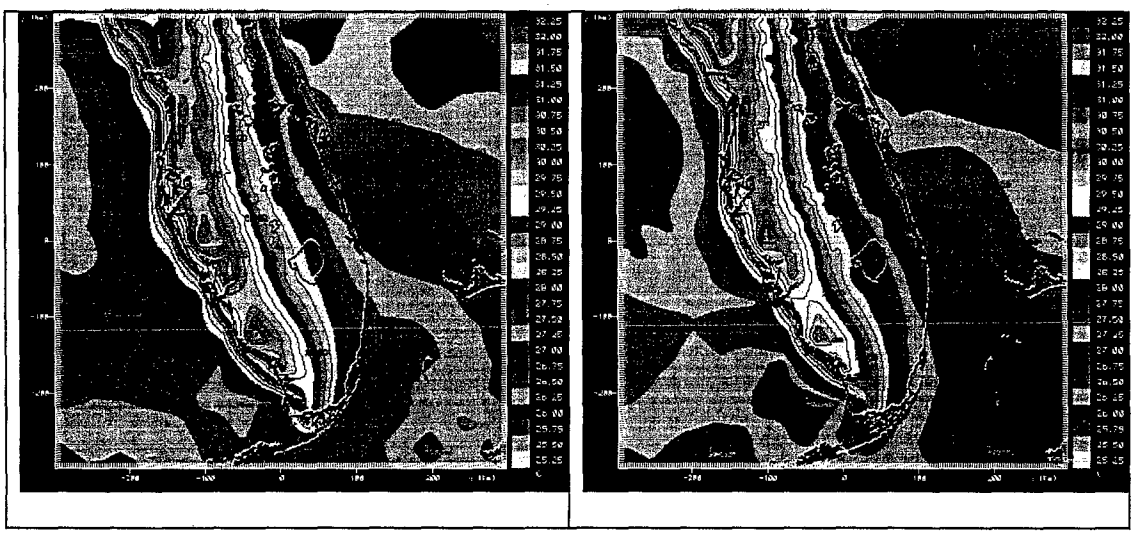

Figure 16: Initial surface temperature from WRF control run (left) and 3DVAR run (right). Contour interval is 0.25° C.

3.3.3.1 Statistical evaluation

WRF gridded forecasts were interpolated to all available METAR observation locations within the model domain (approximately 45 surface stations) for direct comparison.

Figure 17 illustrates domain-wide average forecast temperature, dew point, and wind speed from the control and 3dvar simulations along with observations. Mean absolute error statistics are also displayed.

For each variable, the 3DVAR-initialized simulation showed a small improvement over the control run at the initial time. The small improvement lasted for about nine hours after which time the quality of the control and 3dvar simulations were nearly identical. As with the WRF winter case study results, both the control and 3DVAR runs showed a significant warm temperature and dew point bias especially during the night-time hours. WRF wind speed predictions also showed a high night-time bias, and then wind speed values dropped during the daylight hours at a time when daytime mixing would suggest higher wind speeds that were seen in the observations. Similar inverted wind speed diurnal cycles have been noted in several MM5 projects conducted by ATMET. Overall, any slight improvements resulting from 3DVAR initializations were small.

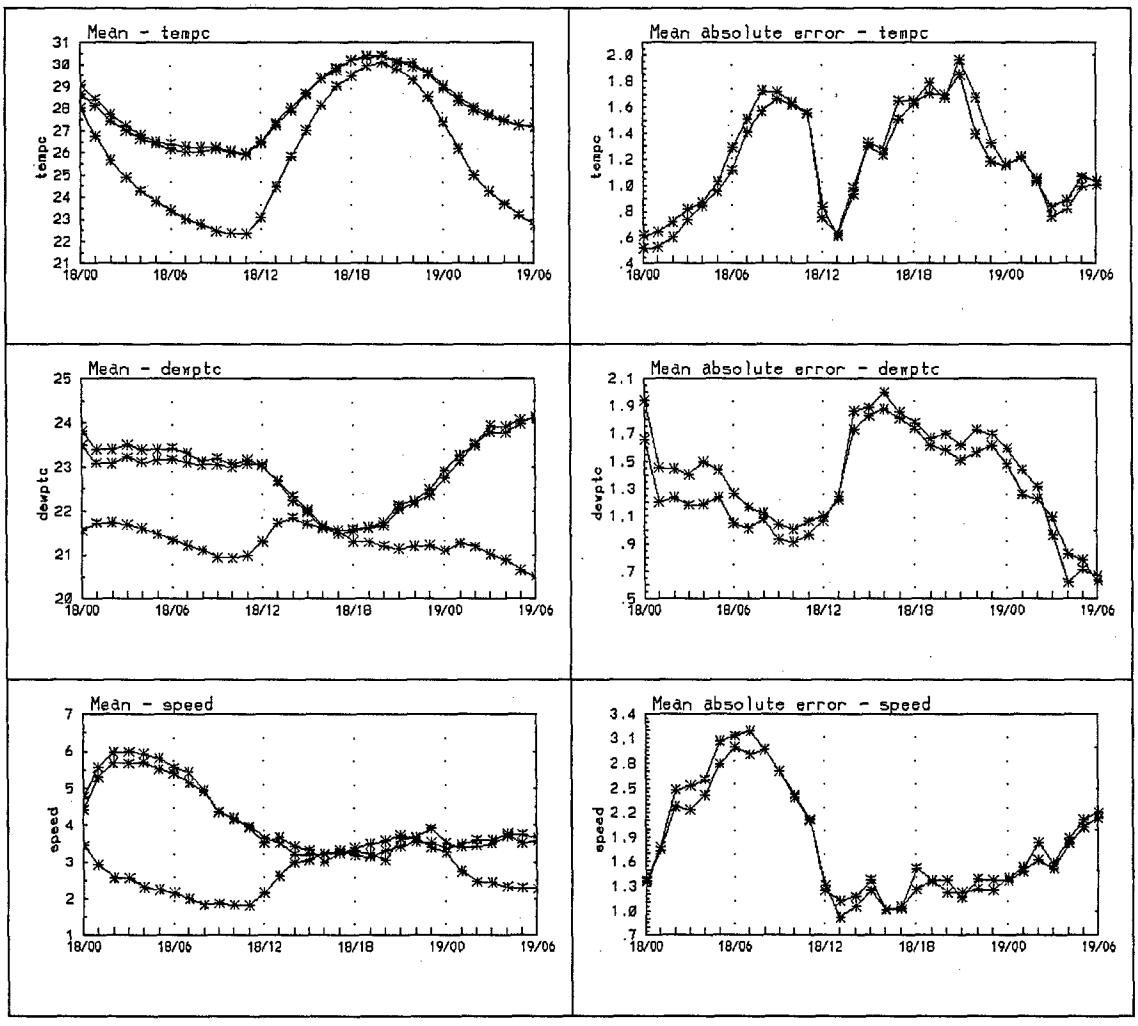

Figure 17: Domain-averaged statistics for temperature (C), dewpoint temperature (C), and wind speed (m/s). The left column contains the domain-averaged quantities; the right column is the mean absolute error. Red lines – average of the observations; blue lines – control run results; brown lines – 3DVAR results.

3.3.3.2 Time series analysis

Time series of WRF surface pressure forecasts at individual locations were examined and compared to results from the winter case study simulations. Simulation results were similar for each of the four different model initialization times with 2-minute increment forecast output; hence, results are only presented from the runs initialized at 0000 UTC 18 June. Figure 18 shows four representative surface pressure time series, positioned at 27N latitude and at 82, 81, 80, 79W longitude, superimposed on the same graph. Significant differences were not easily discerned between the control and 3DVAR simulations. Each time series indicated some small amplitude, high frequency oscillations. Also, a consistent, larger amplitude oscillation with a time period of about 50 minutes was evident in all the simulations and the amplitude dampened with time.

Domain-averaged surface pressure time series using a subset of equally spaced grid points are shown in Figure 19. Even in the domain-averaged time series, the 50-minute frequency oscillation continued to be very obvious in both the control and 3dvar runs. The period of these waves suggest that the oscillation results from a perturbation in the vertical, perhaps a reflection off the model top.

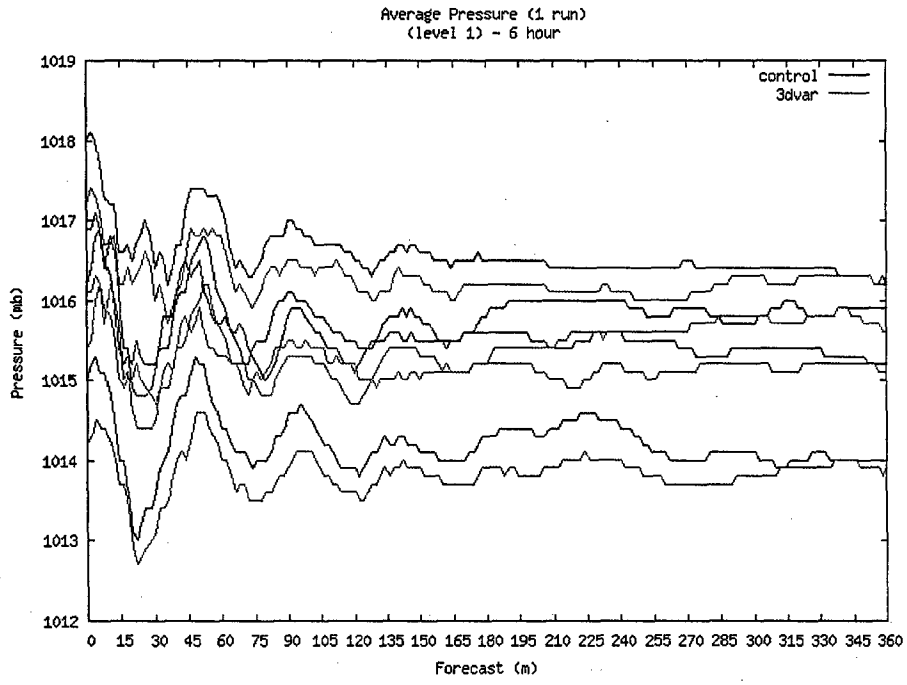

Figure 18: Surface pressure (mb) time series from the WRF control run (red) and 3DVAR run (green) at four representative grid points.

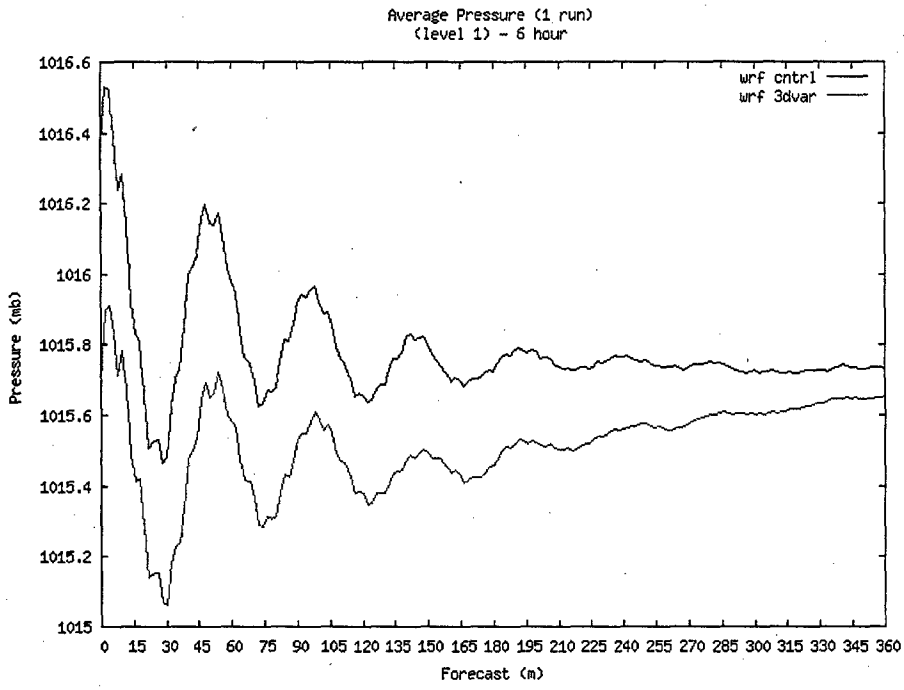

Figure 19: Domain-averaged surface pressure (mb) for the control (red) and 3DVAR (green) runs.

3.3.3.3 Fourier time series analysis

Domain-averaged (using a subset of equally spaced grid points), 6-hour power spectrum results (Figure 20) were generated using the same methodology as used in the winter case study. In general, as in the winter case study, the influence of higher wavenumbers decreased consistently in both the control and 3DVAR simulations with three notable exceptions. Significant power spectrum spikes are observed at wavenumbers 8 and 50, and to a lesser extent at wavenumber 12. Wavenumber 8 over a 6-hour time span relates to about a 45 minute period, which corresponds closely to the 50 minute oscillation period noted in the previous section. It is not entirely clear what the wavenumber 50 oscillation represents. This oscillation can be seen as the small, superimposed perturbations in Figure 19, and may possibly due to the frequency of the radiation scheme updates. When compared to the control simulation, 3DVAR results show a somewhat smaller influence between wavenumbers 5 and 11, a somewhat greater influence below 5 and between 11 and 17, and very little difference above wavenumber 17.

The purpose of the time series analysis is to attempt to determine whether there are imbalances in the initial data analysis that may cause perturbations at the initial time, and

also throughout the simulation if the analysis fields are used in the 4DDA nudging schemes. Unfortunately, because of the status of WRF development (i.e., nudging schemes not implemented, generation of domain-wide perturbations), it is difficult to arrive at a definitive conclusion based on the WRF results at this time.

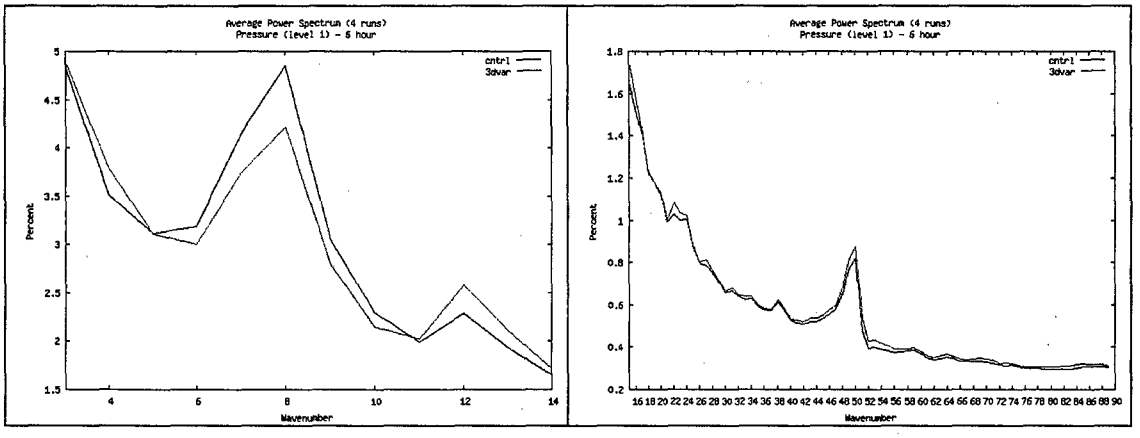

Figure 20: Average power spectrum over the control (red) and 3DVAR (green) runs (see text for details). Left panel shows wavenumbers 3-14, right panel shows 15-90.

3.3.4 RAMS verifications

Figure 21 shows the initial surface temperature analysis from the three RAMS simulations. Comparing the differences between the control and ISAN run, the effect of the Barnes scheme can be clearly seen again. Values from the land observations are being extended over the data poor region of the ocean. The distance that they extend is dependent on the smoothing parameters specified in the RAMS' configuration.

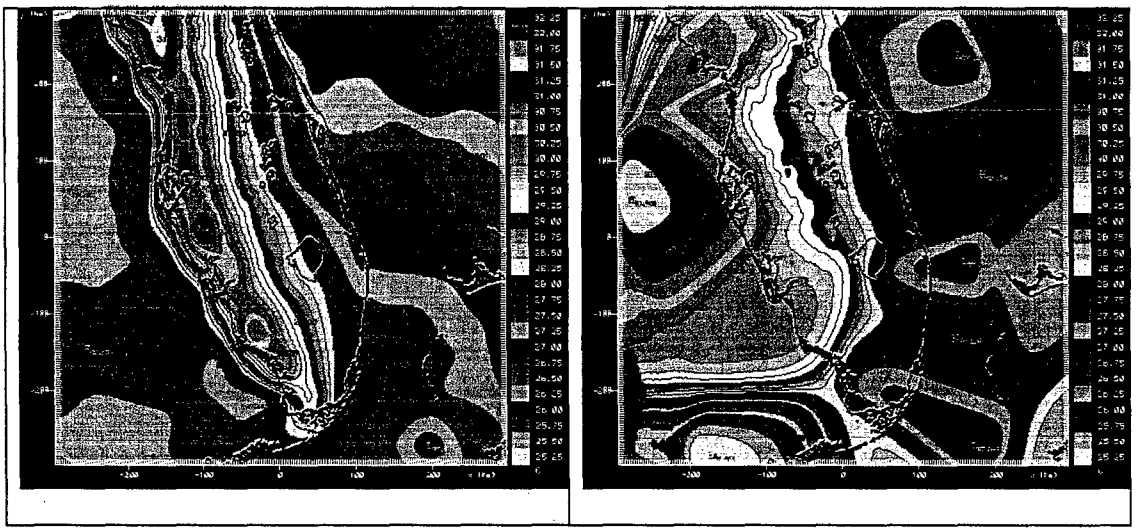

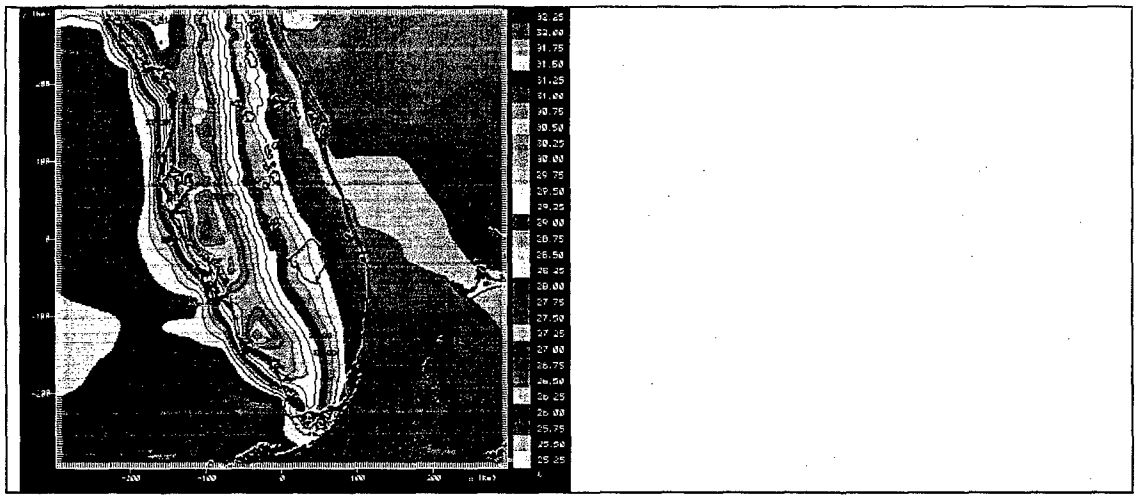

Figure 21: Initial surface temperature from RAMS control run (upper left), ISAN run (upper right), and R3DV run (bottom). Contour interval is 0.25° C.

3.3.4.1 Statistical evaluation

RAMS gridded forecasts were interpolated to all available METAR observation locations within the model domain for direct comparison. Domain-wide average observations and forecasts of temperature, dew point, and wind speed are shown in Figure 22. Mean absolute error statistics are also displayed.

The results are qualitatively similar to the winter run. At the initial time, the ISAN run tended to verify the best for all variables. The 3DVAR analysis pushed the first-guess field closer to the observations as expected, but not as close as the ISAN run. Thereafter, the differences among all runs reduced, especially with wind speed verifications. The mean temperature was virtually the same for all runs, implying that the land surface processes was mostly driving the surface temperatures. The dewpoint temperature showed the most differences. The ISAN run was initialized closest to the observations and maintained that advantage through the nighttime hours. The R3DV run started with a smaller deviation from the first-guess field, and by the early morning hours, had almost the same mean dewpoint as the control run.

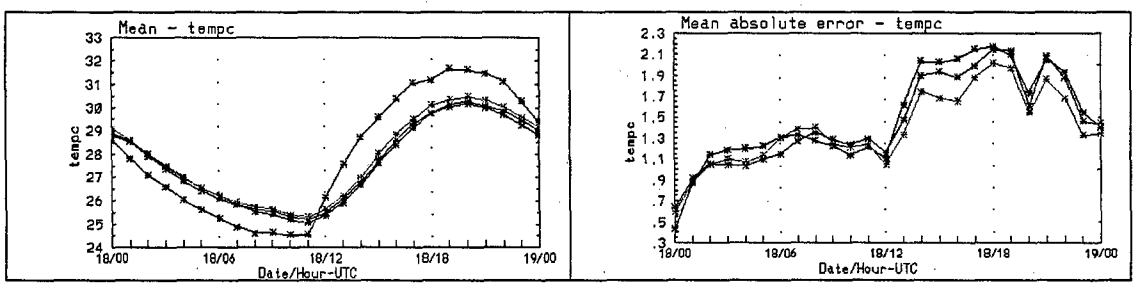

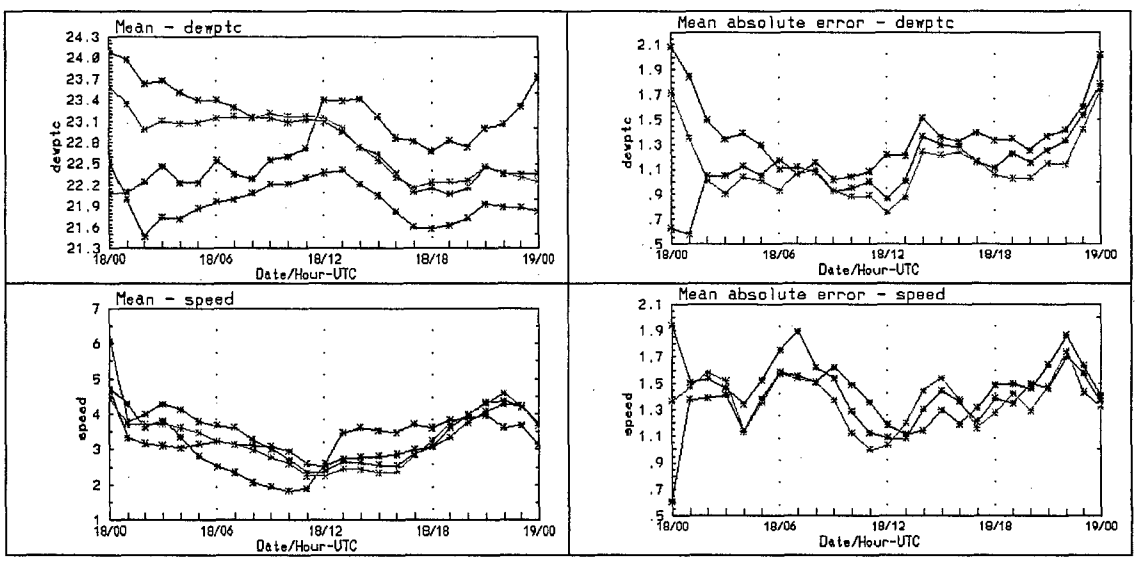

Figure 22: Domain-averaged statistics for temperature (C), dewpoint temperature (C), and wind speed (m/s) from the 3 RAMS runs. The left column contains the domain-averaged quantities; the right column is the mean absolute error. Red line – average of the observations; blue lines – RCTL run results; brown lines – ISAN results; green lines – R3DV results.

3.3.4.2 Time series analysis

Time series of RAMS surface pressures at individual locations were examined and results suggest no significant periodic oscillations at any frequency (not shown). Domain-surface pressure time series, averaged with the same subset of equally spaced grid points used in the WRF analysis, for the control, ISAN, and R3DV simulations are shown in Figure 23. Again, no significant periodic oscillations at any frequency are observed, in contrast to the oscillations noted in the comparable WRF simulations.

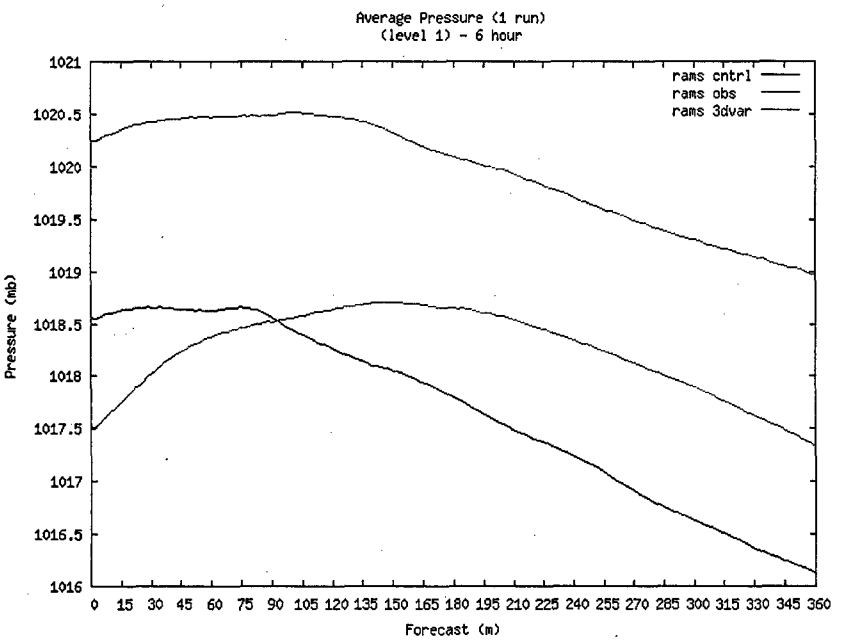

Figure 23: Domain-averaged surface pressure (mb) from the RAMS control run (red), ISAN run (green), and R3DV run (blue).

3.3.4.3 Fourier time series analysis

Domain-averaged, 6-hour power spectrum results (with the same subset of equally spaced grid points used in the WRF analysis) were generated using the same methodology as previously described and are shown in Figure 24.

Results are similar to the RAMS findings from the winter case study with consistent decreasing influence from low to high wavenumbers in both the ISAN and R3DV simulations. The control run has a small increase from wavenumber 4 to 5, but by wavenumber 6, all simulations are very similar. The R3DV results do indicate a greater influence of low wavenumbers when compared to the control runs, and a crossover occurs at wavenumber 8 (compared to wavenumber 9 in the RAMS winter case) after which the higher wavenumbers show less influence in the R3DV results. As suggested previously, the R3DV initialization has acted to remove some of the higher frequency variations. Also, consistent with the RAMS time series analysis, there is no indication of any significant power spectrum spikes that were observed in the WRF simulations. ISAN results indicated power spectrum characteristics that are nearly identical to R3DV at higher frequencies above wavenumber 5 while the ISAN influence is slightly greater than the control at very small wave numbers.

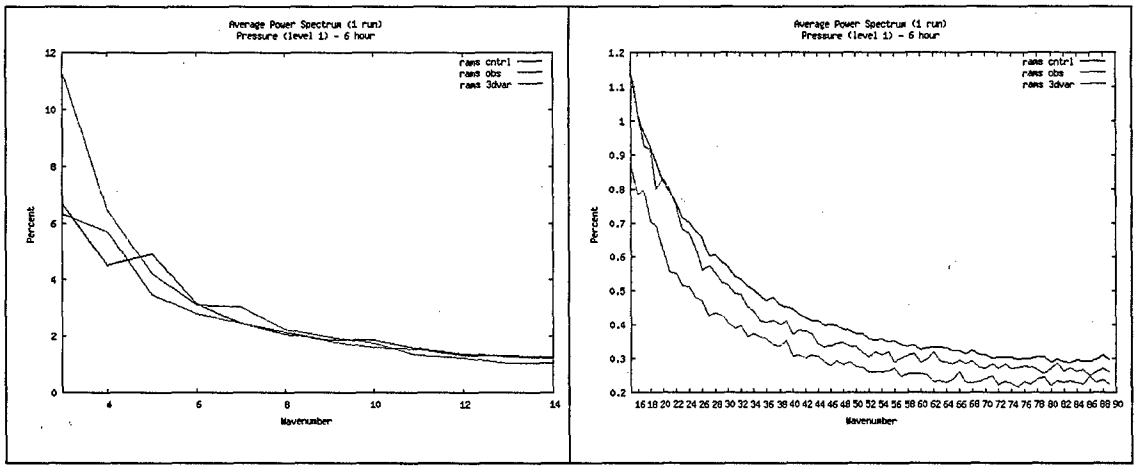

Figure 24: Average power spectrum from the RAMS control run (red), ISAN run (green), and R3DV run (blue) (see text for details). Left panel shows wavenumbers 3-14, right panel shows 15-90.

3.4 Case study summary

Based on the limited number of cases investigated, we can make the following observations about the WRF 3DVAR and its performance with the WRF and RAMS models, along with the WRF model itself:

- WRF 3DVAR improved the statistics of the initial conditions compared to using the first-guess only fields, more so for the coarser grid resolution, winter case than for the higher resolution summer case. This was probably due to the fact that the summer case used the ETA first-guess field, which also uses a 3DVAR scheme that is similar to the WRF scheme. We expected to see some additional improvement due to using higher resolution, but this was not the case.
- The standard RAMS/ISAN initial conditions also showed initial statistical improvement compared to the first guess field, somewhat better than the 3DVAR initial conditions. However, the 3DVAR results were more consistent over data-poor regions than the Barnes objective analysis in RAMS/ISAN.
- Significant improvement in the simulation verifications, in general, only lasted a few hours in both WRF and RAMS. This is consistent with the findings of others using the WRF scheme (McAtee et al. 2005) and the predecessor MM5 3DVAR scheme (Barker et al. 2004). Barker et al. did show a potential improvement throughout a 24-hour forecast period for the u-wind component (they did not show the v-component) on higher resolution grids, but marginal improvements for moisture and temperature at all scales. They also showed that the LITTLE_R scheme (MM5's Cressman analysis) generally had better verifications at the initial time than the 3DVAR scheme.

- It was desirable to make comparisons in WRF comparing the 3DVAR performance versus a "standard" analysis, such as the Cressman technique used in MM5 (as in Barker et al. 2004). However, no capability to use actual observations in creating the initial conditions has been implemented in WRF, aside from the 3DVAR scheme. There is also no capability to use the 3DVAR fields for the lateral nudging boundary conditions.
- Both the WRF summer and winter simulations showed domain-wide pressure oscillations at the start of the model run. The oscillations were of similar magnitude with and without the 3DVAR. This implies that they are caused by boundary conditions or by the WRF numerics and makes perturbation analysis to separate out the effects of the 3DVAR scheme much more difficult.
- The imbalances in the initial fields, as measured by the "noise" in pressure time series at grid points, were not reduced in RAMS by using the 3DVAR fields. Since the 3DVAR cost function does include some balance terms, it was hoped that some noise could be reduced. However, since the terms are for geostrophic (good for large-scale flows) and cyclostrophic (good for hurricanes) wind and pressure balances, the utility of these terms for general mesoscale applications are somewhat limited.
- There is some evidence that the imbalances in RAMS for the 3DVAR initialization may have been somewhat increased in some locations for the winter runs (Figure 11), although there was no qualitative difference in the summer runs. Based on this and the verification results, while more testing should be done, we feel that the interpolation of the WRF 3DVAR fields to the RAMS grid is adequate, at least for this version. If additional balance terms are added, this will need to be reconsidered. More details about this issue are in Section 4.1.

It should be noted again that these observations are based on a limited set of simulations. We did not attempt to vary the "default" set of parameters distributed with the WRF 3DVAR scheme, as there is virtually no documentation or guidance available at this time. Neither did we attempt to change the error matrices, as no code was available to generate other than the default set (the newer WRF version, 2.1, has a program to be used if a history of simulations is available).

While we expect our simulations to be indicative of general performance, especially since they are consistent with other tests in the literature, more tests and cases should be run to verify any conclusions.

4 WRF 3DVAR Scheme Issues

4.1 Use of WRF 3DVAR with RAMS

One of the desires of this research is to have a generalized 3DVAR scheme that can be used for both the WRF and RAMS models. While the WRF scheme should, of course, be compatible with the WRF model, it was less clear how to best utilize the WRF 3DVAR scheme with RAMS. There are three possibilities for the implementation of 3DVAR for WRF and RAMS:

1) Modify the WRF 3DVAR scheme to execute on the RAMS grid structure,
2) Execute the WRF 3DVAR scheme in its current form, then interpolate the results to the RAMS grid structure, or
3) Develop and implement a separate 3DVAR scheme specifically for RAMS.

The first or third option should provide the best results from a technical standpoint. For example, a term which minimizes three-dimensional divergence could be included in the cost function. However, the computation of divergence depends on the actual specific model grid structure. If divergence were minimized on the WRF 3DVAR Arakawa-A grid, then the velocity components were interpolated to the RAMS grid, the divergence may no longer be minimized correctly.

From a practical view, however, the second option has some advantages in that significant modifications to the WRF code will not need to be done. Also, as changes to the WRF scheme might be made by NCAR and others, the interface to RAMS can immediately take advantage of the potential improvements.

As noted in Section 3, for the runs that we have performed herein, we recommend the second option, especially since the WRF 3DVAR implementation is not performed on the native WRF grid anyway. As noted above, this interpolation introduces additional error that is not subject to the minimization procedure of the variational cost function. Interpolation of the optimized initial state to the RAMS grid introduces the same type of error, but this procedure is no less valid than it is for WRF itself.

While it would be possible to implement a separate scheme for RAMS, we feel at this time, since the WRF scheme is under active funding and development, there would be a significant amount of "re-inventing the wheel". We feel it would be better to monitor and use the current and near future developments of the WRF scheme.

Aside from the difficulties in applying 3DVAR in forensic applications, the question still remains as to whether the current development of WRF 3DVAR will meet all of AFTAC's requirements. The current scheme does handle a number of observational platforms, but there may be additional observed variables that are available to AFTAC

(such as satellite radiances) that NCAR will not be implementing. Therefore, we have reviewed the WRF 3DVAR code to assess the difficulties in implementing new observed variables. The following section describes briefly the overall WRF structure and the 3DVAR code.

To summarize our recommendation concerning this question: while the WRF code structure presents some challenges (as noted below) and requires a significant learning curve for many types of implementations, it will be possible to implement new variable types if necessary.

4.2 WRF/3DVAR Code Structure

The WRF 3DVAR scheme has been implemented as one of the WRF "cores". As of this writing, there are three main WRF cores:

- ARW (Advanced Research WRF) – (formerly, the "mass"-core and the "EM"-Eulerian Mass core) – primarily developed by NCAR as the follow-on to MM5. A floating sigma-p vertical coordinate is used again (as in MM4). Most of the physics packages are from MM5.
- NMM (Non-hydrostatic Mesoscale Model) – began development by NCEP (Janjic) as follow-on to the ETA model prior to the WRF concept. It uses a sigma-z vertical coordinate, with most of the physics packages coming from the ETA model.
- 3DVAR (3-Dimensional Variational Scheme) – data assimilation scheme originally developed for MM5 and converted to the WRF code structure.

Following is a short summary of the WRF software structure.

The WRF structure was primarily designed by John Michalakes (Michalakes et. al. 2001) at NCAR. He was also primarily responsible for the two incarnations of the parallelism in MM5 while at Argonne in the 1990's. Figure 25 shows two slides that were taken from the now-standard WRF presentation that has been presented at numerous conferences over the past many years. The slides summarize the design considerations that the developers considered.

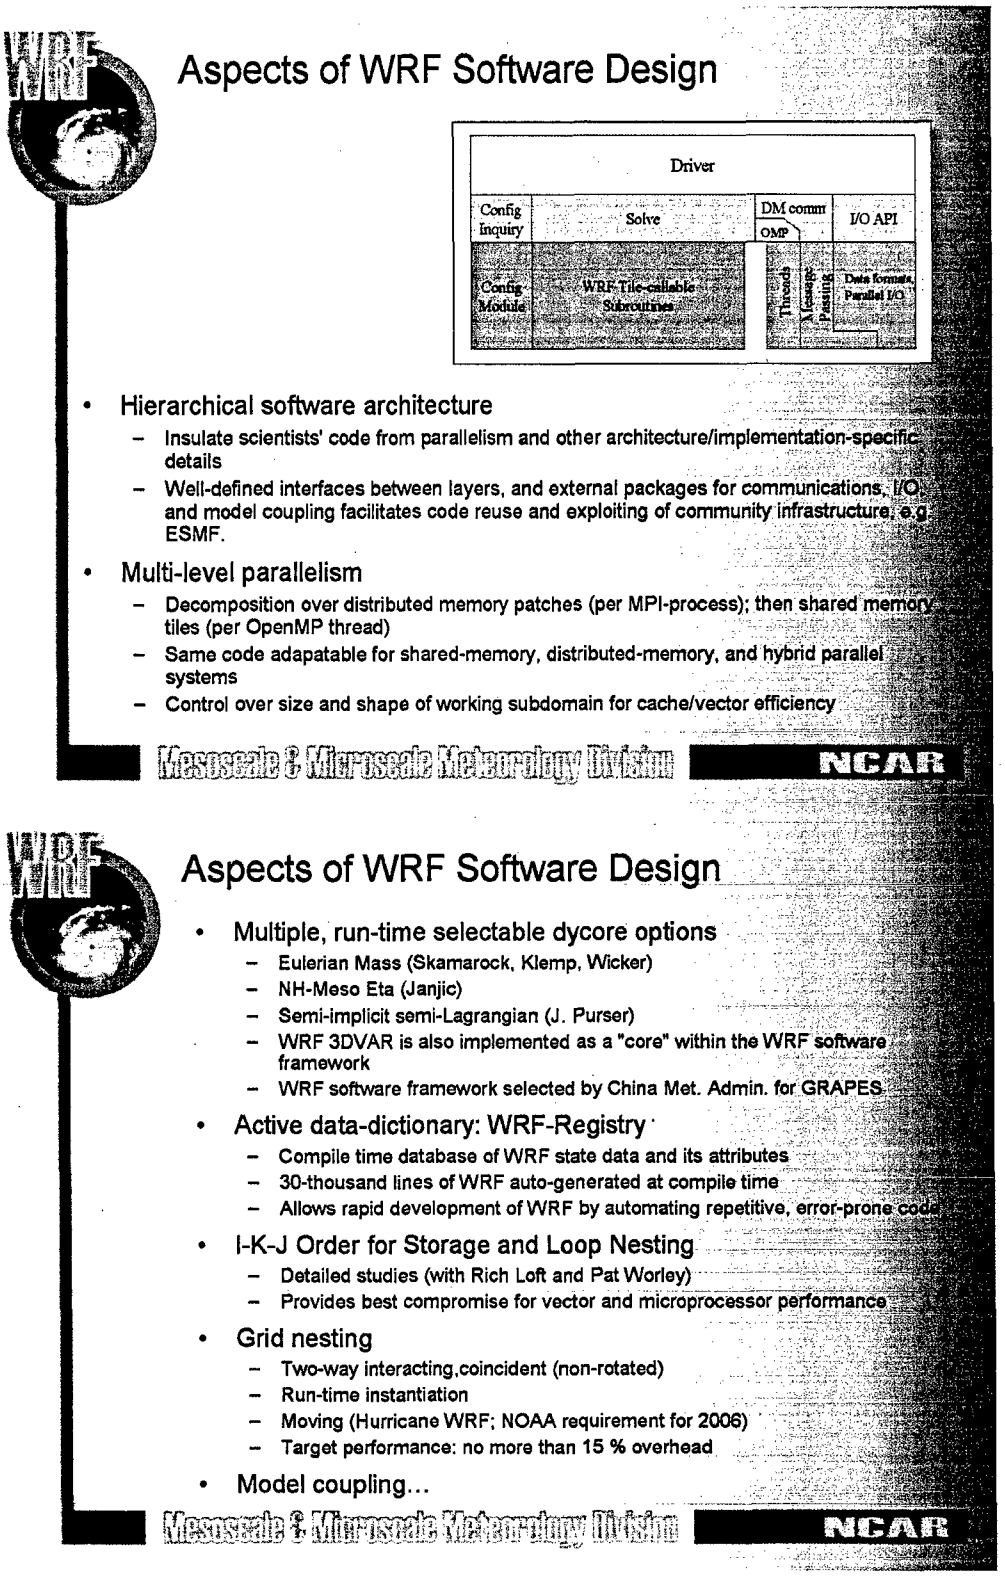

Figure 25: From NCAR/MMM WRF presentation available from their web site.

After examining the WRF code, there were a few of these design points that seemed to be the most important to the developers. In our opinion, these were:

- "Insulate scientist's code from parallelism..."
- "Multiple, run-time selectable dycore options"
- "WRF-Registry"

The WRF memory structure is built around the concept of the "Registry". The concept is rather similar to the RAMS' VTABLE file that was used for a few versions, starting with v4a in about 1992 and ending with v4.4. As with VTABLE, the WRF Registry is a data file which contains a list of all variables that will be allocated for an execution, along with several characteristics for each variable such as order of dimensions, inclusion in the message-passing, etc. The Registry file is processed as the first step in the compilation, which produces numerous "*include*" files. These include files are then inserted into the remainder of the Fortran code in virtually hundreds of places using the C preprocessor. Therefore, the code that the user actually sees and works with can be very different from the code that is eventually compiled. An estimate from one of the WRF presentations had over 40,000 lines of code being inserted from these *include* files. These include files contain not only declaration statements for variables, but also argument lists for CALL and SUBROUTINE statements.

The reliance on the C preprocessor to handle the include files can lead to significant confusion during code development. For example, a developer can have a file in a text editor entering new code. This of course is the file before the preprocessing is performed. He compiles the new code and the compiler reports an error on line 5056. However, the file he is editing may only have 2000 lines. He must then find the expanded file (after the includes were performed) that the compiler actually used, find the appropriate line number in the expanded file, then attempt to cross-reference the line to the actual file he is working with.

WRF was designed with a "standard" computer science-type of hierarchical layer structure. One other example of the use of a layer structure is the IP packet. In WRF, the "Driver" layer sets up the memory, which is passed to the "Mediation" layer through huge argument lists (from one of the *include* files). The Mediation layer configures the I/O and parallelism, then again finally passes the memory, again through long arguments lists, to the "Model" layer, where the actual model code resides.

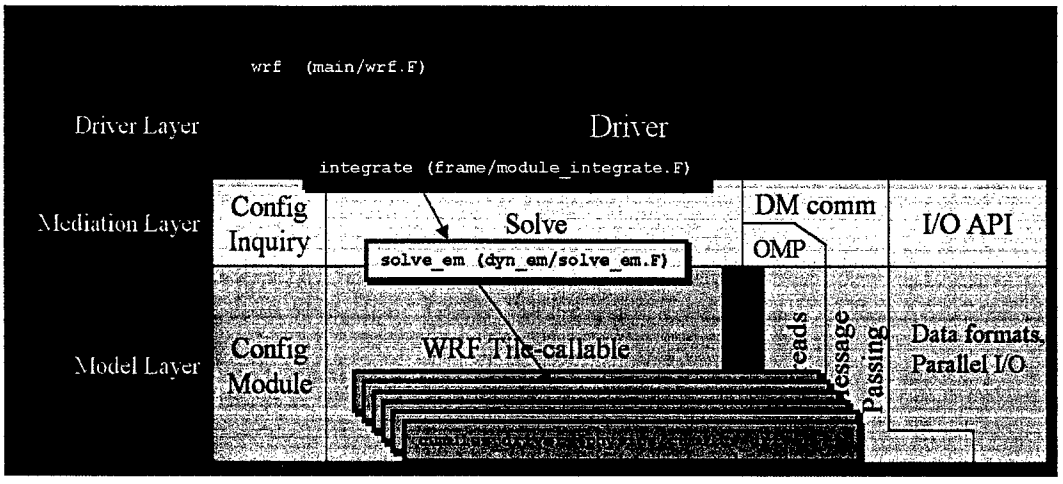

In the early stages of WRF development (circa 1997-8), it was assumed that there would be a single model to be supported in the code. However, with the evolved requirement that WRF support multiple models (or *cores*), this somewhat restrictive structure was implemented. The stated design goal of insulating the scientist from many details was certainly achieved. Through the use of the Registry and passing the memory through the long argument lists, arrays can be named as they are expected by the underlying model.

In some ways, the WRF design has significant flexibility in some parts, but to account for this flexibility, there were sacrifices made in complexity and understandability of the code. As with many codes that were written for flexibility (including some aspects of RAMS), the vast majority of the time, many of these particular features are virtually never or rarely used. An example in the WRF code that comes to mind is the ability to change the array order (e.g., (i,k,j) versus (i,j,k)) through the Registry. And often in practice, maintaining unused features in the code frequently leads to increased chances for errors and lengthens the time required for development.

As long as the model developer is able to follow the WRF structure, and work in the middle of the Model layer (e.g., implementing a convective parameterization), the structure is not that different from, say, developing in MM5. However, if more substantial development is to be done, as evidenced by numerous remarks at the WRF User's Conference, the structure can lead to memory inefficiencies and roadblocks to being able to accomplish new implementations. There has also been concern expressed in other meetings that there are not enough people who actually understand the full complexities of the WRF structure.

A comment about WRF performance in general: we are assisting the USFS with the development and installation of a fire weather forecast system. This system includes MM5, with WRF running in as similar configuration as possible. Basically, WRF has very similar errors and biases to MM5, except it is taking longer to run than MM5. Clifford Mass at the University of Washington reported at the 2005 WRF/MM5 Users' Workshop that they are doing a similar comparison, and finding similar results. This obviously caused significant concern at the Workshop. These results should not be

surprising, however, as WRF (ARW core) has a somewhat modified dynamical framework from MM5, but the physics packages are taken almost directly from MM5.

As mentioned, the WRF 3DVAR code was originally developed for MM5, then converted to the WRF code structure. As of this writing, there is a technical manual (Barker03) available for the scheme, which also goes into some aspects of the usage and code structure. There is also an online tutorial that covers the execution of the code. However, there is not a specific usage document that describes the namelist inputs. In fact, the online tutorial states: "Not all [namelist] values are provided with comments. This is deliberate – we only support changing the values with comments! Feel free to experiment with the others only if you can support yourself by checking the code to see what these other options do!" Unfortunately, none of the namelist variables have any comments in the supplied sample namelist and there are very few comments in the source code itself.

Another difference in the WRF 3DVAR code structure as compared to other geophysical models is the directory and subroutine structure. There is a "standard" structure in which the source code is divided into numerous subdirectories. Each subdirectory contains a Fortran MODULE "shell", with the same name as the subdirectory, with a .F appended. The actual subroutines then are located in numerous .inc files in the subdirectory. Before compilation, all of the .inc files are included into the MODULE shell, again by the C preprocessor. The advantages of this structure are that all subroutines are compiled as members of a MODULE, which allows the compiler to do checking for argument consistency. Also, the subroutines are located in manageable-sized files. The disadvantages are that the line numbers reported by the compiler are not the same as the code files, and the time required for compilation is very long, as with the WRF model itself.

As mentioned, one of the main advantages to variational schemes is the ability to have a variety of observation types, even observations that are not of the state variables. The primary tasks involved with adding a new observation type to the WRF 3DVAR scheme are:

1) read the observation data from input files
2) supply the observation operator H
3) supply the necessary error coefficients

The first task, following the WRF "standard" structure, involves modifications to the Registry and the I/O sublayer. Task 2 would be accomplished in the 3DVAR core, while Task 3 could either be done externally or in the 3DVAR core.

While the WRF code structure does present certain challenges and requires a larger learning curve than many other codes, it will be possible to implement new variable types if necessary.

4.3 Other 3DVAR Schemes

Several other meteorological models and analysis systems also use 3DVAR schemes. Virtually all have been designed and implemented for operational applications. Following is a list of those we found:

4.3.1 UKMO (United Kingdom Met Office) Model

The UKMO was one of the first to implement 3DVAR in their operational cycle. Their scheme (Lorenc et al. 2000) forms the basis for the WRF scheme and several others in use around the world.

4.3.2 ECMWF (European Centre for Medium-range Weather Forecasting)

ECMWF introduced a 3-dimensional variational analysis scheme to their operational system in January 1996, replacing their existing OI scheme. They found that forecasts from 3DVAR for the Northern Hemisphere were of similar quality to the OI forecasts, but forecasts for the Southern Hemisphere tended to be better with 3DVAR. ECMWF replaced the 3DVAR with a 4DVAR scheme in 1997.
(From: http://badc.nerc.ac.uk/data/ecmwf-op/model_changes.html)

4.3.3 HIRLAM (HIgh Resolution Limited Area Model)

HIRLAM is a hydrostatic mesoscale forecast model used by Denmark, Finland, Spain, and others in Europe. The 3DVAR scheme (Gustafsson et al.1999) used is very similar to the WRF scheme; both schemes were based from the UKMO (Lorenc et al. 2000) scheme. The HIRLAM scheme also uses the NMC method for creating the error matrices. Through a series of parallel data assimilation and forecast experiments comparing OI and 3DVAR, they found that 3DVAR forecasts consistently outperformed the OI-based forecasts.

4.3.4 NCEP ETA and GFS Models

NCEP implemented a 3DVAR scheme for the Medium-Range Weather Forecasting model (MRF) in 1991 (Parrish and Derber 1992). The scheme is also known as the SSI (Spectral Statistical Interpolation). The ETA Data Assimilation System (EDAS), used for the ETA model, switch from an OI analysis scheme to a 3DVAR scheme in the mid 1990's.

4.3.5 Japan Meteorological Agency (JMA)

The JMA introduced a three-dimensional variational assimilation method for the Global Spectral Model (GSM) in September 2001 (Tada 2002) for the primary purposes of assimilating satellite and precipitation data.

4.3.6 RUC (Rapid Update Cycle)

RUC (Benjamin, et.al. 2004) is a forecasting and data analysis system, developed by NOAA/FSL and used operationally at NCEP. FSL replaced the OI scheme in 2002 with a 3DVAR scheme. It also follows closely the UKMO scheme and uses the NMC method for creating the error matrices.

4.3.7 LAPS (Local Analysis and Prediction System)

LAPS (McGinley and Smart 2001) has been under development for over 10 years at NOAA/FSL. The most important feature of LAPS is probably the ability to ingest a very wide range of observation types. Unlike the other analysis schemes mentioned above, it uses a combination of variational schemes and standard objective analysis schemes to produce the full three-dimensional analysis. A 3DVAR scheme is used for the wind and streamfunction, while the objective analysis schemes are applied to the temperature, moisture, and cloud fields.

4.3.8 WRF/MM5 applications

Barker et al. (2004) described the original version of the MM5 3DVAR scheme and summarized the implementations at AFWA and the Taiwan Civil Aeronautics Administration. The scheme is also being implemented in Korea for use with a global spectral model. Barker et al. did show a potential improvement throughout a 24-hour forecast period for the u-wind component (they did not show the v-component) on higher resolution grids, but marginal improvements for moisture and temperature at all scales. They also showed that the LITTLE_R scheme (MM5's Cressman analysis) generally had better verifications at the initial time.

McAtee et al. (2005) described an application of the MM5 3DVAR scheme for a local forecasting application (down to 5 km grid spacing) for the Los Angeles basin. By running side-by-side forecasts with and without the use of the 3DVAR scheme, they were able to quantify the impact of the data assimilation. They found, similar to our results in Section 3, that there was some improvement to temperature forecasts in the first few hours, but the differences quickly became small. Surprisingly, their wind forecasts were slightly worse with the 3DVAR on average for the first few hours, and the differences became small again for the remainder of the forecast period.

5 4DVAR Review and Summary

5.1 Overview

3DVAR and similar objective analysis methods are designed to provide initial conditions for a numerical forecast that are optimal according to a chosen mathematical definition. In the specific case of 3DVAR, this means that the initial conditions are chosen to agree as closely as possible with a set of observations and simultaneously with a previous model forecast and/or a set of physical constraints such as hydrostatic and geostrophic balance. Optimal agreement at the start of the forecast does not, however, guarantee that subsequent forecast states will agree more closely with observations than they would had the forecast begun with different initial conditions. In fact, the initial conditions that give rise to the best agreement between forecast and observations over a selected time period (such as 24 hours) are in practice generally not those that would satisfy a 3DVAR optimization at the initial time.

One reason for this is that only observations that apply at the time of the 3DVAR procedure are used directly in the cost function. Although observations at earlier times were previously assimilated to run the model forecast, and these contribute to the cost function through the "first guess" terms, their contribution is less direct as they have been subjected to model errors. A second reason that optimal agreement between model initial conditions and current observations does not necessarily lead to the best agreement between model forecast and future observations is due to the model's own approximations and errors.

These limitations can be reduced to some extent through the application of 4DVAR, which can be described as application of 3DVAR at multiple time levels combined with multiple forward and backward numerical integrations of the forecast model to couple the 3DVAR analyses in time. A 4DVAR forward model integration is not the same as a forward integration with periodic 3DVAR-based corrections to the model state vector. Rather, with 4DVAR, each forward integration proceeds from an initial state with no further assimilation at all, and thus produces a solution whose evolution is subject only to the model equations. Differences between the forward integration and observations over the same time interval are evaluated and stored, but are not applied during the forward integration step. These differences are used to determine a new initial condition for a new forward integration to be made (over the same time interval). The procedure for this consists of backward integration of the adjoint of the model, using a state vector to which is applied the stored differences as the backward integration proceeds until reaching the initial time of each forward integration. The result of the backward integration is a state vector that differs from that used in the previous forward integration, and from the difference between these, a new initial state vector is constructed for the next forward integration. The above procedure is carried out for a number of iterations

until the initial state vector converges to a value that minimizes errors between forward integration and observations over all observation times.

The solution attained at the end of the final forward integration then becomes the initial state vector for a new model forecast. Because this state vector is generated directly from a forward integration of the forecast model, it is well balanced for the model and does not generate the unrealistic oscillations that result from periodic assimilation of observations that is a characteristic of 3DVAR. Thus, the 4DVAR solution is likely to be superior to what could be obtained with either a single or a time series of 3DVAR assimilations.

5.2 Disadvantages of 4DVAR

While 4DVAR schemes can alleviate some of the limitations of 3DVAR, there are a number of important disadvantages to 4DVAR:

- 4DVAR is subject to the same need as in 3DVAR to optimally define covariance matrix coefficients for the minimization of the cost function. This procedure is largely empirical and requires a large number of simulation experiments to be performed, making it also more appropriate for operational applications.

- 4DVAR as compared to 3DVAR is much more complicated and computationally expensive due to the multiple forward and backward integrations. Often these can number in the many tens in order to obtain a solution that is considered sufficiently converged. This means that computational resources expended in the assimilation procedure can be one or nearly two orders of magnitude higher than those used for carrying out the actual forecast. This raises the question as to whether it would be better to spend less of the existing resources on the assimilation procedure and allow more for improving the forecast model itself, for example, by increasing model resolution, implementing more accurate parameterizations, or performing more standard ensembling methods.

  Kalnay et al. (2000) mentioned the example of the ECMWF 4DVAR system: "4D-Var has a large computational cost compared to 3D-Var (typically 10-100 or more iterations are required for convergence, equivalent to about 30-300 model integrations per day). ECMWF, for example, has a powerful supercomputer about 25 times faster than a Cray C90, and has been running a model at a horizontal resolution of T213. Nevertheless, ECMWF had to make several simplifying assumptions in their implementation of 4D-Var (such as using a lower horizontal resolution model of T63 and a short assimilation window) in order to reduce the computational cost."

- The backward (adjoint) model required in 4DVAR is only an approximation to the forward model. It is actually the adjoint to the tangent linear model, which is a linearized form of the forward model. It has also been common practice to form

the tangent linear model from a simplified version of the full nonlinear model (Zhu and Navon 1999). The manner in which it propagates model-observation differences backward in time is therefore not entirely consistent with the behavior of the forward model. However, it would not be possible to run an actual forecast model backwards in time. Such a procedure would be unstable owing to the irreversibility of physical processes such as diffusion.

- Because of the use of forward, tangent linear, and adjoint models, any implementation of a 4DVAR scheme is tied to a specific model. Therefore, if both RAMS and WRF are to be considered, two complete, separate schemes are required. If the WRF NMM core is to be considered, a separate scheme (including tangent linear and adjoint models) will be needed. Also, as the main model versions change, so should the tangent linear and adjoint models. Although software does exist that automatically generates the adjoint of even complicated geophysical models, the resulting adjoint can be very inefficient, resulting in the need for extensive human intervention to produce a practical code (T. Vukicevic, personal communication).

- Kalnay et al. (2000) pointed out that, because the cost functions are being minimized over the entire time frame of a run, predictability issues come into play for longer runs. For example, if a 5 day run tended to stray from observations in day 3, 4DVAR adjustments for days 3-5 would affect how well the scheme performed in days 1 and 2. Thus, 4DVAR may be limited in a practical sense to the types of runs done at operational centers, those where 6-24 hours are run to provide an initial field for the next forecast cycle.

- With all the computational effort, only a single forecast result is obtained. Even assuming that this forecast may be close to the best currently attainable, it does not provide any uncertainty information in the forecast. In this regard, ensemble forecasting is far more useful.

5.3 Implementations of 4DVAR Schemes

Following are some of the locations where 4DVAR schemes are being used or developed:

5.3.1 ECMWF (European Centre for Medium-range Weather Forecasting)

In January 1996, ECMWF introduced a 3-dimensional variational analysis scheme to their operational system in January 1996, replacing it with the first version of a 4DVAR scheme in November 1997. Numerous updates to the scheme were made over the next

five years; most of them involved the inclusion of new observations or changing the way that the observations were processed. The scheme continues to be used.
(From: http://badc.nerc.ac.uk/data/ecmwf-op/model_changes.html)

5.3.2 HIRLAM (HIgh Resolution Limited Area Model)

Development work in 4DVAR started in 1995, with an operational test conducted in 1999. From this test, the forecast verification scores were comparable to the corresponding 3DVAR verification scores. It was unclear from the available literature which of the forecast centers using HIRLAM are currently using the 4DVAR scheme.

5.3.3 WRF

A 4DVAR scheme for WRF is under development at NCAR, funded mostly by AFWA (Air Force Weather Agency) (Huang et al. 2005). As with the 3DVAR scheme, it is based mostly on the 4DVAR code that was originally developed for MM5 (Zou et al. 1995), but was never distributed with the standard MM5 code. For WRF, modifications needed to be done for the tangent linear and adjoint models. Huang et al. estimate that a prototype version will be available in 2005, with a basic version in 2006.

5.3.4 Japan Meteorological Agency (JMA)

JMA implemented a 4DVAR scheme in their limited area Meso-Scale Model (MSM) in March 2002 (Tada 2002).

5.3.5 RAMS (Regional Atmospheric Modeling System)

A 4DVAR scheme based on RAMS was started in the mid 1990's at CSU/CIRA. The original work was based on RAMS version 3b. The efforts have culminated in the development of the RAMDAS system (Zupanski et al. 2005), which is depicted schematically in Figure 26. This scheme uses some aspects of the WRF 3DVAR code, such as the observation operators, and is patterned after the ETA model 4DVAR scheme (Zupanski et al. 2002). The adjoint model is currently based on RAMS v4.2.9. RAMDAS has been applied to assimilation of GOES satellite radiances (Vukicevic et al. 2004).

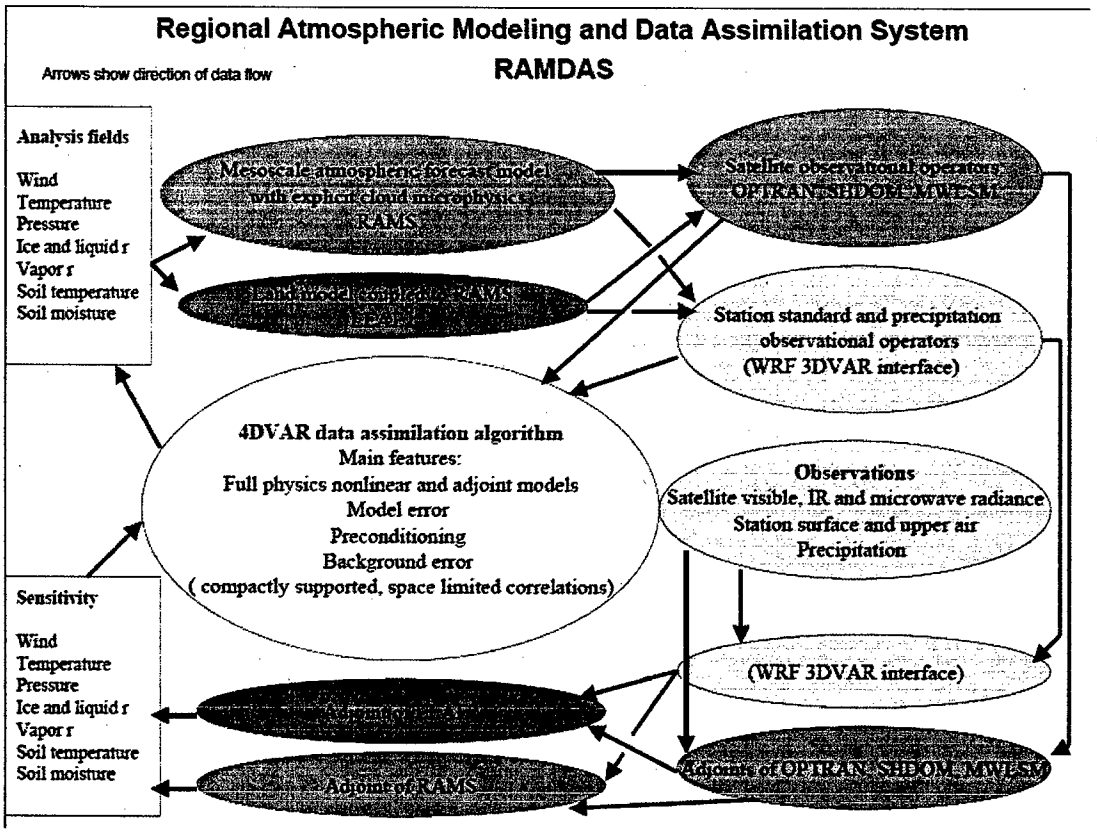

Figure 26: A schematic depiction of the RAMDAS system (from Vukicevic et al. 2004)

6 Ensemble Kalman Filter Review and Summary

6.1 Overview

As noted, the Ensemble Kalman Filter scheme is an approximation to the generalized Kalman Filter. It encompasses several types of procedures, which we will summarize in the following sections.

6.1.1 Kalman Filter

It has been pointed out (Zupanski 2005; Kalnay et al. 2000) that both 3DVAR and 4DVAR, as well as the commonly used data assimilation method of Optimal Interpolation (Daley 1991) are approximations to Kalman filtering theory. In its simplest form, a Kalman filter is a data processing algorithm that combines multiple estimates of a quantity, where each estimate consists of an expectation value and a variance, into a single expectation value and variance. The final single expectation value is a weighted mean of all estimated expectation values in which the relative weights are based on the degree of uncertainty (i.e., the magnitude of the variance) of each original estimate. The final variance is always less than the smallest variance of the original estimates, implying less uncertainty in the final estimate than in any of the originals.

The Kalman filter has often been used in geophysical modeling applications (e.g., Snyder and Zhang 2003, plus other references therein) to optimally blend observational data with numerical forecasts or simulations. In this case, the added complication exists that all model variables at all grid points are interrelated through governing physical constraints. As a consequence, the Kalman filter requires not only variances of each model variable at each grid point but also covariances between each pair. The full covariance matrix containing all these values is of exceedingly large size. For example, a typical global atmospheric or ocean model might contain around 10^7 or 10^8 total state variables (the product of the number of grid cells and the number of prognostic variables), and consequently the covariance matrix would contain 10^{14} to 10^{16} elements. Moreover, the covariance matrix must be propagated forward in time from each data assimilation time to the next. These facts render application of the full Kalman filter impractical, if not impossible, even on today's largest computers, and are the reason why other methods such as 4DVAR have been pursued instead.

However, it is intuitively obvious that the vast majority of covariances would be essentially zero because most pairs of elements are separated from each other by large distances. Hence, many methods have been explored for reducing the size of the covariance matrix to a manageable size without unduly compromising accuracy. Many of these (e.g., Ott et al. 2004) have used the concept of a horizontal radius of influence:

any two variables that are separated by more than a specified cutoff distance are assumed uncorrelated. Such methods have allowed practical application of the Kalman filter.

6.1.2 Ensemble Forecasting

Ensemble forecasting, in simple terms, is the production of multiple, non-identical numerical forecasts for the same geographic region and time period. Differences between forecasts occur because each forecast is initialized with slightly different initial conditions and/or carried out with different numerical model options and parameters. The differences in initial conditions are chosen to represent expected uncertainties in data observations and analysis, and differences in model parameters are chosen to represent expected uncertainty due to model approximations. The primary goals of ensemble forecasting are (1) to obtain a "best" forecast, which is hoped to be close to the mean of forecast results and (2) to obtain an estimate of uncertainty in the forecast, which it is hoped to be indicated by the variation among ensemble members. The first goal is the same as for conventional single forecasts (with any type of data assimilation method such as 3DVAR or 4DVAR). However, the second goal is uniquely addressable in an elegant way through ensemble forecasting, and this constitutes a major advantage of the method.

Within the procedure of ensemble forecasting, the tools for assimilating data into each forecast are in principle the same as for a single forecast. For example, one may initialize each forecast using a 3DVAR procedure, and one means of varying each initial state from all others could be to alter the covariance matrix coefficients each time. Alternatively, one could produce from 3DVAR or even 4DVAR a single "best" initial state vector for one forecast in the ensemble, and then add unique perturbations to that state vector for each additional member of the ensemble.

6.1.3 EnKF

Another method of obtaining covariance information is to estimate it from covariances between members of an ensemble of numerical model forecasts or simulations. This technique, known as the Ensemble Kalman Filter (EnKF) is an elegant means of obtaining not only a practical and relatively inexpensive blending of modeled and measured quantities, but also a direct estimate of forecast or analysis uncertainty that is an inherent advantage of ensemble methods.

To illustrate how the EnKF obtains an improved analysis by combining model forecast and covariance information with observations, we consider the following simple example from Snyder and Zhang (2003) in which radial velocity measurements from a single Doppler radar are blended with an ensemble of numerical model forecasts. We shall focus on a single Doppler measurement of radial velocity, V_r, at a particular location and time. This is represented in Figure 27a as a bold arrow (with a value of 14 m/s) along the Y-axis. Expected measurement error is depicted as a Gaussian distribution with the

arrow. A 50-member ensemble of numerical forecasts, initialized at some earlier time, is integrated forward until reaching the time in which the V_r measurement is made. At this time, the individual forecasts of radial velocity at the same time and location of the Doppler measurement are denoted by the 50 dots on the figure, and range from near 0 to about 19 m/s. We also consider one additional quantity that is forecast by the model but is not measured, say, a vertical velocity, W, at some different location. The distribution of W, over the 50 ensemble members is illustrated on the horizontal axis in Figure 27a. Ensemble-mean values of V_r and W are shown on the Y and X axes, respectively, by thin arrows.

The important thing to recognize here is that although predictions of V_r and W vary widely over the ensemble, a clear correlation (covariance) between them is indicated. Possible reasons for this covariance are many and varied, but the details are unimportant. The fact alone that the 50-member ensemble indicates a strong correlation implies, with a high level of statistical confidence, that some physical constraint between them exists in the model for the particular conditions of the present forecast. Thus, the covariance matrix that is determined from the forecast ensemble contains a nonzero element representing the (V_r, W) pair of members of the model state vector.

Next, the Kalman filter is applied to combine the Doppler V_r measurement with the ensemble forecasts to obtain improved estimates for each ensemble member. This procedure is represented in the following linear equation, applied to each ensemble member:

$$x_a = x_f + P_f H^T \left(H P_f H^T + R\right)^{-1} \left[y - H(x_f)\right]$$

where x_a is the model state vector resulting from the analysis, i.e., the set (V_r, W) in our simple example, x_f is the model state vector prior to the analysis, P_f is the model error covariance matrix, y is the observation (of V_r in our example), R is the observation error covariance, and H is a transformation matrix that relates the model state vector to the observation (i.e., it would include appropriate spatial interpolation coefficients when the observation location is not exactly collocated with a grid cell). The above equation applies to the Kalman filter in general, but in the specific case of the EnKF, P_f is estimated from the ensemble. Many successful algorithms have been developed for this procedure (e.g., see the extensive review given by Evensen 2003).

The result is depicted by the black dots in Figure 27b. For comparison, the gray dots illustrate the ensemble values prior to application of the filter. Thin arrows indicate updated ensemble means of V_r and W and gray arrows show the previous means for comparison. The spread of corrected V_r values is now much smaller than before. This is because the expected error of the Doppler measurement is much smaller than the original

spread over the ensemble, indicating that higher confidence is placed in the measurement than the forecast ensemble.

Although W was not measured by Doppler, the ensemble distribution of W values is also narrower following application of the EnKF. This results entirely from the positive covariance indicated by the ensemble, and illustrates the important role that the covariance matrix plays in the EnKF.

Schematic Radar Example

Update w given an observation of v_r

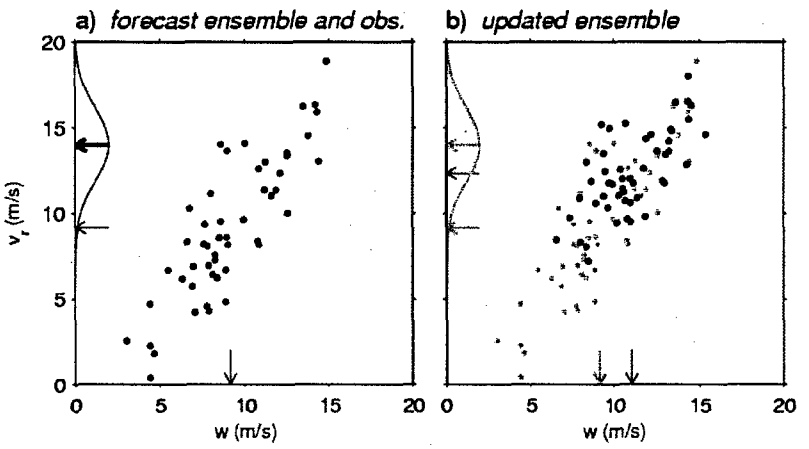

Figure 27: Idealized schematic of an application of an Ensemble Kalman Filter scheme (from: Snyder and Zhang 2003). See text for details.

One of the main differences in the various EnKF schemes is how the ensemble members are chosen. Szunyogh et al. (2004) divide the schemes into two groups. The first group uses random perturbations of the observations, based on the observation error estimates, to generate the ensembles. An example of this is the system developed at the Canadian Meteorological Service (Mitchell and Houtekamer 2004). The second group of schemes is described as the Kalman square-root filter schemes, where the analysis is only done once to generate a mean analysis and the error covariance matrix. The ensemble perturbations are used to then generate analysis perturbations that are limited to some subset of the ensemble. Since there are virtually an infinite number of ways to define this subset, many different types of EnKF schemes can be devised following this technique.

Further refinement of the EnKF is still an active research topic (e.g., Snyder and Zhang 2003; Ott et al. 2004, Evensen 2003; Zupanski 2005). For example, the technique described above of applying a cutoff radius of influence to reduce the size of the covariance matrix is used with the EnKF, but the distance chosen and how it may or may not depend on the local environment are generally chosen arbitrarily. Anderson (2004) has made progress toward developing a generic, self-adjusting method of determining this distance through a Monte Carlo approach. Although considerable improvements are

still to be gained through future development, the EnKF has already demonstrated considerable success in many geophysical modeling applications in other fields along with meteorology.

Another attractive aspect of the EnKF is that it does not require development and maintenance of the numerical model adjoint as does 4DVAR. The EnKF employs only the numerical model in forward integration, utilizing its full nonlinearity and complexity of physical parameterizations, to generate the best possible numerical forecast or analysis with ensemble-estimated uncertainty.

6.2 Development and application of EnKF Schemes

Following are some of the locations where EnKF schemes are being developed or applied:

6.2.1 Colorado State University

M. Zupanski at CSU/CIRA has worked on EnKF theory for several years. His recent work has focused on modification of the basic EnKF scheme, which he calls the Maximum Likelihood Ensemble Filter (Zupanski 2005). The MLEF scheme combines some ideas from 3DVAR, 4DVAR, and EnKF, using an iterative minimization algorithm to reach the maximum likelihood state estimate. They have also investigated techniques to allow non-Gaussian error distributions (Fletcher and Zupanski 2005).

6.2.2 University of Maryland

Researchers at the University of Maryland are developing EnKF schemes in conjunction with the NCEP GFS model. Szunyogh et al. (2004) describe the development of a scheme based on the Kalman square-root filters, where the assimilation of the observations is local to every grid point.

6.2.3 NCAR

Snyder et al. (2005) presented recent work describing their prototype EnKF for WRF. The prototype uses the square-root filter approach similar to the work at the University of Maryland. They also compared EnKF with the WRF 3DVAR and found comparable or better results with the EnKF schemes.

6.2.4 University of Washington (UW)

A real-time EnKF system, focusing on the analysis over the Pacific Ocean has been implemented at the University of Washington (Torn et al.; see http://www.atmos.washington.edu/~enkf/enkfpy.cgi). The system was started in December 2004.

6.2.5 Meteorological Service of Canada (MSC)

Mitchell and Houtekamer (2004) presented the EnKF that has been developed at the Meteorological Service of Canada (MSC) to provide the initial conditions for their operational medium-range ensemble prediction system. Their EnKF includes the input of real observations from the standard observational network as well as microwave radiances from the AMSU A and B instruments. The assimilation is being done with their forecast model that includes the standard operational set of physical parameterizations. The system was started in January 2005.

7 Suggestions for Future Research and Development

With the model testing, literature reviews, and investigations that were performed on this project, we have a clearer understanding of the data assimilation field and are able to provide the following suggestions for follow-on work.

7.1 3-D Variational Schemes

We arrived at the following conclusions concerning 3DVAR schemes and the WRF scheme in particular:

- 3DVAR schemes are much more in a "research mode" than we anticipated at the beginning of the project.
- Significant investigation still is required to determine the best ways to develop the error covariance matrices, which are very important to the accuracy of the scheme.
- It is easiest to develop the error matrices for operational applications, where a static domain and location are configured and can produce a history of simulations. We did not find any literature specifically addressing using the schemes for forensic re-creation of past events.
- For the WRF 3DVAR scheme and RAMS, it is probably adequate to run the scheme on the WRF grid structure and interpolate the results to the RAMS grid, although more testing is required.
- The WRF 3DVAR scheme is under active development, being funded by AFWA, Korea, and Taiwan. Since the WRF scheme is under active development, attempting to coordinate any new development we perform for AFTAC's applications with the development that NCAR performs for their applications would be problematic, at best.

***Recommendations*:**

1) Continue to monitor the WRF 3DVAR development and test new versions.
2) Investigate methods for specifying error matrices for forensic simulations.
3) Perform additional comparisons of 3DVAR versus RAMS/ISAN initialization/FDDA for use in forensic applications. It is possible the WRF 3DVAR scheme could serve as an adequate option for RAMS, as an alternative for ISAN and the Barnes scheme, although other packages or simpler schemes could suffice also.
4) Do not pursue specific 3DVAR development at this time, unless there will be specific observation types that AFTAC has access to that will not be considered by NCAR, or if further testing shows improvements can be easily made.

7.2 4-D Variational Schemes

We arrived at the following conclusions concerning 4DVAR schemes:

- 4DVAR schemes have the important disadvantages of 3DVAR schemes (the requirement of empirical tuning of error covariance matrices), especially for forensic applications.
- In addition, they have the very significant added complexities of the need for tangent linear and adjoint models.
- Significantly more computing resources are required than for 3DVAR.

***Recommendations*:**

1) Continue to monitor the WRF and RAMS 4DVAR development.
2) Do not pursue 4DVAR development at this time.

7.3 Ensemble Kalman Filter Schemes

We arrived at the following conclusions concerning the EnKF schemes:

- EnKF schemes are relatively new to the meteorological field. They show a great deal of promise, but still require research and testing.
- EnKF scheme eliminate two significant problems of 3DVAR and 4DVAR: 1) no adjoint models are required, and 2) pre-specification of error covariance matrices are not needed.
- Uncertainty information can be produced since an ensemble is run.
- More computing resources are needed than for 3DVAR, but less than 4DVAR.

***Recommendations*:**

1) Continue to monitor the WRF and RAMS EnKF development
2) Determine which EnKF algorithms in current existence are most suitable for forensic analysis applications based on accuracy, flexibility, and efficiency.
3) Develop, test, and implement these algorithms for RAMS applications.
4) As part of the algorithm investigation for the RAMS EnKF, take into consideration the efforts at NCAR on the WRF EnKF scheme. Determine the feasibility of coordinating development with NCAR researchers, and work with them if advantageous.

7.4 Other Model Development

While data assimilation is one way to attempt to improve model accuracy, there are other model developments and implementations that can also help. It is often unclear as to which techniques will make the biggest improvements. Given the extra computational time that many data assimilation schemes require, we can pose the question as to whether it would be better to spend less of the computer resources on the assimilation procedure and give more resources for improving the simulation results (e.g., by increasing model resolution or implementing more accurate parameterizations).

Following is a partial list of improvements that could be implemented to improve general simulation results, especially for AFTAC's forensic-type of applications:

- Use of simpler variational techniques – Rather than full 3-dimensional variational schemes, the use of 2-D or partial 3-D (such as in LAPS) schemes could be experimented with. Or even the full implementation of a package like LAPS or RUC could be implemented to provide 3-dimensional analyses. These can serve as a replacement for the current successive correction schemes (e.g., Barnes) and not require large amounts of extra computer resources.

- Observation quality control – A standard implementation of a quality control package would help both the initial data analysis and the use of the observational-nudging scheme. We have developed a package in the past to handle surface and upper air observations, but updating and improvements are needed.

- More sophisticated observation-nudging scheme – With the implementation of the observation-nudging data assimilation scheme in RAMS, observational data can be inserted in a very flexible manner. However, various improvements can be made to the scheme, such as the use of non-circular horizontal influence functions.

- Improved physics parameterizations:

    - Treatment of urban areas – The characteristics of urban areas are rather grossly approximated in virtually all mesoscale models. There have been two partial implementations of a sophisticated urban canopy scheme, TEB (Town Energy Balance), in RAMS at CSU and the University of Sao Paulo in Brazil. TEB was developed by Dr. Valery Masson at Meteo-France and we have received permission from Meteo-France to distribute the scheme with RAMS if it were officially implemented. There are portions of the scheme that may also be able to assist with other types of sub-grid complexities, such as valleys, rock outcroppings, etc.

    - Radiation schemes – Other radiative transfer schemes are available which may be able to provide more accurate radiative fluxes than the existing schemes. The RRTM (Rapid Radiative Transfer Model) has been

developed by AER, Inc. and is one example. MM5 and WRF include the RRTM longwave scheme; a newer shortwave scheme has also been developed. Other schemes from the climate modeling community are also available.

- Generalized diffusion – As computer power becomes cheaper, model resolutions will continue to increase. We will soon be routinely passing a similar threshold as we passed in the 1970's and 1980's with regards to convective parameterizations. Regarding diffusion schemes, RAMS has schemes that work reasonably well at 1km grid spacing and larger, and schemes that work well at less than 100m grid spacing. However, as with convective parameterization, there is a zone where the current schemes can have problems. In conjunction with colleagues in Italy, we have designed the theoretical framework of a generalized scheme that would provide a smooth transition from the mesoscale to the LES (and smaller) scales.

- Direct RAMS building simulation – With the recent addition of the shaved-ETA-type coordinate in RAMS v6.0, the dynamics of the model can handle very small-scale simulations routinely. However, there are various aspects of the model physics that still need modification, such as surface fluxes from vertical walls, different types of surface materials, and non-vertical radiative fluxes.

- Integration of dispersion/chemistry modules – While RAMS has been interfaced to various dispersion models (e.g., HYPACT, CALPUFF) and photochemical models (e.g., CAMx, CMAQ), there are sometimes advantages to having an integrated capability. Generally, there is a one-way interaction between the meteorological simulation and the dispersion. With integration, there can be a two-way interaction. Various physics, such as wet deposition and aqueous-phase chemistry, can be much more accurate by including them at the regular timestep resolution of the meteorological model. We have partially completed the integration of HYPACT in RAMS, but more work is needed. We also have had discussions with ENVIRON concerning the issues of an integration of CAMx chemistry into RAMS.

- Shared-memory parallelism – While RAMS uses MPI for distributed-memory parallelism, which is efficient on many shared-memory machines, we expect additional parallel efficiency would be gained by implementing shared memory parallel constructs in the code to work in conjunction with MPI. New CPU architectures, such as multi-core chips, may make this very beneficial for the near-future PC cluster, along with shared-memory nodes in machines such as SGI and IBM.

- Unstructured grids – RAMS and almost all other models use a *structured* grid system, meaning that at a given grid point, the neighboring grid points can be

referenced by incrementing or decrementing an array index. With RAMS v6.0, we have started to remove the global simulation capabilities, that were partially implemented, to a companion global model called OLAM (Ocean Land Atmosphere Model), under development at Duke University by R. Walko. OLAM uses the same physics parameterizations as RAMS, but employs an *unstructured* grid system. The unstructured grid in OLAM was primarily designed for the model "equatorial" regions and the nesting scheme, but could be extended to other uses, such as providing high-resolution following diagonal coastlines, rivers, or regions of complex topography. It is a similar concept to the triangular mesh implemented in the OMEGA model by SAIC, but using the more sophisticated RAMS physics and numerics.

8 References

Anderson, J.L, 2004: A Hierarchical Ensemble Filter for Data Assimilation. Submitted to *Mon. Wea. Rev.*

Barker, D. M., A. Bourgeois, Y.-R. Guo, W. Huang, and Q.N. Xiao, 2003: A Three-Dimensional (3DVAR) Data Assimilation System For Use With MM5. NCAR Tech Note, Submitted.

Barker, D.M., W. Huang, Y. -R. Guo, A.J. Bourgeois, and Q. N. Xiao, 2004: A three-dimensional variational data assimilation system for MM5: Implementation and initial results. *Mon. Wea. Rev.*, **132**, 897-914.

Benjamin, S., and Coauthors, 2004: An hourly assimilation–forecast cycle: The RUC. *Mon. Wea. Rev.*, **132**, 495–518.

Daley, R., 1991: *Atmospheric Data Analysis*. Cambridge University Press, 457 pp.

Dharssi, I., A. C. Lorenc, and N. B. Ingleby, 1992: Treatment of gross errors using maximum probability theory. Quart. J. Roy. Meteor. Soc., 118, 1017-1036.

Evensen G., 1992: Using the extended Kalman filter with a multilayer quasi-geostrophic ocean model. J Geophys Res, 97, 905–924.

Evensen, G., 2003: The Ensemble Kalman Filter: theoretical formulation and practical implementation. Ocean Dynamics, 53, 343–367.

Fletcher, S. J., and M. Zupanski, 2005: A framework for data assimilation which allows for non-Gaussian errors. Submitted to *Proc. Royal Soc. of London A*.

Gustafsson N. et al.,1999: Three dimensional variational data assimilation for a high resolution limited area model (HIRLAM), HIRLAM Technical Report, 40, January 1999, 74 pp.

Huang, X. and Co-authors, 2005: The Weather Research and Forecasting Model based 4-dimensional variational data assimilation system. WRF/MM5 Users' Workshop, Boulder, Colorado.

Kalnay, E., S.K. Park, Z.-X. Pu, and J. Gao, 2000: Application of the quasi-inverse method to data assimilation. *Mon. Wea. Rev.*, **128**, 864-875.

Liu, D. C., and J. Nocedal, 1989: On the limited memory BFGS method for large-scale optimization. *Math. Program.*, **45**, 503-528.

Lorenc, A. C., 1995: Development of an Operational Variational Assimilation Scheme. UKMO Forecasting Research Division Technical Report No. 116.

Lorenc, A. C., S. P. Ballard, R. S. Bell, N. B. Ingleby, P. L. F. Andrews, D. M. Barker, J. R. Bray, A. M. Clayton, T. Dalby, D. Li, T. J. Payne, and F. W. Saunders, 2000: The Met. Office global three-dimensional variational data assimilation scheme. *Quart. J. Roy. Meteor. Soc.*, **126**, 2991-3012.

McAtee, M.D., J.F. Drake, and L.O. Belsma, 2005: Impact of Three Dimensional Data Assimilation on High Resolution Weather Forecasting in the Los Angeles Basin. Presentation at the Ad-Hoc Met. Modeling Workshop, USEPA and LADCO, Lakewood, Colorado, June 2005.

McGinley, J.A. and J.R. Smart, 2001: On providing a cloud-balanced initial condition for diabatic initialization. Preprints, 18th Conf. on Weather Analysis and Forecasting, Ft. Lauderdale, FL, Amer. Meteor. Soc.

Michalakes, J., S. Chen, J. Dudhia, L. Hart, J. Klemp, J. Middlecoff, and W. Skamarock, 2001: Development of a next-generation regional weather research and forecast model. Developments in Teracomputing: Proceedings of the 9th ECMWF workshop on the use of high performance computing in meteorology, 269-276.

Mitchell, H.L. and P.L. Houtekamer, 2004: Atmospheric Data Assimilation with an Ensemble Kalman Filter. Workshop on Ensemble Methods, UKMO, Exeter, UK, October, 2004.

Ott, E., B.R. Hunt, I. Szunyogh, A.V. Zimin, E.J. Kostelich, M. Corazza, E. Kalnay, D.J. Patil, and J.A. Yorke, 2004: A local ensemble Kalman filter for atmospheric data assimilation. Tellus, 56A, No. 5, 415-428.

Parrish, D.F., J. Derber, 1992: The National Meteorological Center's spectral statistical interpolation analysis system. Mon. Wea. Rev., 120, 1747-1763.

Snyder, C., and F. Zhang, 2003: Assimilation of Simulated Doppler Radar Observations with an Ensemble Kalman Filter. Mon. Wea. Rev., 131, 1663-1677

Snyder, C., A. Caya, and D. Barker, 2005: An ensemble KF for WRF and comparison with WRF 3DVAR. WRF/MM5 Users' Workshop, Boulder, Colorado.

Szunyogh, I., E.J. Kostelich, G. Gyarmati, B.R. Hunt, E. Ott, A.V. Zimin, E. Kalnay, D. Patil, and J.A. Yorke, 2004: A local ensemble Kalman filter for the NCEP GFS model. 20th Conference on Weather Analysis and Forecasting/16th Conference on Numerical Weather Prediction, AMS, Seattle, WA, January 2004.

Tada, H., 2002: Present and future uses of satellite observations for tropical cyclone forecasting and research. Fifth International Workshop on Tropical Cyclones, Cairns, Australia, 3-12 Dec 2002.

Vukicevic, T., T. Greenwald, M. Zupanski, D. Zupanski, T. Vonder Haar, and A. Jones, 2004: Method for explicit 3D cloud analysis by direct assimilation of visible and IR cloudy radiance. *Mon. Wea. Rev.*, **132**, 3066-3077.

Zhu, Y, and M. Navon, 1999: Impact of Parameter Estimation on the Performance of the FSU Global Spectral Model Using Its Full-Physics Adjoint. *Mon. Wea. Rev.*, **127**, 1497-1517.

Zou, X., Y.-H. Kuo, and Y.-R. Guo, 1995: Assimilation of atmospheric radio reflectivity using a nonhydrostatic mesoscale model. Mon. Wea. Rev., 123, 2229-2249.

Zupanski, D., M. Zupanski, E. Rogers, D. Parrish, and G. DiMego, 2002: Fine resolution 4DVAR data assimilation for the Great Plains Tornado Outbreak. *Wea. Forecasting.*, **17**, 506-525.

Zupanski, M., D. Zupanski, T. Vukicevic, K. Eis, and T.I. Vonder Haar, 2005: CIRA/CSU Four-Dimensional Variational Data Assimilation System. Mon. Wea. Rev., 133, 829-843.

Zupanski, M., 2005: Maximum likelihood ensemble filter: Theoretical aspects. Mon. Wea. Rev., 133, 1710–1726.