5.4 THE NCSC-PSU NUMERICAL AIR QUALITY PREDICTION PROJECT: INITIAL EVALUATION, STATUS, AND PROSPECTS

John N. McHenry[1]*, Nelson Seaman[2], Carlie Coats[1], Dave Stauffer[2], Annette Lario-Gibbs[2],
Jeff Vukovich[1], Elizabeth Hayes[3], Neil Wheeler[4]

[1]Environmental Programs
MCNC–North Carolina Supercomputing Center
Research Triangle Park, North Carolina

[2]Department of Meteorology, The Pennsylvania State University, University Park,
Pennsylvania

[3]Silicon Graphics Inc.,
Environmental Sciences Group,
North Carolina Supercomputing Center, RTP, NC

[4]Sonoma Technology, Inc.,
1360 Redwood Way, Suite C, Petaluma, CA

1. INTRODUCTION*

Recently, “third generation” air quality modeling systems such as MCNC’s Multiscale Air Quality Simulation Platform (MAQSIP, Odman and Ingram, 1996; Srivastava et al., 1994) and EPA’s Models-3 Community Multiscale Air Quality Model (CMAQ, EPA, 1999) have been developed to address multiscale, multipollutant problems. In the case of MAQSIP, atmospheric scientists and chemists have worked closely together with computational scientists and mathematicians to create a modular, flexible modeling system.

In the U.S., the paradigm for using such systems has been to apply them in retrospective historical simulations in order to contribute to policy assessments and foster scientific progress. However, the progress that can be made using such a paradigm is ultimately constrained by a limited number of test cases that are chosen for their “representativeness.” Unlike numerical weather prediction, which by its very nature involves multiple runs of multiple models at multiple centers every day of the year, air quality modeling has not been able to benefit from the wide application and scientific advancements that result from repeated forecast applications of the available state-of-science models.

This “discipline of forecasting” has been cited in a report by the National Research Council’s Board on Atmospheric Sciences and Climate (BASC: Entering the 21st Century, 1998) as a *necessary next step* for continuing to advance the state of the environmental sciences as we enter the new millennium.

For the past two years, we have been able to exercise this discipline through the joint NCSC-PSU Numerical Air Quality Prediction (NAQP) Project. In this project, we have coupled MAQSIP, the Sparse Matrix Operator Kernel for Emissions (SMOKE, Coats, 1996), and the PSU/NCAR MM5 Mesoscale Model (Grell et al., 1994) into a quasi-operational real-time atmospheric chemical forecast system. During the summers of 1998 and 1999, this system was used to produce once daily, experimental, real-time 48-hour NAQPs on a nested 108-36km grid covering the CONUS (MCNC, 1999; Penn. State, 1999). The 108-36km forecasts are produced on a 4-processor Silicon Graphics, Inc. Power Challenge computer at the Penn. State University Dept. of Meteorology. The system was expanded during 1999 by adding additional forecast domains of 45-15km resolutions running at the North Carolina Supercomputing Center.

We first reported on this system at the 1999 Annual Meeting of the AMS (McHenry et al., 1999), having demonstrated that it was possible to complete a 48-hour forecast in a timely manner. Further, a simple forecast example was used to show the potential that the coupled model has measurable predictive skill.

In this paper, we review the system’s component models, model coupling, and model initialization. We then present an initial performance analysis conducted on four forecast periods from simulations run at Penn State during summer 1998. We conclude with a summary of our findings and an outlook for the future.

2. COMPONENT MODELS

MM5 has been run in real-time at Penn State for nearly a decade (Warner and Seaman, 1990). It has gone through significant development since that time including the transition from a hydrostatic to a fully compressible non-hydrostatic model. Initial and boundary conditions are constructed after downloading

---

Corresponding author address: John McHenry, MCNC, P.O. Box 12889, 3021 Cornwallis Rd., Research Triangle Park, NC, 27709; E-mail: mchenry@mcnc.org

32km, grid-104 ETA GRIB forecast files from the National Center for Environmental Prediction (NCEP). MM5's coarse (108km) domain covers most of North America from the arctic circle south to Panama, and includes large portions of the eastern Pacific and western Atlantic. The fine 36km domain covers the CONUS. Both MM5 domains use a Lambert conformal map projection with 30 sigma-coordinate layers in the vertical. Model top is fixed at 100mb. The MAQSIP domains (Figure 1.) are windows into these MM5 domains and share this common vertical structure.

MM5 is configured with a 1.5 order TKE boundary layer scheme (Gayno, 1994), a deep convection scheme (Kain and Fritsch, 1993), a simple ice phase explicit moisture scheme, an atmospheric radiation scheme (Dudhia, 1989), and its standard land-surface model. A shallow convection scheme has been developed under another project (Deng et al., 1998) and is undergoing 3-D tests; it is designed specifically to support the needs of air quality models.

MAQSIP is a modular modeling system which, for this application, is configured for prediction of lower tropospheric ozone. This configuration uses a modified version of the Carbon Bond IV (Gery et al., 1989) chemical mechanism, a flux-form advection scheme (Bott, 1989), a K-theory scheme for turbulent vertical re-distribution of pollutants (Alapaty et al., 1997), and a dry deposition scheme (Walcek, et al., 1986). Clear-sky photolysis rates are calculated following Madronich (1987). It also uses an innovative cloud processes package which represents the effects of deep and shallow convection, grid-scale stratiform/cirruform layer clouds, and additional sub-grid scale layer clouds. Both the cloud package and turbulence parameterization were discussed in McHenry et al., 1999.

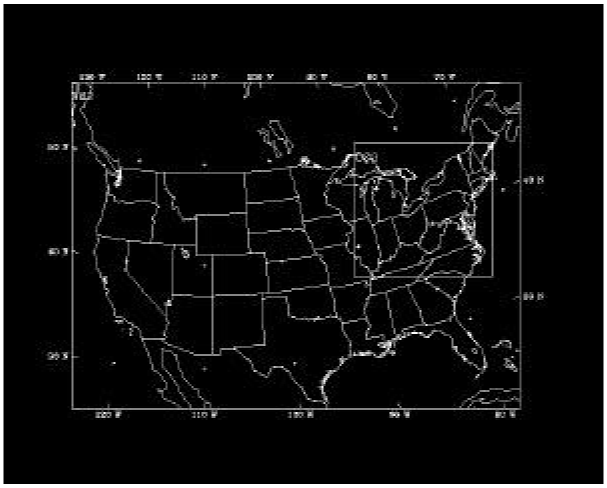

Figure 1. MAQSIP Real-time Domains: Coarse Domain-108km; Fine (Nested) Domain--36km

SMOKE is a high-efficiency emissions processing system. It features several hundred-fold computational performance enhancements due to its sparse-matrix algorithms for factor-based computations. It consumes less than 10 minutes of execution time on MAQSIP's 108-36km nested grid. The real-time simulations use SMOKE to process an improved Ozone Transport Assessment Group (OTAG) 1995 emissions inventory (Pechan, 1996). For real-time application, no attempt is made to project the inventory toward the future, but an emissions projection model or an updated inventory may be needed (see below).

The system uses real-time *forecast* meteorology to supply SMOKE with data needed to model meteorologically sensitive phenomena such as biogenic and vehicular emissions, and point-source plume rise. Because of the higher vertical resolution in third-generation models, SMOKE's discretization of plume rise *for every point source* is used.

3. MODEL COUPLING

Figure 2 shows the high-level dataflow diagram for the real-time system. An output module, MCPL (Coats et al., 1998), was developed within MM5 to generate outputs compatible with SMOKE and MAQSIP, which both use an I/O Applications Programming Interface (I/O API, Coats, 1995). The present I/O API implementation is layered on top of NetCDF from the National Center for Atmospheric Research, although the I/O API was designed from the beginning so that this lower layer could be easily changed.

MCPL fits into MM5 with minimal effort, and is callable at a variety of time scales from the advection-step frequency on up. It provides functionality generalizing the present MM5 subroutine OUTTAP.

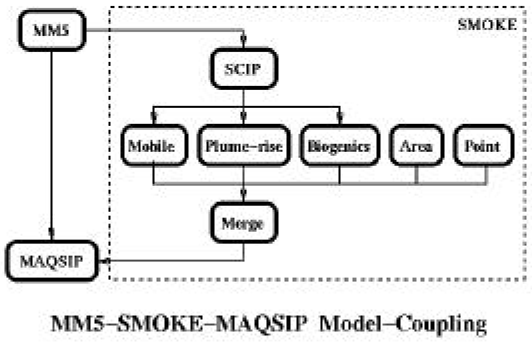

Figure 2. High Level Data Flow Diagram

MCPL is extremely flexible and configurable, allowing selected output variables, output formats, and windows into MM5 nests, all through the use of environment variable flags. For this application, MCPL is called every simulation hour for both the coarse and fine MAQSIP domains.

As Figure 2 indicates, MM5 (top left) is first run on both its coarse (108km) and fine (36km) domains. Then, a small interface processor (SCIP) is used to build a database of Kain-Fritsch deep convection soundings for use in MAQSIP. SCIP also constructs cloud fractional coverage and liquid water contents using algorithms in MAQSIP's new cloud package. This provides SMOKE with the self-consistent cloud data needed to account for attenuation of photosynthetically active radiation (PAR), which drives biogenic emissions. Once SMOKE finishes, all data is prepared and MAQSIP runs, first on the coarse domain, then on the fine domain.

MAQSIP was extensively re-written to run efficiently on microprocessor-based parallel machines. Run-times on the target machine were reduced substantially. Typically, the 48 hour MAQSIP forecast is now completed by 8AM EDT–barring initial meteorological data problems–in time to provide numerical guidance for this afternoon's and tomorrow afternoon's ozone outlook throughout the northeast quadrant of the US.

4. MODEL INITIALIZATION

There is little *real-time* atmospheric chemical measurement data suitable for use in initializing an NAQP model. Ground level ozone data are now being assembled through EPA's "AIRNOW" Web-site (EPA-OAQPS, 1999), but a coordinated effort would be needed to make it usable for *real-time* initialization. Real-time ground level primary precursor data (NO_X, VOC) are collected, but here again, assembly for use in real-time presents geopolitical problems. Further, there is no regular upper air chemical sounding data. Thus, establishing initial and boundary conditions for MAQSIP is more problematic than for meteorological models.

For boundary conditions, clean tropospheric background concentrations are set on the boundaries of the coarse domain, including a climatological vertical ozone profile as a function of latitude. The fine domain utilizes time-dependent chemical boundary conditions provided by the coarse domain MAQSIP run. Within the coarse domain, the emissions database extends to just east of the Rockies allowing for adequate build-up of polluted air upwind of the fine domain boundary.

Initial conditions for both domains are determined by using the 24 hour forecast concentrations from the previous day's forecast (a "warm start"); or, in the case when the model needs to be restarted, from a "cold start" file. We are currently experimenting with ingest of limited real-time surface chemical observations over a small sector of our 1999 Texas domain. In general, though, without adequate spatial and temporal surface and upper air atmospheric chemical measurements, the establishment of initial conditions will remain one of our largest sources of forecast uncertainty.

5. FORECAST PERFORMANCE ANALYSIS

To describe the forecast performance, we have chosen 20 "warm start" forecast days starting after the modeling system was frozen in mid-July, 1998. Because of occasional interruptions, these days are not all in sequence. They divide into four episodes, July 27-July 31; Aug 2-4; Aug 15-20, and Aug 27-31. This sample of 19 is a subset of the 57 warm start forecasts completed in July, August, and September for which verification data were saved. Thus, the statistics presented comprise about 1/3 of all forecasts and thus should be reasonably representative of model performance.

A. MM5 Performance and Implications

Because errors in the meteorological forecast will contribute to errors or even biases in both the SMOKE-emissions and MAQSIP-chemistry forecasts, we first characterized MM5's performance over the 19 chosen days. Figure 3 summarizes MM5's bias errors (BE) for the key variables wind-speed, wind direction, temperature, and water vapor mixing ratio, averaged over the depth of the mean (80-1700m) daytime PBL. With the exception of wind-speed, there is some bias increase in all variables from first to second day forecasts; i.e., between the 24-hour and 48-hour forecast times.

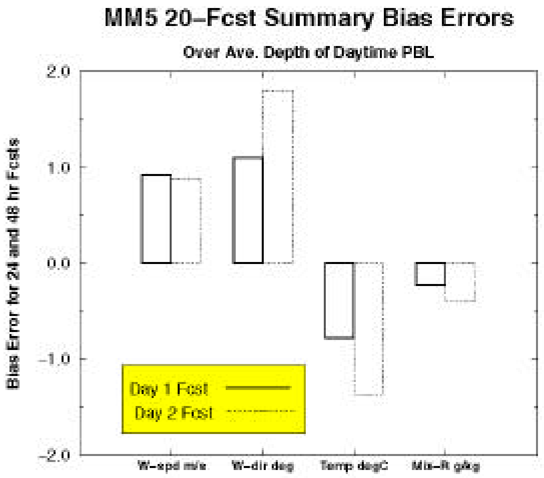

Figure 3. MM5 20-Fcst Bias Error Statistics

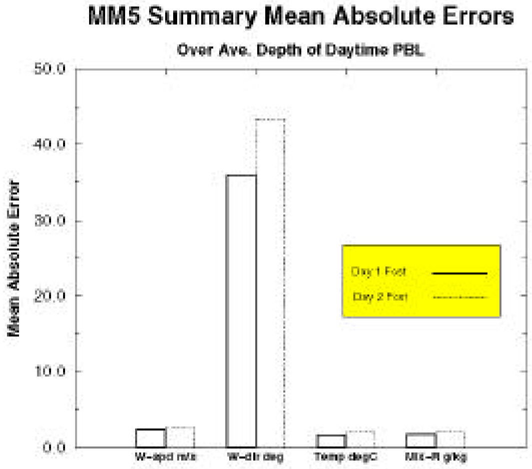

Figure 4. MM5 20-Fcst Mean Absolute Error Statistics

Figure 4 provides a summary of the mean absolute errors for the same key variables. Statistics were also computed for the surface layer (not shown).

Key points to note are the following: (1) though some of the speed bias can be explained, a true model error, somewhat less than 1 ms^{-1}, may tend to overestimate ventilation and thus contribute to an underestimate of chemical concentrations; (2) very small BEs for wind direction suggest that the opportunity for direction errors to impact plume transport over a 48-hour forecast is, on average, very small; (3) rather large MAEs for wind direction are typical for mesoscale models, and the variance implied is similar in magnitude to that observed in nature. Thus, the effect of mesoscale plume meandering should actually be represented reasonably well; (4) the small temperature and mixing ratio biases could have a small impact on biogenic emissions and chemical reaction rates that are sensitive to T and Q_v, as well as plume rise, which is sensitive to T.

B SMOKE Performance and Implications

There is little real-time emissions data with which to compare forecast emissions; thus, it was not plausible to attempt such a comparison. However, because we used the 1995 OTAG inventory, it is fairly safe to assume that it represents an underestimate of actual emissions in 1998, especially for urban areas characterized by rapid growth in vehicle miles traveled and power consumption. This could contribute to an underestimate of the NO_x pool available for production of ozone during daylight hours, together with an underestimate available for titration during nighttime.

C. MAQSIP Performance Analysis

To analyze the 19-forecast subset, data was extracted from EPA's Aerometric Information Retrieval System (AIRS) database after release in January, 1999. Data associated with monitors not coincident with those that voluntarily reported ozone observations through the AIRNOW Website during summer 1998 were filtered, so as to provide the most realistic comparison with "available" real-time observations. The monitor density per MAQSIP 36km grid cell is shown in Figure 5.

In order to perform the comparison, MAQSIP ozone forecasts for model layer 1, about 40 meters AGL, were bi-linearly interpolated to the latitude/longitude locations of the monitors. Hourly ozone values from local 1PM to local 6PM were then paired from the observation and interpolated model data sets for each of the 391 monitor locations. This resulted in approximately 2346 data pairs per afternoon for each of the 19 days examined. Pairs which included missing monitor data were discarded. Hour-by-hour pairs were constructed, rather than, for example, afternoon averages, in order not to obscure the "raw" performance of the model.

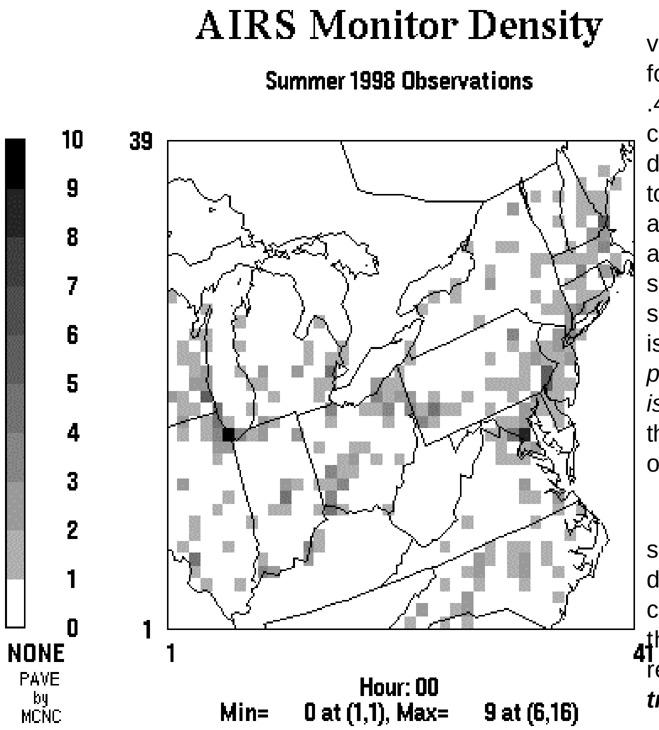

Figure 5. No. of Monitors per Grid Cell in 36-km Domain per Summer 1998 AIRNOW WebSite.

Figure 6 depicts a resulting "typical" scatter plot of MAQSIP forecast versus observed ozone for 1 of the 19 forecast days, including regression analysis. For July 30, the figure shows that there was a clear linear relationship between predicted and observed ozone values. The linear regression resulted in an R^2 of.297 for the second day forecast valid 7/30, and an R^2 of .416 for the first day forecast valid 7/30. These correspond to correlation coefficients r_{corr} of .54 for 2nd day and .644 for 1st day respectively. MAQSIP is seen to over-predict ozone concentrations at values below about 40PPB, and under-predict ozone concentrations at values above about 80 PPB, as evidenced by the slopes of the regression lines. The improvement in slope and regression statistics in the 1st day forecast is apparent. Note that the R^2 statistic specifies the *proportion of the variability in the observation data that is linearly accounted for by MAQSIP predictions.* Thus, the 1st day forecast linearly accounted for 11.9% more of the observed variability in this example.

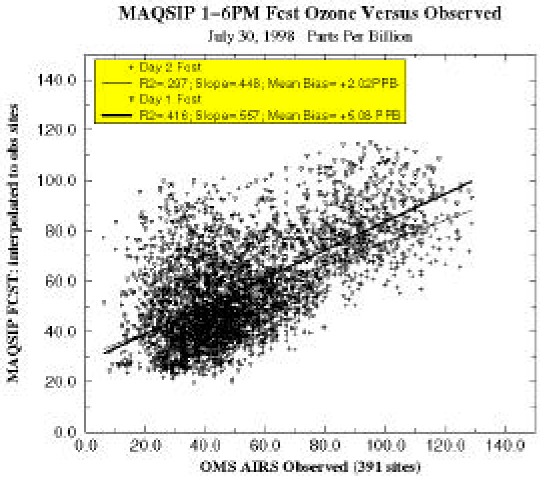

Figure 6. Representative MAQSIP Scatter Diagram: July 30, 1998

Figure 7 shows the complete set of regression statistics for all 19 analyzed days. For convenience, the days are plotted sequentially even though there are calendar date separations between the four episodes, the abscissa labels depicting this. For example, "j27" refers to July 27 whereas "a4" refers to August 4. ***No trend information between episodes can be inferred.***

Figure 7 shows that the R^2 and correlation coefficients are consistently higher for 1st day forecasts than for 2nd day forecasts, but not by a wide margin. The second day, 48-hour forecasts still exhibit a linear correlation and thus retain predictive information. With the exception of one uncorrelated (August 4th) and one weakly correlated (August 27th) forecast, all days reveal at least modest to at most quite good linear correlation between the predicted and observed hourly afternoon ozone concentrations. Because the sample size is large, we would expect that the 95% confidence limits for r_{corr} would be narrow and thus closely representative of the actual population r_{corr} between the model and observations for each day. The variability of

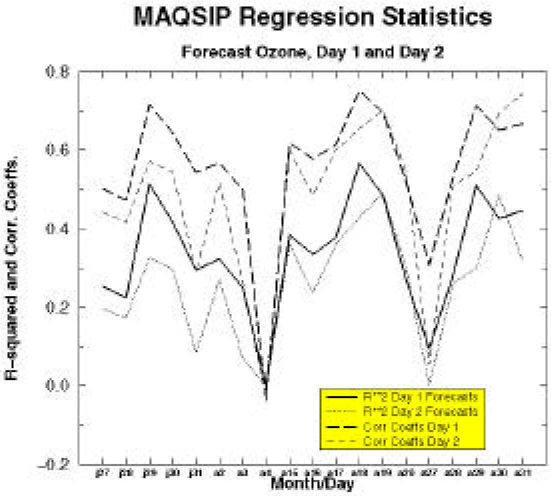

Figure 7. MAQSIP Regression Statistics for 19 Day-1 and Day-2 Forecasts, Afternoon 1-6PM

the regression statistics from day to day reveal that certain "types of days" may be more "predictable" than others, all other things being equal. Some of this variability is attributable to MM5's representation of the atmosphere as opposed to the "true" atmosphere, although MM5's statistics suggest that its bias errors will not induce large errors in the mean. However, it is well known that the atmosphere's own predictability is governed by chaotic nonlinear dynamics that induce greater or lesser predictability on any given day. See Rao, et al., 1997, for some insightful background here.

In Figure 8, the MAQSIP forecast bias error for first-day forecasts is plotted as a function of simulation day for all observations and for three observation bins: <50PPB ozone, between 50-80PPB ozone, and above 80PPB ozone. Note that the bias errors peak on the two days when forecasts were uncorrelated with observations. On August 4th, the modeling system over-predicts observed concentrations <50PPB by about 40PPB, and under-predicts observed concentrations >80PPB by about 20PPB. This suggests a displacement of weather features such as a frontal location in MM5. This is the only forecast considered to be a true "bust" among the 19 studied so far.

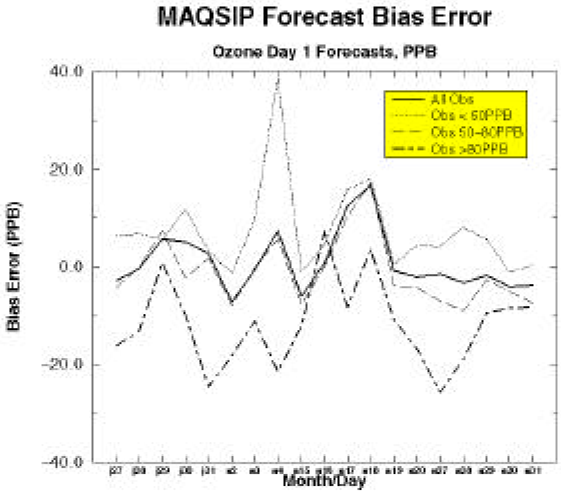

Figure 8. MAQSIP Bias Error Statistics for 19 Day-1 Forecasts

On August 27th, the first-day forecast bias is most evident as an under-prediction of observations >80PPB. However, we note that on this day only 8 of 2210 monitors reported ozone concentrations greater than 80PPB for any of the 6 afternoon hours considered, so that the primary bias results from a very small monitor sample. Further, though the scatter diagram (not shown) reveals a rather random pattern of predictions versus observations, they are clustered between 50PPB and 80PPB. Thus, despite a lack of "accuracy," MAQSIP correctly forecasts a rather "clean" day: note that the other "bins" are essentially unbiased. Thus, in this instance, the predictive value of the model is rather good despite a lack of correlation with the observations, since ozone forecasts rely on *thresholds* and *ranges of values* more than they do on specified concentrations. (See below.)

In general, Figure 8 *reveals that, taken over all observations, MAQSIP's first-day ozone concentration forecasts are relatively unbiased.* Corresponding to this is a modest under-prediction of concentrations above 80 PPB and a modest over-prediction of concentrations under 50PPB. Thus, most regression line slopes are similar to those shown in Figure 6.

There is a sequence of two forecast days that do not fit this general pattern. On August 17 and 18 all "bins" were biased modestly high, while at the same time exhibiting good linear correlation. This suggests that there may have been subtle meteorological phenomena at work that led to concentrations on the order of 20PPB lower in nature than predicted by MAQSIP. If this is the case, it is possible that the lack of a real-time chemical observation re-initialization capability in MAQSIP could be responsible. Forced to reinitialize on an over-forecast on the 17th, MAQSIP could have retained too much near-surface and mixed layer ozone and precursors through the 18th. Had it been available, an objective analysis system for chemical observations could have corrected these initial conditions for the 00z 8/18 run.

By August 19th, the model "catches up," when the forecast bias returns to nominal and the mean absolute error, shown in Figure 9, reduces to around 10PPB. In fact, the monitor data show that the number of sites experiencing concentrations >80PPB dropped from an average of about 150 on the 17th and 18th, to 10 monitors on August 19th. It is likely then, that synoptic forcing—possibly a cold front—provided a domain wide cleansing effect that the model was able to respond to on the 19th.

Figure 9 shows the MAEs for MAQSIP surface layer first-day ozone forecasts for the 19 days analyzed. It reveals large MAEs for each of the concentration bins on August 4th, but in general, the MAEs are clustered fairly tightly around 12 PPB. The MAEs for observations above 80PPB are a little higher , 16.3PPB on average, indicating that MAQSIP has marginally more error in forecasting these higher concentrations. Further, since the forecast bias error is negative for this "high concentration" bin, a majority of the errors in forecasting high concentrations are concentration under-forecasts vis-à-vis the monitors.

It should be noted that, for forecasting purposes, EPA as well as most state agencies issue ozone action "alerts" based on predictions that have ranges at least as large as MAQSIP's MAEs. For example, forecasts are usually given in terms of an AQI (Air Quality Index) code of green, yellow, orange, or red, where green represents forecast concentrations in a range 1-64PPB, yellow in the range 65-84PPB, orange in the range 85-104PPB, and red at 105PPB or above. Thus, the ability of a forecast model to hit a range of expected values rather than a specific number is relevant to the current public alert approach. MAQSIP's 1998 average bias of about –15PPB for observed values >80PPB could lead, on average, to an under-forecast of one "code level" about three-quarters of the time when observed concentrations are in the range 84-119PPB. Above 119PPB, this bias wouldn't affect the code level at all: code red would still be code red even if the observed level was 150PPB and MAQSIP forecast 135PPB. Thus, it is in the intermediate range that current operational forecasts would be affected by the bias. We will summarize some reasons that we think contribute to this low bias above 80PPB below.

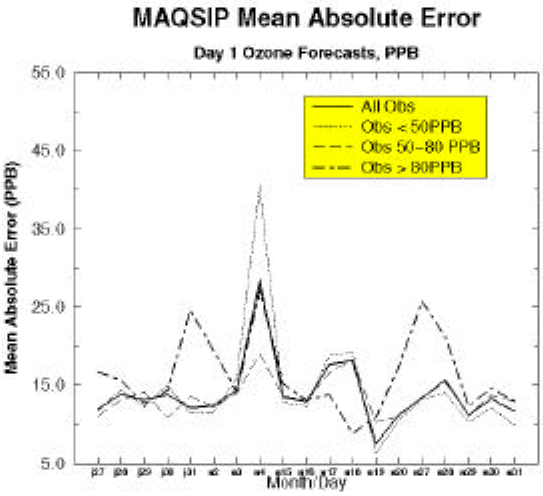

Figure 9. MAQSIP Mean Absolute Error Statistics for 19 Day-1 Forecasts

Because ozone forecasting is so reliant on thresholds and ranges of concentrations, it is appealing to consider the *threat score* as a statistic that may be useful for understanding the predictive skill of the model, even though no one statistic alone can be conclusive. It is used extensively in quantitative precipitation forecast (QPF) analysis, and is defined as:

$$TS = C/(F + R - C) \qquad (1),$$

where C is the number of stations correctly forecast to receive greater than a threshold amount of precipitation, F is the number of stations actually forecast to receive that amount, and R is the actual number of stations observing that amount (Anthes, 1983). The threat score measures the relative overlap or collocations between predictions and observations of the given threshold value or above: a perfect score is 1, and no overlap yields 0.

For ozone forecasting, we applied this to two thresholds by the following rule: if the observations showed that >5% of all monitors reported values >80PPB, then 80PPB was the chosen threshold. These days were broadly classified as "dirty." If the observations showed that <5% of all monitors exceeded 80PPB for the 1-6PM afternoon hours, then we chose 60PPB as the threshold, and these days were classified as "clean." For the 19 days studied, there were 16 "dirty" days and only 3 "clean" days.

Results of the threat score calculation are shown in Figure 10. *Again, no trend information can or should be inferred between episodes*. It is seen that for most days, the first-day forecast TSs are better than the second-day forecast TSs, consistent with r_{corr}, R^2, and MAE (not shown) results. The average threat score for all first-day forecasts is .249; the average for all second-day forecasts is .195. A threat score of .25 says that half of all monitors exceeding the threshold were actually forecast to exceed it, and an identical number were forecast to exceed it that didn't. Such a result can occur because of phase errors in the meteorological forecast timing—we will be using a structure function to examine this possibility in the near future (results not available). Too, in our case, elimination of the low bias above the 80PPB threshold would no doubt improve the scores.

To put these results in context, it is useful to borrow some insights from the QPF community. Anthes (1983) points out that subjectively prepared threat scores of NMC predicted precipitation for 0-24 hour forecasts between the years 1960-1980 had an annual average of about 0.20. Between the years 1976 and 1983, the old LFM model mean monthly threat scores were on the order of .35-.40. We point out that NAQP, which uses a complex atmospheric chemistry modeling system, is still in its infancy, and note that our TSs are not dissimilar to NMC reported "success" at QPF in the 1970's.

Based on all of these analyses, we have devised a tentative, "work-in-progress" system for ranking the objective skill of a given forecast. We define a forecast to have substantial objective skill if it meets the following criteria: (1) threat score > .20; (3) R^2 > .2; (4) MAE <15 PPB for monitors reporting above threshold; and (4) absolute bias error <7PPB. Under these criteria, out of the 19 analyzed forecasts, 10 have substantial skill, and two more are nearly that good. Five others have lesser degrees of skill, meeting at least two of the criteria, and only two, Aug 2nd and Aug 4th, fail three out of the four tests.

5. DISCUSSION AND CONCLUSIONS

In this paper we have described the NCSC/PSU Numerical Air Quality Prediction Project and presented results from analyses of 19 forecasts made during summer 1998. We have identified some major sources of error in the forecasts, the principal one being the lack of an objective atmospheric chemical initialization package. We have shown that MAQSIP tends to under- forecast ozone concentrations at monitor values exceeding 80PPB by, on average, about 15PPB. Further, we have discussed several shortcomings of the modeling system that may contribute to this bias, in order of importance: (1) an outdated emissions inventory; (2) over-ventilation of the PBL by MM5 winds; (3) small negative temperature biases in MM5 leading to slower reaction rates.

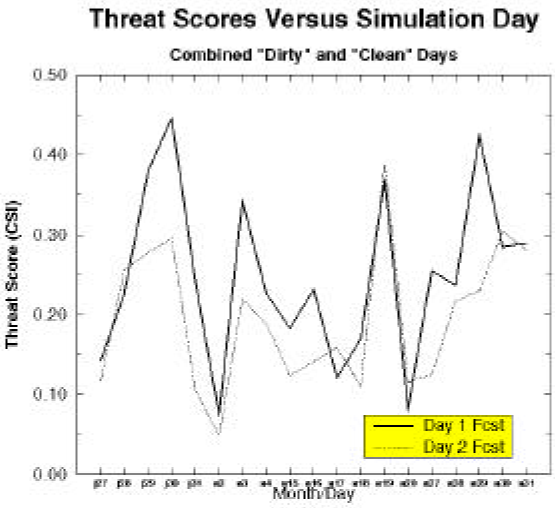

Figure 10. MAQSIP Threat Score Statistics for 19 Day-1 and Day-2 Forecasts

A fourth reason, we feel, is that the 36km grid is too coarse to resolve finer detail that influences high monitor readings on a more local scale. As pointed out by Rao et al., 1997, the level of spacing needed to capture all of the ozone exceedances in a monitoring network would be about 5km. We expect, thus, that 5km would be an upper limit on model resolution required to resolve all of the fine-scale variability contained in the monitor data.

Despite the shortcomings of our current system, we concluded that, objectively, 17 of 19 forecasts have some measurable forecast skill, with ten of those 17 meeting four different skill criteria. Further, we showed that the underprediction bias at 80PPB and above would affect operational forecasts only in the intermediate range between 84 and 119PPB, and that this would lead to an underforecast of one public alert "code color level" about 75% of the time for this range of values. Given the known, correctable shortcomings of this early operational system, these results provide great hope for NAQP improvement in the future, including eventual acceptance as an important operational forecasting tool. During summer 1999, we undertook NAQP on a 15-km resolution fine domain over the state of Texas. Qualitative results suggest a greater dynamic range in the predictions as compared to the 36km results. Analysis is in process. Current plans call for an expansion of forecast capabilities during summer 2000.

Disclaimer. Although the research described in this article has been funded wholly or in part by the United States Environmental Protection Agency through Grants #R825210 and #R825211 to MCNC and Grant #R825254 to Penn State University, it has not been subjected to the Agency's required peer review and policy review and therefore does not necessarily reflect the views of the Agency and no official endorsement should be inferred. Mention of trade names/products does not imply official endorsement.

REFERENCES

Alapaty, K., J.E. Pleim, S. Raman, D.S. Niyogi, and D.W. Byun, 1997: Simulation of atmospheric boundary layer processes using local- and nonlocal-closure schemes. *J. App. Met.* **36**,3,214-233.

Anthes, R., 1983: Review: Regional models of the atmosphere in middle latitudes. *Mon. Wea. Rev*., 111, 1306-1335.

Bott, A., 1989: A positive definite addiction scheme obtained by nonlinear renormalization of the advective fluxes. *Mon. Wea. Rev*., **117**, 1006-1015.

Coats, C.J, 1995: The EDSS/Models-3 I/O applications programming interface. *http://envpro.ncsc.org/products/ioapi/index.html* MCNC Env. Programs, RTP NC.

Coats, C.J, 1996: High performance algorithms in the sparse matrix operator kernel emissions (SMOKE) modeling system. *Ninth AMS Joint Conference on Applications of Air Pollution Meteorology with A&WMA*. Amer. Meteor. Soc., Atlanta, 584-588.

Coats, C.J., J.N.McHenry, A. Lario-Gibbs, and C.D.Peters-Lidard, 1998: MCPL(): A drop-in MM5-V2 module suitable for coupling MM5 to parallel environmental models; with lessons learned for the design of the weather research and forecasting (WRF) model. *The Eighth PSU/NCAR Mesoscale Model Users' Workshop.* MMM Division, NCAR, Boulder, Co., 117-120.

Deng, A., N.L.Seaman, and J.S.Kain, 1998: A parameterization for detraining cumulus-cloud mass into dissipating layer clouds. *Twelfth Conference on Numerical Weather Prediction.* Amer. Meteor. Soc., Dallas, 120-123.

Dudhia, J., 1989: Numerical study of convection observed during the winter monsoon experiment using a mesoscale two-dimensional model. *J. Atmos. Sci.,* **46**, 3077-3107.

EPA-Office of Research and Development, 1999: Science Algorithms of the EPA Models-3 Community Multiscale Air Quality (CMAQ) Modeling System. *Technical Report EPA-600/R-99/030*, ORD, Washington, DC.

EPA-Office of Air Quality Planning and Standards (OAQPS), 1999: Ozone mapping project, environmental monitoring for public access and community tracking. *http://www.epa.gov/airnow/*

Gayno, G.A., 1994: Development of a higher order, fog producing boundary layer model suitable for use in numerical weather prediction. M.S. Thesis. Dept. of Meteorology, Penn State University.

Gery M.W., G.Z. Whitten, J.P. Killus, and M.C. Dodge, 1989: A photochemical kinetics mechanism for urban and regional scale computer modeling. *J. Geophys. Res*., **94**, 12925-12956.

Grell, G.A., J. Dudhia, and D.R.Stauffer, 1994: A description of the fifth-generation Penn State/NCAR mesoscale model (MM5). NCAR Technical Note TN-398, Bldr. Co.

Kain, J.S., and J.M. Fritsch, 1993: Convective parameterization for mesoscale models: the Kain-Fritsch Scheme. **Meteor. Mono. 24, No. 46**, Amer. Meteor. Soc., Boston.

National Research Council, Board on Atmospheric Sciences and Climate, 1998: The Atmospheric Sciences Entering the 21st Century. National Academy Press, Washington, DC.

MCNC, 1999: Numerical air quality prediction (NAQP) web-site. *http://envpro.ncsc.org/NAQP*

Madronich, S., 1987: Photodissociation in the atmosphere. 1. Actinic flux and the effects of ground reflections and clouds. *Jou. Geophys. Res*. **92**, D8, 9740-9752.

McHenry, J.N., N.L. Seaman, C. Coats, A. Lario-Gibbs, J. Vukovich, N. Wheeler, and E. Hayes, 1999: Real-time nested mesoscale forecasts of lower tropospheric ozone using a highly optimized coupled model numerical prediction system. *Symposium on Interdisciplinary Issues in Atmospheric Chemistry*. Amer. Meteor. Soc., Dallas, 125-128.

Odman, M.T., and C.L. Ingram, 1996: Multiscale Air Quality Simulation Platform (MAQSIP): Source Code Documentation and Validation. *Technical Report ENV-96TR002-V1.0*, MCNC-North Carolina Supercomputing Center, RTP, NC.

Pechan, E.H., and Assoc., 1996: Ozone transport assessment group emissions inventory development report, volume II, version 1.0. *Pechan Report No. 68-D3-0035*, Durham, NC.

Penn. State, 1999: NAQP web-site: http://mm5.met.psu.edu/mm5/MM5maps.html

Rao, S.T., I.G. Zurbenko, R. Neagu, P.S. Porter, J.Y. Ku, and R.F. Henry, 1997: Space and time scales in ambient ozone data. *Bull. Amer. Meteor. Soc*., 78, 10, 2153-2165.

Srivastava, R.K., D. S. McRae, and M.T. Odman, 1994: Governing Equations of Atmospheric Diffusion. *Technical Report ENV-94TR001*, MCNC-North Carolina Supercomputing Center, RTP, NC.

Walcek, C.J., R.A. Brost, J.S. Chang, and M.L. Wesely. 1986. SO_2, sulfate, and HNO_3 deposition velocities computed using regional land-use and meteorological data. *Atmos. Environ.* **20**:949-964.

Warner, T.T., and N.L. Seaman, 1990: A real-time weather-prediction system used for research, teaching, and public service at the Pennsylvania State University. *Bull. Amer. Meteor. Soc.* **71**, 6, 792-805.