Climate Dynamics (1997) 13:35–43

Climate Dynamics

A new general circulation model: formulation and preliminary results in a single- and multi-processor environment

L. M. Leslie[1], **K. Fraedrich**[2]

[1] School of Mathematics, UNSW, Sydney, 2052 Australia
[2] Meteorologisches Institut, University of Hamburg, Hamburg, Germany

Received: 28 February 1996 / Accepted: 30 July 1996

Abstract. This article describes a new general circulation model (GCM) developed jointly by The University of New South Wales (UNSW) and the University of Hamburg. The model is versatile in that it can be run as a medium-range (1 to 15 days) global numerical weather prediction (NWP) model; as an extended range (15 to 30 days) NWP model; and as a GCM for periods extending from seasons, through annual and decadal periods, and beyond. The model can be coupled with ocean models that vary in complexity from simple "swamp" oceans to complex ocean GCMs. The atmospheric GCM also has a number of novel features, particularly in the numerical integration scheme which is a high-order, mass-conserving, semi-implicit semi-Lagrangian scheme, thereby removing the stability restriction on the time-step and allowing efficient long-term integrations. The emphasis here will be on demonstrating that the new model performs effectively on the usual measures of skill (statistics such as mean errors, root-mean-square errors and anomaly correlations) in several standard applications upon which new models usually are assessed. These applications include medium range weather forecasts out to 10 days on a daily basis over a one year period; a limited 10-year simulation climatology, prediction of atmospheric anomalies using SST anomalies in an El Nino year; and an alternative two-way approach to regional modelling (the "down-scaling problem") made possible because the unconditional stability of the semi-implicit, semi-Lagrangian formulation permits large variations in grid spacing without changing the time step size. Finally, the model is run on a variety of parallel computing platforms and it is shown that near-linear speed-up can be attained. This is significant for both medium range NWP and very long-term GCM integrations.

1 Introduction

A collaborative general circulation (GCM) modelling effort has been established between The University of New South Wales (UNSW), Sydney, and the University of Hamburg. The modelling is part of a larger planned program of prediction and diagnosis and is intended to provide forecasts/simulations over a large range of time scales. There are three main aims of the present study. The first aim simply is to describe the salient features of the new model, particularly any features that can be regarded as being different from existing GCMs. The second aim is to demonstrate the performance of the new model in a range of applications regarded as necessary for a credible GCM. These applications involve assessing the model: as a medium range forecast model providing daily 10-day forecasts using operational global data; as a climate model, integrated over a 10-year period, and then compared with global analyses for the same period; as an atmospheric model driven by a regional SST anomaly field over a twelve month period that was an El Nino year; and in a down-scaling application in which a new, more efficient procedure requires only one forecast and permits two-way interaction between the GCM and the region of interest. The third and final aim is to compare the efficiency of the model on both single processor and parallel computing platforms.

Returning to the first aim, in Sect. 2 the new GCM is described briefly. It is the end result of a number of years of research on developing increasingly accurate and efficient semi-implicit, semi-Lagrangian schemes. This work has been carried out almost entirely in collaboration with the National Meteorological Center (NMC), now a National Center for Environmental Prediction (NCEP), Washington DC. The attractiveness of the semi-implicit, semi-Lagrangian schemes has

---

Correspondence to: L. M. Leslie, e-mail: L.Leslie@unsw.edu.au

This paper was presented at the Third International Conference on Modelling of Global Climate Change and Variability, held in Hamburg 4–8 Sept. 1995 under the auspice of the Max Planck Institute for Meteorology, Hamburg. Editor for these papers is L. Dümenil

been manifold: there are no formal (linear) stability restriction applies to the time step size (a favourable feature for long-term/climate simulations), it has excellent phase and amplitude properties for semi-Lagrangian schemes third-order and above, it is easy to incorporate two-time level formulations thereby avoiding the generation of spurious computational modes, and it has small computational dispersion. For a review of the semi-Lagrangian procedure, see Staniforth and Côté (1991), and for a selection of the work leading up to the current model, see Leslie and Purser (1991), Purser and Leslie (1995) and Leslie and Purser (1995). On the negative side it is well known that the semi-Lagrangian schemes formally do not conserve domain integrals, such as total mass, that are conserved by the governing equations. This is not a problem in short-range NWP (1–2 days) but has serious implications for long integrations, rendering them unusable. Fortunately, in recent years there has been considerable progress in redressing this problem, and the present model uses the procedure developed by Leslie and Purser (1995) who demonstrated that global conservation of mass could be built into the semi-Lagrangian interpolations in a direct, "local", manner.

The second aim of this study, presented in Sect. 3, is to demonstrate how well the new GCM performed when it was applied to a range of problems. Three problems are chosen to demonstrate that the new model performed at an acceptable level of skill. The first performance indicators are the anomaly correlations for days 1 to 10, calculated each day over the period January 1, 1994 to December 31, 1994 using archived operational 00 UTC global analyses. The second aim is to produce a brief model climatology, based largely on time-averaged fields of zonal means of temperature and wind, and of SLP and precipitation, for the December to February and June to August periods. The third application of the GCM is to simulate the atmospheric pattern over the Pacific Ocean in which the SST anomalies were provided. Here it was hoped that the GCM could exhibit skill in predicting atmospheric anomalies over eastern Australia and the adjacent south Pacific Ocean. The fourth and final application of the model, as described in Sect. 4, introduces a new procedure for driving a high resolution regional simulation from a GCM, in which the idea is to concentrate the model resolution over the region of interest such that the resolution slowly degrades away from the region of interest to match that of the GCM.

In order to meet the ever-increasing needs of GCMs, much faster computers will be introduced and the GCMs must be able to take advantage of this increased capacity. It is anticipated that most of these developments will come from highly parallel computers. The GCM described here has been ported to such a computing environment, namely a 16 processor computer that has a peak performance of about 4 Gflops and a sustained performance of around 1 Gflop. Details of the implementation procedure are given in Sect. 5, along with performance indicators.

2 The model

In this section a short outline will be given of the model in terms of the numerical schemes used for the temporal and spatial differencing. Following that, the representation of the various physical processes such as vertical turbulent fluxes, radiation, cumulus convection, etc. will also be described briefly. Most of the novel features of the model are in its "numerics", while the "physics" is quite standard.

Numerics

As mentioned already, the model is based on the semi-implicit, semi-Lagrangian regional/global model of Leslie and Purser (1991) and Leslie and Purser (1995). The most advanced form of the model employs 6th-order differencing and forward trajectories. However, the version used in this study is that which has been used most heavily in previous research applications. It is a two-time level, bi-cubic semi-Lagrangian scheme, with third-order accurate back-ward trajectory calculations (Leslie and Purser 1991). The governing equations are in conservation form wherever possible, and the local mass-conserving scheme of Leslie and Purser (1995) also has been incorporated so that the model can be run as a climate model. The novel features of the numerics compared with other GCMs are the high-order accurate differencing and the mass-conserving formulation of the semi-Lagrangian scheme. Together, these features enable very long-term integrations of the model to be carried out with only a small change in the global total mass and total energy. The high-order differencing (greater than or equal to third order) is made possible by the "cascade" interpolation procedure which divides the multi-dimensional interpolation into a sequence of one-dimensional interpolations. The cascade interpolation, which produces enormous computational savings for medium and large grid dimensions, is described in full detail by Purser and Leslie (1991). The cascade interpolation procedure also is used for the exact, conservation of mass and tracers by working with the cumulative distributions of the variable to be conserved, in this case mass and tracer quantities (see Leslie and Purser 1995). Unlike the a posteriori semi-Lagrangian conservation schemes (Priestley 1993; Gravel and Staniforth 1994), which in principle can allow spurious local corrections to be made to ensure global conservation, the direct conserving scheme employed here ensures that a high level of local accuracy is maintained.

Physics

The representation of physical processes in the model is summarized next. The vertical fluxes of momentum, heat and moisture in the boundary layer are the same as those used by Louis (1979), while shallow convection is the same as that described by Geleyn (1987).

Large-scale precipitation falls if the relative humidity exceeds 95%. Cumulus convection is based on the mass flux scheme of Tiedtke (1989). Gravity wave drag is obtained from the formulation of Palmer et al. (1986). The model uses the modified Schwarzkopf and Fels (1985, 1991) scheme for the long-wave radiation and the Lacis and Hansen (1974) scheme for the short wave radiation. Cloud cover diagnosis for the radiation schemes is calculated in three layers: low, middle and high, using a modification of the scheme developed by Slingo (1987). The surface temperatures are calculated using a prognostic equation based on a surface energy budget, with a three-layer soil model (Leslie et al. 1985). Finally, a crude ice model formulated by Parkinson and Washington (1979) is used to compute ice growth.

Ocean model coupling

Coupling with an ocean model is still in test mode and at this point the only calculations completed have been those in which SST anomalies were provided for the Pacific ocean. Elsewhere the SST field has been specified from the new Reynolds monthly climatologies, which became available in 1994. The Reynolds climatologies have been combined with the analyzed SST anomalies from the Australian Bureau of Meteorology's operational archives to obtain the actual SST field used in the model simulations.

3 Applications and results

The model was applied to a variety of problems in an assessment of its performance as both a NWP model and a GCM. It has already been shown to have a high level of skill as a short range (1 to 2 days) NWP model in a number of studies (Leslie and Purser 1995). As a GCM the model is run at a resolution of about 300 km, and has 19 levels in the vertical. A time step of 30 minutes is used in all the forecasts/simulations.

Medium range forecasting experiments

The model was tested as a medium range forecast model by assessing its performance for the entire year January 1, 1994 to December 31, 1994. The measure of skill adopted for the model was the commonly used anomaly correlation coefficient. The model was run out to 10 days, once per day, using archived global analyses at 00 UTC. On the parallel workstation cluster the 10-day forecast took approximately one hour for each 10 day run, at 1.5 degrees resolution. The relatively small CPU time was due in large measure to the large time step size of 30 minutes. For higher resolution runs, even more efficiency is possible as the 30 min step was chosen as the maximum value that could be "risked" in order not to cause difficulties with the representation of physical processes. At higher resolutions of, say, 0.5° the efficiency gains are greater as the same time step can be maintained in the semi-implicit, semi-Lagrangian framework.

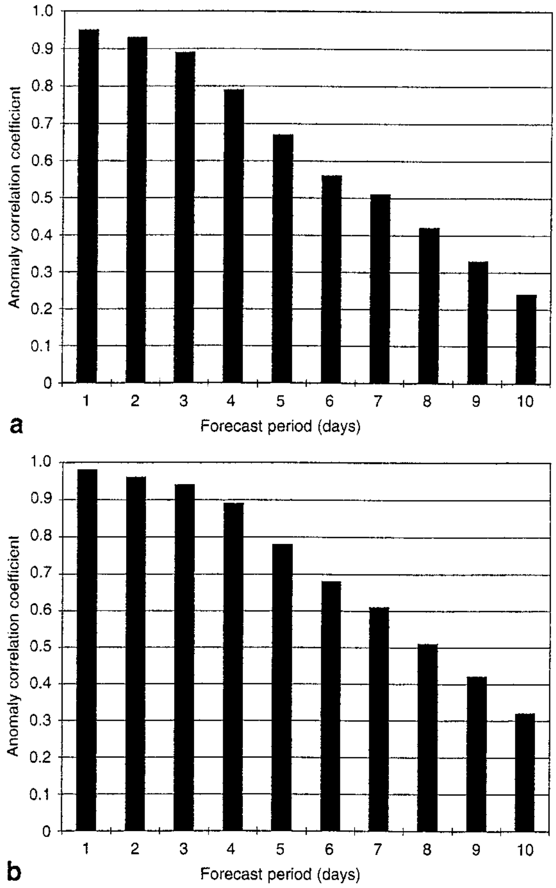

Fig. 1a, b. Anomaly correlation (S. Hem.). **a** Southern Hemisphere anomaly correlation coefficients for days 1 to 10 of the new GCM, averaged over the calendar year January 1, 1994 to December 31, 1994; **b** as in a except for Northern Hemisphere

The anomaly correlation coefficients for the 500 hPa height fields over the Southern Hemisphere for the 10-day forecasts are shown in Fig. 1a and are comparable with those from other modelling centres around the world. The anomaly correlation is above 0.8 at day 3, and does not drop below 0.6 until after day 5. Beyond day 6 there is still skill that extends even out to day 10 in particular situations. However, techniques such as ensemble forecasting procedures (Toth and Kalnay 1993) are required to help determine which forecasts have skill at these longer time periods. Over the Northern Hemisphere the anomaly correlation coefficients (Fig. 1b) remain above 0.6 out to day 7, a full day or so longer than for the Southern Hemisphere.

Brief climatology of the model

This section is included simply to verify that the model basically is performing as might be expected, by computing time averages of a number of key variables and comparing them, in the usual manner, with numerical

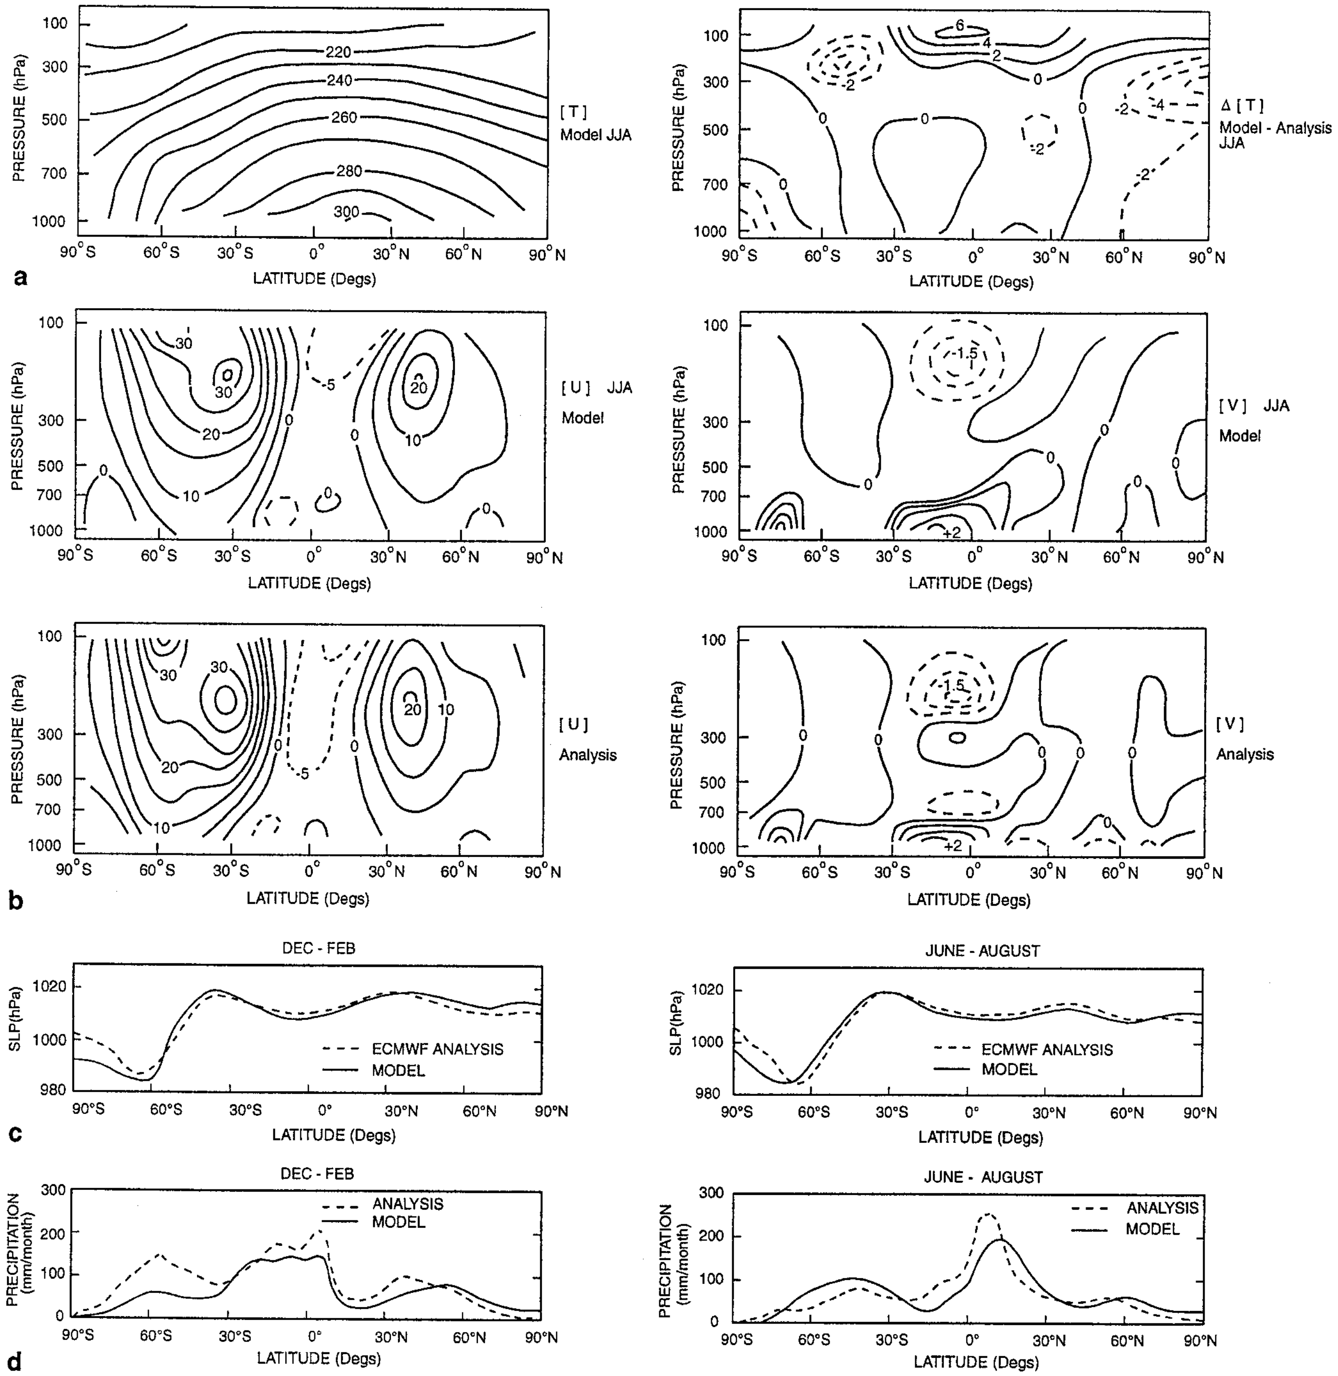

Fig. 2. **a** Analyzed and predicted zonal mean temperature for the 10 year period 1981–1990; **b** as in **a** except for u and v components of wind; **c** as in **a** except for SLP; **d** as in **a** except for precipitation

analyses. The ECMWF analyses have long been taken as the standard analyses of temperature, wind and SLP, and we use them here also. Mean global precipitation was obtained from the widely used climatology developed by Jaeger (1976). The distributions of zonal mean temperature and zonal mean wind components are shown in Fig. 2a, b. The corresponding SLP and precipitation zonal means are given in Fig. 2c, d. Close inspection of the model simulated climatologies reveals that they are very similar in quality to the earlier ECHAM2/ECHAM3 model climatologies as reported by Roeckner et al. (1992). This comparison with an established set of GCM simulations suggests that the new model is performing at a satisfactory skill level as a GCM.

SST anomalies in an El Nino year

The predictability of the El Nino/Southern Oscillation (ENSO) provides a good test of a coupled atmosphere-ocean model on a time scale of months to years. The

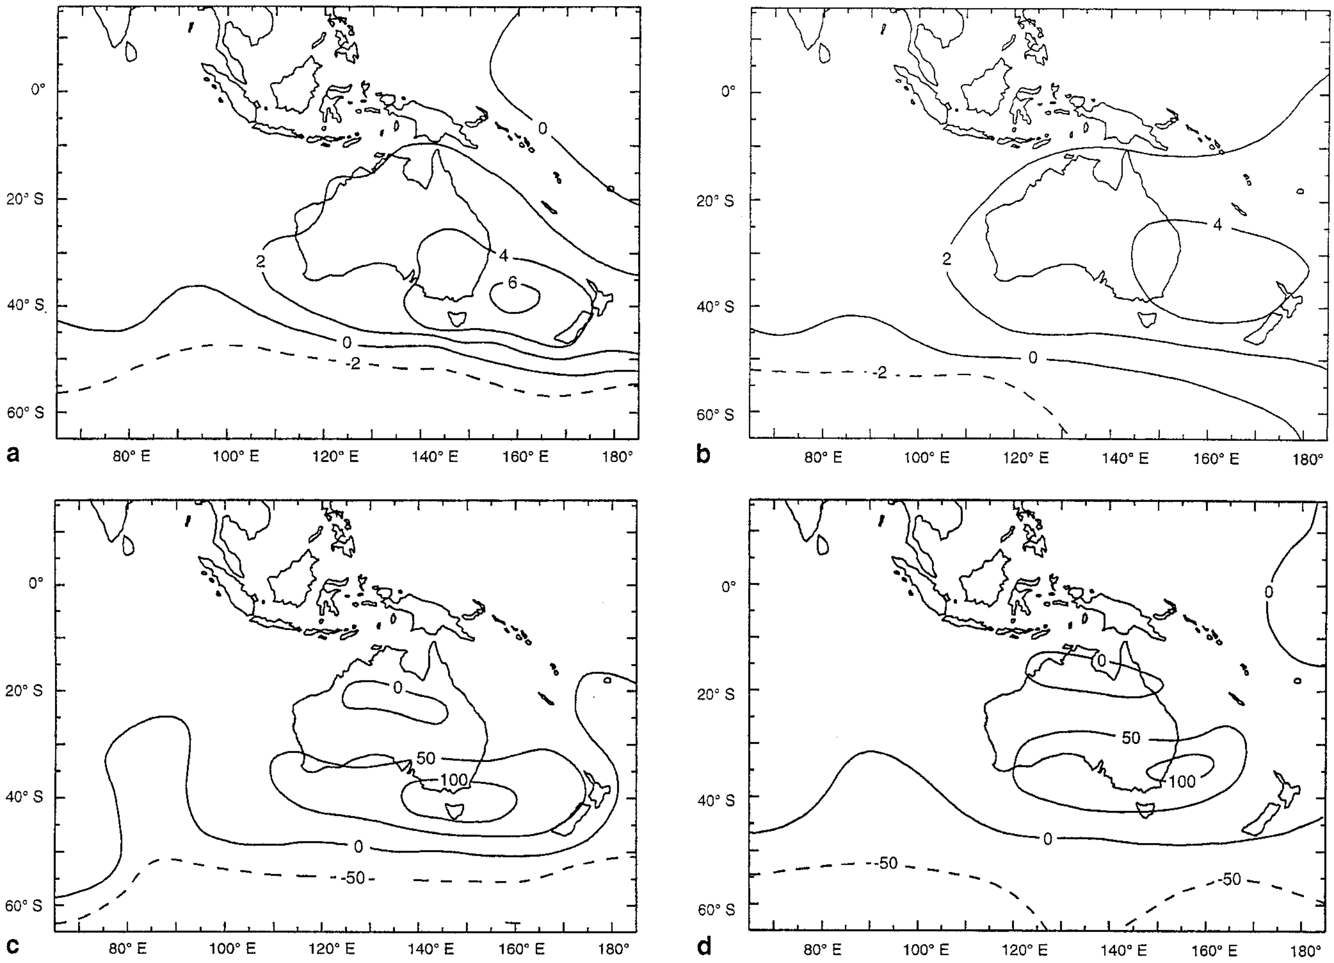

Fig. 3. a Observed SLP anomaly (hPa) for November-December 1994; **b** predicted SLP anomaly (hPa) for November-December 1994; **c** as in **a** except for 500 hPa geopotential height anomaly (metres); **d** as in **b** except for 500 hPa geopotential height anomaly (metres)

interest in ENSO is largely in the capability of coupled ocean-atmospheric models to simulate it, and ultimately to predict its structure and occurrence over seasons or even years ahead. The aim here is much more modest, being determined to a large extent by the fact that the authors are still working on the coupled atmospheric-oceanic model. Therefore, since only the atmospheric response is possible, we used an SST comprising the new Reynolds monthly climatological SST field plus an anomaly SST field over the South Pacific Ocean provided by the Australian Bureau of Meteorology. The simulation therefore was a simple one in that the climatological forcing was provided solely by the SST distribution, so that the experiment was similar to that carried out by Smith (1994). The main differences from the Smith (1994) simulations resulted from the fact that the aims were different so that the simulation period was much shorter, being over the El Nino year January 1 to December 31, 1994, and that analyzed SST anomalies provided by the Australian Bureau of Meteorology were used, rather than just climatological monthly means. The period November-December 1994, which were effectively the 11 and 12 month predictions, was one in which there were large positive pressure anomalies in the SLP and 500 hPa fields over eastern Australia and the adjacent south Pacific Ocean. In this simple example of atmospheric modelling-SST coupling, the predicted and observed anomalies are shown in Fig. 3. Figure 3a, b shows the observed and predicted SLP mean anomalies (hPa) for November-December 1994, and the results are in good agreement given the constraints of the numerical experiment. The areas of positive anomalies are very similar over eastern Australia and the South Pacific, however the intensity is underforecast. Figure 3c, d shows the corresponding observed and simulated 500 hPa height anomalies and the results are again in good agreement.

Regional simulations

An additional bonus, apart from speed, provided by the parallel computer was the ease with which an alternative approach to "down-scaling" could be implemented. Simply defined, down-scaling is the transfer of model information from larger scales (for example, relatively coarse resolution GCM runs) to finer scales, such as specific regions (for example, Australia or Europe). A number of down-scaling schemes have been

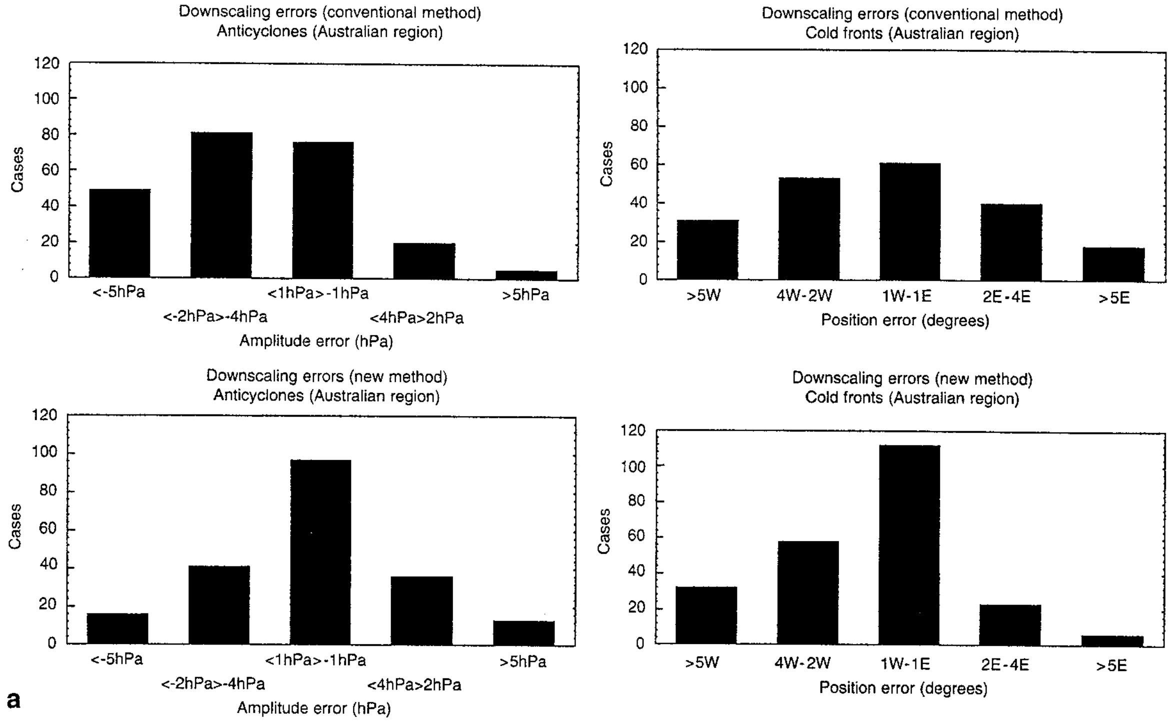

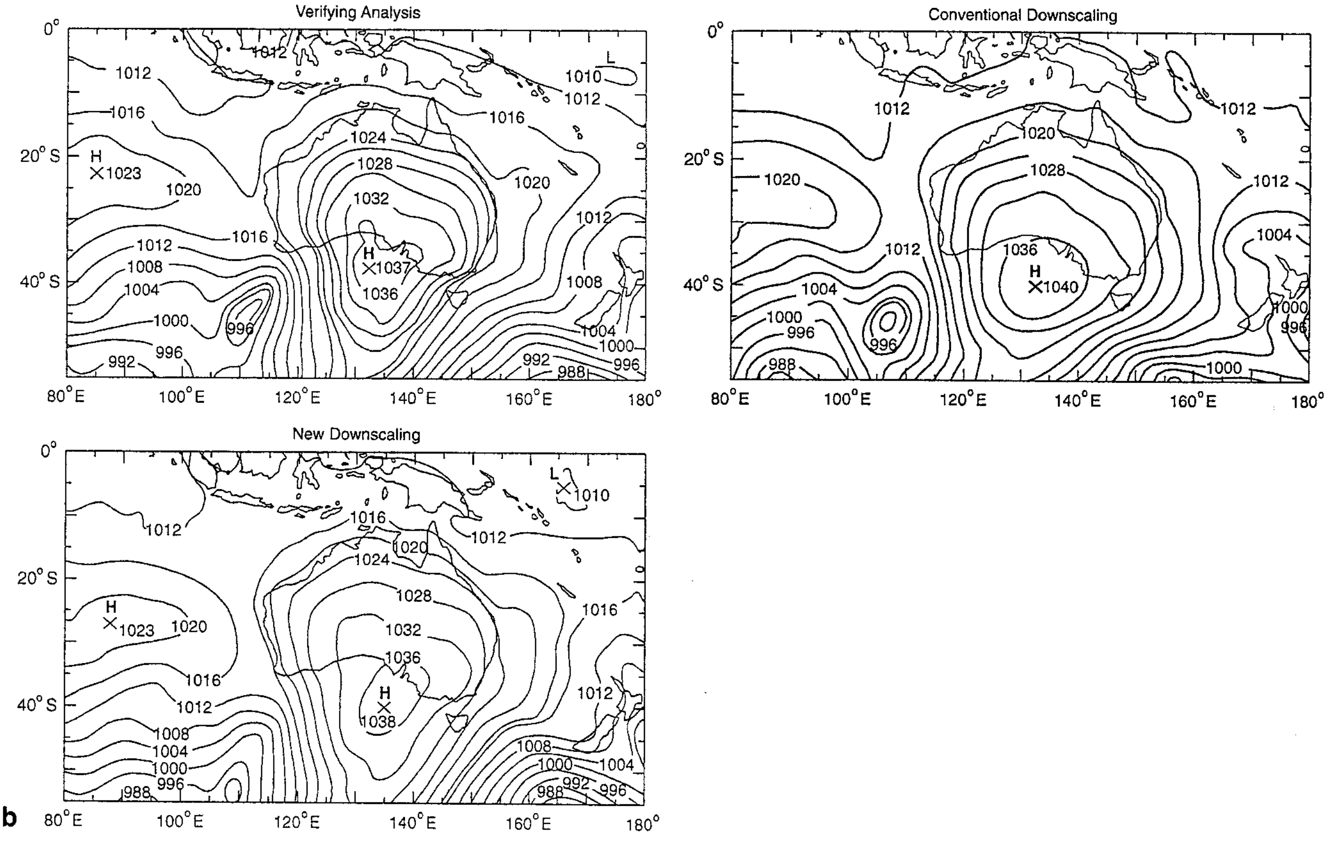

Fig. 4. a Histograms of the location errors in the standard downscaling procedure for Australian region anticyclones and cold fronts (top two panels) and the new procedure (*bottom two panels*); **b** verifying analysis valid at 00 UTC August 5, 1994 (*left top panel*) and the two 72 h forecasts over the Australian region using standard numerical down-scaling (*middle panel*) and the new procedure (*bottom panel*)

devised and these may broadly be divided into two main classes: statistical down-scaling and numerical down-scaling. In statistical down-scaling, a search is made for correlations between the large-scale flow and regional variables such as rainfall patterns. The method can be quite powerful when the correlations are large and statistically significant. In the case of numerical down-scaling, the GCM is run for some pre-determined period and the simulations are then used to drive a much higher resolution regional model. There are numerous examples of numerical down-scaling, such as the work of Giorgi and Mearns (1991). Also, there are two main problems. GCMs are typically run at resolutions of around 500 km, thereby imposing completely artificial lower limits on the resolvable length scales. Such an approach greatly reduces the possibility of physically realistic energy fluxes taking place. Moreover, as the GCM is run first, and then used to drive the regional numerical model, the information exchange is only one-way, hence the literal meaning of the term "down-scaling". As a result, there is considerable debate over whether or not additional information is being provided by the regional model (see, e.g. Anthes 1985). That is, extra detail is not necessarily extra information as the larger scales determine the scales that are resolved by the regional model.

In the alternative procedure employed here, the GCM is run at standard resolutions everywhere except over the region of interest, where a graded mesh is used which is set to be constant over the region of interest and is slowly increased (typically by a factor of 10 per cent per grid point) over an intermediate region surrounding the region of interest until it matches the GCM resolution at some suitably chosen distance from the centre of the region of interest. The grid spacing over the region of interest itself is set at the same value (50 km) as in the conventional procedure. There are three advantages in this approach: two-way interaction is allowed between the GCM and the region of interest; the model needs to be run only once; and the graded mesh eliminates any problem of boundary conditions that normally cause difficulties when two models are involved. Finally, the approach fits in very well with parallel architectures.

A simple test of the procedure has been carried out over Australia. The GCM was run for the 2 year period January 1, 1993 to December 31, 1994, the first year being regarded as a "spin-up" period, and the second year used for an assessment of the performance of the model in positioning of two of the most significant and readily identifiable features of the flow patterns over Southern Australia, namely the subtropical ridge (the quasi-permanent anti-cyclone centres) and the long-wave troughs. The Australian region poleward of about 30°S is dominated by these systems. Statistics were collected for prediction errors in the position of these systems by tabulating the predicted positions of the ridges and troughs forecast by the model and comparing these with the positions of the observed systems. To avoid confusion, it is emphasized that this was not a case of comparing the deterministic forecasts with observations of individual high- and low-pressure cells over the 12 month period, but rather comparing the mean forecast locations of systems that were predicted by the model over the 12 month period with the mean observed positions of the subtropical ridge and the long-wave trough. The conventional down-scaling procedure was used with a regional version of the GCM being run at 50 km resolution and the statistics are shown in the two upper panels of Fig. 4a. The corresponding mean position errors for anticyclones and cold fronts using the new regional climate down-scaling procedure are shown in the two bottom panels of Fig. 4a. It is clear from the naked eye, and verified at greater than the 99% confidence level using a standard t-test that the new approach has improved the prediction of the location of the systems. A case study illustrating the gains obtained from the new procedure over the Australian region is presented in Fig. 4b in which 72 h forecasts verifying at 00 UTC August 5, 1994 are compared using the conventional and new down-scaling procedures. The upper panel of Fig. 4b shows the verifying analysis for the 72 h forecast; the middle panel the regional forecast obtained from the conventional down-scaling; and the lower panel the corresponding forecast from the new two-way down-scaling procedure. Close inspection of the forecasts show that the conventional down-scaling procedure produces obvious phase lags in the systems to the west, southwest and southeast of the Australian continent. There are other gains from the new down-scaling procedure that can be noted by compering the forecasts with the verifying analysis. In summary, it is obvious that the forecasts from the conventional down-scaling procedure lag those from the new procedure, both in the interior of the domain and in the boundary regions.

4 Model implementation in a parallel computing environment

Numerical models of the atmosphere and the oceans have been a major driving force behind the development of increasingly fast scalar and vector computers. This is true not only for mesoscale models, but also for very long-term integrations of coupled atmospheric-oceanic GCMs. Whereas vector supercomputers have dominated the field during the 1980s and early 1990s, parallel computers of various architectures are beginning to emerge as a preferred platform for these applications. The main reason for the shift to parallel computing platforms is that it appears to provide the best chance of achieving the projected requirements of the late 1990s for teraflop performance and for petaflop machines in the early years of the next century.

The challenge for numerical modellers, therefore, is either to adapt existing models or to introduce new models that can exploit fully the emerging parallel machines. The remainder of this section will discuss the implementation of the semi-implicit semi-Lagrangian

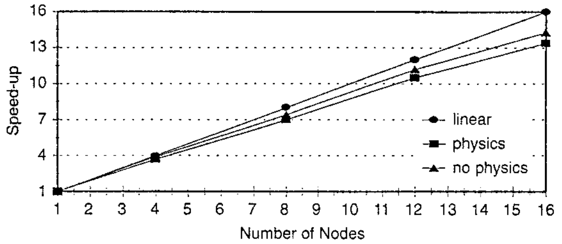

Fig. 5. Computing scalability. IBM SP-2 (16 nodes). Speedup factors obtained on 30 ten-day forecasts. Physics and no-physics performance is compared with linear speedup, for 1, 4, 8, 12 and 16 processors

model on a distributed-memory multiple-instruction multiple data (MIMD) parallel computer (Wightwick and Leslie 1995), with 16 processors connected by a high-performance switch (HPS). Each workstation has a performance rating of about one-quarter of a Cray C-90 processor, based on the results of the NASA Numerical Aerodynamic Simulation (NAS) benchmarks. The parallel architecture and HPS provided the topology, low-latency and high-bandwidth that enabled the parallel computer to achieve high levels of performance as described below.

The performance, or scalability, of the parallel computations were assessed using speedup measured relative to a single processor. The model chosen was run out to 48 h at a range of resolutions, from 2.5° down to 0.5°. The results for 0.5° are shown in Fig. 5. Results for other resolutions are not shown as there was little change in the speedup factors obtained. Results were obtained for one, four, eight, twelve and sixteen processors. Finally, the scalability of the code was evaluated with and without representation of physical processes included. It was found that with four processors, close to linear speedup was achieved with little fall-off using eight processors, with speedup factors of 7. 4 and 7.0 obtained for the " no physics" and "physics" code, respectively. With 12 processors, the performance began to drop away at the full complement of 16 processors the factors were 14.3 and 13.4, respectively. These numbers are very acceptable, given that there is room for further improvement because of the inherent complexity of parallelizing semi-implicit, semi-Lagrangian code.

5 Summary and conclusions

The main purpose of this study was to introduce a new general circulation model (GCM) that is intended to form the basis of a major collaborative research effort between The School of Mathematics, University of New South Wales, Sydney, and The Meteorologisches Institut, Hamburg. The new GCM produced an encouraging set of results in a number of applications. It produced anomaly correlations out to 10 days over the Australian region that were very competitive with existing operational medium range NWP models; it produced an acceptable 10-year model climatology; it quite accurately reproduced SST-driven atmospheric anomalies over eastern Australia; a new down-scaling feature produced a statistically significant improvement in the forecasting of anti-cyclone and cold front location over that predicted by conventional numerical down-scaling; and, finally, the model was ported to a parallel computer with near-linear speedup obtained on a 16 node machine.

Acknowledgements. This study benefited greatly from the help received from the Australian Computing and Communications Institute (ACCI) in Melbourne, and from Glenn Wightwick of IBM in particular. The authors are grateful to Andrew Bennett of Oregon State University, Corvallis and to Russel Morison of The University of New South Wales, Sydney, for valuable suggestions. The paper was in initiated as a result of the Max Planck Prize (Alexander von Humboldt Stiftung).

References

Anthes RA (1985) The general question of predictability. Mesoscale meteorology and forecasting. Ray PS (ed) American Meteorological Society Boston MA, USA

Geleyn J-F (1987) Use of a modified Richardson number for parametrizing the effect of shallow convection. In: Short- and medium-range numerical weather prediction. Matsuno T (ed) Meteorol Soc Japan, Tokyo, 831 pp

Giorgi F, Mearns LO (1991) Approaches to the simulation of regional climate change: a review. Rev Geophys 29:191–216

Gravel S, Staniforth A (1994) A mass-conserving semi-Lagrangian scheme for the shallow-water equations. Mon Weather Rev 122:243–248

Jaeger L (1976) Monatskarten des Niederschlages für die ganze Erde. Ber D Wetterdienstes 139:1–38

Lacis AA, Hansen JE (1974) A parametrization for the absorption of solar radiation in the Earth's atmosphere. J Atmos Sci 31:118–133

Leslie LM, Purser RJ (1991) High-order numerics in a three-dimensional time-split semi-Lagrangian forecast model. Mon Weather Rev 119:2492–2498

Leslie LM, Purser RJ (1995) Three-dimensional mass-conserving semi-Lagrangian scheme employing forward trajectories. Mon Weather Rev 123:2551–2566

Leslie LM, Mills GA, Logan LW, Gauntlett DJ, Kelly GA, Manton MJ, McGregor JL, Sardie J (1985) A high resolution primitive equations NWP model for operations and research. Aust Meteorol Mag 33:11–35

Louis J-F (1979) A parametric model of vertical eddy fluxes in the atmosphere. Bound-Layer Meteor 17:187–202

Palmer TN, Shutts GJ, Swinbank R (1986) Alleviation of a systematic westerly bias in general circulation and numerical weather prediction models through an orographic gravity wave drag parametrization. Q J R Meteorol Soc 112:1001–1040

Parkinson CL, Washington WM (1979) A large-scale numerical model of sea ice. J Geophys Res 84:311–337

Priestley A (1993) A quasi-conservative version of the semi-Lagrangian advection scheme. Mon Weather Rev 121:621–629

Purser RJ, Leslie LM (1996) Generalized Adams-Bashforth time integration schemes. Q J R Meteorol Soc (in press)

Roeckner E, Arpe K, Bengtsson L, Brinkop S, Dümenil L, Esch M, Kirk E, Lunkeit F, Ponater M, Rockel B, Sausen R, Schlese U, Schubert S, Windelbland M (1992) Simulation of the present-day climate with the ECHAM Model: impact of model physics and resolution. Rep 93, Max-Planck-Institut für Meteorologie, Hamburg

Schwarzkopf MD, Fels SD (1985) Improvements to the algorithm for computing carbon dioxide transmissivities and cooling rates. J Geophys Res 90:10541–10550

Schwarzkopf MD, Fels SD (1991) The simplified exchange method revisited: an accurate, rapid method for computation of infrared cooling rates and fluxes. J Geophys Res 96:9075–9096

Slingo JM (1987) The development and verification of a cloud prediction scheme for the ECMWF model. Q J R Meteorol Soc 113:899–927

Smith IN (1994) A GCM simulation of global climate trends: 1950–1988. J Clim 7:732–744

Staniforth A, Côté J (1991) Semi-Lagrangian integration schemes for atmospheric models: a review. Mon Weather Rev 119:2206–2223

Tiedtke M (1989) A comprehensive mass flux scheme for cumulus parametrization in large-scale models. Mon Weather Rev 117:1779–1800

Toth Z, Kalnay E (1993) Ensemble forecasting at NMC: the generation of perturbations. Bull Am Meteorol Soc 74:2317–2330

Wightwick GR, Leslie LM (1996) A new numerical weather prediction model for the IBM SP2 parallel computer. IEEE Trans Comput (in press)