

**Current Status,
Preliminary Validation Study
of the
Australian Air Quality Forecasting System**

Prepared by Project Management Committee
Chair: Dr Peter Manins
CSIRO Atmospheric Research
PMB 1 Aspendale 3195
Australia



Contact: peter.manins@dar.csiro.au

Table of Contents

1. Current Status of AAQFS	2
Emissions Inventory	2
Meteorological modelling	2
Chemical transport model	2
Validation	3
Data display/storage/communication	3
Conclusions	3
References	3
2. Comparison of the AAQFS with Other Air Quality Forecasting Systems	5
3. Preliminary Validation Study: Meteorological Component	7
Introduction	7
Model Performance: Overall	8
Model Performance: Overall	9
Model Performance: Case Study	11
Conclusions	16
4. Preliminary Validation Study: Chemical Transport Model Performance	18
Introduction	18
Victoria: Observed and forecast air pollutant time series	19
Victoria: Residual analysis	31
Delivery performance: Victoria	33
NSW: Observed and forecast air pollutant time series	36
NSW: Residual analysis	47
NSW: Delivery performance	49
Conclusions	50
5. Acknowledgments	51

1. Current Status of AAQFS

The AAQFS consists of five components: an emissions inventory component, a meteorological modelling system, a chemical transport model, a validation system, and a data display, storage and communications component.

Since program inception in 1998, considerable work has been undertaken on all five components in order to develop a state-of-the-art system with the potential for generating accurate, high-resolution air pollution forecasts for regions in NSW and Victoria.

The progress to date has been fully documented in six Progress Reports to Environment Australia, in conference proceedings, and in published research papers. A list of these is included here.

Emissions Inventory

Both the NSW and the Victorian emissions inventories have been updated for AAQFS. In the case of NSW, the 1992 Metropolitan Air Quality Study (MAQS) inventory has been upgraded to the year 2000 and substantially enhanced. The inventory is available in two spatial resolutions. The regional-scale inventory has a resolution of about 5 km and covers the greater MAQS study region (includes Newcastle, Sydney and the Illawarra). Urban scale inventories (1 km resolution) have been developed for Sydney, Newcastle and Wollongong. Note however, that the urban-scale forecasts are currently only undertaken for Sydney. The regional-scale Victorian inventory covers the State at a resolution of approximately 5 km. An urban-scale inventory (1 km) covers the Port Phillip Bay region (includes Melbourne and Geelong).

A module for meteorologically dependant emissions modelling has been developed and implemented. This includes biogenic emissions of reactive organic compounds and sea-salt aerosol. A module for near-road sub-grid-scale emissions has also been developed and implemented, but has not yet been tested within the context of the full system.

Methodologies for including wind blown dust and bushfire emissions have been developed but have not yet been implemented into the system.

Meteorological modelling

The Bureau of Meteorology's operational forecasting system (LAPS) has been substantially enhanced for the AAQFS project.

The LAPS model horizontal resolution has been increased from 25 km to 5 km for the AAQFS project. The number of vertical levels has been increased from 19 to 29 (which includes 9 levels within the lowest kilometre). Data sets for vegetation, land use and soil texture at 5 km resolution have been implemented along with high-resolution topography and sea surface temperature. The physics has been updated to include advanced algorithms for land surface interaction and boundary-layer representation.

These modifications make it possible to resolve mesoscale features such as bay breezes, sea breezes, topographically-forced eddies, all of which may have a significant impact on air quality.

A new method for computing the diagnostic 10 m wind, screen temperature and screen dewpoint temperature has also been developed but is not yet implemented.

Chemical transport model

A new chemical transport model (CTM) has been developed for the project. The CTM treats the emission, dispersion and chemical transformation of photochemical smog and precursors, aerosols and some air toxics. The CTM software has been specifically designed to enable rapid execution on the BoM/CSIRO supercomputer, thus enabling air quality forecasts to be rapidly undertaken for multiple, large computational domains.

A revised, condensed chemical mechanism (GRS2) has been developed but not yet implemented. A near-road sub-grid-scale dispersion model has also been developed and awaits implementation.

Validation

Meteorological methods for validation have been developed and implemented for the 10 m wind, screen temperature and screen dewpoint temperature using the METAR/SYNOP network, and for upper winds and temperature using AMDAR data.

Data from vertical profilers are now becoming available and a methodology for incorporating these data has been developed but has not yet been implemented. Daily values of ventilation based on measurements are also compared to predicted values. Trajectories are now computed to aid in the interpretation of model results and validation.

Methods for validation of air quality variables have been developed and implemented. These include graphical display of observed and forecast air pollutant concentration time series, and time series of the ratios of concentrations, and statistical methodologies such as residual analysis.

Following sections of this report discuss a preliminary validation study done for the four month period 13–27 November 2000.

Data display/storage/communication

The spatial fields of ground-level air pollutant concentrations are archived each day. A comprehensive Web site display system has been developed and implemented. The site includes spatial plots of air quality concentrations and meteorological conditions, air quality indices, validation measures and resource documents.

A system of emails to the EPA's showing plots of the highest concentration of key pollutants has been developed and implemented. A data display system has been developed and upgraded to include both meteorology and air quality variables. This is currently being further enhanced and will be provided to the participating EPAs so that investigations of archived events are easily done.

The methodology for plotting the ground-level air pollutant fields is being refined. Once completed, the daily forecast products will be available at an earlier time, and will also be substantially faster to download from the Web site.

Conclusions

The Demonstration System is substantially complete and has been operating since August 2000. We will be progressively implementing the remaining modules into the full system over the next few weeks and will be conducting further validation of the System during the summer photochemical smog season.

References

Many of the following papers and reports are available on the Web at <http://www.dar.csiro.au/info/Images/lib/pubsearch.html>.

- Manins, P. C., (Chair of Committee). (1999). Air quality forecasting for Australia's major cities: 1st progress report [to Department of the Environment and Heritage] / prepared by Project Management Committee: Environment Protection Authority (Victoria), Environment Protection Authority of NSW, Bureau of Meteorology Research Centre, CSIRO Energy Technology and CSIRO Atmospheric Research CAR: SB/1/407. Aspendale, Vic.: CSIRO Atmospheric Research. 29, [27] p.
- Manins, P. C., (Chair of Committee). (1999). Air quality forecasting for Australia's major cities: 2nd progress report [to Department of the Environment and Heritage] / prepared by Project Management Committee: Environment Protection Authority (Victoria), Environment Protection Authority of NSW, Bureau of

- Meteorology Research Centre, CSIRO Energy Technology and CSIRO Atmospheric Research CAR: SB/1/407. Aspendale, Vic.: CSIRO Atmospheric Research. 35, [6], 5, 16 p.
- Manins, P. C., (Chair of Committee). (2000). Air quality forecasting for Australia's major cities: 3rd progress report [to Department of the Environment and Heritage] / prepared by Project Management Committee: Environment Protection Authority (Victoria), Environment Protection Authority of NSW, Bureau of Meteorology Research Centre, CSIRO Energy Technology and CSIRO Atmospheric Research CAR: SB/1/407. Aspendale, Vic.: CSIRO Atmospheric Research. 41, [6], 5, 14 p.
- Manins, P. C., (Chair of Committee). (2000). Air quality forecasting for Australia's major cities: 4th progress report [to Department of the Environment and Heritage] / prepared by Project Management Committee: Environment Protection Authority (Victoria), Environment Protection Authority of NSW, Bureau of Meteorology Research Centre, CSIRO Energy Technology and CSIRO Atmospheric Research CAR: SB/1/407. Aspendale, Vic.: CSIRO Atmospheric Research. 86 p.
- Manins, P. C., (Chair of Committee). (2000). Air quality forecasting for Australia's major cities: 5th & 6th progress report [to Department of the Environment and Heritage] / prepared by Project Management Committee: Environment Protection Authority (Victoria), Environment Protection Authority of NSW, Bureau of Meteorology Research Centre, CSIRO Energy Technology and CSIRO Atmospheric Research CAR: SB/1/407. Aspendale, Vic.: CSIRO Atmospheric Research. 56 p.
- Cope, M.E., Manins, P.C., Hess, G.D., Mills, G.A., Puri, K., Dewundege, P., Tilly, K. and M. Johnson (1998) 'Development and application of a numerical air quality forecasting system', Proc. 14th International Clean Air and Environment Conference, Melbourne, 18–21 October: Clean Air Society of Australia and New Zealand; 1998, 353–358.
- Hess, G. D., Cope, M. E., Lee, S., and Tory, K. (1999). LAPS and the Australian Air Quality Forecasting System. In: Parallel computing in meteorology and oceanography: abstracts of presentations at the eleventh annual BMRC Modelling Workshop, BMRC, Melbourne, D. J. Jasper, and P. J. Meighen (editors) (BMRC Research Report, 75) . Melbourne: Bureau of Meteorology Research Centre. p. 35–40.
- Hess, G. D., and Cope, M. E. (1999). The Australian Air Quality Forecasting System. In: Annual BMRC-DAR Research Discussion Meeting: meeting report, Melbourne, Vic. Melbourne: Bureau of Meteorology. p. 5.
- Manins, P.C., Cope, M.E., Hurley, P.J., Noonan, J.A. and W.L. Physick (1999) 'Regional prognostic climatology. In: Proceedings of fifth International Joint Seminar on the Regional Deposition Processes in the Atmosphere, Seoul, Korea. SU Park (editor). [Seoul] Department of Atmospheric Sciences, Seoul National University, 141–146.
- Manins, P.C., Cope, M.E., Hurley, P.J., Noonan, J.A. and W.L. Physick (1999) 'Prognostic urban climatology. [electronic publication] In: Proceedings of the 15th International Congress of Biometeorology & International Conference on Urban Climatology. [CD ROM] : Wesley Conference Centre, Sydney, R.J. de Dear and J C Potter (editors). NSW Macquarie University. p. ICUC17.1.
- Cope, M. E., Hess, G. D., Lee, S., Azzi, M., Carras, J. N., Wong, N., and Young, M. A. (1999). Development of the Australian Air Quality Forecasting System: current status [electronic publication]. In: Proceedings of the 15th International Congress of Biometeorology & International Conference on Urban Climatology [CD-ROM], Wesley Conference Centre, Sydney, R. J. de Dear, and J. C. Potter (editors). N.S.W.: Macquarie University. p. ICUC20.5.
- Manins, P. C. (2000) Systems for prognostic air quality forecasting. In Air quality forecasting for Australia's major cities : 5th and 6th progress report / prepared by Project Management Committee ; chair: Peter Manins. Aspendale, Vic. : CSIRO Atmospheric Research, 2000. p.48–56.
- Hess, G.D., Cope, M.E., Lee, S., Manins, P.C., Mills, G.A., Puri, K. and K. Tory (2000) 'The Australian Air Quality Forecasting System', AMOS Bulletin **13**, 67–73.

2. Comparison of the AAQFS with Other Air Quality Forecasting Systems

Forecasting air quality is a relatively new concept. However there are a number of groups around the world providing such air quality forecasts. It is worthwhile to see how the Australian Air Quality Forecasting System (AAQFS) compares to these other systems.

- The most common ozone forecasting systems are based on statistical regression for predictants such as screen temperature, mean sea-level pressure, solar radiation, wind speed, and boundary-layer height. Examples of groups practicing this method include EPA-VIC, EPA-NSW, and the University of Maryland which forecasts for Washington, DC, Baltimore, MD, Philadelphia, PA and Virginia. There are many other groups using this technique around the world.
- The next level of complexity is a box-model forecast. The UK Department of Environment, Transport and the Regions uses this methodology to forecast for ozone, nitrogen dioxide, sulfur dioxide, carbon monoxide and PM10.
- The Cambridge Environmental Research Consultants use the UK Meteorological Office's wind fields, a Gaussian dispersion model that they have refined, and the condensed chemistry developed by CSIRO (GRS) to produce suburban-scale forecasts for York, Leicester, South Gloucestershire, Bristol, North Somerset, Bath and North East Somerset.
- GMD FIRST performs forecasts for Berlin using a simplified meteorological model called REGOZON. The model is restricted to fair weather, stagnant air flow, and relatively flat terrain without strong orographic features or strong horizontal temperature gradients (no land/sea breezes). It uses a comprehensive chemistry, CBIV, and has about 2 km horizontal resolution.
- The Air Resources Laboratory performs (unverified) forecasts for ozone for the Eastern half of the United States and for the region near Houston, Texas. They employ a hybrid Lagrangian puff model and a comprehensive chemistry, CBIV, with 50 km horizontal resolution. The wind fields are calculated by Eulerian models from the US National Weather Service.
- The Danish Atmospheric Chemistry Forecast System model computes backward trajectories from receptor points, and then retraces the paths using a comprehensive chemistry EMEP MSC-W. It has 50 km horizontal resolution. The wind fields are calculated by the Eulerian model HIRLAM.
- Another Danish system, called DMU-ATMI-THOR uses an Eulerian meteorological model at 25 km horizontal resolution and a comprehensive chemistry, CBIV. Air quality predictions are made to high resolution, including street canyons, using the same 25 km windfield.
- At the North Carolina Supercomputing Center they use a meteorological model called MM5 and an Eulerian chemical model MAQSIP with a comprehensive chemistry, CBIV. They perform (unverified) forecasts for the Eastern half of the United States at 36 km horizontal resolution.
- The Canadian Environment Service is developing an air quality forecasting system and is using AAQFS as a model of what is possible.
- **AAQFS** uses the LAPS model with a horizontal resolution of 5 km and an Eulerian chemical model with a horizontal resolution of 1 km over urban areas and 5 km over rural areas. The chemical transport model has 25 pollutants and includes the condensed chemistry, GRS. A comprehensive chemistry, CB99, is available for offline use. Extensive verification is performed. The domains of application demand a high level of modelling skill because of the

presence of topographically forced flow (drainage winds, lee eddies), bay and sea breezes, and synoptic-scale/mesoscale interactions.

The World Meteorological Organization has initiated a project called GURME. The GURME project arose in response to the requests for assistance by many National Meteorological Services dealing with urban issues, and in recognition that the management of urban environments requires special attention. A major need that has been identified is for WMO to assist countries in providing air quality services of high quality, particularly the capabilities to provide meteorological and air quality forecasts of urban environments. This need also entails measurement efforts that support operational and verification aspects of forecasting, and performed in co-operation with appropriate agencies. The status of the program can be found on the Web at <http://www.cgrer.uiowa.edu/people/carmichael/GURME/GURME.html>.

Following the presentation of a paper about AAQFS at the Millennium NATO/CCMS International Technical Meeting on Air Pollution Modelling and Applications in the U.S. in May this year, the AAQFS team was invited by Greg Carmichael (co-chairman GURME- Scientific Advisory Group) to participate in a WMO/GURME workshop in order to provide and participate in a series of lectures on urban air quality forecasting (held in Kuching, Malaysia 14–17 August 2000). Two members of the AAQFS team (Kamal Puri and Peter Manins) responded to the invitation and presented a series of lectures to the assembled group. GURME has a made available on the Web site copy of the presentation by Peter Manins at <http://www.cgrer.uiowa.edu/people/carmichael/GURME/prediction.pdf>. There are also links there to AAQFS reports and other CSIRO modelling information.

The outstanding message from the above, reinforced by the GURME objectives and the promotion by GURME of AAQFS, is that the AAQFS is world-leading in its integration of high resolution air quality and meteorological forecasting into one system that brings together the environment agencies and the meteorological agency of a country. The focus on verification, an integral component of AAQFS, is also unusual but essential for quality results. AAQFS is the benchmark against which other groups are beginning to measure their approaches and systems.

As further evidence of these statements, recently Bruce Hicks, Director of the Air Resources Laboratory, US National Oceanographic and Atmospheric Administration, has been in touch with us. He asked us to send him all of our published papers about the AAQFS, because he is organising the US Government's effort to develop an urban air quality forecasting system along similar lines.

In summary, AAQFS is a world-class air quality forecasting system. It is leading the way in model resolution, advanced algorithms to deal with complex meteorology and chemistry, its treatment of an extensive range of pollutants, including air toxics and PM_{2.5}, and its detailed verification program.

3. Preliminary Validation Study: Meteorological Component

G. D. Hess, K. Tory and K. Puri

Bureau of Meteorology Research Centre, Melbourne

Introduction

The period, 13–27 November 2000 has been chosen to carry out a preliminary validation study of the Australian Air Quality Forecasting System (AAQFS). This period comes at the end of the first AAQFS Demonstration Period, which ran between 7 August to 4 December 2000. Here we document the performance of the meteorological component of the system.

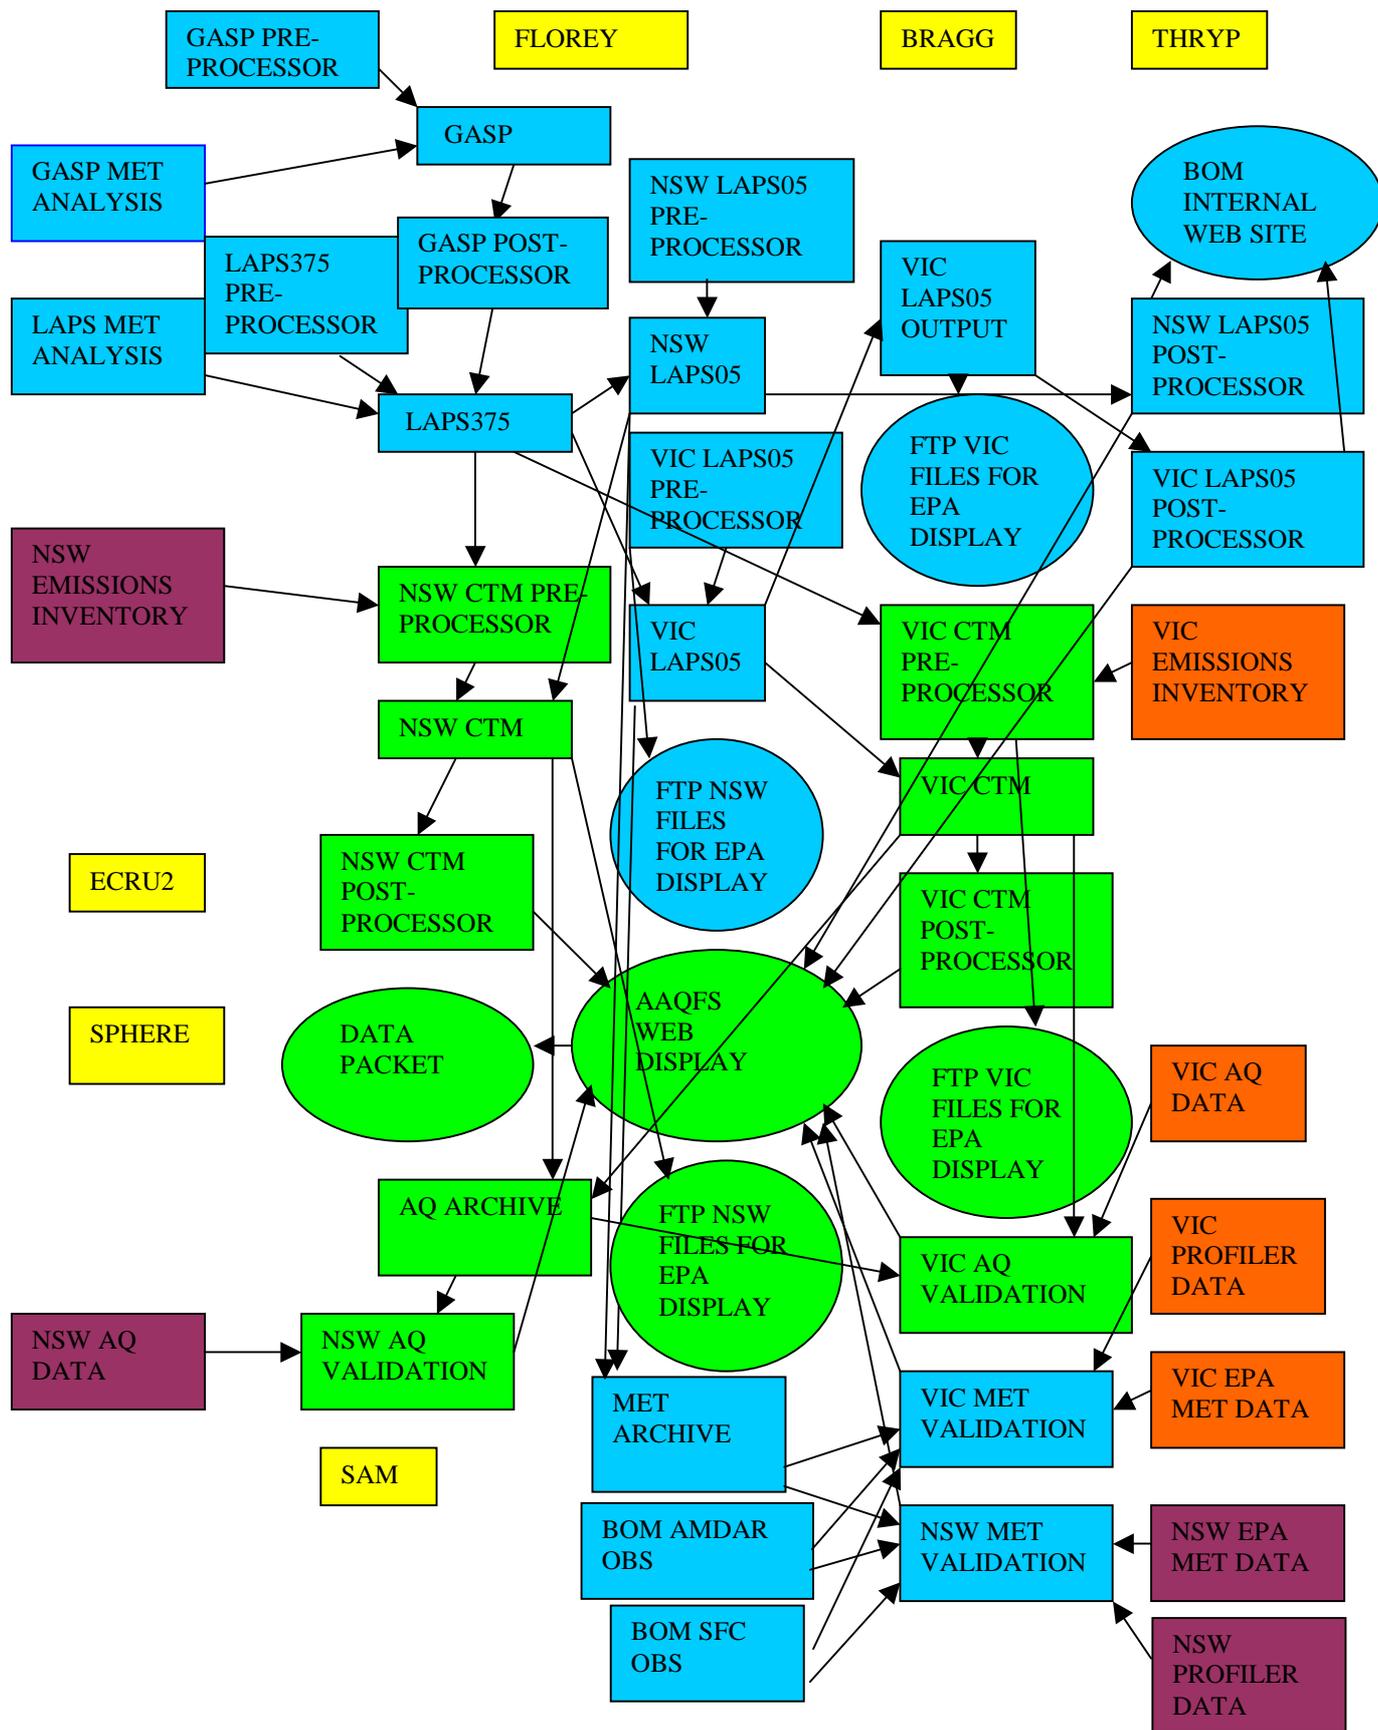
The AAQFS is a complex numerical forecasting system requiring collaborative input from the Bureau of Meteorology (BoM), CSIRO Atmosphere Research (CAR), the Environment Protection Authority of Victoria (EPA-VIC) and the Environment Protection Authority of New South Wales (EPA-NSW) for the daily operations. The computations and communications are carried out on a variety of computers ranging from supercomputers to workstations to personal computers. The coordination of the logistics of this network is extremely involved. Figure 3.1 shows a flow diagram for the system, which was current during the Demonstration Period up to the end of the Sydney Olympics. The blue regions indicate processes carried out by BoM, the green regions indicate those carried out by CAR, the orange regions indicate input by EPA-VIC and the plum regions indicate input by EPA-NSW. The yellow boxes indicate the computers involved, where Bragg and Florey are supercomputers, and Thryp, Sam, Sphere and ECRU2 are workstations. The boxes indicate computing processes and circles and ellipses indicate displays and output.

Since the end of the Olympics, the configuration of the system has changed dramatically. The NEC SX4 supercomputer (bragg), shared jointly by BoM and CAR, and the CRAY J-90 supercomputer (not shown in the flow diagram), owned by BoM, have been phased out. Consequently this has meant that computing and users on those machines have migrated to florey, greatly increasing the work load and congestion there. Adding to the problems, the NEC system scheduler for florey is still in a state of development, and new archival retrieval system for BoM has not yet been installed. This has meant delays or sometimes failure in jobs starting and files copying. Recently, reorganisation of the computer logistics has improved the situation. In a few weeks time the second NEC supercomputer node will come online and this will substantially increase our computing capacity.

In a fully operational system, we expect a much simplified flow with fewer computers and much easier logistics.

Here we examine the performance of the meteorological component during the preliminary validation period. We do this in two ways: by studying the overall performance statistics and by examining the model performance in a case study of a complex meteorological event, involving three mesoscale circulations that are particularly conducive to high concentrations of air pollutants.

Figure 3.1: A simplified flow diagram for the AAQFS during the Demonstration Period up until the end of the Olympics.



Model Performance: Overall

We begin by looking at the statistics of the overall model performance. These daily statistics are computed for all observing stations and all hours over the domain. For the Victorian domain this comprises about 2000 daily observations; for New South Wales it comprises about 1500 daily observations.

In Figure 3.2 and Figure 3.3 we present daily variations of the bias, root-mean-square (RMS) error and correlation for the 10 m u-component and v-component of the wind, the screen temperature and the screen dewpoint temperature for the Victorian domain; Figure 3.4 and Figure 3.5 show the results for the New South Wales domain. The horizontal lines indicate the mean values.

A summary of the model performance is given in Table 3.1. Overall the LAPS model is performing very well. The wind bias is small and the correlations of the model data interpolated to the locations of the observing stations are 0.5–0.6. The screen temperature and dewpoint show an absolute mean bias of between 0.5 and 1°C with a correlation approaching 0.9 for the temperature and 0.6–0.8 for the dewpoint temperature.

Figure 3.2: The bias (blue), RMS error (red) and correlation (green) for (left) the 10 m u-component, (right) the 10 m v-component, of the wind for the Victorian domain.

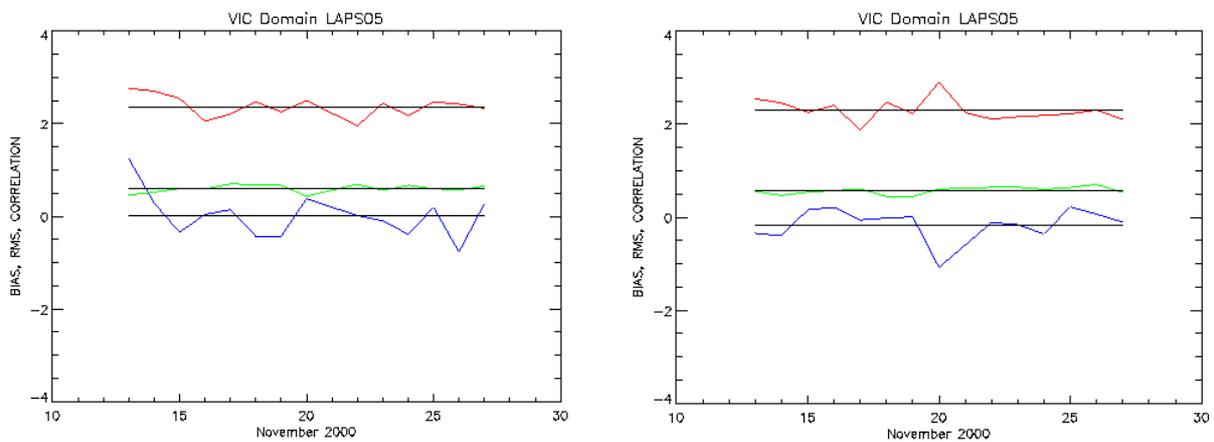


Figure 3.3: The bias (blue), RMS error (red), and correlation (green) for (left) the screen temperature and (right) the screen dewpoint temperature for the Victorian domain.

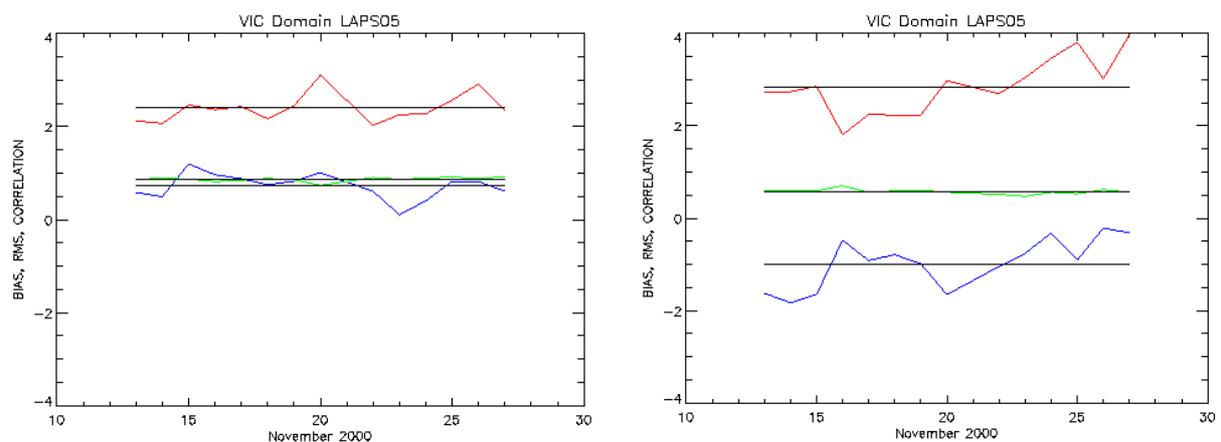


Figure 3.4: The bias (blue), RMS error (red), and correlation (green) for (left) the 10 m u-component, (right) the 10 m v-component, of the wind for the New South Wales domain.

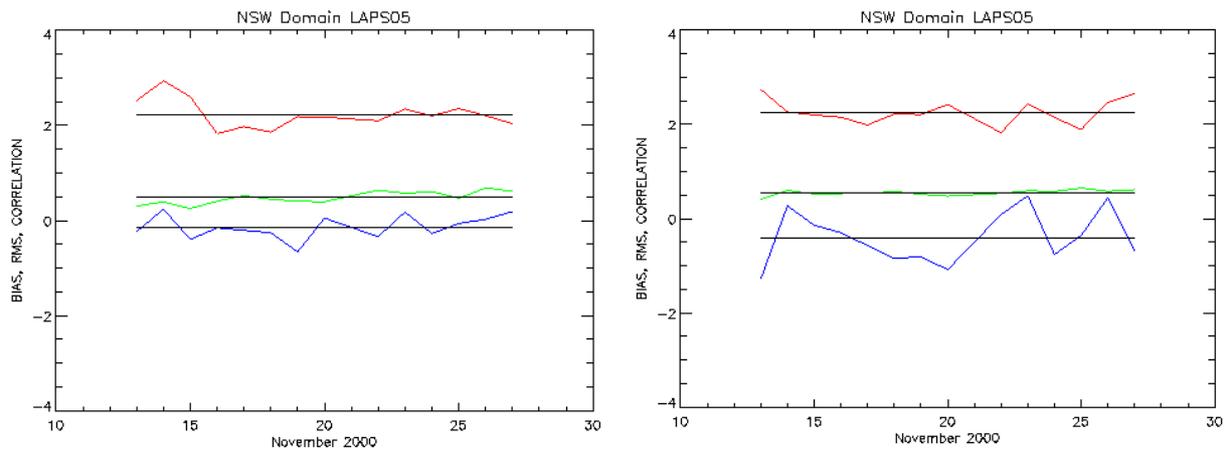


Figure 3.5: The bias (blue), RMS error (red), and correlation (green) for (left) the screen temperature, (right) the screen dewpoint temperature, for the New South Wales domain.

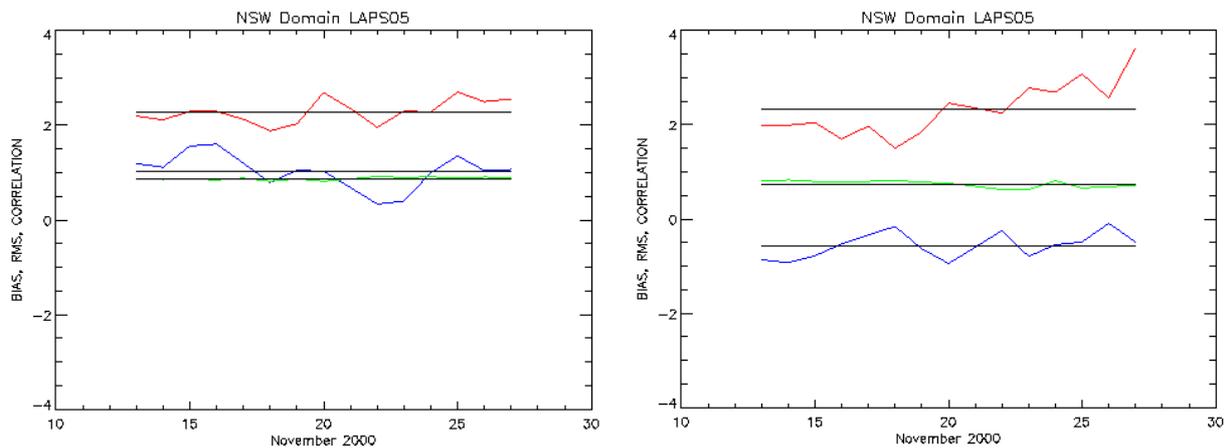


Table 3.1: Summary of the LAPS model performance for the period 13–27 November 2000. The mean bias and RMS error for the wind components are in m s^{-1} and for the temperatures are in $^{\circ}\text{C}$.

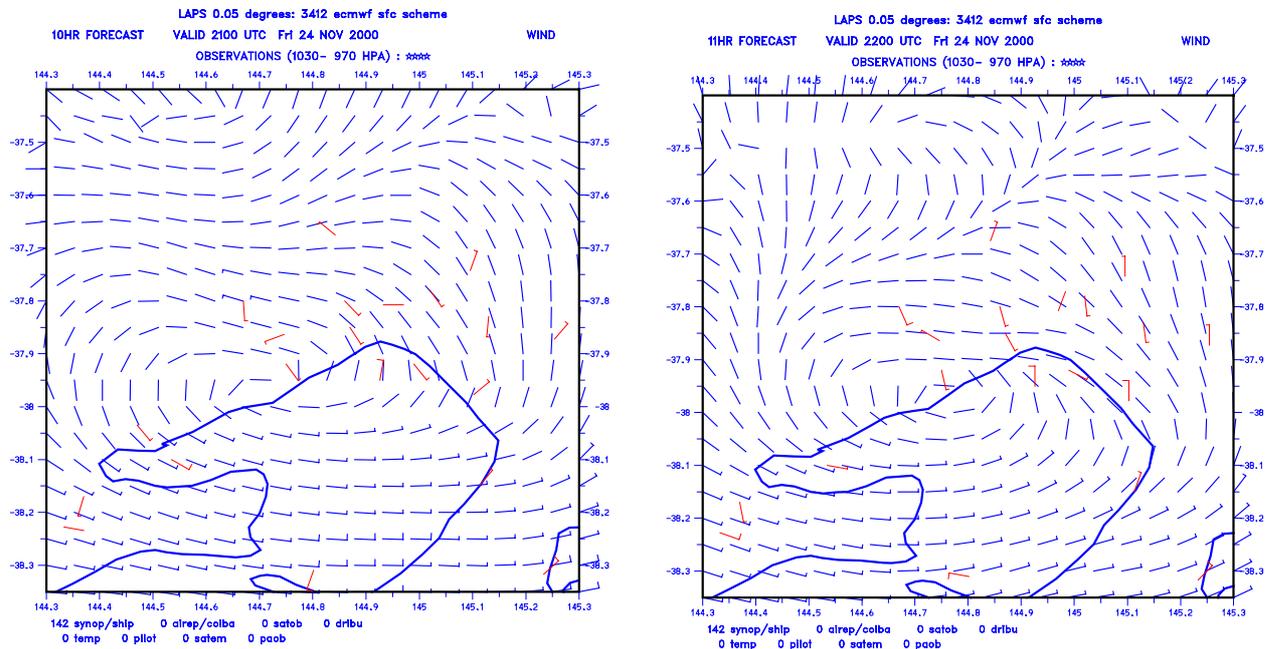
	Victoria	New South Wales
	10 m u-Component of Wind	
MEAN BIAS	0.022	-0.136
MEAN RMS-ERROR	2.372	2.233
MEAN CORRELATION	0.598	0.483
	10 m v-Component of Wind	
MEAN BIAS	-0.170	-0.403
MEAN RSM-ERROR	2.303	2.251
MEAN CORRELATION	0.577	0.550
	Screen Temperature	
MEAN BIAS	0.727	1.029
MEAN RSM-ERROR	2.413	2.288
MEAN CORRELATION	0.867	0.876
	Screen Dewpoint Temperature	
MEAN BIAS	-0.990	-0.562
MEAN RSM-ERROR	2.848	2.326
MEAN CORRELATION	0.576	0.746

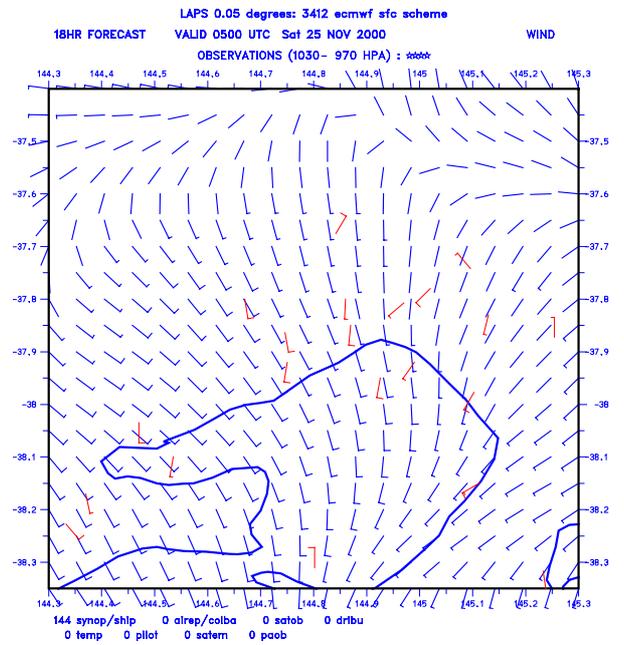
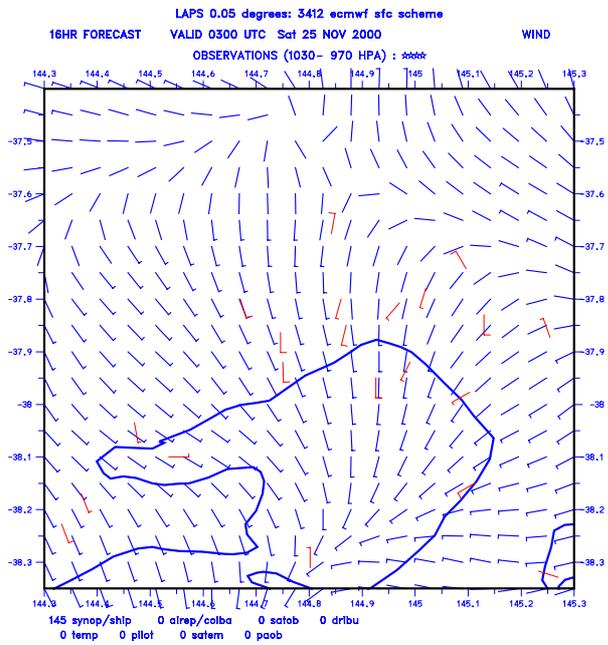
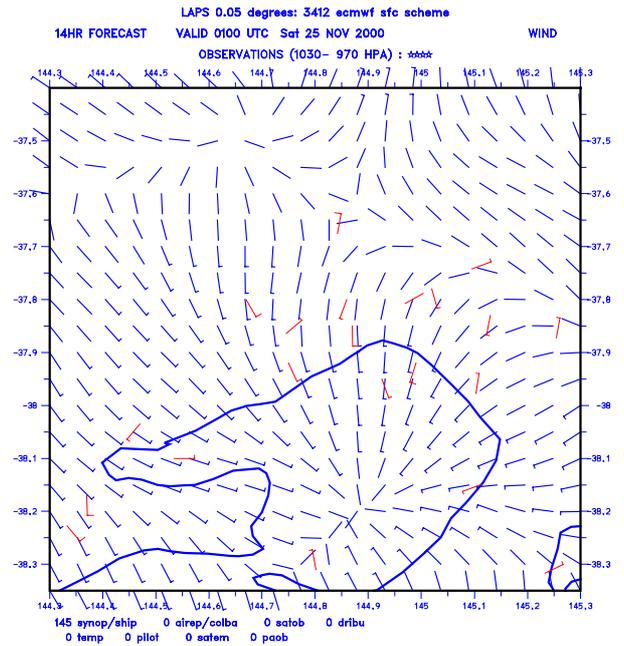
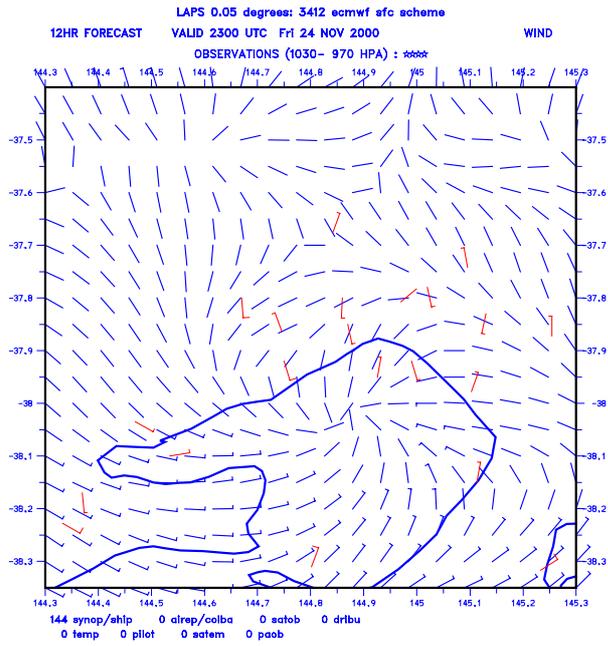
Model Performance: Case Study

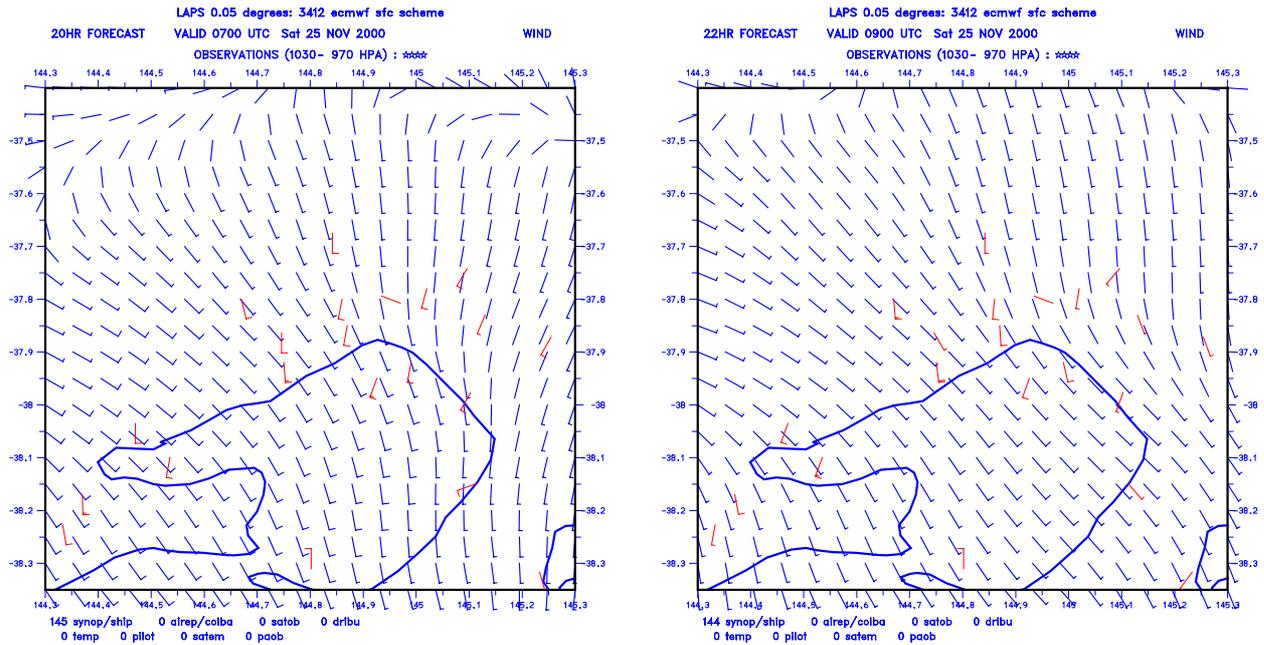
We now turn our attention to a case study for Victoria, starting 1100 UTC (10 pm Australian Eastern Daylight Saving Time – AEDST) on 24 November 2000. This day exhibited three mesoscale circulation patterns in turn, all potential contributors to high concentrations of pollutants: the Melbourne eddy, the bay breeze and the sea breeze. Figure 3.6 show a comparison of the evolving circulation patterns with observations. The Melbourne eddy forms over night in the wake of topography. The eddy is a result of the stable stratification of the air that forces the air to flow around the topography instead of over it. It transports NO_x and VOCs from the western industrial areas over the city and Port Phillip Bay. As the morning heating begins, turbulent mixing breaks up the eddy circulation and it is replaced by the bay breeze. The strength of the onshore flow associated with the bay breeze is proportional to the land/bay temperature gradient. This gradient increases most rapidly with time over the northern shores of the bay, leading to the southerly migration of the bay breeze divergence point. After this point passes the heads, the bay breeze begins its transition to a sea breeze. The model captures the timing of the onset of the bay breeze on the western side of the bay, but is one hour earlier than indicated by the observations at Fawkner Beacon near the top of the bay.

(The actual timing difference is unknown because the model output frequency is hourly and the observational frequency is also hourly: the timing difference could be less than an hour.) In general the modelled flow pattern agrees with the observations.

Figure 3.6: Comparison of the flow pattern for 10 m winds with observations for the forecast beginning 1100 UTC (10 pm AEDST) 24 November 2000. Note the presence of the Melbourne eddy in the top two panels. The bay breeze replaces this circulation in the next two panels. The bay breeze makes a gradual transition to a sea breeze in the remaining panels.







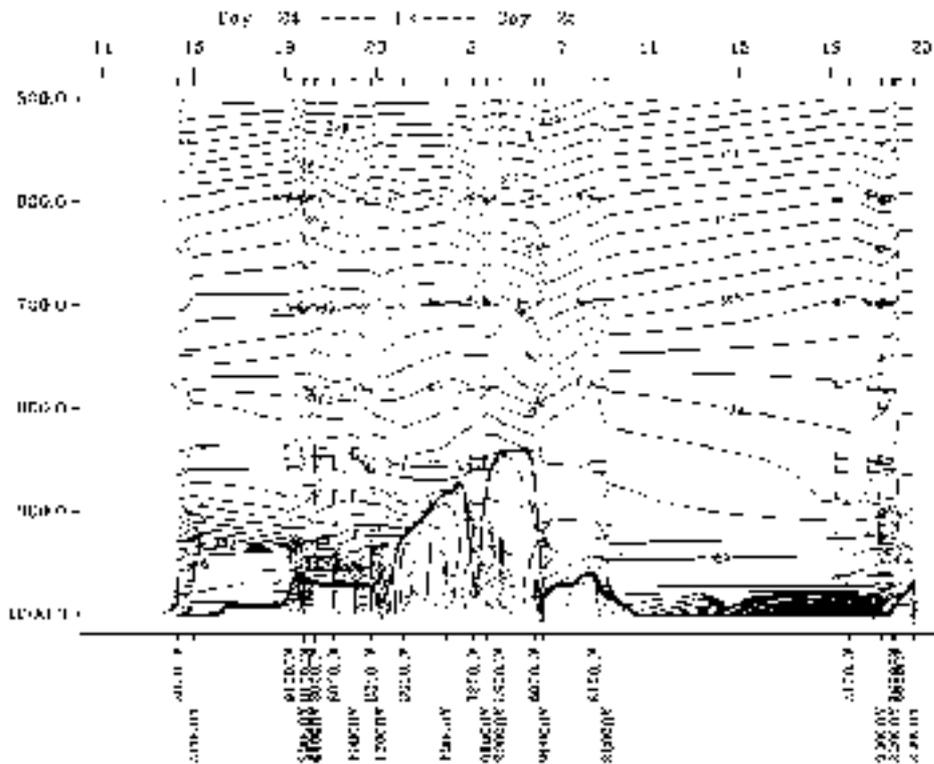
Another method to investigate the flow is by the use of trajectories. In Figure 3.7 we show multiple trajectories starting at 1900 UTC, 2300 UTC on 24 November 2000 and 0900 UTC 25 November 2000. These times represent the times of the Melbourne eddy, the bay breeze onset and the sea breeze onset. The larger left panels in each row of Figure 3.7 have starting locations along a longitude line on the eastern side of the bay that are spaced 0.05° apart. Stars indicate the starting locations and the symbols along the trajectories indicate 6-hour time ticks (at fixed hours). The trajectories show the path that a parcel of air would take over a 12 hour period. The larger right panels in row of Figure 3.7 have starting locations along a latitude line across the bay that are spaced 0.10° apart. The smaller panels show the vertical movement of the trajectories. If the trajectories stay close together, it indicates that the flow is deterministic; if the trajectories spread widely apart, then the flow is highly sensitive to initial conditions, or is said to be 'chaotic'.

In Figure 3.7 (top) most of the air parcels from the eastern side of the bay are transported to the western side by the Melbourne eddy. Exceptions occur for the northern-most trajectory and the eastern-most trajectory which are caught by the onset of the bay breeze and transported back to the eastern side. In Figure 3.7 (middle) the effect of the bay breeze and the sensitivity of where a parcel of air finally end up is clearly seen (note the radial spread of trajectories). Air parcels can be transported to the western side of the bay (see the lower panel), provided the starting location is not too far from the western coast. Under the sea breeze regime as shown in Figure 3.7 (bottom) the flow has returned to be more coherent and uniform over the bay region. Normally one would expect the flow to continue in a SE/NW direction until it reached opposing northerly flow. However in this case a frontal wind change moved north-northeastward from Cape Otway and limited the western transport. These Figures show that the model can carry pollutants to the western suburbs where the highest concentrations of ozone are generally observed, but accurate simulations of the final position are highly sensitive to spatial location and timing.

Further model verification is performed using data collected by aircraft as they take off from Melbourne airport (AMDAR data). Time-height profiles of the observed atmosphere are constructed using these data (Figure 3.8), and the equivalent time-height profiles of the modelled atmosphere from LAPS are also constructed (Figure 3.9). Comparisons of Figure 3.8 and Figure 3.9 give an indication of the model's success in predicting the diurnal development of the boundary layer, as well as the arrival time of mesoscale phenomena at the airport, such as sea and bay breezes, and cold fronts. Figure 3.8 shows the observed atmosphere from about 1400 UTC 24 November to almost 2300 UTC 25 November (1 am 25 November to 10 am 26 November, AEDST). A shallow inversion (approximately 950 hPa) was present until about 2200 UTC 24 November (9 am). Within the next two hours surface heating led to enhanced turbulent mixing and the break down of the inversion. Figure 3.9 shows that the stable layer and its erosion were well predicted by the model. This result provides added confidence in the model's ability to predict the transition from Melbourne eddy to bay breeze, illustrated above, which is largely dependent on the surface heating and breakdown of the nocturnal stable stratification. Comparisons also show that the model's mixed-layer development was very well predicted, with peak heights at 850 hPa. The arrival of the bay breeze near 6 UTC 25 November (5 pm) brought a drop in temperature at low levels (below 950 hPa) and an intensification of the southerly wind component. Again comparisons of the Figures show that the timing of the bay breeze arrival was very well predicted by the model. (Figure 3.6 shows that by the time the bay breeze reached the airport, it was virtually indistinguishable from the sea breeze. Thus, all further reference to the bay breeze at the airport includes the afternoon hybrid bay/sea-breeze.)

Minor discrepancies in the wind speed and direction are evident during the day between about 2200 and 600 UTC (9 am to 5 pm). These are related to the very delicate balance that exists between the bay-breeze blowing from the south and the larger scale northerly flow. In the model the balance was tipped more in the favour of the bay breeze, which meant that for much of the daytime period the convergent zone was located north of Melbourne airport. (Compare the convergence line position in Figure 3.6 with the Melbourne airport location, given by the most northerly red wind barb). Close inspection of the observed wind barbs shows the observed convergent zone was stationary between 2300 and 500 UTC, after which it accelerated rapidly northward beyond the airport in a manner consistent with a bay- or sea-breeze front. In the model the location of the convergent region was well predicted at 2300 UTC in the eastern half of the domain, illustrated in Figure 3.6. Further west the model convergent region was located a few kilometres further north. This discrepancy is likely to be due to subtle differences in the remnant Melbourne eddy. After 2300 UTC, Figure 3.6 shows the model convergent region, with light southerly winds behind, slowly propagated northward beyond the airport. This explains the early arrival of light southerly flow in Figure 3.9, compared with the continued northerly flow in the observations (Figure 3.8) due to the stationary observed bay breeze to the south of the airport. Although the model's southerly flow arrived relatively early at the airport, the actual bay breeze arrival did not occur until 0600 UTC (5 pm), when the temperature below 950 hPa dropped noticeably and the southerly winds intensified. As mentioned above, the model very well forecast the timing of this change.

Figure 3.8: Time-height profile of AMDAR (observed) potential temperature (contour interval 1 K) and winds (full barb = 10 knots; flag = 50 knots), for the 36 hour period following 11 UTC 24 November 2000. The thick line marks the approximate position of the top of the PBL, determined by the level at which the atmosphere is first 1 K warmer than the 10 m potential temperature.

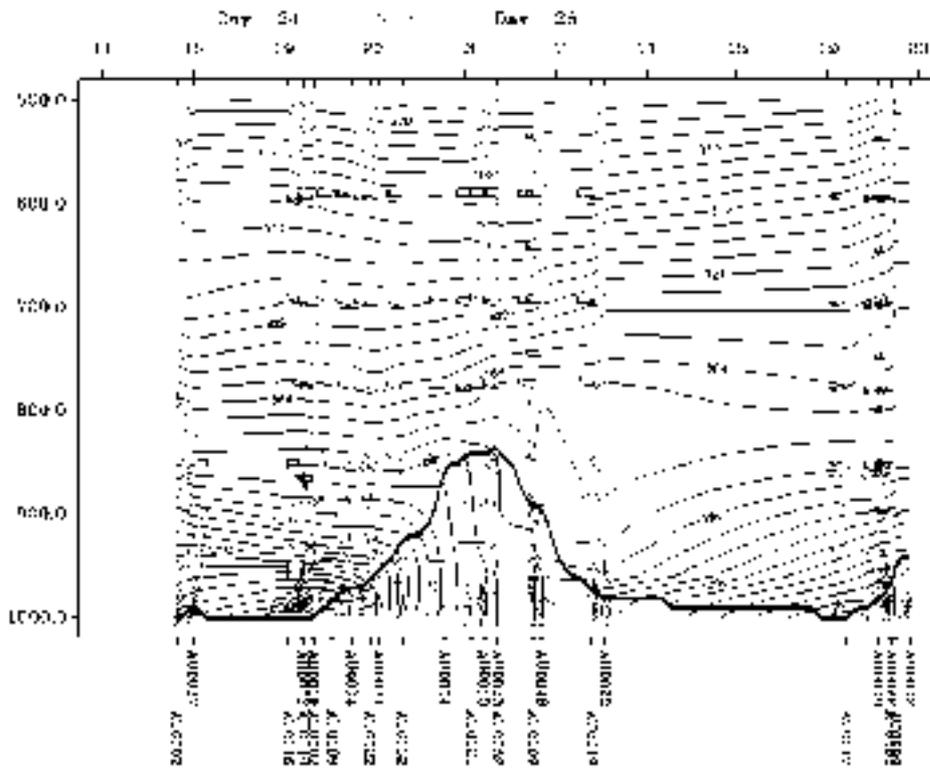


Conclusions

We have shown that the model statistics for all stations and all hours for bias, RMS error, and correlation indicate that the LAPS model performed well during the Preliminary Validation Period. Furthermore the case study beginning 1100 UTC 24 November 2000 shows that the LAPS model can simulate complex meteorological conditions including the Melbourne eddy, the bay-breeze and the sea-breeze flow patterns. Each of these mesoscale circulations has the potential to contribute to high concentrations of air pollutants. Thirdly we have shown by computing trajectories that the morning traffic emissions can impact on the western suburbs of Melbourne where the highest values of ozone concentration are measured, however the ability to simulate them accurately will depend on the accuracy of simulating the flow over the bay. Fourthly we have shown good agreement between measured (AMDAR) and predicted (LAPS) upper air temperatures and winds at Melbourne Airport and along flight paths.

This Validation Period occurred early in the photochemical smog season when concentration levels did not exceed the standards. The ability of the LAPS model to perform under conditions in which high concentrations of ozone and other pollutants are measured needs further investigation, and these studies will be carried out over the coming summer.

Figure 3.9: Same as Figure 3.8, except constructed from LAPS model data.



4. Preliminary Validation Study: Chemical Transport Model Performance.

Martin Cope and Sunhee Lee

CSIRO Atmospheric Research and Energy Technology

Introduction

The Australian Air Quality Forecasting system (AAQFS) has been operating in Demonstration Mode for more than 3 months. During this period, we have been generating twice-daily forecasts of hourly air pollution concentrations for 25 pollutants for 5 km spaced grids covering Victoria and NSW and 1 km spaced grids covering Melbourne/Geelong and Sydney. Spatial distributions of ground-level pollutant concentrations have been generated from the forecasts and placed on the CSIRO AAQFS web site for review by the Environment Protection Authorities (EPA's).

An important component of the AAQFS forecasting is verification. This is particularly critical for the inventory and chemical transport model (CTM), because the databases and software for the daily inventory generation and for the CTM are new and thus need to be carefully scrutinised for errors.

System verification is a multi-pronged approach. Examples of how the meteorological forecasts are verified are given in Section 3 of this report. Inventory verification is more difficult, but does commence with careful querying and analysis of the data and comparison with past and contemporary inventories. Further insight may be gained through comparison of observed and modelled pollutant ratios, and through the use of data from specialist measurement studies which may, for example, involve the use of instrumented aircraft to measure pollutant fluxes or in-tunnel pollutant measurements for motor vehicle emissions. Verification of the CTM software is undertaken through checking, at subroutine level, that the software description of each key physical or chemical process has been correctly implemented. This may be undertaken through comparison of process-level results with analytic solutions (where available), or through comparison with other (previously verified) numerical systems, or through comparison with specialist data sets. This process of software verification for the CTM is still in progress. Verification of the CTM in isolation using ambient air quality observations is complicated by the fact that the predicted concentrations are a function of the meteorology and the emissions inventory, in addition to the pollutant transformation processes that are embodied in the CTM. As such, a comparison of observed and modelled air quality time series should be considered a verification of the combined modelling system.

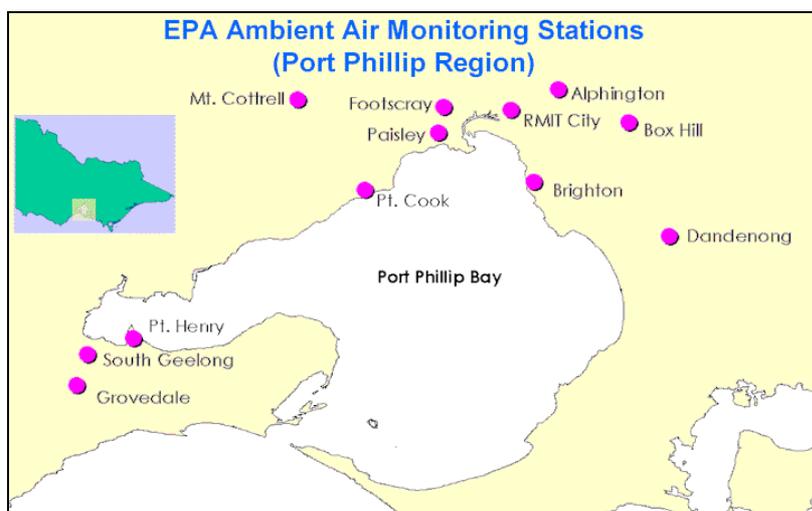
The results of a preliminary verification of AAQFS for the trial period 13–27 November 2000 are described in the following sections. Here we have compared observed and modelled hourly concentration time series for oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), carbon monoxide (CO), particulate matter with aerodynamic diameter $< 10 \mu\text{m}$ (PM10), and $< 2.5 \mu\text{m}$ (PM2.5). We also compare modelled and observed ratios of $\text{NO}_2:\text{NO}_y$, $\text{PM}_{2.5}:\text{PM}_{10}$, $\text{PM}_{10}:\text{NO}_y$, $\text{CO}:\text{NO}$ and $\text{CO}:\text{NO}_y$. Species ratios are included because they are able to provide additional insights into the underlying emissions and/or chemical transformation processes. Note that we have not included ozone in the verification study because the CTM has not yet been fitted with latest version of the Generic Reaction Set chemical transformation mechanism (GRS2, personal communication — Merched Azzi, CSIRO Energy Technology) and it is already known that the use of GRS1 can lead to under predictions of peak ozone concentration under some conditions (as described in earlier AAQFS Progress Reports).

Following the discussion of predictive performance, we consider system performance in terms of delivering images of the air quality forecast to the AAQFS web site in time for the 3 pm EPA air quality forecast.

Victoria: Observed and forecast air pollutant time series

Observed and modelled air pollutant time series are presented in Figures 4.2–4.10 for selected monitoring stations in the Melbourne/Geelong region (see Figure 4.1 for localities) for the period 13–27 November 2000. Figure 4.2 shows the time series of observed and forecast concentrations of NO_y and NO_2 together with the forecast and observed ratios of $\text{NO}_2:\text{NO}_y$ for Alphington monitoring station. Note that the observed concentrations have not yet been subject to EPA Victoria’s full validation process. Note also that the forecast concentrations and ratios correspond to those within a radius of ± 1 grid cell, which most closely match the observed value for the current hour. For the 1 km model grids, this corresponds to selecting a forecast concentration within a radius of 1–2 km of the monitoring station, and thus is considered to be representative of a suburb-level forecast.

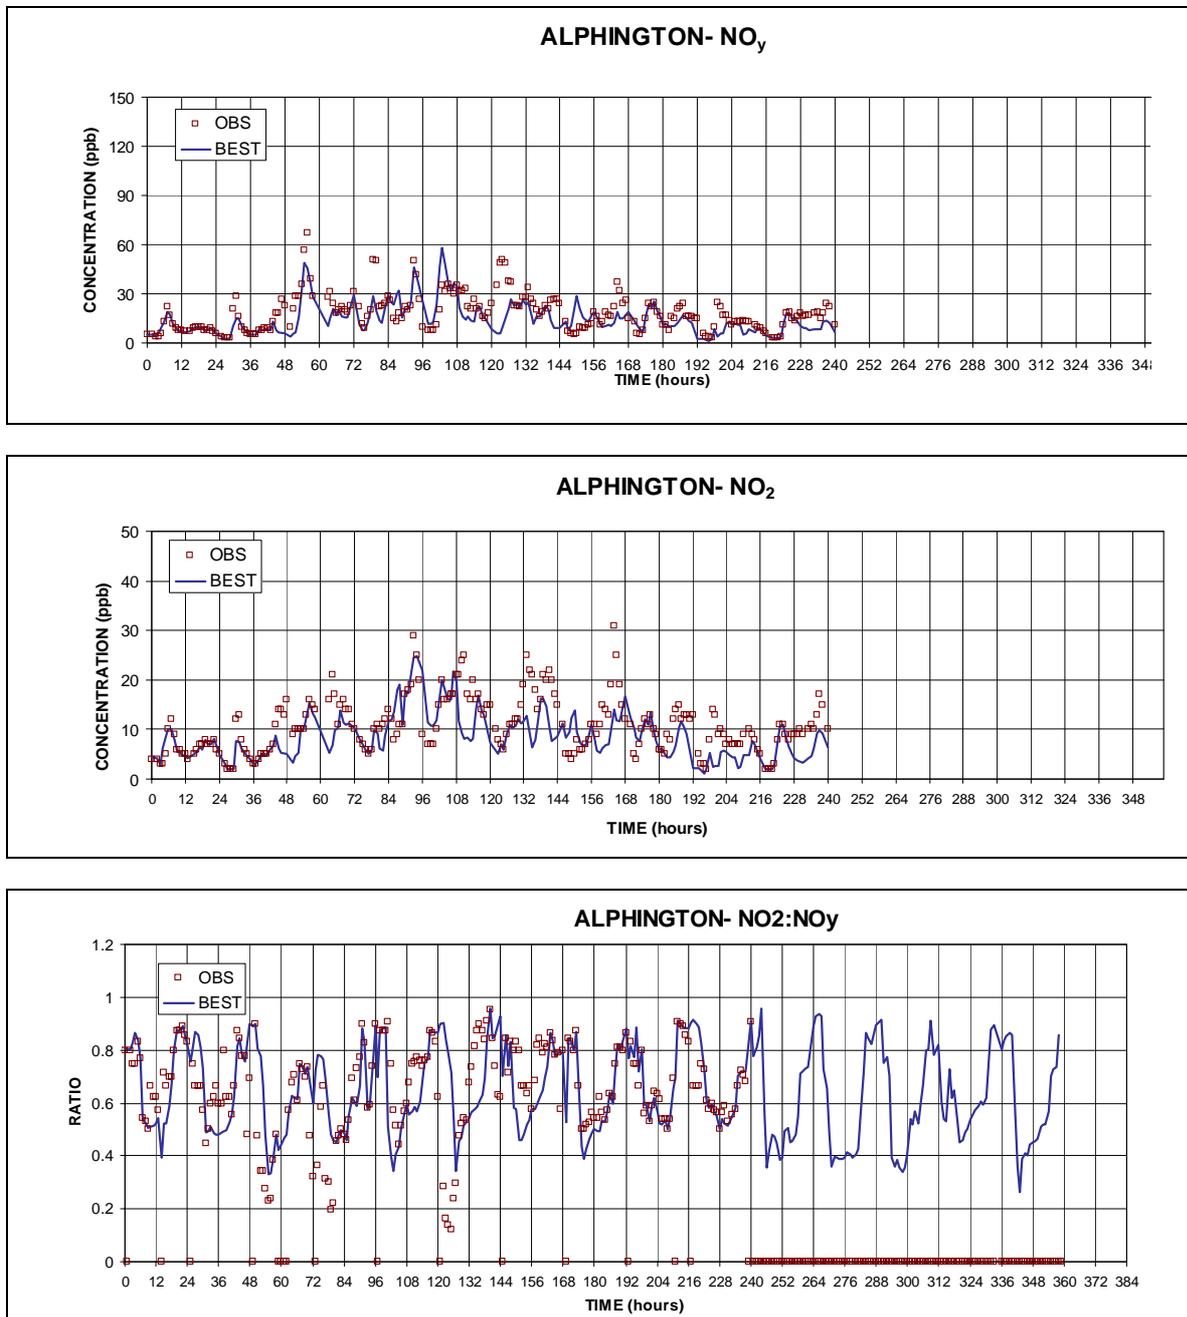
Figure 4.1: Location of EPA Victoria monitoring network for Melbourne/Geelong region (image courtesy of EPA Victoria).



Observations of NO_y were only available from Alphington for 10 days of the 15 day trial period. However, for the period when data were available, it can be seen (Figure 4.2) that the observed concentrations of NO_y are low (peaks generally < 60 ppb) and that the observations have generally been matched with good accuracy by the suburb-level forecast. The same comment can be made with respect to NO_2 with the exception of the peaks observed just past hour 132 and just prior to hour 168. The model has also generally matched the observed ratios of $\text{NO}_2:\text{NO}_y$ with good skill. Matching of the observed $\text{NO}_2:\text{NO}_y$ ratios indicates that the model is able to reproduce the observed diurnal distribution of turbulent mixing and the titration of fresh nitric oxide emissions (e.g., for motor vehicles, which are the dominant source of NO_y in the Melbourne/Geelong region, $> 90\%$ of NO_y is emitted as nitric oxide on a molar basis). Thus under conditions of low mixing, it may be expected that the $\text{NO}_2:\text{NO}_y$ ratio will approach a lower limit of about 0.1. Under conditions of strong turbulent mixing and/or photochemistry, it may be expected that the $\text{NO}_2:\text{NO}_y$ ratio will approach an upper limit of 1.0.

For example, periods between hours 48–60, 72–84 and 120–132 likely correspond to times where fresh and relatively undiluted motor vehicle emissions have impacted Alphington monitoring station. For these cases it would appear that the System has overestimated the dilution rates.

Figure 4.2: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), and $\text{NO}_2:\text{NO}_y$ ratio for Alphington monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx=1$ km).



Observed and suburb-level forecast time series concentrations of PM_{10} , $\text{PM}_{2.5}$ and the ratio $\text{PM}_{2.5}:\text{PM}_{10}$ are shown in Figure 4.3 for Alphington monitoring station. Note that we currently restrict our consideration to 1-hour rather than 24-hour PM_{10} (the latter is required for the calculation of Air Quality Index [AQI]) because we are interested in the System's ability to reproduce the significant diurnally varying processes. Moreover, our AQI methodology parameterises aerosol backscatter using 1-hour $\text{PM}_{2.5}$.

From a consideration of Figure 4.3 it can be seen that the System tends to underestimate PM10 and PM2.5 at Alphington monitoring station. However, it is encouraging to see that the System has correctly forecast the occurrence of peaks in PM10 and PM2.5 between hours 96–18, 288–300 and 312–324, although again, the magnitudes of the peaks have been underestimated. The System has forecast the day-to-day variation of the ratio of PM2.5:PM10 with good skill. Periods when the observed and modelled PM2.5:PM10 ratio lies around 0.6 likely corresponds to the presence of motor vehicle emissions. Periods when the ratio is lower to correspond to conditions of higher ventilation and/or the presence of alternative aerosol source groups.

Time series plots of observed and forecast concentrations of CO, the ratio CO:NO_y and the ratio PM10:NO_y are presented in Figure 4.4 for Alphington. It can be seen that the System has matched the observed concentrations of CO with good skill for the second half of the trial period, but has tended to underpredict concentrations during the first half of the trial period. The System has matched the observed ratio of CO:NO_y with good skill between hours 48 and 180. With respect to the ratios of PM10:NO_y it can be seen that the System tends to match the lower observed limit of about 0.2 with good skill, but can often underpredict the upper peaks. The lower bound ratios likely correspond to the presence of motor vehicle emissions under conditions of limited dilution. Peaks of observed high PM10:NO_y may correspond to the presence of inefficient combustion sources (i.e., wood fires), to the presence of wind blown dust, to sea-salt aerosol, or to aged or well diluted air masses. Note that the version of the CTM used during the trial did not include emission estimates for wind blown dust or sea salt aerosol, thus some of the discrepancies between model and observed may be due to the omission of these natural source groups.

Figure 4.3: Observed vs modelled concentration time series of PM10, PM2.5, and PM2.5:PM10 ratio for Alphington monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).

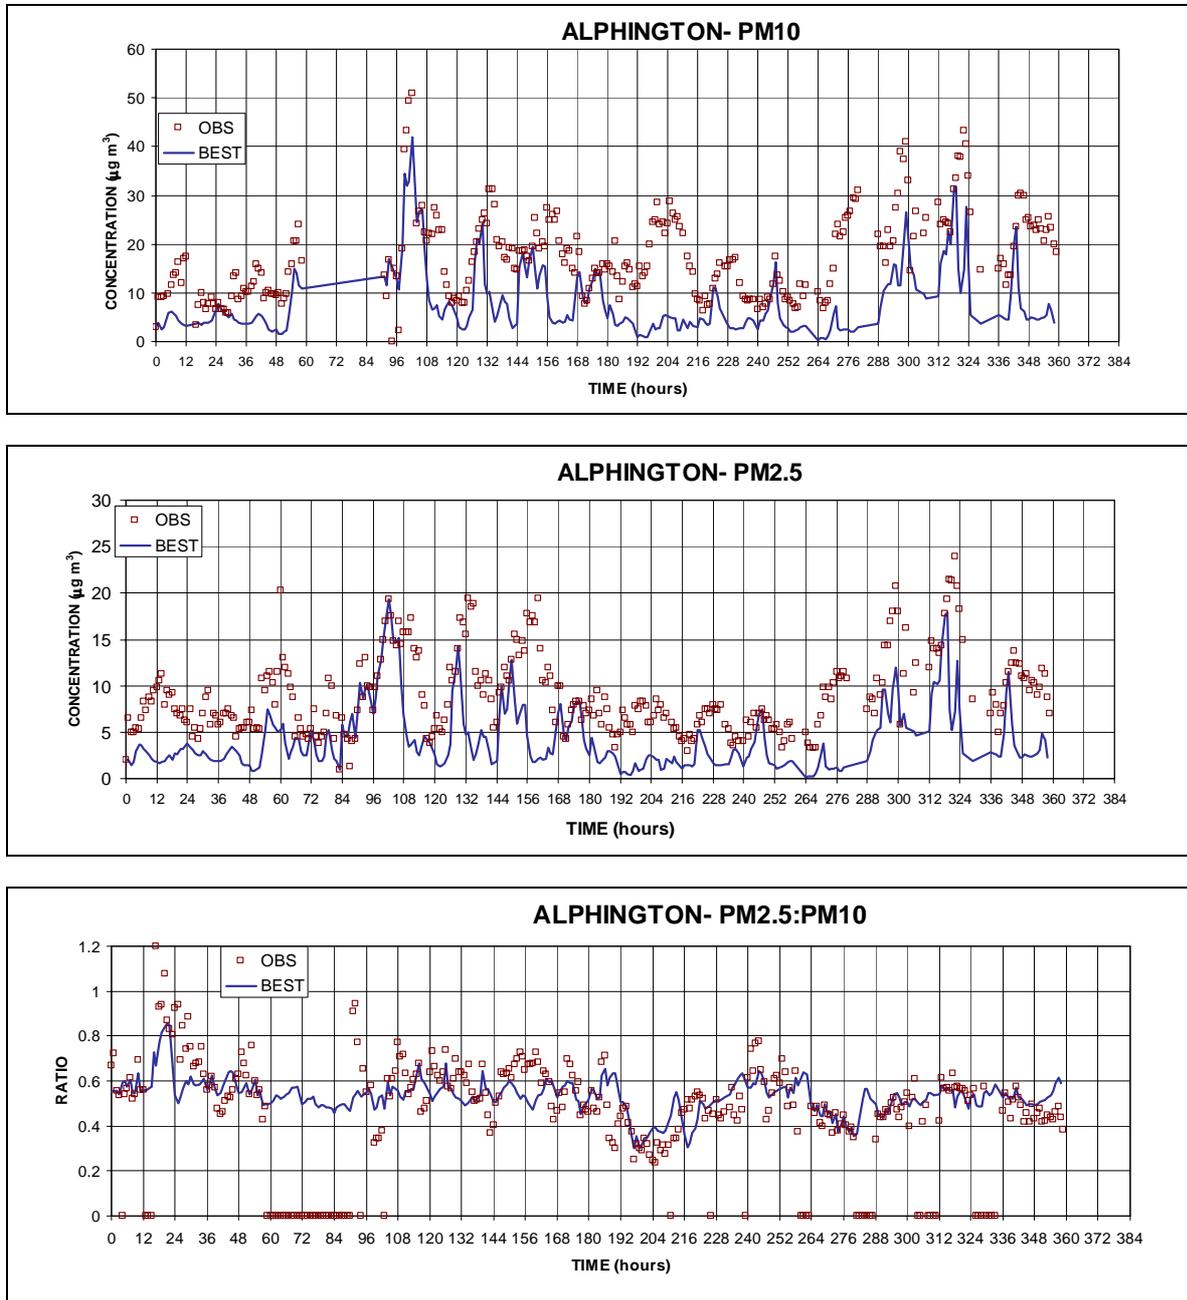
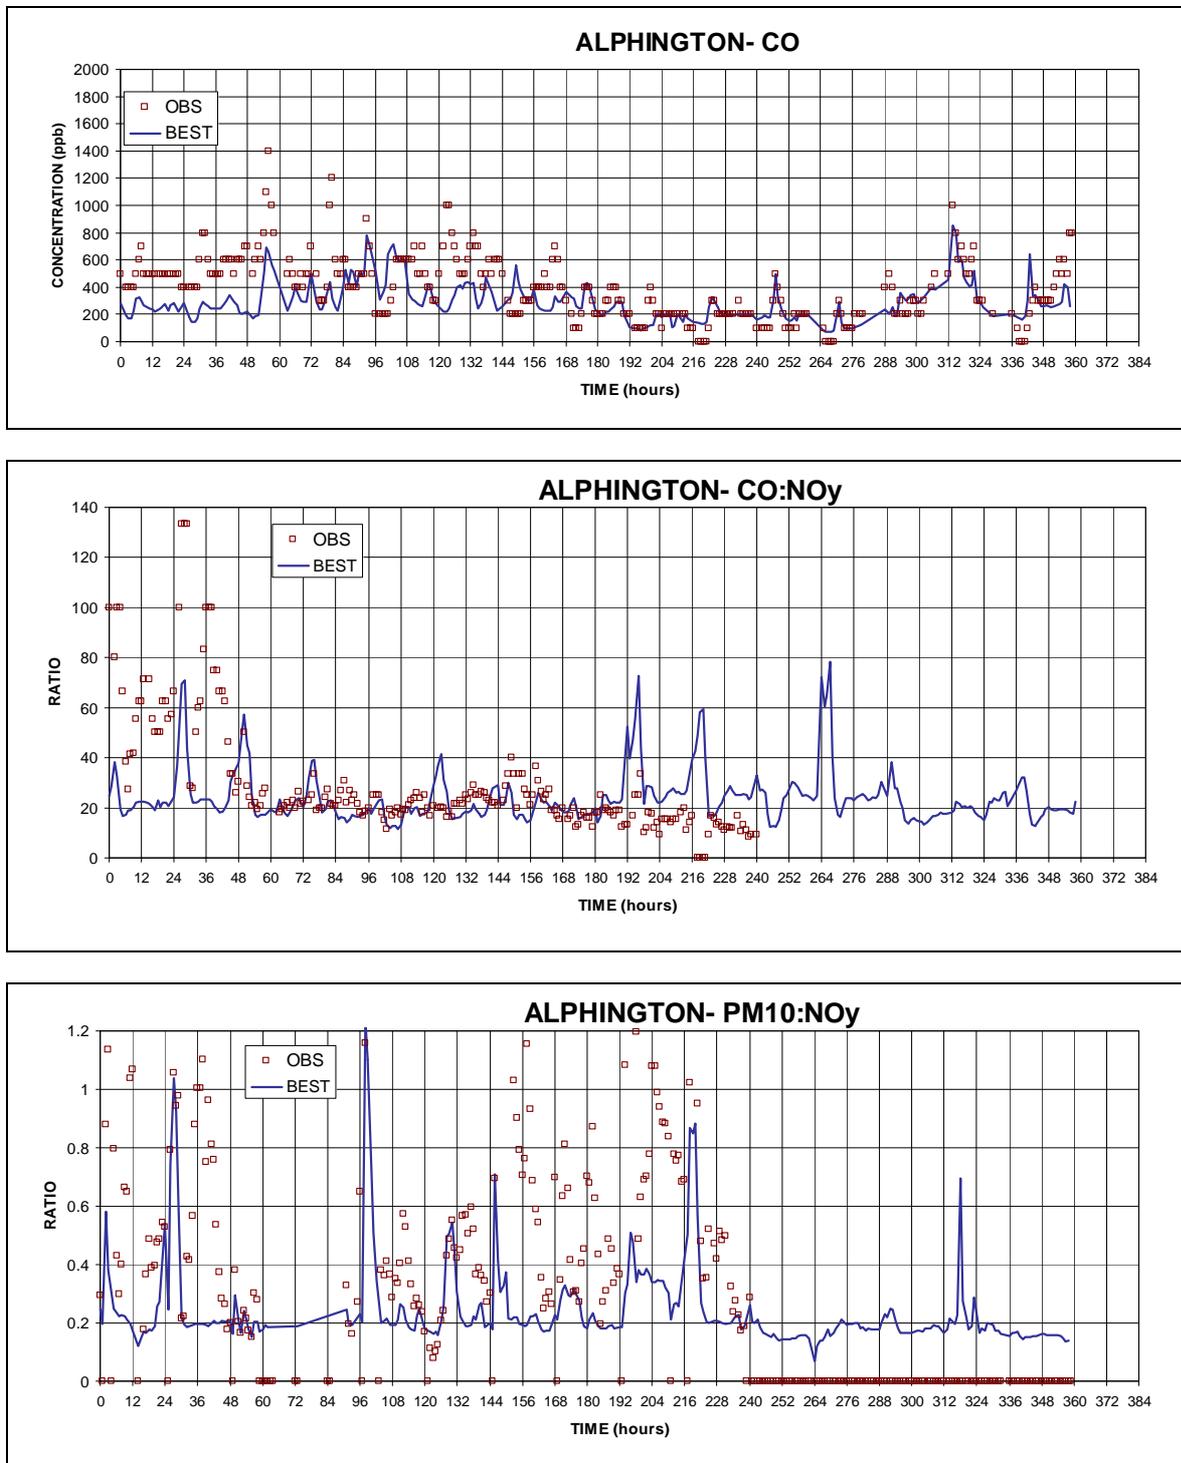


Figure 4.4: Observed vs modelled concentration time series of CO, CO:NO_y ratio and PM10:NO_y ratio for Alphington monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site (dx = 1 km).



Time series plots of observed and forecast concentrations of NO_y , NO_2 and ratios of $\text{NO}_y:\text{NO}_2$ are shown in Figure 4.5 for Brighton monitoring station. The System has forecast the presence of NO_y peaks observed between hours 288–300 and hours 312–324. However, it has underpredicted the peak observed between hours 48 and 60. With regard to NO_2 , the System has done particularly well at matching the NO_2 peak observed at hour 300 and between hours 312 and 324. However, the peak observed prior to hour 60 has been underestimated. The System has been able to reproduce the observed ratios of $\text{NO}_2:\text{NO}_y$ with good skill.

Observed and modelled concentrations of PM_{10} , $\text{PM}_{2.5}$ and $\text{PM}_{2.5}:\text{PM}_{10}$ for Brighton monitoring station are shown in Figure 4.6. It is pleasing to see that the System has correctly predicted the presence of concentration peaks for hours 48–60, 84–108, 132–156 and 312–324. On the other hand concentration peaks observed between hours 192–216 and 288–312 have been underpredicted by the System. Its interesting to note that the period of elevated aerosol concentrations observed between hours 192–204 correspond to a minimum in the $\text{PM}_{2.5}:\text{PM}_{10}$ ratio (both modelled and observed). In this case, the elevated aerosol concentrations may correspond to wind blown dust or sea salt aerosol (neither of which was included in the CTM during the trial period).

Plots of observed and forecast $\text{PM}_{10}:\text{NO}_y$ ratio for Brighton and observed and forecast CO concentration and $\text{CO}:\text{NO}_y$ ratio, both for RMIT (as an alternative for Brighton where CO observations were not available) are presented in Figure 4.7. The $\text{PM}_{10}:\text{NO}_y$ ratios for Brighton are similar to those for Alphington in that the lower bound ratios are generally reproduced while the upper bound ratios are often underpredicted. The System has generally forecast the observed CO concentrations with good skill for the RMIT site (see Figure 4.1 for location), although some of the peak observed concentrations have been underpredicted. The results for the $\text{CO}:\text{NO}_y$ time series are interesting in that the System tends to systematically underpredict the observed ratios. This suggests that either the emissions of NO_y are overestimated or the emissions of CO are underestimated upwind of this site.

A comparison of observed and modelled concentration time series for NO_y , NO_2 and $\text{NO}_2:\text{NO}_y$ ratio for Footscray monitoring station is given in Figure 4.8. It can be seen that the System has done an excellent job in reproducing the observations, although it should also be noted that some of this performance can be attributed to the presence of strongly inhomogeneous emissions around the region. Thus the suburb-level search algorithm has generally been able to find a good match between observation and System predictions.

The time series plots, for Footscray, of observed and forecast PM_{10} , $\text{PM}_{2.5}$ and the ratios of $\text{PM}_{2.5}:\text{PM}_{10}$ are shown in Figure 4.9. It can be seen that an improved match has been achieved for the Footscray monitoring site in comparison to the other sites. The System has done particularly well in forecasting the presence of $\text{PM}_{2.5}$ peaks observed between hours 96–108, 300–312 and 212–324. The suburb-level forecasts of $\text{PM}_{2.5}:\text{PM}_{10}$ also match the observations with good skill.

Ratios of observed and forecast $\text{PM}_{10}:\text{NO}_y$ for Footscray are presented in Figure 4.10. In the absence of CO measurements at Footscray, we have also includes time series plots of observed and forecast CO and $\text{CO}:\text{NO}_y$ ratios for Box Hill monitoring station (see Figure 4.1 for location). It can be seen that, once again, the observed lower-bound ratios of $\text{PM}_{10}:\text{NO}_y$ are well matched by the System, while the upper bound ratios are often underpredicted. Thus this outcome appears to hold in both the eastern and western suburbs of Melbourne. The observed CO concentration time series for Box Hill is generally well forecast by the System, however, it is notable that that the $\text{CO}:\text{NO}_y$ ratio is overpredicted during periods of low observed CO concentration. This may be caused by the prescription of a CO boundary condition in the model which is slightly too high (150 ppb in the boundary layer, decreasing to 100 ppb in the free troposphere above), or truncation error in the CO observations (which can be seen in the time series plots).

Figure 4.5: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), and $\text{NO}_2:\text{NO}_y$ ratio for Brighton monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx=1$ km).

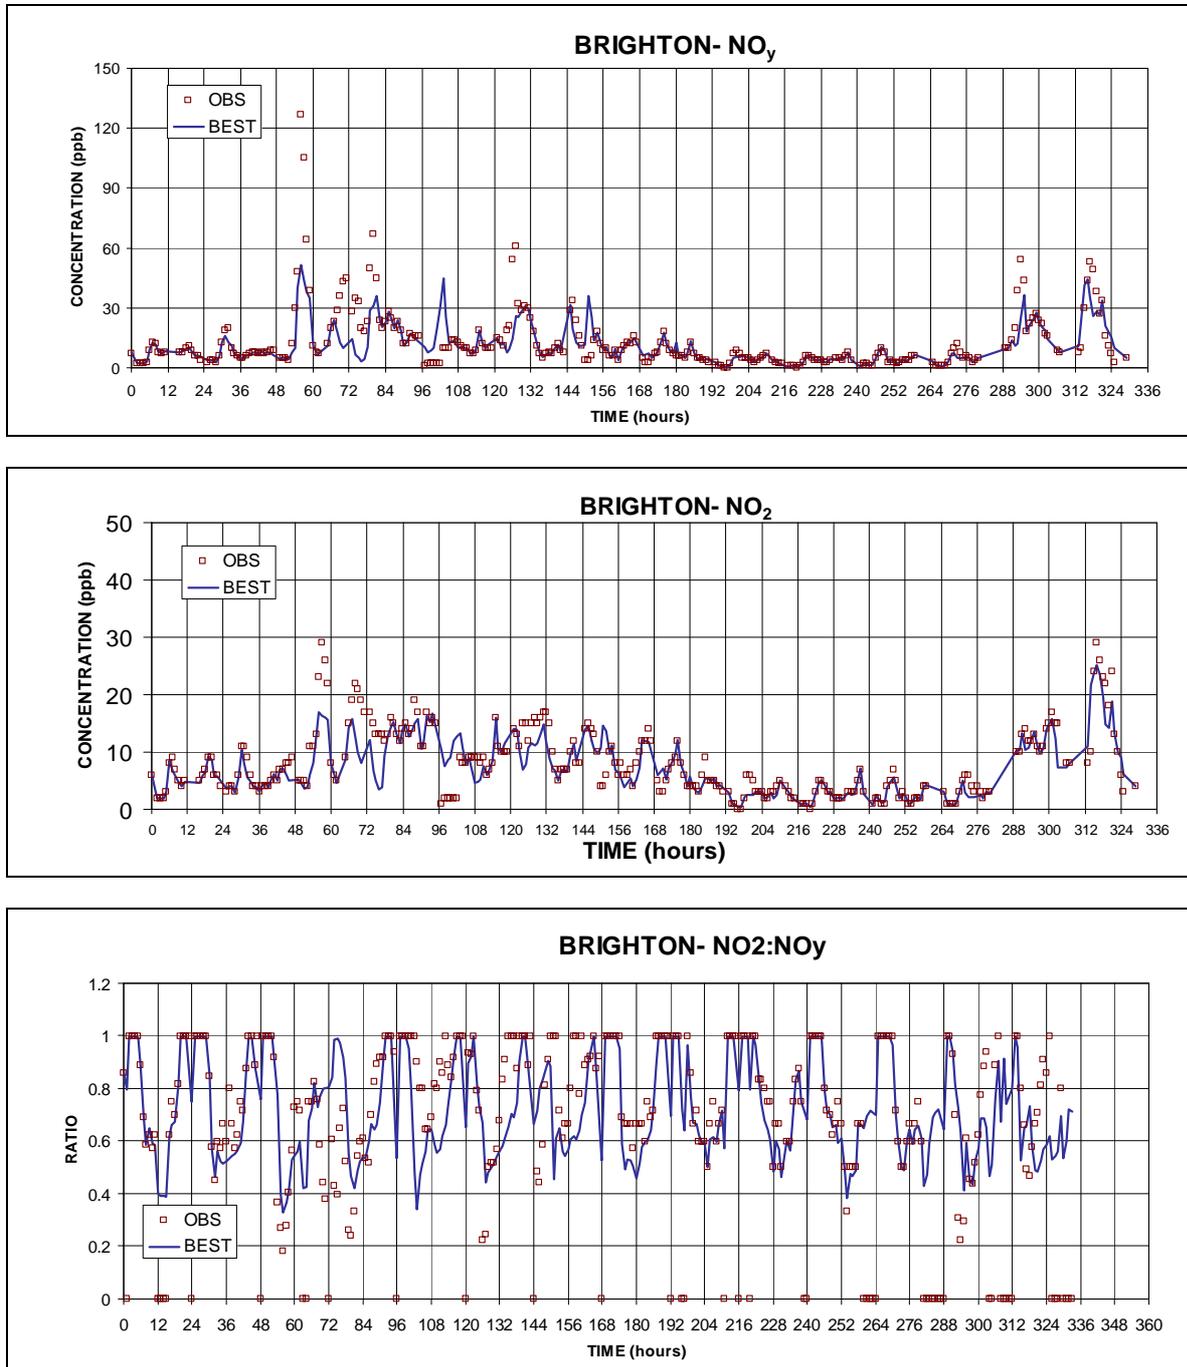


Figure 4.6: Observed vs modelled concentration time series of PM10, PM2.5, and PM2.5:PM10 ratio for Brighton monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).

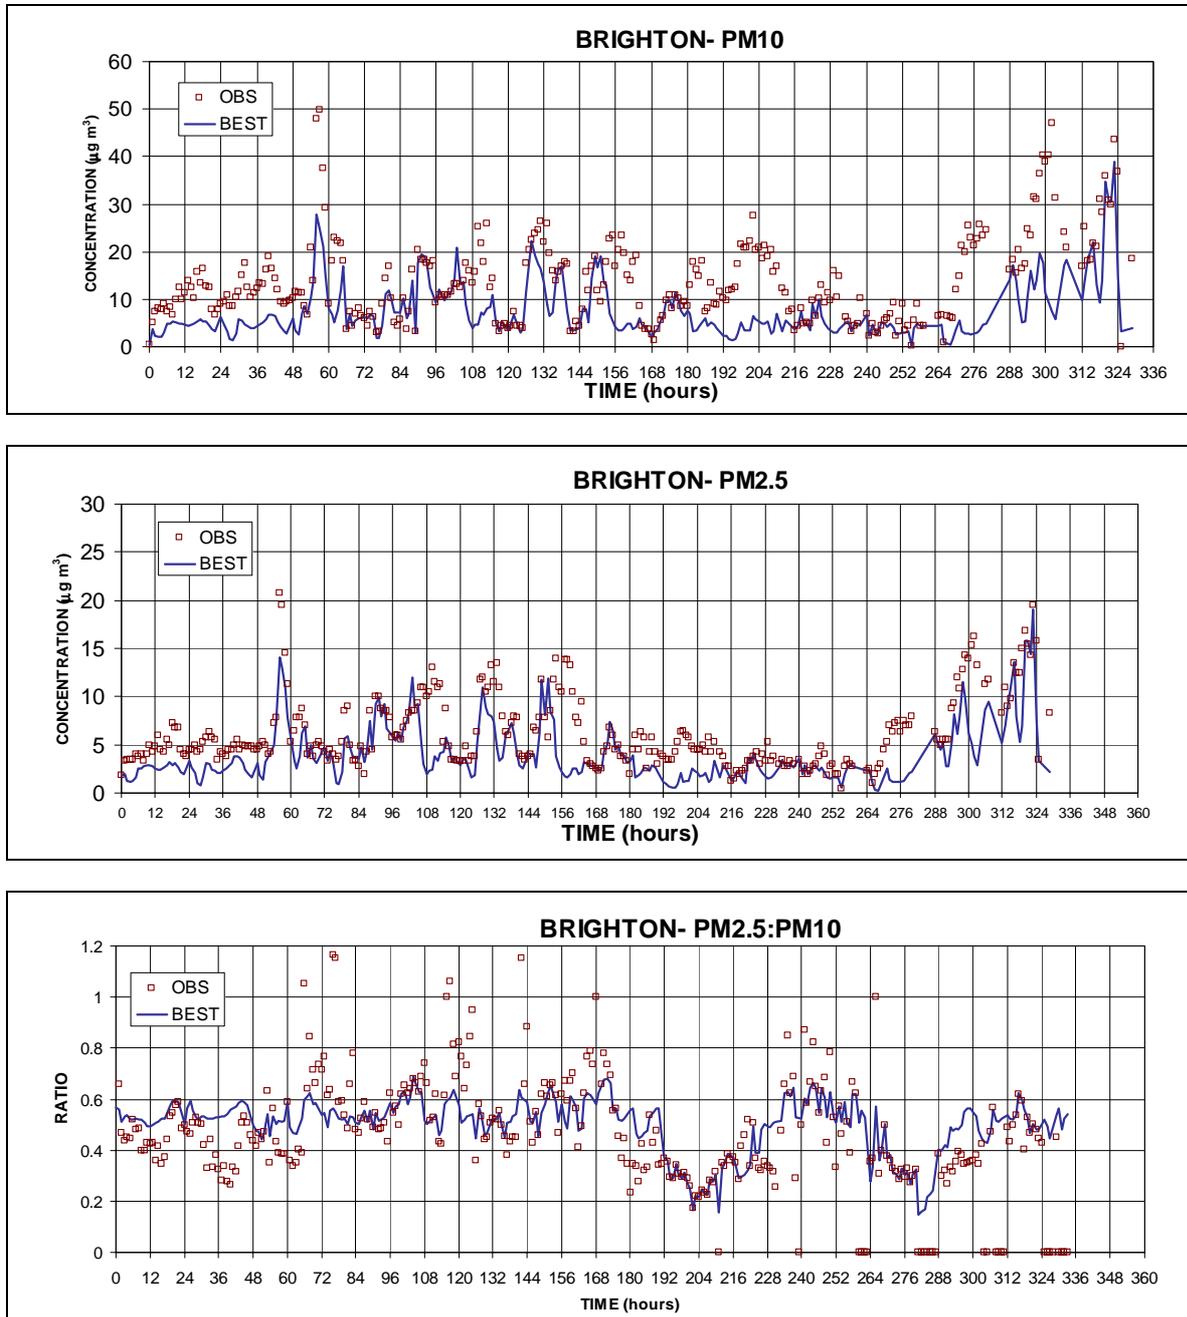


Figure 4.7: Observed vs modelled concentration time series of PM10:NO_y ratio for Brighton monitoring station and CO and CO:NO_y ratio for RMIT monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site (dx = 1 km).

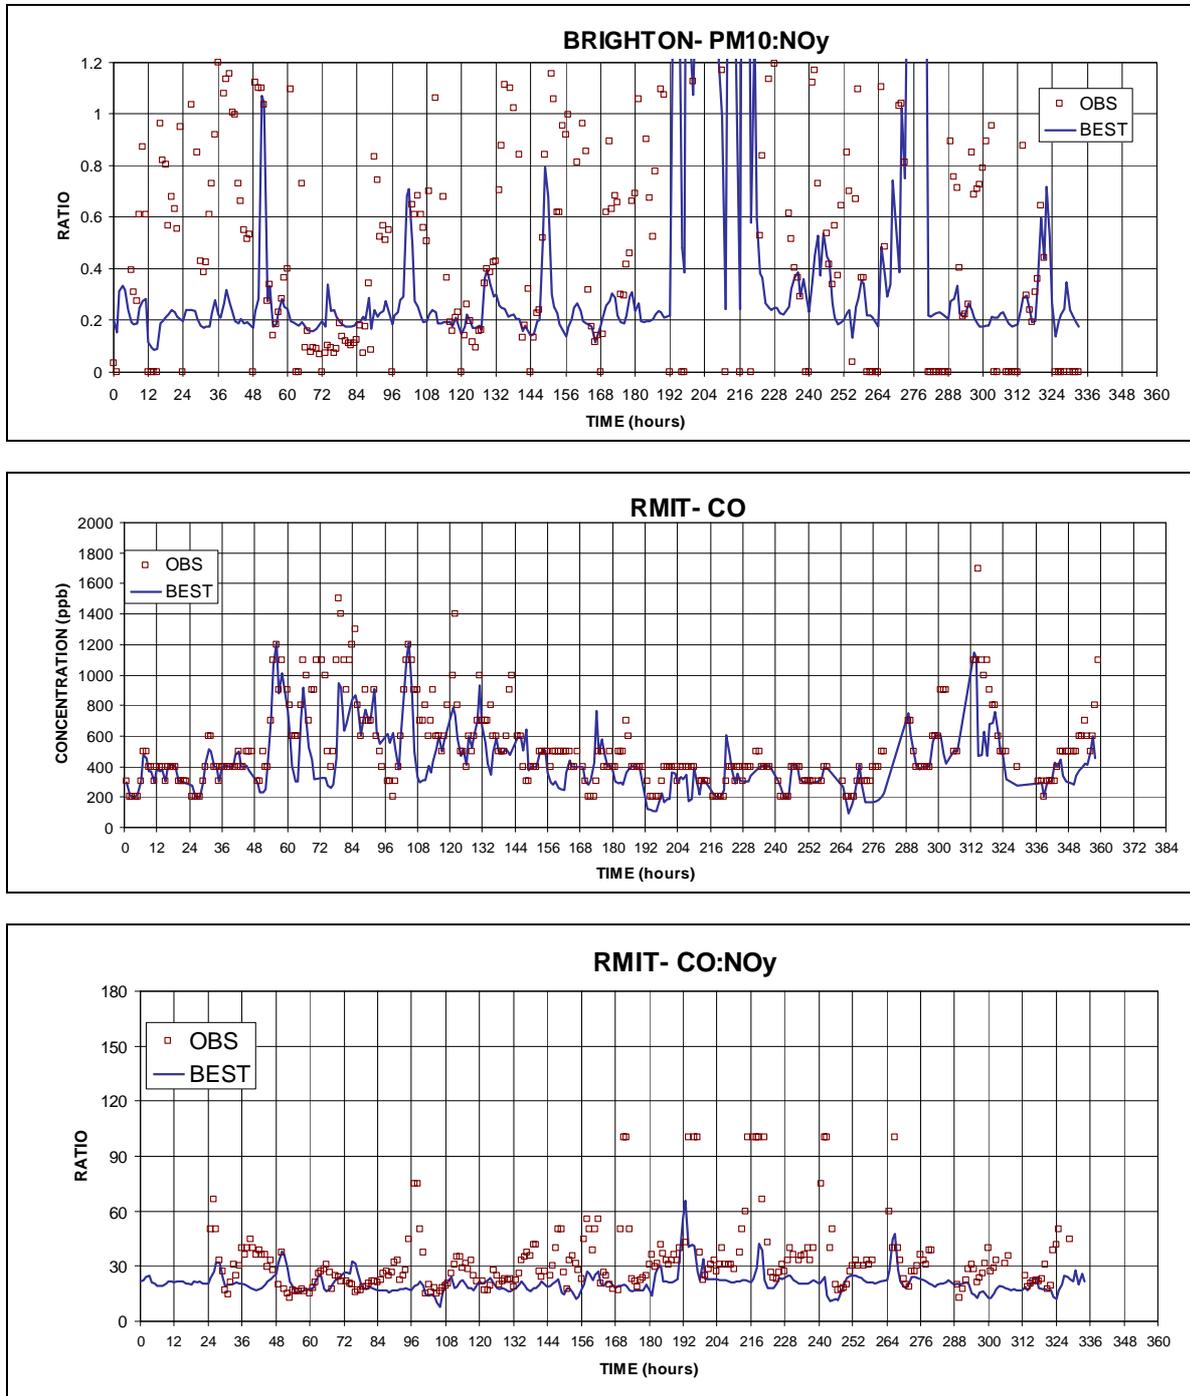


Figure 4.8: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), and $\text{NO}_2:\text{NO}_y$ ratio for Footscray monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx=1$ km).

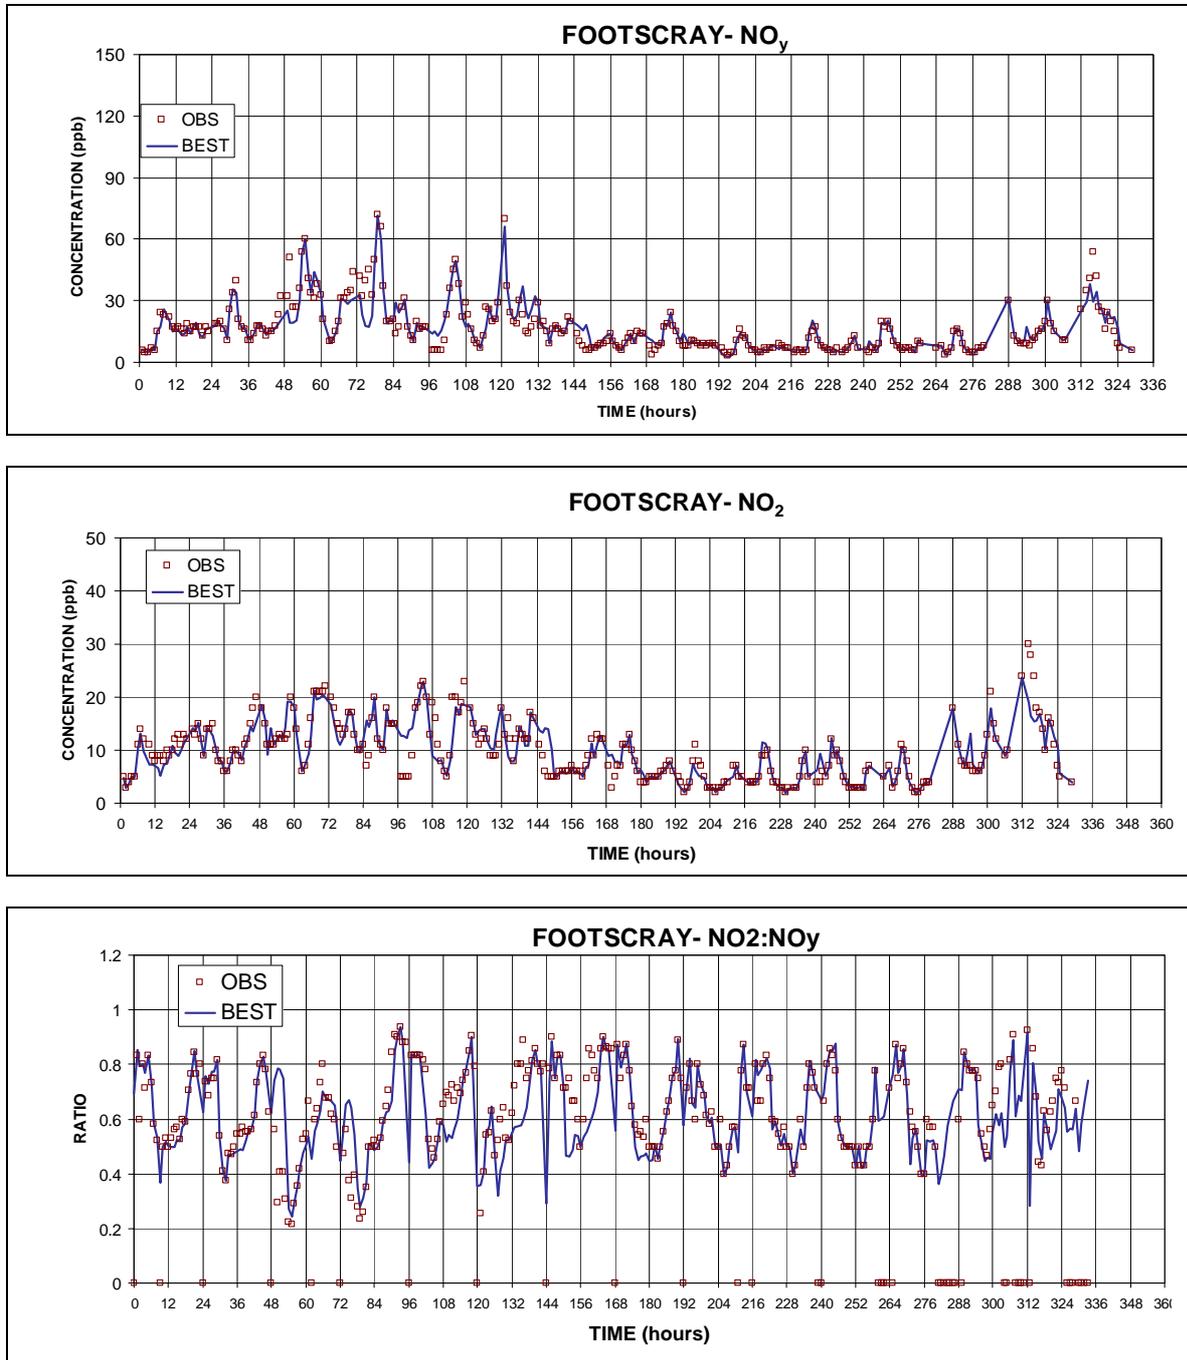


Figure 4.9: Observed vs modelled concentration time series of PM10, PM2.5, and PM2.5:PM10 ratio for Footscray monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).

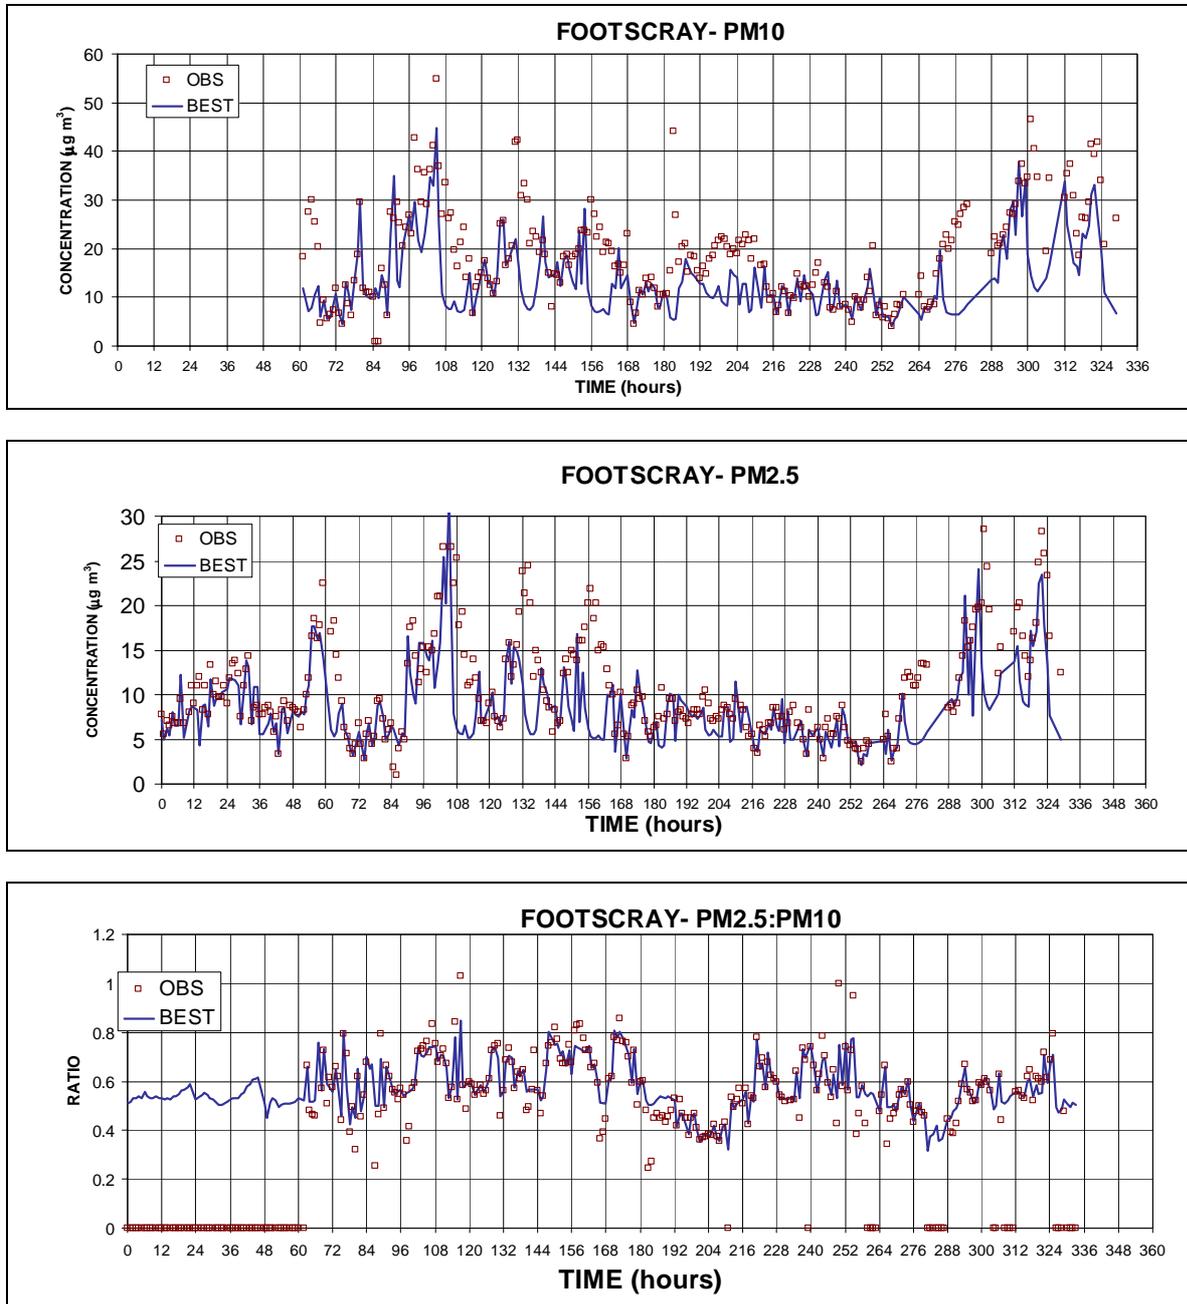
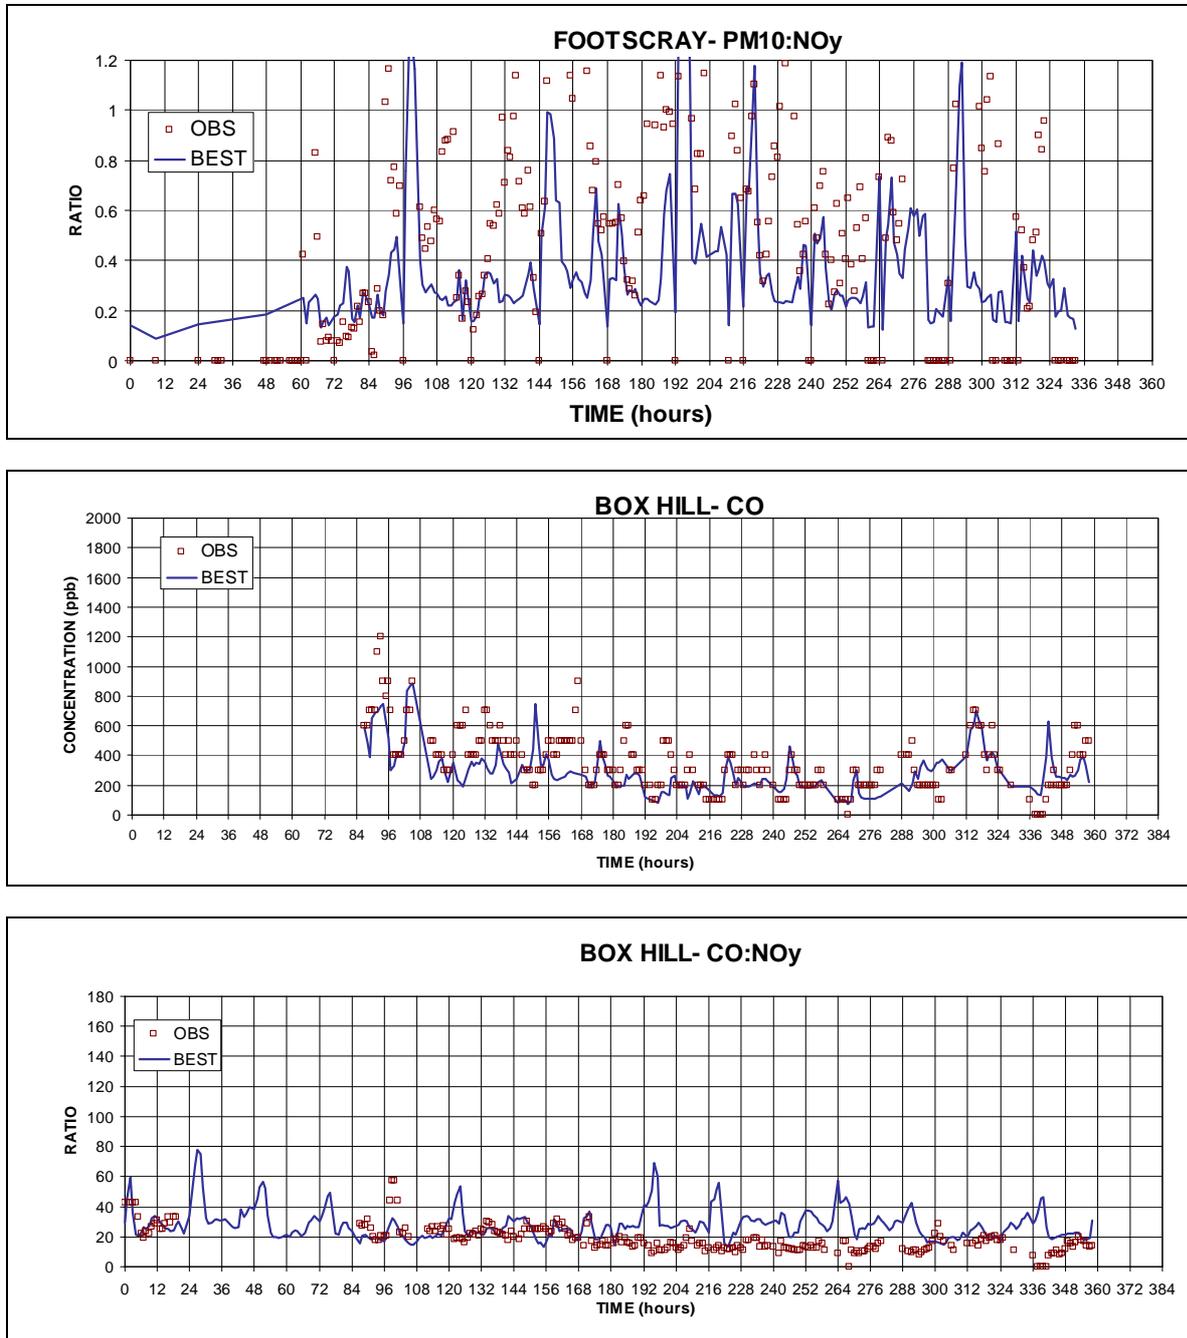


Figure 4.10: Observed vs modelled concentration time series of PM₁₀:NO_y ratio for Footscray monitoring station and CO and CO:NO_y ratio for Box Hill monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site (dx = 1 km).



Victoria: Residual analysis

The results of the time series comparison that was qualitatively discussed in the previous section can be formalised using the range of statistical measures presented in Table 4.1. Note that we calculate these statistical measures for four time- and space-coupling alternatives between model and observations:

1. the observed concentration at a given hour is paired to the modelled concentration, which is interpolated to the monitoring site from the four nearest grid points. All modelled concentrations are for the current hour. In this case, observed and modelled concentrations are considered to be matched in time and space (ST).
2. the observed concentration at a given hour is paired to the best matching modelled concentration taken from the four nearest grid points. All modelled concentrations are for the current hour. In this case observed and modelled concentrations are matched in time but the spatial matching is relaxed (T). This is equivalent to our suburb-level forecast.
3. the observed concentration at a given hour is paired to the best matching, interpolated modelled concentration within ± 1 hour of the current hour. In this case, observed and modelled concentrations are matched in space but the temporal matching is relaxed (S).
4. the observed concentration at a given hour is paired to the best matching, modelled concentration taken from the four nearest grid cells, and from within ± 1 hour of the observation hour. For this case, the matching is slightly relaxed in both time and space (U).

Table 4.1: Statistical measures used to assess the performance of the AAQFS chemical transport model.

$D = \frac{1}{N} \sum_{i=1}^N (C_e(\hat{x}_i, t) - C_o(\hat{x}_i, t))$ $d_i = (C_e(\hat{x}_i, t) - C_o(\hat{x}_i, t))$	<p><u>Mean Bias.</u> This is the mean difference between modelled and observed concentrations. Large negative or positive values of D are indicative of bias in model estimates. However, it should also be noted that a small bias may also be caused by the cancellation of errors through large positive and negative values of D. Thus D should also be used in combination with the other metrics described in this table.</p>
$D^* = \frac{1}{N} \sum_{i=1}^N \frac{(C_e(\hat{x}_i, t) - C_o(\hat{x}_i, t))}{C_o(\hat{x}_i, t)}$	<p><u>Normalised Mean Bias.</u> Here the bias has been normalised by the observed concentration. The normalised bias is a particularly useful metric for inter-study comparisons.</p>
$S_d^2 = \frac{1}{N-1} \sum_{i=1}^N (d_i - D)^2$	<p><u>Variance.</u> The variance of the residual distribution is indicative of the scatter between observed and modelled concentrations. Thus a small mean bias, coupled with a large variance, is indicative of the cancellation of errors in the mean bias calculation.</p>
$E = \frac{1}{N} \sum_{i=1}^N C_e(\hat{x}_i, t) - C_o(\hat{x}_i, t) $	<p><u>Gross Error.</u> The metric is the mean of the absolute difference between observed and modelled concentrations. A large gross error is indicative of substantial differences between modelled and observed concentrations. Unlike the mean bias, this metric is not affected by the cancellation of errors.</p>
$E_d^* = \frac{1}{N} \sum_{i=1}^N \left \frac{C_e(\hat{x}_i, t) - C_o(\hat{x}_i, t)}{C_o(\hat{x}_i, t)} \right $	<p><u>Normalised Gross Error.</u> As with the mean normalised bias, this is a particularly useful metric for the comparison of results from other studies, where the concentration range considered may be substantially different.</p>

Results from the residual analysis for all monitoring stations in Melbourne and Geelong (for which data were available) for NO_y, CO, PM10 and PM2.5, for the entirety of the trial period, are presented

in Table 4.2 and Table 4.3. Furthermore, plotted in Figure 4.11 and Figure 4.12 are the daily variation in normalised bias for NO_y, CO, PM10 and PM2.5. Data are presented for the four methods of coupling the observed and modelled concentration pairs as discussed above.

Turning our attention to NO_y (Table 4.2, Figure 4.11), it can be seen that the mean bias is negative and ranges between 5 ppb and 8 ppb underprediction for the various coupling alternatives. This translates into a mean level of underprediction of between 22 and 32% for NO_y. Thus when trying to forecast hourly average NO_y concentrations at a given location and hour, we can expect to underestimate concentrations by 32%, on the average. However, if we instead consider the case in which the forecast may be taken from any grid point within a radius of 2 km, and within a time band of ± 1 hour, it can be seen that such a forecast may, on the average, be expected to lie within 22% of the observed NO_y concentration. This is an excellent result for an air quality forecasting system. From a consideration of Figure 4.11 it can be seen that there is a considerable day-to-day variation in the normalised bias for NO_y, with cases of smallest bias often occurring when NO_y concentrations are elevated. Again, this is a very promising outcome for a forecasting system where a major goal is to predict the presence of peaks in the air pollution concentration fields. As may be expected, the normalised gross error is larger than the bias (see Table 4.1 for an explanation), however the uncoupled (U) result of 26% is still very acceptable. Note how the greatest reduction in gross error occurs when the spatial coupling is relaxed (rather than the temporal coupling), this suggests that it is more difficult to forecast the spatial variation in pollutant concentrations than the temporal variation.

The normalised bias for CO lies within 25% for all coupling alternatives (Table 4.2). This is also an excellent result. However, the slightly larger normalised gross error for some of the coupling alternatives indicates that the System may, on occasions, slightly overpredict hourly CO concentrations. This is apparent from Figure 4.11 (bottom) where it can be seen that the System has matched or overpredicted CO concentrations on 3 of the 15 days of the trial period.

System performance for PM10 and PM2.5 is presented in Table 4.3 and Figure 4.12. It can be seen that PM10 is, on the average, underpredicted by 5–9 $\mu\text{g m}^{-3}$ and PM2.5 by 2–5 $\mu\text{g m}^{-3}$. This translates into a normalised negative bias of 33 to 55% (PM10) and 28–62% (PM2.5). Again, the best forecast performance corresponds to a suburb-level forecast with a time bandwidth of ± 1 hour. In this case the System is, on the average matching the observations to within 33% (PM10) and 28% (PM2.5). Note that the level of underprediction for the aerosol is larger than that seen for NO_y and CO, and may partially result from an absence of the natural particle source groups (wind blown dust and sea salt aerosol) in the CTM during the trial period. This hypothesis is supported by the fact that the System performs better for PM2.5 than PM10, and the knowledge that the majority of the particulate mass for wind blown dust and sea salt resides in the coarse fraction.

Table 4.2: Performance measures for NO_y and CO for Melbourne/Geelong monitoring stations for 13–27 November 2000.

Statistical Measure ¹	NO _y				CO			
	Pairing type				Pairing type			
	ST	T	S	U	ST	T	S	U
Bias (ppb)	-8	-6	-7	-5	-144	-143	-142	-123
Normalised Bias (%)	-32	-27	-32	-22	-21	-25	-23	-22
σ of residuals (ppb)	13	9	10	7	228	166	187	147
Gross error (ppb)	14	8	11	6	233	170	191	140
Normalised gross error (%)	58	34	45	26	57	40	45	32

¹Based on the average of daily values for 13-27 November 2000.

Table 4.3: Performance measures for PM10 and PM2.5 for Melbourne/Geelong monitoring stations for 13–27 November 2000.

Statistical Measure ¹	PM10				PM2.5			
	Pairing type				Pairing type			
	ST	T	S	U	ST	T	S	U
Bias (ppb)	-9	-6	-9	-5	-5	-2	-4	-2
Normalised Bias (%)	-54	-40	-52	-33	-62	-35	-56	-28
σ of residuals (ppb)	8	6	7	6	4	3	4	3
Gross error (ppb)	12	8	10	6	6	4	5	3
Normalised gross error (%)	72	47	62	37	70	42	62	32

¹Based on the average of daily values for 13-27 November 2000.

Delivery performance: Victoria

In addition to generating realistic air quality forecasts, it is critically important that the System be able to generate and disseminate the forecasts in time to provide input into the EPA's daily forecasting process. In the case of both the NSW and Victorian EPAs this requires that the images of ground-level pollutant concentrations and air quality index be available for download from the AAQFS web site before 3 pm each day, and furthermore, that a second forecast be available by 9 am the following morning.

Experience has shown that the 3 pm deadline is the more difficult to achieve because it requires that the System be run during a period of peak demand on the NEC-SX5 supercomputer (i.e., between 11 am and 2 pm). System performance for achieving the 3 pm deadline is summarised in Figure 4.13 where we have plotted the commencement and completion times for the pre-processing, CTM integration, and post-processing stages. The plot is for the period 23 October 2000 to 5 December 2000 and thus encompasses the trial period. Note that the commencement time corresponds to the time at which data files become available from the LAPS05 36 hour forecast (thus the elapsed time doesn't include the time taken to run the numerical weather prediction system). The pre-processing stage includes the generation of day-specific emissions, initial and boundary concentrations and the creation of a NetCDF data packet. This stage typically takes 24 minutes to execute (34% of total time), with the majority of the time taken in the area source module of the inventory generation software. The CTM integration stage takes about 14 minutes to complete (24 hour integration, 350 K grid points, 25 species — 20% of total time) and the supercomputer-based plot file generation takes about 2 minutes — 3%). Approximately 650 plots are generated which are then converted to GIF format, merged with a background map and copied to the web server. The conversion and copying process takes about 30 minutes (43%). Thus the 24-hour PM Victorian forecast takes about 110 minutes to generate (not including the time taken to run the numerical weather prediction system).

Figure 4.11: Time series of daily normalised bias for NO_y (top) and CO (bottom) the period 13–27 November 2000. ST – model and observations paired in time and space; T – paired in time, smallest bias within ± 1 cell ($dx = 1$ km); S – paired in space, smallest bias within ± 1 hour of observation time; U – smallest bias within ± 1 cell and ± 1 hour.

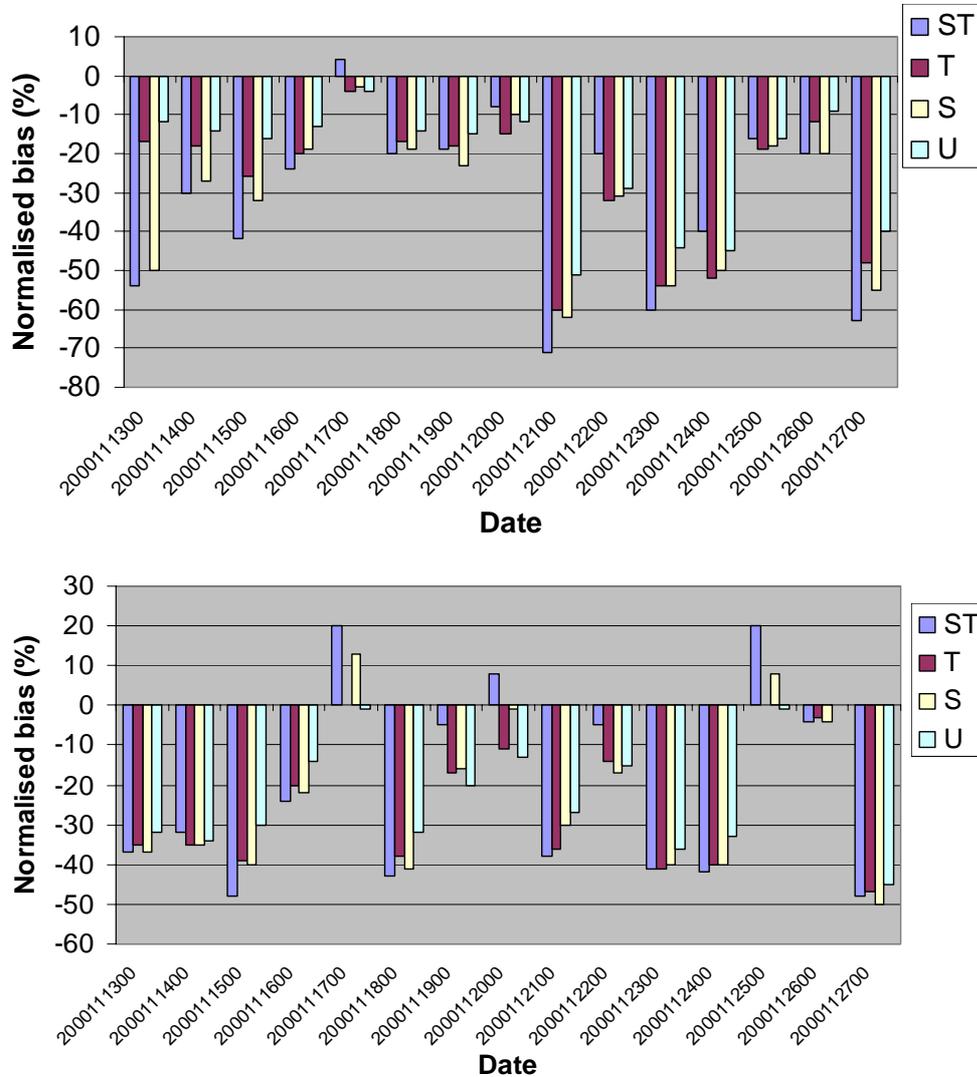
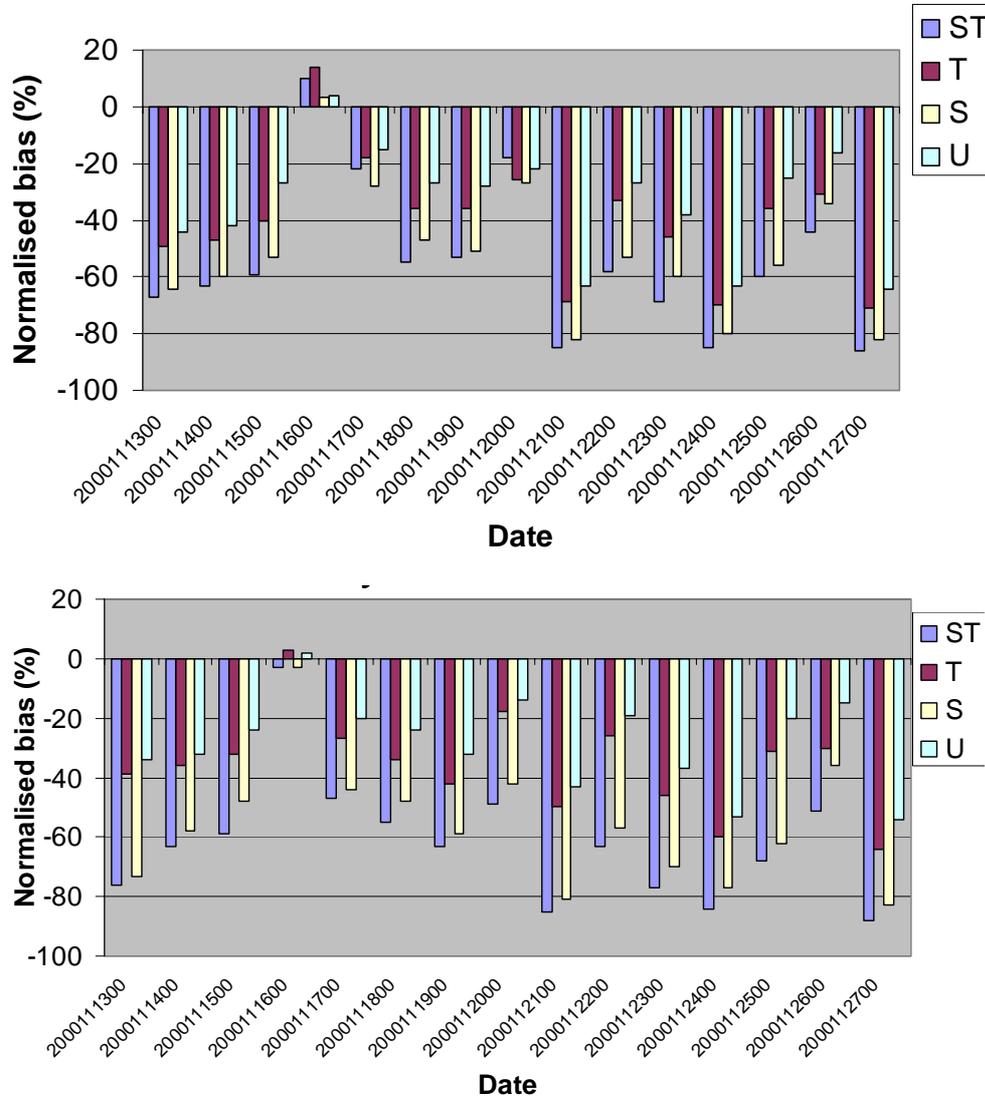
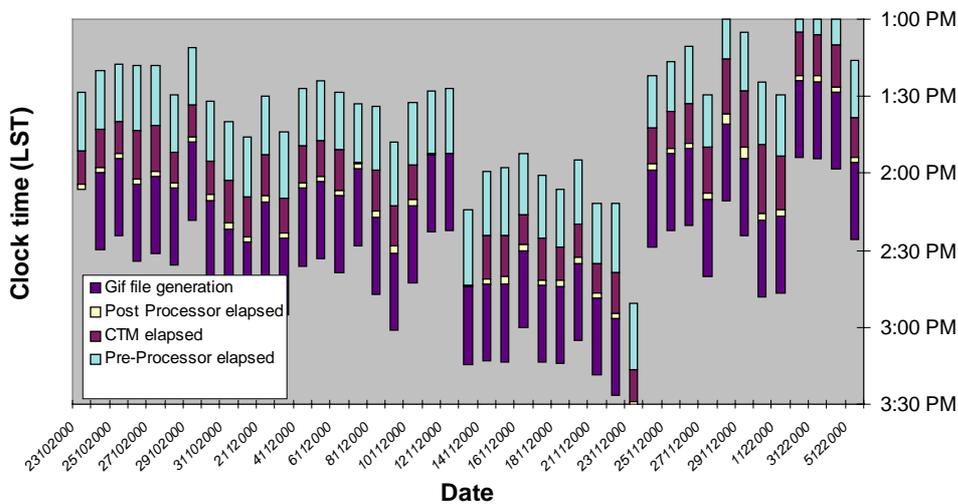


Figure 4.12: Time series of daily normalised bias for PM10 (top) and PM2.5 (bottom) the period 13–27 November 2000. ST – model and observations paired in time and space; T – paired in time, smallest bias within ± 1 cell ($dx = 1$ km); S – paired in space, smallest bias within ± 1 hour of observation time; U – smallest bias within ± 1 cell and ± 1 hour.



From Figure 4.13 it can be seen that the 3 pm deadline was generally achieved until 12 November, after which there were ten occasions when the post-processing was not completed until after the 3 pm (note however, that the images of selected air quality species were placed on the AAQFS web site prior to completion of the entire post-processing step). This corresponded to the period in which one of HPCCC's two supercomputers was decommissioned and its users began to share resources on the NEC-SX5 (see discussion in Section 3). This exacerbated an existing scheduling problem, which led to delays in the starting time of all components of the AAQFS (including the operational weather forecasting system). This problem has since been fixed and (as can be seen from Figure 4.13), the pm air quality forecast for Victoria is now available well before the 3 pm deadline.

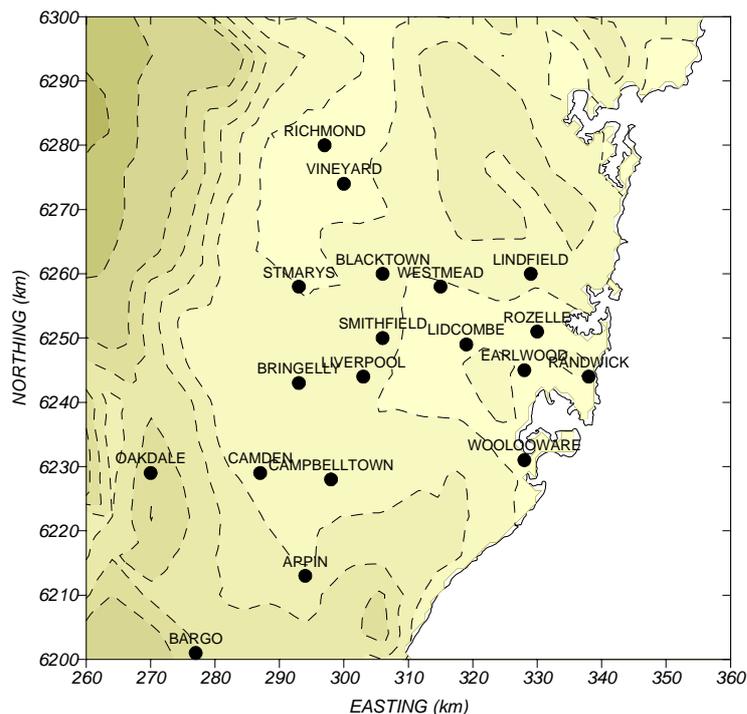
Figure 4.13: Performance data for run times and delivery times of the pre-processing, post-processing and AAQFS chemical transport integration for the 3 pm Victorian air quality 24-hour forecast.



NSW: Observed and forecast air pollutant time series

Plots of observed and forecast air pollutant time series for representative monitoring stations in the Sydney region (site locations are indicated in Figure 4.14) are shown in Figure 4.15 to Figure 4.23. As for the Victorian case, consideration is limited to oxides of nitrogen, nitrogen dioxide, PM₁₀, PM_{2.5}, carbon monoxide and, for two sites, sulfur dioxide. Again we note that model performance for photochemical oxidants will be postponed until the enhanced version of the Generic Reaction Set photochemical mechanism has been installed and tested. Consideration is primarily given to model performance within the vicinity of monitoring stations at Earlwood, Lidcombe and Richmond, thus representing conditions in the coastal, eastern and central regions of the Sydney airshed.

Figure 4.14: Location of key EPA NSW monitoring stations in the Sydney region.



Observed and predicted time series concentrations of NO_y , NO_2 and the ratio $\text{NO}_2:\text{NO}_y$ are given in Figure 4.15 for Earlwood monitoring station for the period 13–27 November 2000. Note that the modelled results have been taken from the 1-km Sydney grid. Also, as discussed in relation to Melbourne, note that the model results refer to the best match between model and observation (for the same hour) within a radius of ± 1 grid cell of the cell that encompasses the monitoring site. This is our suburb-level forecast and is equivalent to the Forecaster stating that the forecast air quality at any given location will fall within the range predicted within a radius of ~ 2 km of the location. From Figure 4.15 it can be seen that the System has predicted the temporal variation of low and moderate concentrations of NO_y with good skill, but has tended to underpredict the highest concentrations (i.e., between hours 96–120, 216–228 and 240–252). The System has matched the observed concentrations of NO_2 with good skill for the entirety of the concentration range, although it should be noted that a significant fraction of the observed variation is due to titration of emitted nitric oxide with the background ozone field (and thus represents a reasonably robust predictive quantity). The time series ratio of $\text{NO}_2:\text{NO}_y$ is a useful variable for comparison because (as discussed in relation to Melbourne), it is indicative of how much the dilution/titration/photochemical ageing of NO_y has occurred since being emitted. NO_y freshly emitted from a motor vehicle typically has a $\text{NO}_2:\text{NO}_y$ ratio of < 0.1 , while aged air may have a ratio which approaches unity. For example, if we consider the observed and modelled ratios for the period between hours 216 and 252 (when the System underpredicted the peak concentrations of NO_y), it can be seen that the observed ratio of $\text{NO}_2:\text{NO}_y$ approaches 0.1, thus implying that Earlwood was downstream of a relatively undiluted plume of NO_y . The System, on the other hand, predicted ratios in the range 0.3–0.5, which suggests that the rate of dilution may have been overestimated during these times.

Observed and modelled concentration time series for PM_{10} , $\text{PM}_{2.5}$ and the ratio $\text{PM}_{2.5}:\text{PM}_{10}$ are shown Figure 4.16 for Earlwood monitoring station. It can be seen that the System has generally matched or slightly underpredicted observed PM_{10} concentrations prior to hour 192, and more strongly underpredicted concentrations after hour 252. The correspondence between the observed and the suburb-level forecast of $\text{PM}_{2.5}$ is very good for all days of the trial period. The plot of observed and modelled $\text{PM}_{2.5}:\text{PM}_{10}$ ratio suggests that the System is consistently underpredicting the concentrations of PM_{10} (as evidenced by the observed/modelled PM_{10} time series discussed above). Noting that the observed ratio of $\text{PM}_{2.5}:\text{PM}_{10}$ was generally well predicted for the Melbourne monitoring stations (e.g., see Figure 4.3), indicates that either the particle emission distributions are substantially different for Sydney, or the prevailing meteorological conditions have led to the generation of higher concentrations of wind-blown dust or sea salt aerosol (neither of which are currently in the model). With regard to the former, it should be noted that paved road dust is not currently included in the Sydney AAQFS emissions inventory. Current estimates of paved road dust made for Melbourne suggest that this is a significant source group, contributing more than 20% of anthropogenic particle mass on an annual basis (motor vehicles are estimated to contribute 10%) to the coarse fraction of the aerosol distribution (PM_{10} - $\text{PM}_{2.5}$).

Observed and forecast ratios of $\text{PM}_{10}:\text{NO}_y$ are shown in Figure 4.17 for Earlwood. Also shown are time series plots of observed and forecast CO and the ratio of $\text{CO}:\text{NO}_y$, but for Rozelle monitoring station (CO is not measured at Earlwood). The $\text{PM}_{10}:\text{NO}_y$ traces follow a similar trend to that observed for Melbourne with both System and observations being constrained to a lower bound of approximately 0.2 (which likely corresponds to the fleet-average ratio for motor vehicles). On the other hand, while observed upper bound ratios are frequently observed to exceed unity, the forecast $\text{PM}_{10}:\text{NO}_y$ ratios do not exceed an upper bound of 0.6. It was suggested in the discussion for Melbourne that large ratios of $\text{PM}_{10}:\text{NO}_y$ may correspond to an aged air mass, or to a non-vehicular source such as wood smoke combustion, industrial, or possibly wind blown dust or aerosol. Upper bound ratios of $\text{PM}_{10}:\text{NO}_y$ are likely to be underestimated by the System in Sydney because the inventory does not include paved road dust and because wind blown dust and sea salt aerosol were not present in the System during the trial period.

Observed and modelled concentrations of CO for Rozelle monitoring station Figure 4.17 (middle) are generally in good agreement. The only significant discrepancy is the underestimate of a peak observed between hours 108 and 120. The modelled ratio of $\text{CO}:\text{NO}_y$ is also in general

agreement with the observed. This suggests that the mean ratio of CO:NO_y for motor vehicle emissions is well represented in the AAQFS emissions inventory for Sydney.

Figure 4.15: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO₂), and NO₂:NO_y ratio for Earlwood monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site (dx= 1 km).

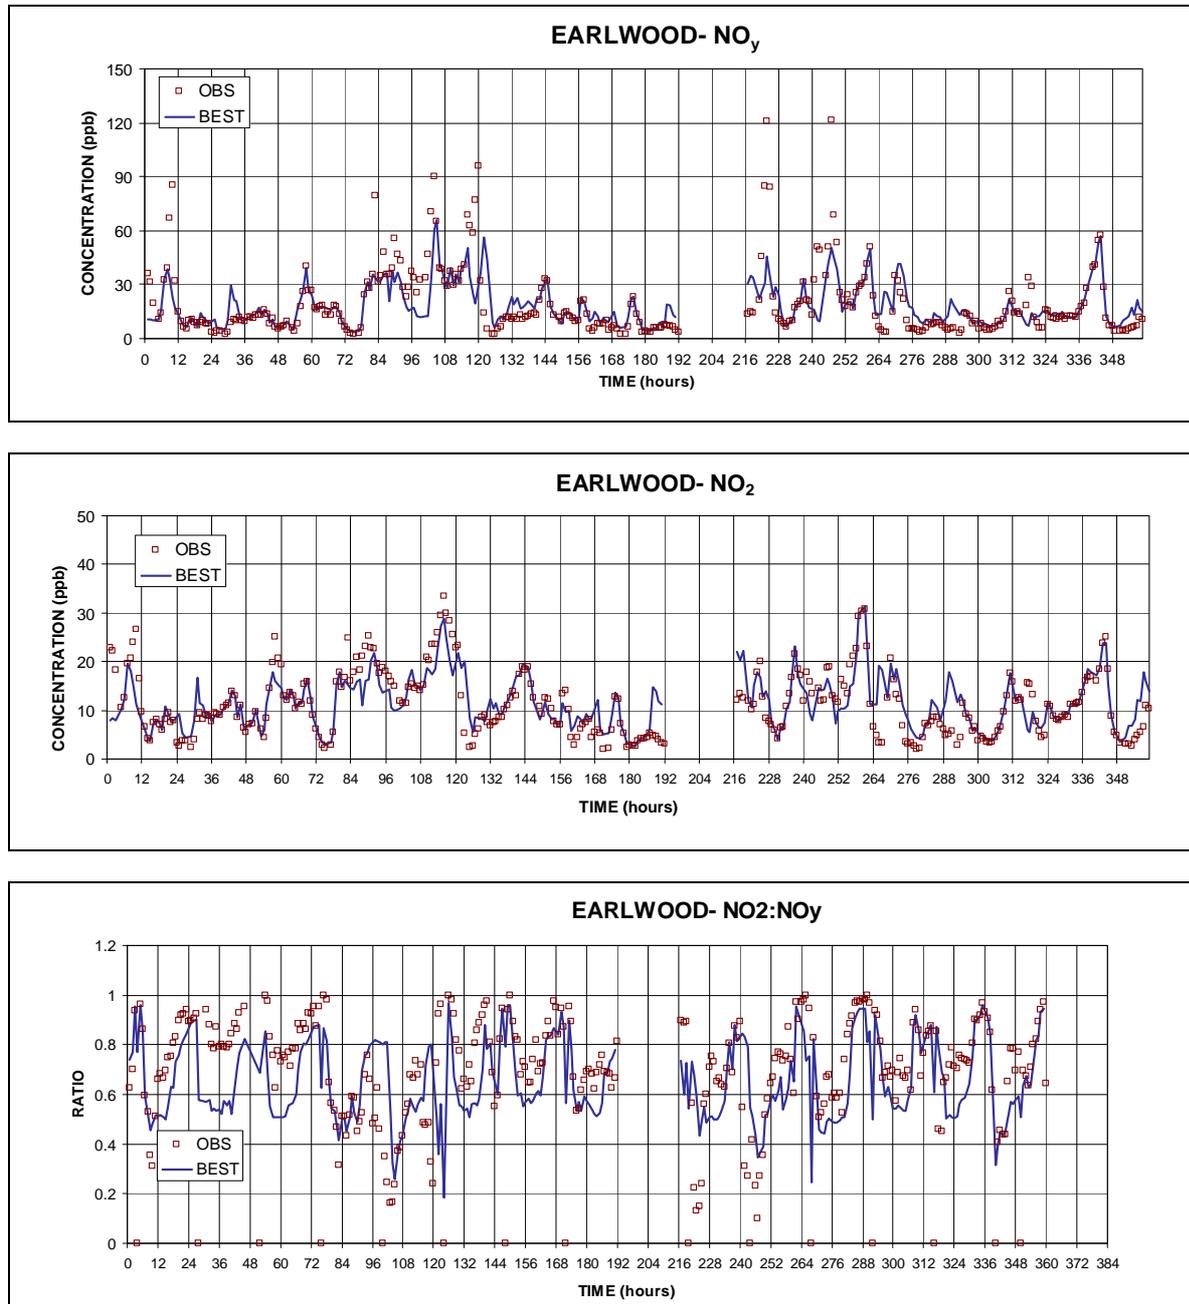


Figure 4.16: Observed vs modelled concentration time series of PM10, PM2.5, and PM2.5:PM10 ratio for Earlwood monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).

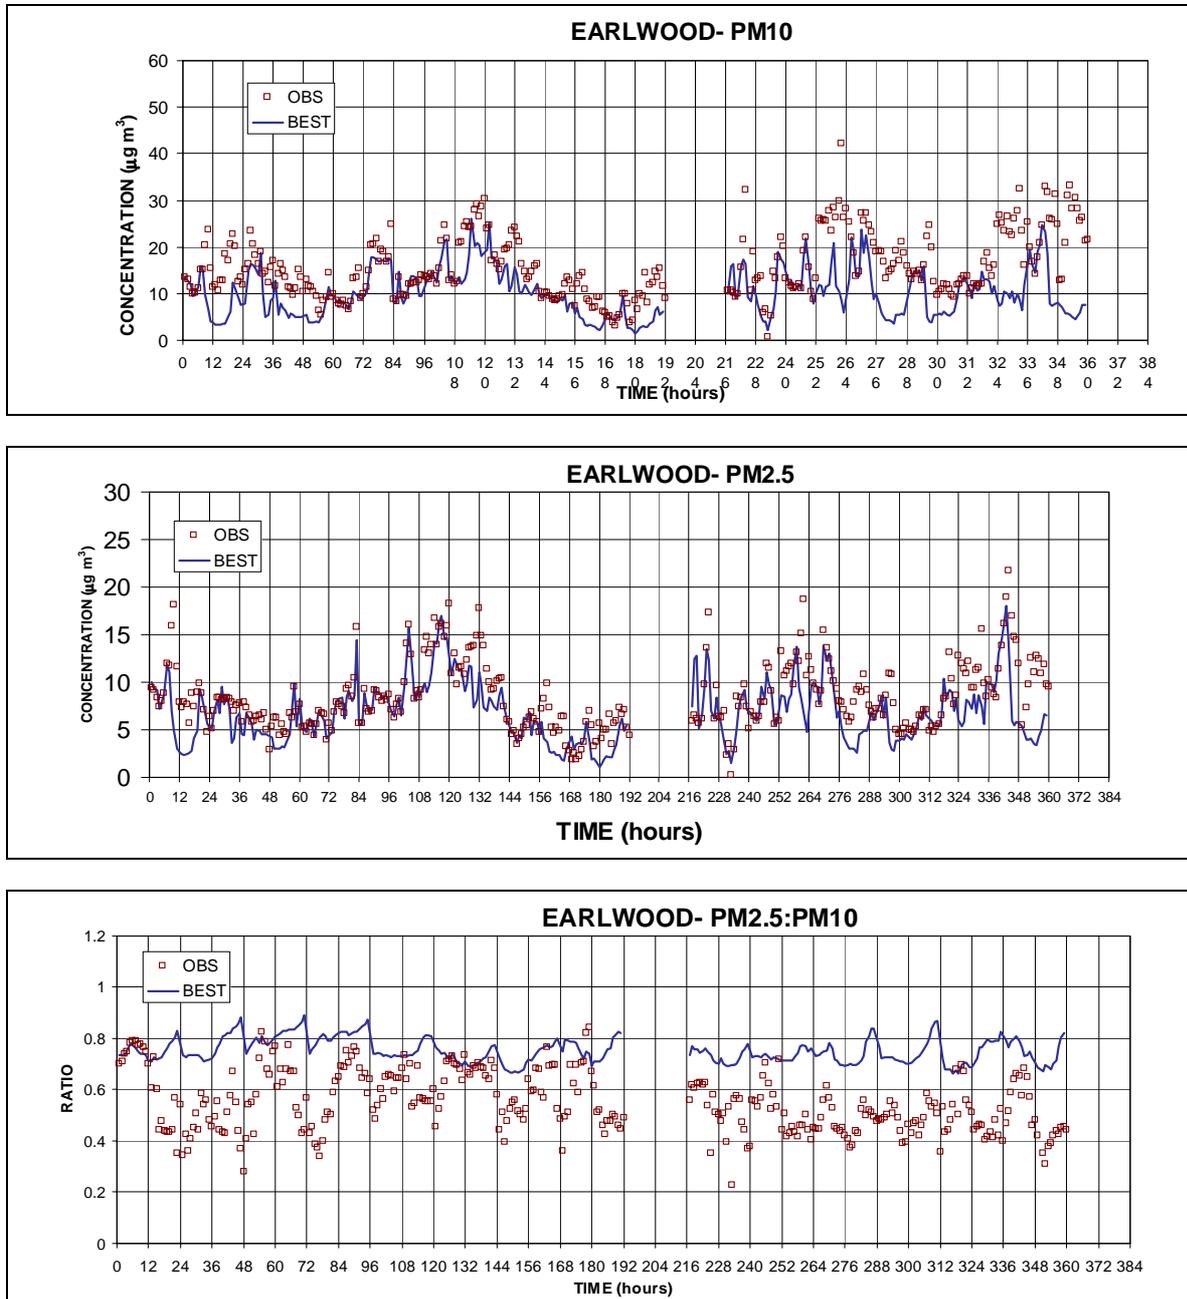
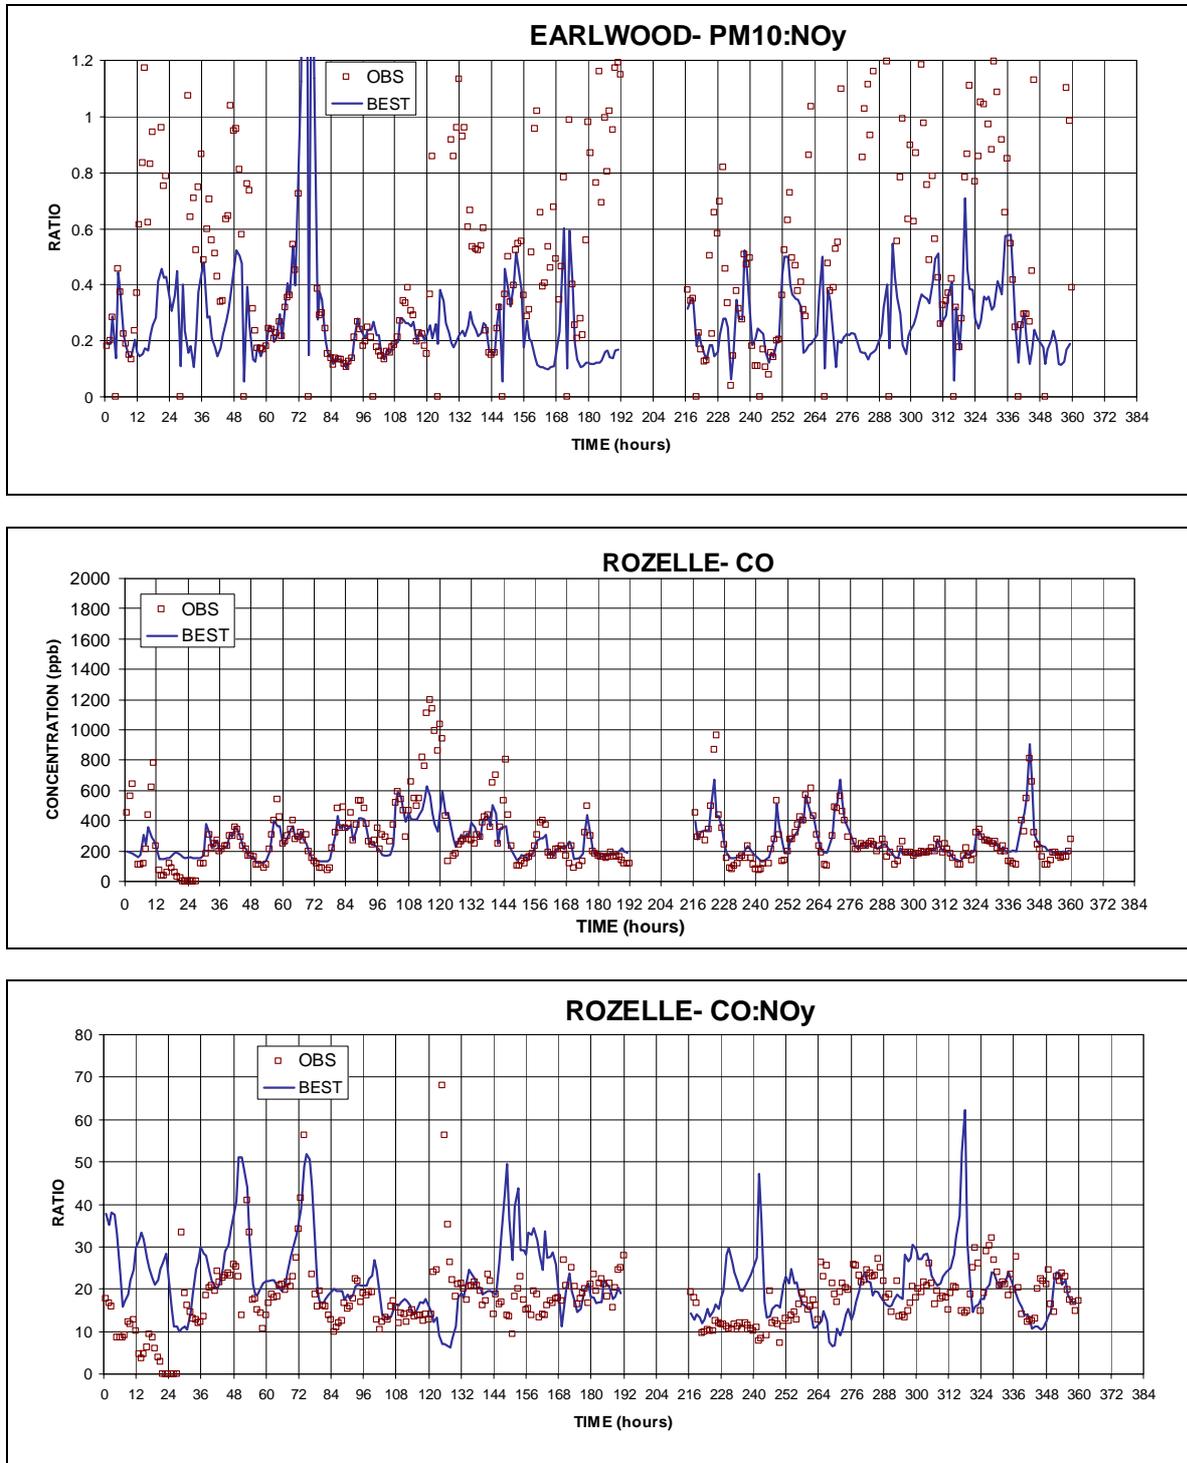
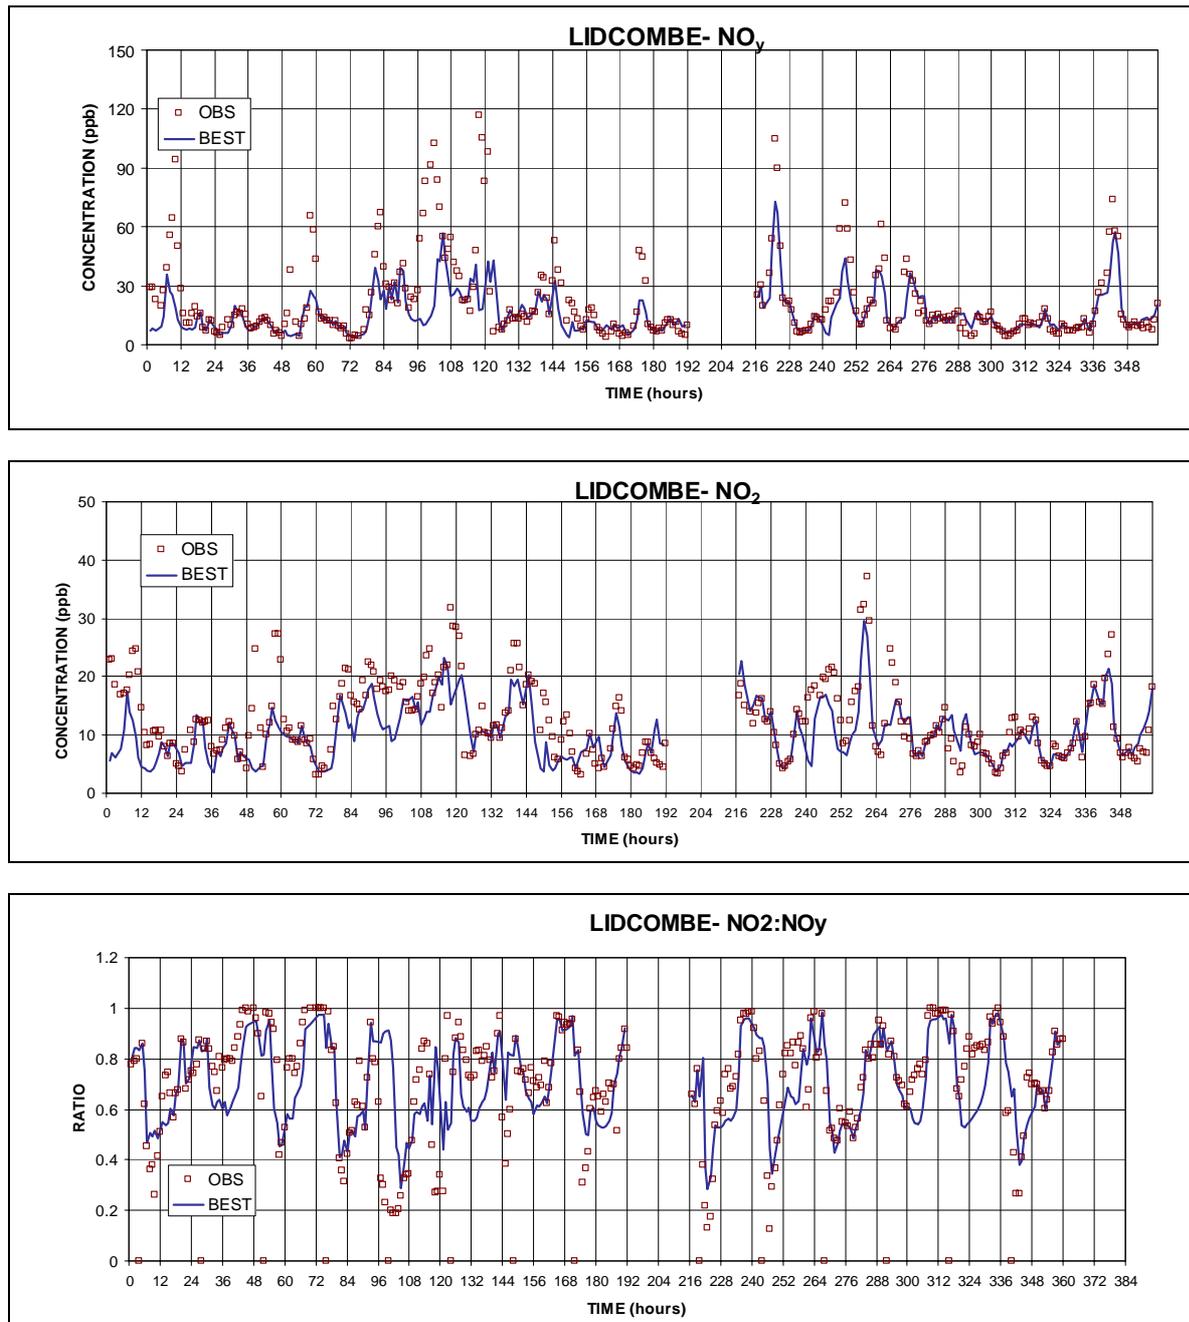


Figure 4.17: Observed vs modelled concentration time series of PM10:NO_y ratio for Earlwood monitoring station and CO, CO:NO_y ratio for Rozelle, 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site (dx = 1 km).



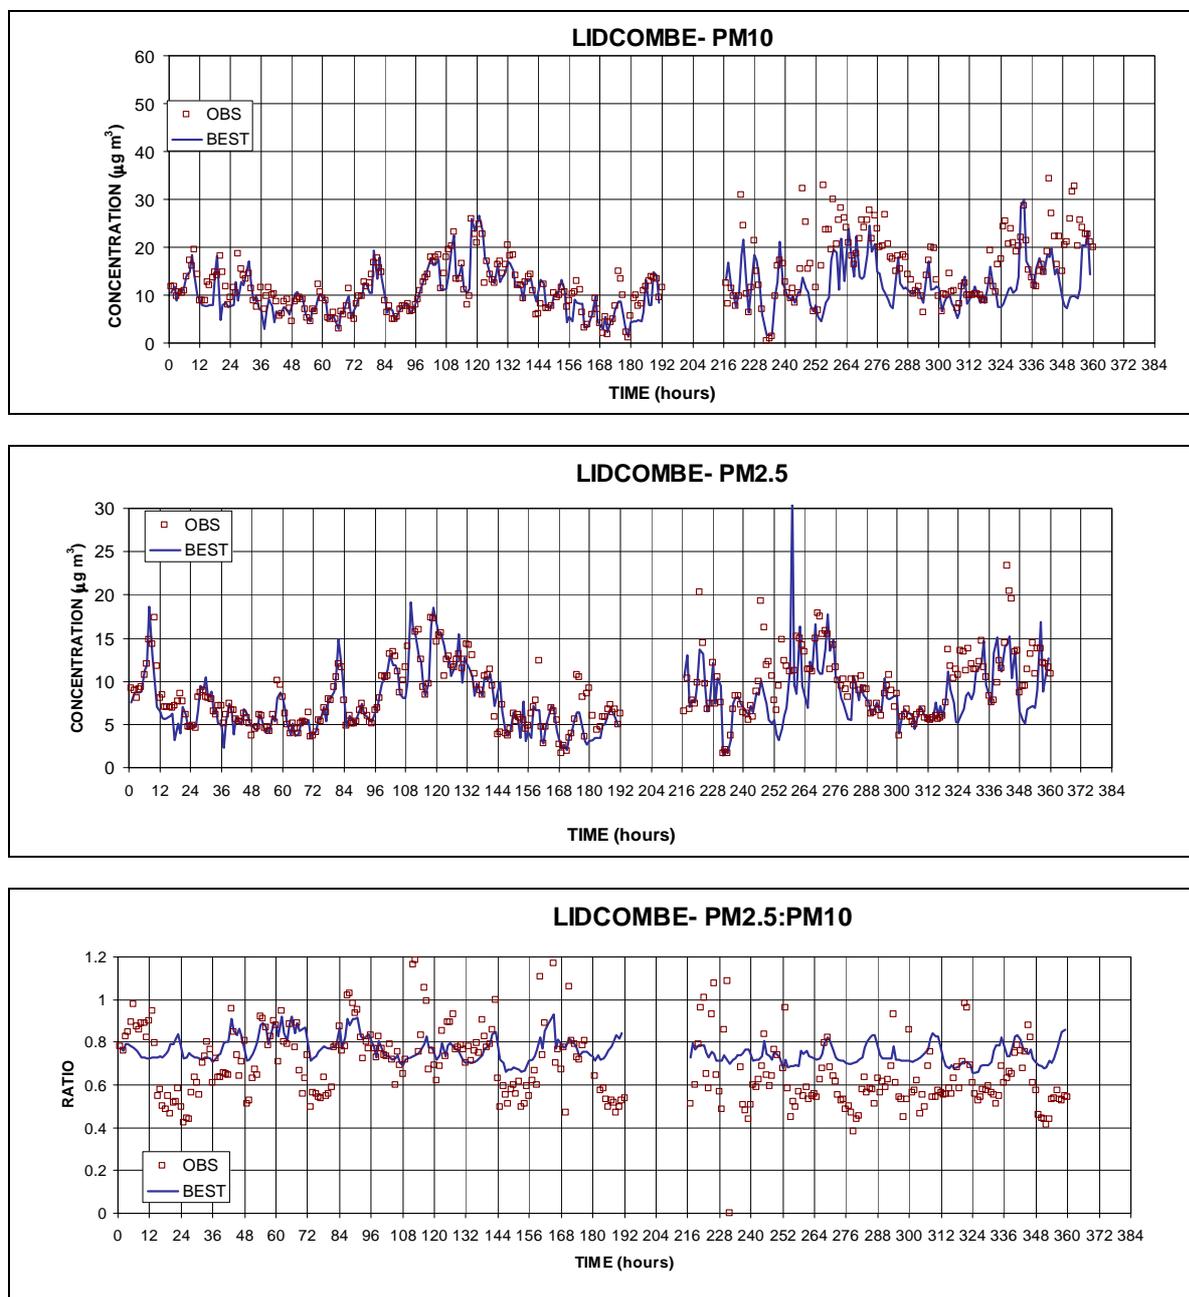
Time series plots of observed and forecast NO_y , NO_2 and $\text{NO}_2:\text{NO}_y$ ratio are shown in Figure 4.18 for Lidcombe monitoring station (see site map — Figure 4.14) for 13–27 November 2000. It can be seen that the System is again able to match the low and mid-level NO_y concentrations with good skill, but often underestimates the peak concentrations. Consideration of the $\text{NO}_2:\text{NO}_y$ ratios suggests that the System may be overestimating the degree of dilution during the period of the observed NO_y peaks.

Figure 4.18: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), and $\text{NO}_2:\text{NO}_y$ ratio for Lidcombe monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).



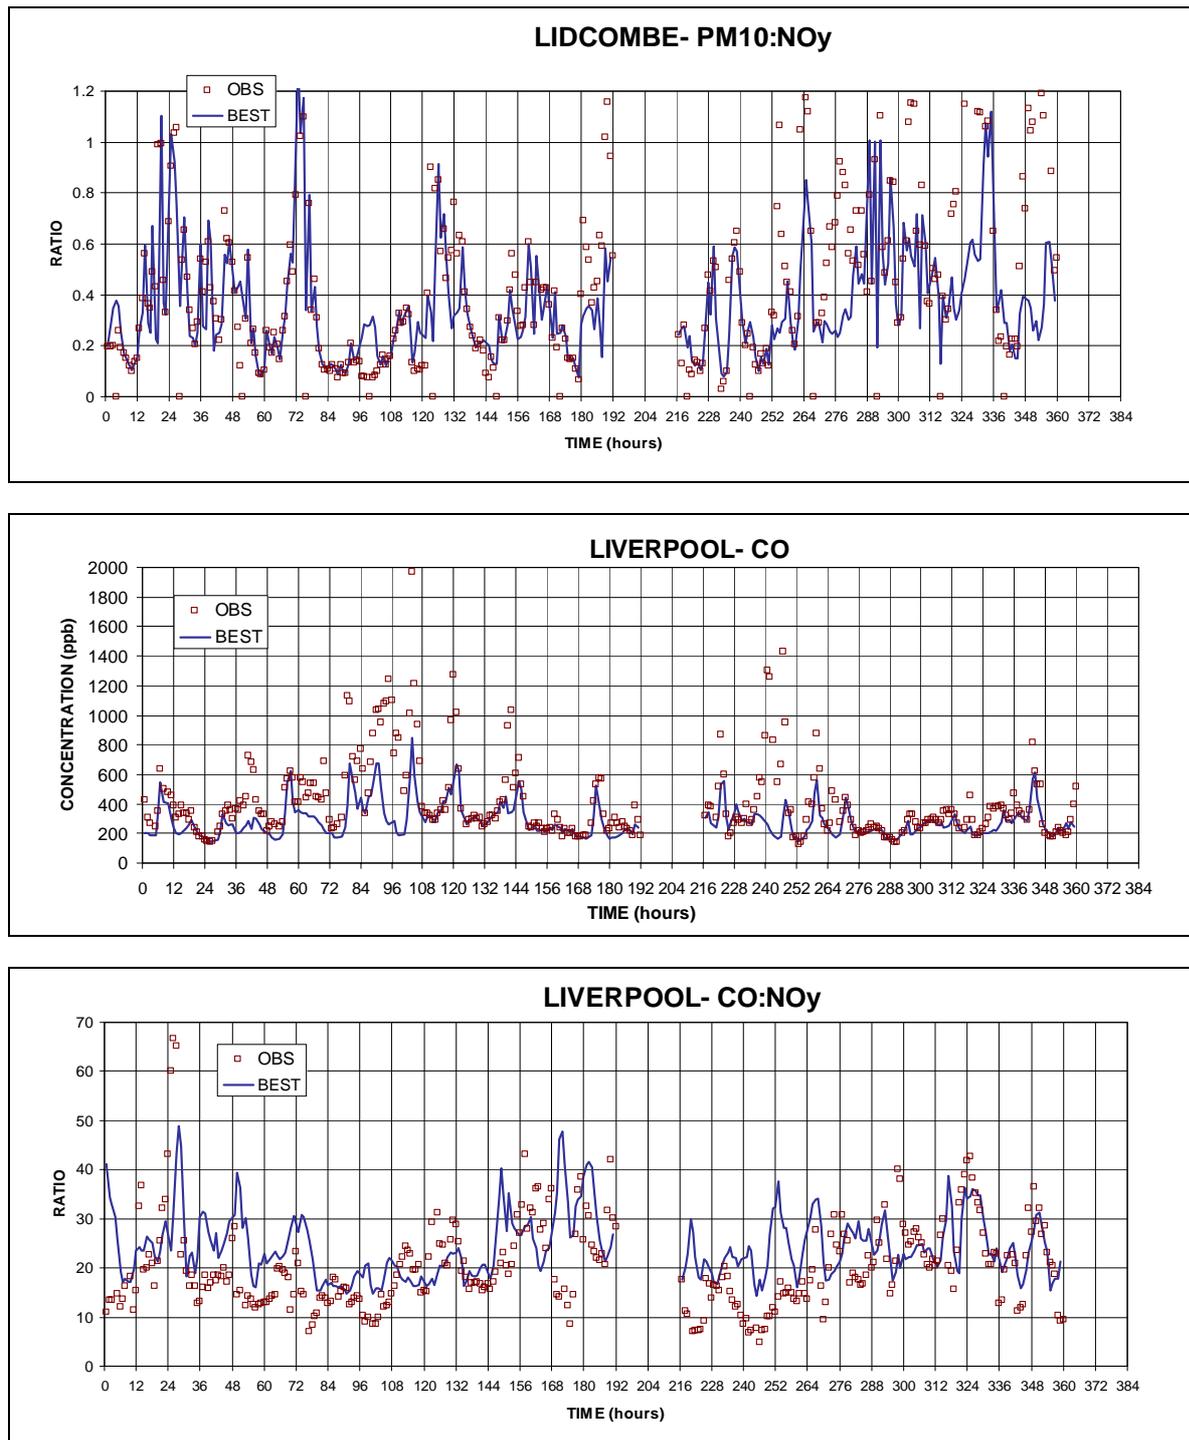
Plots of observed and suburb-level forecasts of PM₁₀, PM_{2.5} and the ratio PM_{2.5}:PM₁₀ are given in Figure 4.19 for Lidcombe. The level of agreement between the observed and the System is very good for both PM₁₀ and PM_{2.5}. The observed ratio of PM_{2.5}:PM₁₀ is either matched or systematically overpredicted by the System. Time series plots of PM₁₀:NO_y ratio for Lidcombe are shown in Figure 4.20. It can be seen that the suburb-level search algorithm has been able to generate a good level of agreement between System and observations during the first half of the trial period. However peak ratios of PM₁₀:NO_y are underpredicted during the second half of the trail period. This, and the PM_{2.5}:PM₁₀ plots discussed earlier, indicate that there is an observed trend of increasing mass in the coarse fraction of the Sydney aerosol over the course of the study period.

Figure 4.19: Observed vs modelled concentration time series of PM₁₀, PM_{2.5}, and PM_{2.5}:PM₁₀ ratio for Lidcombe monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).



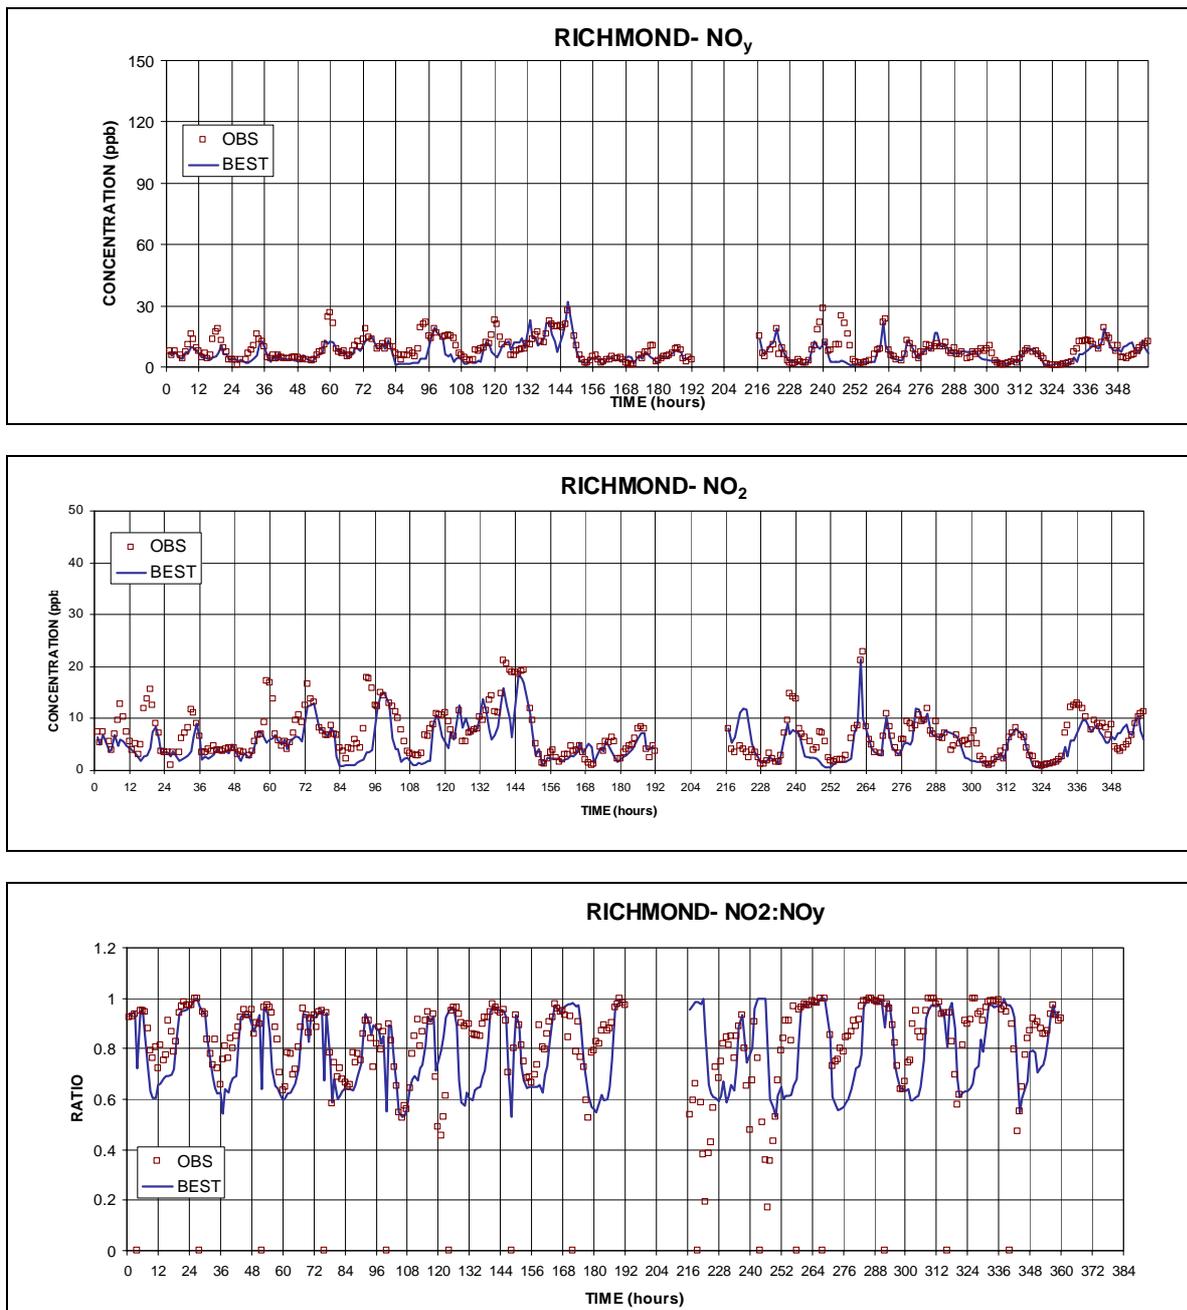
In lieu of CO measurements at Lidcombe, we have plotted a comparison of observed and modelled CO and CO:NO_y ratio for Liverpool monitoring stations (Figure 4.20 — middle and bottom). Again, low and mid-level concentration peaks have been well matched by the System while upper bound peaks are underestimated. From Figure 4.20 (bottom) it can be seen that there is a general good level of agreement between observed and modelled CO:NO_y ratio. Both of these results are consistent with what has been previously discussed.

Figure 4.20: Observed vs modelled concentration time series of PM10:NO_y ratio for Lidcombe monitoring station and CO, CO:NO_y ratio for Liverpool, 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).



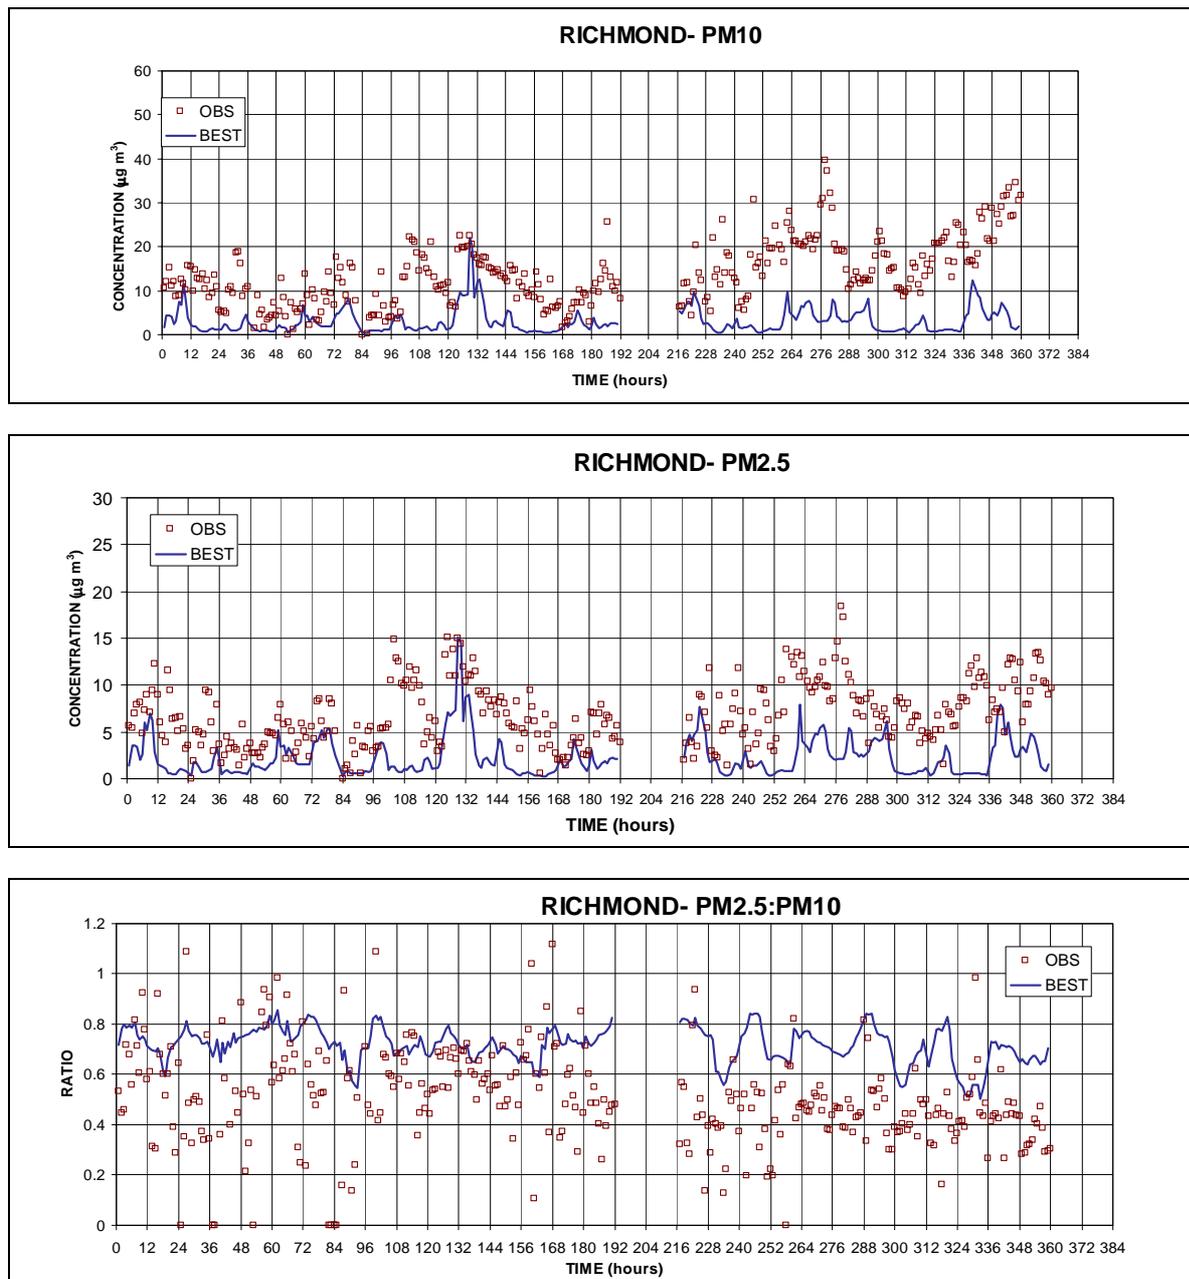
Richmond monitoring station is the final site for which a comparison of observed and modelled pollutant concentration time series are reported. It can be seen (Figure 4.21) that low concentrations of NO_y are both observed and forecast. Moreover, the System has done an excellent job in matching the observed diurnal variation of NO_2 (e.g., note the good correspondence between observed and modelled peak at hour 263). Observed and modelled time series of $\text{NO}_2:\text{NO}_y$ ratio are also generally well matched, although it can again be seen that two periods of low observed $\text{NO}_2:\text{NO}_y$ ratio have been underestimated by the System.

Figure 4.21: Observed vs modelled concentration time series of oxides of nitrogen (NO_y), nitrogen dioxide (NO_2), and $\text{NO}_2:\text{NO}_y$ ratio for Richmond monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx=1$ km).



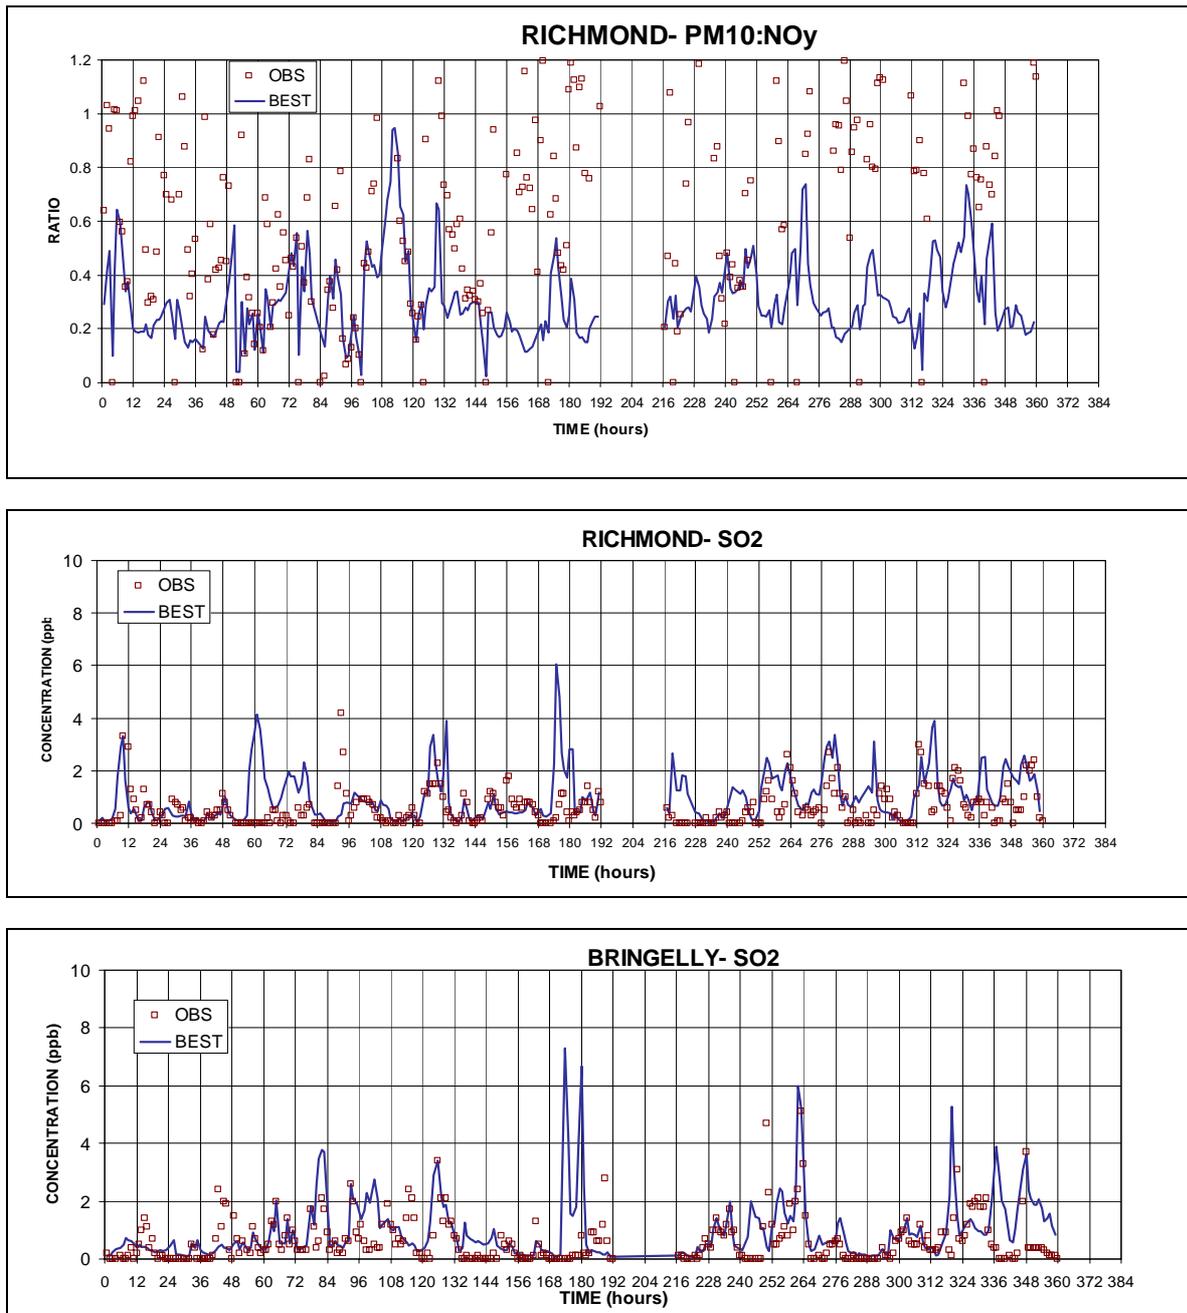
The System does less well at matching the observed time series of PM₁₀, PM_{2.5} and PM_{2.5}:PM₁₀ ratio (Figure 4.22). Here it can be seen that the System has underpredicted the concentrations of PM₁₀, PM_{2.5} and has overpredicted the PM_{2.5}:PM₁₀ ratio, particularly during the second half of the trial period. Moreover, the observed time series of PM₁₀:NO_y is generally underpredicted by the System (see Figure 4.23). This suggests that Richmond monitoring station may have been influenced by a particle source that has not yet been included in the System (wind blown dust and possibly wood combustion).

Figure 4.22: Observed vs modelled concentration time series of PM₁₀, PM_{2.5}, and PM_{2.5}:PM₁₀ ratio for Richmond monitoring station 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).



We complete the concentration time series comparisons with plots of observed and modelled sulfur dioxide concentration (SO_2) for Richmond and Bringelly monitoring stations (Figure 4.23 — middle and bottom). Given that SO_2 is primarily emitted from isolated industrial point sources, it can be seen that the System has generally performed quite well in matching the diurnal variation and magnitude of SO_2 peaks observed at both monitoring stations.

Figure 4.23: Observed vs modelled concentration time series of $\text{PM}_{10}:\text{NO}_y$ ratio for Richmond monitoring station and SO_2 for Richmond and Bringelly monitoring stations 13–27 November 2000. Forecast concentrations are the best (closest matching) to the observed within ± 1 grid cell of the observation site ($dx = 1$ km).



NSW: Residual analysis

A description of the statistical measures used to quantify model performance errors are given in Table 4.1 and discussed further in relation to the Melbourne air quality forecasting performance. Performance statistics for Sydney for NO_y, PM10 and PM2.5 (there are insufficient data to perform the residual analysis for carbon monoxide) are given in Table 4.4 and Table 4.5. Plots of the daily variation in the normalised bias are presented in Figure 4.24 and Figure 4.25.

Table 4.4: Performance measures for NO_y for Sydney monitoring stations for 13–27 November 2000.

Statistical Measure ¹	NO _y			
	Pairing type			
	ST	T	S	U
Bias (ppb)	-13	-10	-12	-8
Normalised Bias (%)	-41	-30	-38	-24
σ of residuals (ppb)	18	14	15	13
Gross error (ppb)	18	12	14	9
Normalised gross error (%)	58	36	46	27

¹Based on the average of daily values for 13–27 November 2000.

Table 4.5: Performance measures for PM10 and PM2.5 for Sydney monitoring stations for 13–27 November 2000.

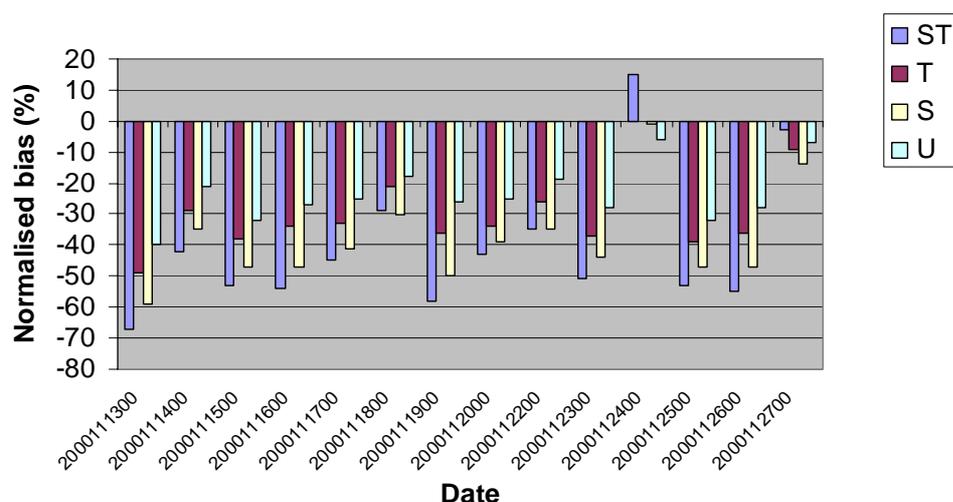
Statistical Measure ¹	PM10				PM2.5			
	Pairing type				Pairing type			
	ST	T	S	U	ST	T	S	U
Bias (μg m ⁻³)	-9	-7	-8	-6	-4	-3	-3	-2
Normalised Bias (%)	-59	-48	-56	-42	-53	-36	-48	-30
σ of residuals (μg m ⁻³)	9	8	8	8	4	3	4	3
Gross error ((μg m ⁻³)	11	8	10	7	5	4	5	3
Normalised gross error (%)	73	52	64	45	62	41	54	33

¹Based on the average of daily values for 13–27 November 2000.

From Table 4.4 and Figure 4.24 it can be seen that, similar to Melbourne, the System is negatively biased with respect to the forecasting of 1-hour NO_y concentration. For this species, the normalised mean bias varies from -41% for time-space paired comparisons to -24% for observed-modelled pairs where the best concentration prediction within ± one grid cell and ± one hour of the observation location and time are chosen. Although the latter is a very acceptable result for a forecast system, it nevertheless represents a slightly larger degree of underprediction than was observed in Melbourne

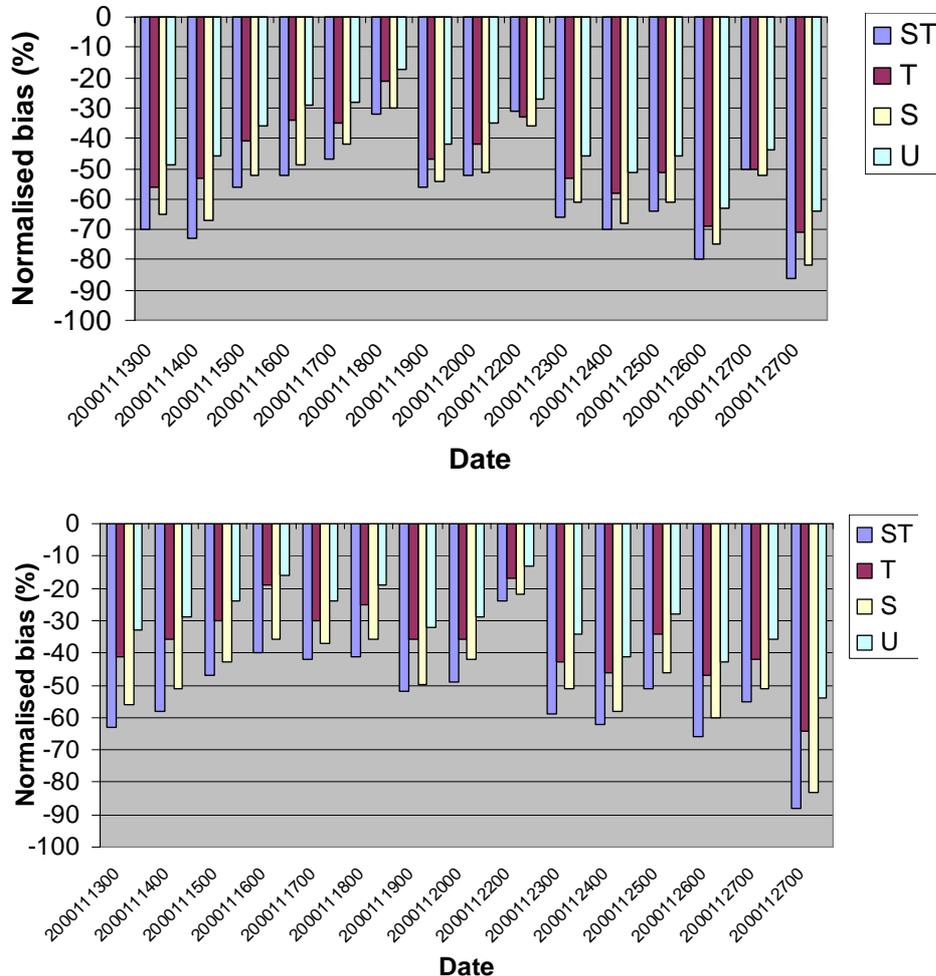
(-32% to -22%). A similar result holds in the case of the normalised gross error (Melbourne: 58% (ST), 34% (T), 45% (S), 26% (U); Sydney: 58% (ST), 36% (T), 46% (S), 27% (U)). Similar to Melbourne, the greatest level of improvement occurs when the spatial coupling is relaxed. Thus the model is better able to forecast fine-scale temporal variations than fine-scale spatial variations in the primary species. This is perhaps not surprising when it is recalled that a significant amount of the temporal variation is a result of strong diurnal variations in vehicular emissions convoluted with strong variations in boundary-layer depth.

Figure 4.24: Time series of daily normalised bias for NO_y for the period 13–27 November 2000. ST – model and observations paired in time and space; T – paired in time, smallest bias within ± 1 cell ($dx = 1$ km); S – paired in space, smallest bias within ± 1 hour of observation time; U – smallest bias within ± 1 cell and ± 1 hour.



Performance statistics and plots of normalised bias for PM10 and PM2.5 are shown in Table 4.5 and Figure 4.25. It can be seen that the normalised bias and normalised gross error for Sydney PM10 are somewhat larger than was calculated for Melbourne, while the PM2.5 statistics for Sydney show a comparable level of performance for the System. Thus, in the case of the latter, a suburb-level forecast for PM2.5, with a temporal uncertainty of ± 1 hour will, on the average, match observed PM2.5 concentrations by about 30%. This can be considered a very promising result given that we plan to introduce additional particle source groups into the System.

Figure 4.25: Time series of daily normalised bias for PM10 (top) and PM2.5 (bottom) the period 13–27 November 2000. ST – model and observations paired in time and space; T – paired in time, smallest bias within ± 1 cell ($dx = 1$ km); S – paired in space, smallest bias within ± 1 hour of observation time; U – smallest bias within ± 1 cell and ± 1 hour.

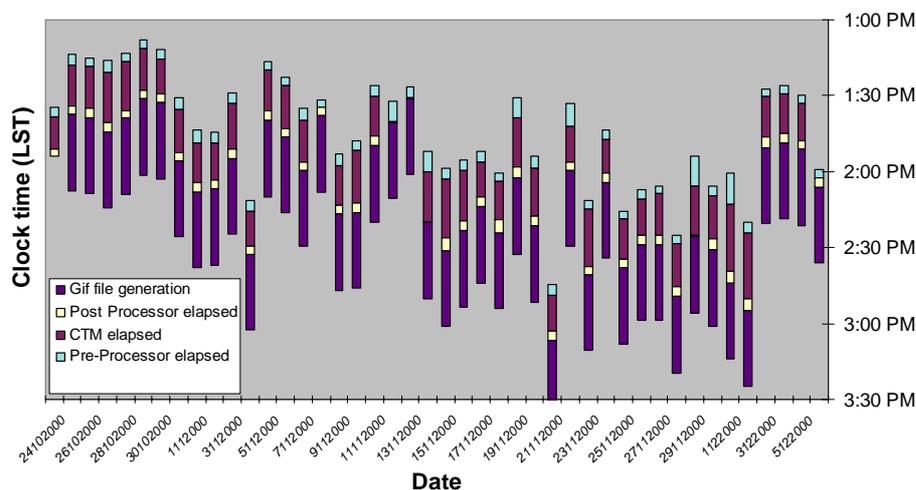


NSW: Delivery performance

As discussed in relation to Melbourne, the 3 pm forecast is the most difficult to achieve with the AAQFS and thus delivery performance is limited by this forecast. The commencement and completion times for the NSW/Sydney pm air quality forecast are shown in Figure 4.26 for the period 23 October 2000 to 5 December 2000. Note that, in contrast to the Victorian case, the forecast is for 36 hours. This is because the generation of a 36-hour forecast for NSW takes 53 minutes, which amounts to less than half the time taken to generate the 24-hour Victorian forecast. In part this is because the Victorian domain is 30% larger (353K vs. 241K grid points), but also stems from the fact that the pre-processor stage takes up to ten times longer to complete for the more detailed Victorian emissions database. Typical elapsed times for each stage of the CTM system execution for NSW are: pre-processing- 4 minutes (9% of the total); CTM integration- 16 minutes (29%); plot file generation- 3 minutes (6%); gif file conversion and copying- 30 minutes (56%). With reference to Figure 4.26 it can be seen that, similar to the Victorian case, delivery performance was very respectable until a reduction in the total supercomputing capacity, coupled with job scheduling problems led to some occasions where the

NSW air quality forecasts were not completed in time for the 3 PM deadline. However, it can also be seen that these scheduling problems have now been overcome, leading to early delivery of the forecasts for the period 2-5 December 2000.

Figure 4.26: Performance data for run times and delivery times of the preprocessing, postprocessing and AAQFS chemical transport integration for the 3 pm NSW 36-hour air quality forecast.



Conclusions

This preliminary validation study for the Victorian and NSW air quality forecasts highlights the potential for the AAQFS to provide accurate suburb-level forecasts for a range of air pollutants. The system has generally been able to forecast the day-to-day variation of the NO_y , NO_2 and CO with good skill, although there is currently a trend to underpredict concentrations. Aerosol concentrations are also predicted with reasonable skill, although the degree of underprediction is larger than for the gaseous species. One of the reasons for the System's tendency to underpredict aerosol concentrations may be the absence of wind blown dust and sea salt aerosol in the model during the trial period. Paved road dust is also a significant source that is not currently included in the AAQFS inventory for NSW. With regard to these source groups, sea salt aerosol has now been included into the CTM and a software description of an algorithm for generating wind blown dust is currently being developed and tested. Subject to EPA-NSW approval, EPA-VIC has committed to including paved road dust into the NSW inventory.

The general trend of underprediction for the primary gaseous species has three potential causes: 1: dilution rates are slightly overpredicted by the meteorological model; 2: emission fluxes are slightly low; 3: turbulent dispersion processes are not correctly represented in the CTM. For example, in the case of the latter, the stability of the nocturnal urban boundary layer is forced to neutral. Use of this parameterisation of the urban heat island effect may not be appropriate for all regions of the Sydney basin. It is recommended that each of the potential causes of underprediction listed above be examined by the responsible groups before further System validation is undertaken. Note also that the improved version of the Generic Reaction Set photochemical mechanism will be installed into the CTM. Following this, system performance with respect to the forecasting of ozone events will be examined.

5. Acknowledgments

The Australian Air Quality System is supported by funds from the Natural Heritage Trust component of *Air Pollution in Major Cities Program*, approved by the Minister for the Environment as one of the *Clear the Air* projects on 15 May 1998, and managed by Mr Paul Kesby, Director of the Air Quality Section in Environment Australia. The project is legally organised as a contract between the Commonwealth and CSIRO Atmospheric Research, with subcontracts to collaborators including: CSIRO Energy Technology, Bureau of Meteorology Research Centre, Environment Protection Authority (Victoria) and Environment Protection Authority of New South Wales.