25.5 DEVELOPMENT OF FULLY PARALLELIZED REGIONAL SPECTRAL MODEL AT NCEP

Jongil Han*
RS Information Systems, Inc. at Environmental Modeling Center, NCEP, Washington D.C.

and

Hann-Ming Henry Juang
Environmental Modeling Center, NCEP, Washington D.C.

1. INTRODUCTION

The National Centers for Environmental Prediction (NCEP) regional spectral model (RSM) developed by Juang and Kanamitsu (1994) is a limited area spectral model for regional weather and climate prediction and operationally being used as a member of short-range ensemble forecast system at NCEP. The RSM was designed to have the same model structure, dynamics, and physics as the NCEP global forecast system (GFS) which is a global spectral model for global weather and climate prediction. While in the GFS, total fields are predicted using spherical harmonic functions for spectral computation, in the RSM, perturbation fields from coarse-resolution base fields (usually global analysis or forecast fields) are represented by double Fourier trigonometric functions, which are relaxed to zero at lateral boundaries.

The efforts for improving weather forecast have drawn the model to be run at higher resolution with more sophisticated physics, which requires higher model performance in computation speed. Furthermore, recent development of ensemble forecast system and extensive activities in regional climate simulations using limited area model needed a significant improvement in model performance. Toward this effort, most of up-to-date operational weather forecast models take advantage of massive parallel processor computer to improve their performance. The current parallel version of RSM at NCEP (hereafter old RSM-MPI) is based on 1D decomposition using massage passing interface (MPI) which enables data exchange among distributed memories. However, the performance of the old RSM-MPI falls short of NCEP's current demands for the model to be run at higher resolution (about 30km over North America). This is because the old RSM-MPI has a limitation in use of the number of working processors (limited by model vertical layers) as well as in memory usage related to increased resolution. In this study, we present development of a fully parallelized RSM-MPI code based on 2D decomposition on IBM-SP, a distributed-memory parallel supercomputer.

2. FULL PARALLELIZATION OF RSM WITH MPI

A spectral model essentially requires a transformation from spectral to grid-point space or vice versa using the Fast Fourier Transform (FFT) algorithm. Because the FFT algorithm requires all of grid-point values within the row or column, 3D decomposition for parallelization in a spectral model is almost impossible while 1D or 2D decomposition is possible.

Figure 1 shows the schematic diagram for 1D decomposition used in the old RSM-MPI with an example of 6 processors. First, the regional domain in horizontally spectral space is partitioned into 6 vertical layer bands (upper left cube in Fig. 1). Then, transformation from spectral to grid-point space for each partial horizontal domain is conducted as shown schematically in upper cubes in Fig. 1. Then, the vertical layer bands are transposed into latitude bands using the MPI routines (upper right and lower cubes in Fig. 1) and all nonlinear model dynamics such as nonlinear advection and model physics that entails vertical data dependency are computed in parallel to those latitude bands. The transformation back to the spectral space is conducted in the reverse way denoted as the reverse arrows in the Fig.1. The computation of the linear dynamics such as semi-implicit and time filter is performed in the spectral space (upper left cube in Fig. 1).

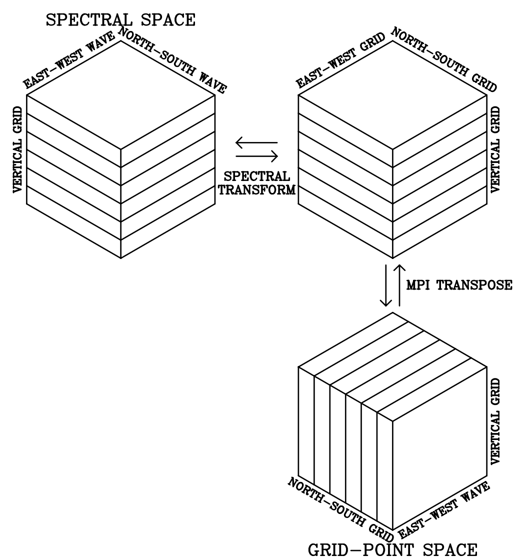

Figure 1. Schematic diagram of 1D decomposition with 6 processors for transformation between spectral to grid-points space using MPI transpose.

The disadvantage of 1D decomposition in Fig. 1 is that the number of working processors is limited by the vertical layers which are usually much smaller than horizontal grid points in NWP models. In the old NCEP-RSM, furthermore, the number of vertical layers assigned in each processor is designed to be the same and thus, it has to be a factor of the number of the whole vertical layers.

The new 2D decomposition which we developed in this study has the advantage to use more processors than the old 1D decomposition because there are much more grid points in 2D to be partitioned. Figure 2 shows the schematic diagram for 2D decomposition used in the new RSM-MPI code with an example of 12 processors. Each

* Corresponding author address: Jongil Han, RSIS/EMC/NCEP, Rm. 207, 5200 Auth Rd., Camp Springs, MD 20746; e-mail: Jongil.Han@noaa.gov

stage is characterized by full dimension in one of three directions. As mentioned earlier, in a spectral model a full dimension in one of two horizontal directions (East-West [X] and North-South [Y] directions in Fig. 2) is needed for the FFT computation. Similar method has been introduced in several previous studies (Oikawa, 2001; Juang et al., 2001; Foster and Worley, 1997; Barros et al, 1995; Skalin and Biorge, 1997). First, rows along X and Y directions in the spectral space are divided depending on the number of working processors while the vertical column has a full dimension (upper left corner cube in Fig. 2). Then, for transformation from spectral to grid-point space in Y direction using FFT algorithm, rows along X direction and vertical columns are divided while Y direction has a full dimension (upper middle and right corner cubes in Fig. 2), for which the transpose (rearrangement of partial and full dimensions as shown in upper left corner and middle cubes in Fig. 2) is performed using the MPI routines. Then, for transformation from spectral to grid-point space in X direction, rows along Y direction and vertical columns are divided while X direction has a full dimension (lower right corner and middle cubes in Fig. 2), again using the MPI routines for the transpose (upper and lower right corner cubes in Fig. 2). After the last transpose (lower middle and left corner cubes in Fig. 2), we have partial dimensions along the horizontal directions but a full dimension along the vertical direction in grid-point space (lower left corner cube in Fig. 2) where nonlinear model dynamics and model physics that entails vertical data dependency are computed. The transformation back to the spectral space is conducted in the reverse way denoted as reverse arrows in the Fig.2. The computation of the semi-implicit time integration is conducted in the spectral space with partial dimensions along the horizontal directions but a full dimension along the vertical direction (upper left corner cube in Fig. 2).

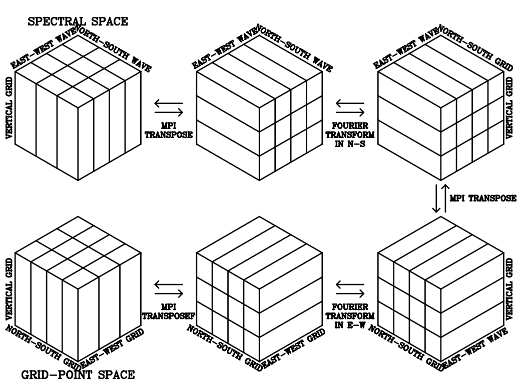

Figure 2. Same as Fig. 1 but for 2D decomposition with 12 processors.

3. PERFORMANCE

In this section, the performance of the model with MPI is shown in terms of speedup. The speedup, $S_p(n)$, is defined as the ratio of the time spent by single processor run and the one by n processor run. For the perfect parallelization, therefore, $S_p(n)=n$. Practically, however, $S_p(n)$ is always less than n because there are not only sequential codes that cannot be parallelized but also there is a time spent for communication. Assuming that the time spent in communication is negligible, the speedup can be expressed as (e.g., see Juang et al., 2003)

$$S_p(n) = \frac{s + p}{s + p / n} \qquad (1)$$

where s is the time spent by the sequential portion of the code in the computation and p the time spent by the parallel portion of the code in computation. Therefore, *Eq. (1)* indicates a maximal speedup achievable theoretically for a percentage of parallelization for given n processors. This theoretical speedup curve is used to evaluate the speedup performance of the new RSM-MPI code.

The speedup performance of the 24-hour RSM runs at NCEP IBM-SP machine with 4, 8, 16, 32, 42, and 64 working processors is shown in Fig. 3 for two different dimensions. Note that the results for the old RSM-MPI with processors more than 42 are not available because the number of working processors in the old RSM-MPI code is limited by the number of the vertical grid points. For the smaller dimension of (X,Y,Z)=(120,121,42) (Fig. 3a), the speedup of the new RSM-MPI is in between those of the theoretically maximal values [i.e., Eq. (1)] achievable from 95% and 99% of parallelization, while the old RSM-MPI shows less speedup than the theoretical one from 95% of parallelization. For the larger dimension of (X,Y,Z)=(240,241,42) (Fig. 3b), the speedup of the new RSM-MPI is about same as the theoretical one achievable from 99% of parallelization. Compared to the old RSM-MPI, overall speedup performance of the new RSM-MPI is about twice larger.

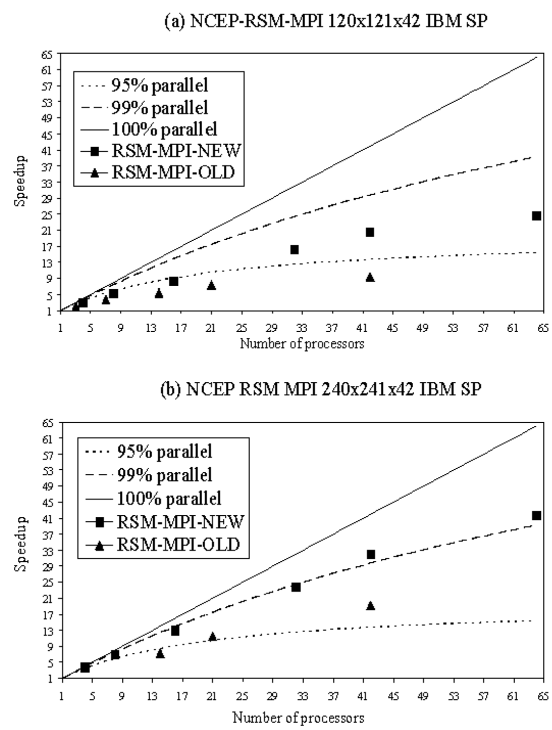

Figure 3. Performance of parallel RSM.

In the MPI code, the memory usage generally decreases with increasing working processors because each processor has smaller dimension with increasing working processors. Note that for the larger dimension (Fig. 3b), the runs of the old RSM-MPI with 4 and 8 working processors were failed because of too much large dimension associated with smaller number of working processors. However, the runs with the new RSM-MPI were successful even for much larger dimension. This is because for the same working processors, each processor has smaller dimension in 2D decomposition than in 1D decomposition. Although not shown in a figure, our tests showed that for a limited computer resource, the new RSM-MPI could use about factor of 5 larger dimensions than the old RSM-MPI, indicating that the memory usage is much smaller in the new RSM-MPI than in the old RSM-MPI.

4. SUMMARY

A new fully parallelized RSM using MPI based on 2D decomposition has been developed, which shows significant improvement in speedup and memory usage compared to the previous RSM-MPI that is based on 1D decomposition. While in the old RSM-MPI the number of working processors is limited by vertical layers, in the new RSM-MPI it is limited by horizontal dimension, which is usually much larger than vertical dimension in numerical weather models. Compared to the old RSM-MPI, the new RSM-MPI is about twice faster and can use about factor of 5 larger dimensions. To improve model performance further, parallelization in model input/output is being developed.

5. REFERENCES

Barros, S.R.M., D. Dent, L. Isaksen, G. Robinson, G. Mozdynski, and F. Wollenweber, 1995: The IFS model: A parallel production weather code. *Parallel Computing*, **21**, 1621-1638.

Foster, I., and P.H. Worley, 1997: Parallel algorithms for the spectral transform method. *SIAM J. Sci. Comput.,* **18(2)**, .806-837.

Juang, H-M. H., and M. Kanamitsu, 1994: The NMC nested regional spectral model. *Mon. Wea. Rev.,* **122**, 3-26.

Juang H.-M. H., and M. Kanamitsu, 2001: The Computational Performance of the NCEP Seaonal Forecast Model on Fujitsu VPP5000, Ninth ECMWF Workshop on the use of high performance computing in meteorology, 13-17 November, 2000, Published by World Scientific under the title: Developments in teracomputing. Ed. by W. Zwieflhofer and N Kreitz. Singapore, ISBN 981-02-4761-3, 384pp.

Juang, H-M. H., C.-H. Shiao, and M.-D. Cheng, 2003: The Taiwan Central Weather Bureau regional spectral model for seasonal prediction: multiparallel implementation and preliminary results. *Mon. Wea. Rev.,* **131**, 1832-1847.

Oikawa, Y., 2001: Performance of parallelized forecast and analysis models at JMA. Ninth ECMWF Workshop on the use of high performance computing in meteorology, 13-17 November, 2000, Published by World Scientific under the title: Developments in teracomputing. Ed. by W. Zwieflhofer and N Kreitz. Singapore, ISBN 981-02-4761-3, 384pp.

Skalin, R., and D. Bjorge, 1997: Implementation and performance of a parallel version of the HIRLAM limited area atmospheric model. *Parallel Computing*, **23**, 2161-2172.