International Journal of Geographical Information Science, 2013
Vol. 27, No. 4, 765–784, http://dx.doi.org/10.1080/13658816.2012.715650

Using adaptively coupled models and high-performance computing for enabling the computability of dust storm forecasting

Qunying Huang[a], Chaowei Yang[a]*, Karl Benedict[b], Abdelmounaam Rezgui[a], Jibo Xie[c], Jizhe Xia[a] and Songqing Chen[d]

[a]*Center of Intelligent Spatial Computing for Water/Energy Sciences and Department of Geography, GeoInformation Sciences, George Mason University, Fairfax, VA, USA;* [b]*Earth Data Analysis Center, University of New Mexico, Albuquerque, NM, USA;* [c]*Center for Earth Observation and Digital Earth, Chinese Academy of Sciences, Beijing, China;* [d]*Department of Computer Science, George Mason University, Fairfax, VA, USA*

(*Received 11 September 2011; final version received 9 July 2012*)

Forecasting dust storms for large geographical areas with high resolution poses great challenges for scientific and computational research. Limitations of computing power and the scalability of parallel systems preclude an immediate solution to such challenges. This article reports our research on using adaptively coupled models to resolve the computational challenges and enable the computability of dust storm forecasting by dividing the large geographical domain into multiple subdomains based on spatiotemporal distributions of the dust storm. A dust storm model (Eta-8bin) performs a quick forecasting with low resolution (22 km) to identify potential hotspots with high dust concentration. A finer model, non-hydrostatic mesoscale model (NMM-dust) performs high-resolution (3 km) forecasting over the much smaller hotspots in parallel to reduce computational requirements and computing time. We also adopted spatiotemporal principles among computing resources and subdomains to optimize parallel systems and improve the performance of high-resolution NMM-dust model. This research enabled the computability of high-resolution, large-area dust storm forecasting using the adaptively coupled execution of the two models Eta-8bin and NMM-dust.

Keywords: parallel computing; Cyber GIS; atmospheric modelling; nested models; computing intensity; applied sciences; geospatial platform; spatiotemporal thinking and computing

1. Introduction

Dust storms are the result of strong turbulent wind systems entraining particles of dust into the air, reducing visibility down from miles to several meters (Goudie and Middleton 1992). Global climate change has driven up the frequency and intensity of dust storms in the past decades with negative consequences on the environment, human health, and assets. For example, dust storms (1) contain marine nutrients, such as active iron and phosphorus, which can result in algal blooms over the ocean surface when decomposing into the ocean water (Dulac *et al*. 1996); (2) act as a pollutant which reduces air quality and affects public health by causing allergies, respiratory diseases, and eye infections (Nickling and Gillies

*Corresponding author. Email: cyang3@gmu.edu

1993); and (3) impact both the environment and climate from regional to global scale (Sokolik and Toon 1996, Gong *et al*. 2003) by causing the cooling of oceans through reflecting solar radiation back to space (Gong *et al*. 2003) and contributing to global aerosol mass load and optical thickness (Gong *et al*. 2003), where a dust-laden atmosphere with an average optical thickness of 0.5 would cause a net radiative forcing of $+20$ to $+40$ W/m^2 over arid regions and -5 to -15 W/m^2 over the ocean (Sokolik and Toon 1996).

The severe impacts of dust storms on our environment have motivated scientists to develop dust models for (1) predicting dust storms; (2) understanding dust processes; (3) quantifying the global dust cycle; and (4) re-constructing past climates (Shao and Dong 2006). Since the late 1980s, several research groups have developed dust models that can correctly predict spatiotemporal patterns, evolution, and order of magnitude of dust concentration, emissions, and deposition (e.g. Westphal *et al*. 1988, Gong *et al*. 2003, Shao *et al*. 2007, Han *et al*. 2004).

However, utilizing those existing models to predict high-resolution dust storms poses several critical challenges:

(1) Simulating dust storm phenomena is very complex and computing intensive (Xie *et al*. 2010). Such a periodic phenomenon simulation requires the iteration of computing-intensive numerical equations for many times. For a given domain size, the computing cost of an atmospheric model is a function of n^4, where n is the grid dimension, including two horizontal dimensions, one vertical dimension, and one time dimension (Baillie *et al*. 1997). Therefore, doubling the geographic domain on the horizontal direction would result in a fourfold increase in the computing cost. Halving the spatial resolution only could result in an eightfold increase in the computing cost because it would also require halving the time-step to keep the model accuracy (Baillie *et al*. 1997).

(2) Dust storm forecasting is a time-critical task that requires a limited computing time. For example, a 2-hour computing limit is recommended for 1-day forecasting to make the results useful (Drake and Foster 1995). Limited geographic domain and/or resolution forecasting is usually performed to complete the simulations within the time limit (Wolters *et al*. 1995). However, a postcode-level resolution is needed for dust storm forecasting to support decision-making, such as preparing medications in public health (Yang *et al*. 2008). Figure 1 shows the computing time required for a 24-hour forecasting over different domain sizes on the horizontal directions and with the same vertical layers (37 layers) and spatial resolution (3 km). More than 7.5 hours are needed to forecast a 10×10 degree domain size using cutting-edge hardware configuration (Yang *et al*. 2011a). Based on the computing cost of an atmospheric model and computing requirement trends (Figure 1), forecasting the whole southwest United States with a domain size of 37×20 degree would require about 60 hours. Such a computing performance is not acceptable because we would be forecasting yesterday.

The limitation of computing power precludes a direct solution to complete the southwest United States high-resolution forecasting within reasonable time constraints (Huang 2011), and more efficient computing strategies should be explored (Kuligowski and Barros 1999). One solution is to embed a finer-resolution subdomain or succession of subdomains within the entire domain with lower resolution to enhance the resolution over specific areas of interest while moderating the required computing cost (Anthes

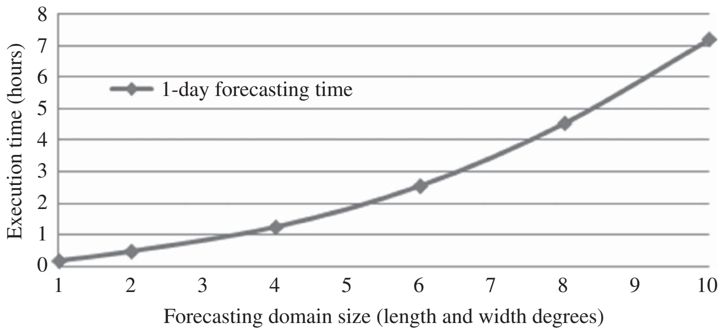

Figure 1. Execution time for different geographic domain forecasts.

1983, Kuligowski and Barros 1999). Such an approach is called nested models or, more commonly, high-resolution limited area models.

Generally, nested models require *a priori* knowledge of where to place the high-resolution subdomains inside the entire domain (Constantinescu *et al*. 2008). For a real-time dust storm forecasting system, users are not aware of where the dust storm will occur, and therefore traditional tightly coupled nesting approach is not suitable. In addition, finer-scale module and coarse-scale module are tightly coupled to execute concurrently and extensive efforts are required to enable the two modules to be coupled together (Michalakes *et al*. 1998). We propose a loosely adaptively coupled solution to overcome this issue. In the adaptive loosely coupled approach, a coarse dust storm model (Eta-8bin) is first executed with low spatial resolution (22 km) to identify potential dust event areas. Afterwards, a finer model, NMM-dust, will run with higher resolution (3 km) for those potential areas on different groups of computing resources with each group performing one subdomain forecasting. This way, simulation of each subdomain with much smaller size can be completed with much less computing cost and without too much extra efforts to modify the two models.

In this article, NMM-dust model is used to produce high-resolution results for weather forecasting executed in parallel. High-performance computing (HPC, Armstrong *et al*. 2005) or grid computing (Yang *et al*. 2008) have been used to address the increasing computing need in geographic science problems and enable a lot of research and geospatial applications, such as dust storm simulation (Xie *et al*. 2010) and digital elevation model (DEM) interpolation (Huang and Yang 2010). However, in parallelized HPC or grid computing environments, communication (Drake and Foster 1995, Sterling *et al*. 1995), synchronization, and load balance (Yang *et al*. 2011a) are potential bottlenecks in the overall performance.

One possible direction for reducing communication overhead of HPC systems is to optimize the configuration and scheduling of the parallel system based on the spatiotemporal relations and principles among the computing resources and phenomena (Huang and Yang 2010). Yang *et al*. (2011a) provided a general guide about how to discover the spatiotemporal principles with several examples to improve the performance of dust storm simulation. This article builds upon but significantly enhances that research to better explore and improve HPC by using spatiotemporal patterns and constraints.

Section 2 reviews related research on loosely adaptively coupled models and HPC support for dust storm forecasting. Section 3 introduces the methodology, including adaptive loosely coupled models and HPC performance improvement strategies based on spatiotemporal principles. Section 4 presents and analyses the experiments of loosely coupled nested models, and scheduling strategies of computing resources (CPU and storage) to demonstrate how to improve the performance based on these strategies. Finally, in Section 5, we demonstrate that our approach can provide a potential solution to complete the computation of dust storm forecasting for large areas and high resolution within acceptable time frame.

2. Related work

2.1. Nested models

Nested models are able to properly simulate and predict spatial features and resolve processes with small scales on subdomains of a big domain with higher resolution. Therefore, to produce the regional high-resolution simulation results without the enormous computing cost of a global model at the high resolution, nested models are used for a variety of research and operational applications (Ramón *et al*. 2002). For example, nested models are greatly used in the research on chemical transport and decomposing process over the atmosphere (Pleim *et al*. 1991), regional air pollution (Constantinescu *et al*. 2008), regional and global weather and climate pattern analyses (Giorgi *et al*. 1996, McGregor 1997, Fennessy and Shukla 2000, Ramón *et al*. 2002, Jasper *et al*. 2002).

However, most models are developed and nested in a 'static tightly coupled' approach (Sela 1980, Giorgi *et al*. 1996, McGregor 1997, Fennessy and Shukla 2000, Nellson *et al*. 2005). Within the static tightly coupled approach, a high-resolution module could be one way, two way, or even triple way nested with a low-resolution module and both are executed at the same time (Fennessy and Shukla 2000). For example, a subdomain, high-resolution atmospheric module (with 80 km resolution) was nested in the Center for Ocean-Land-Atmosphere Studies (COLA) global general circulations module (GCM; Sela 1980) (with 1.88×2.88 degree spatial resolution) for seasonal climate prediction over North America (Fennessy and Shukla 2000). Extensive modifications of both modules are required to make them ready for nesting because of inconsistent subroutine interfaces, definition of physical constants, data structures, and other problems (Michalakes *et al*. 1998). In addition, nesting of finer subdomain into coarser domain requires *a priori* knowledge of where inside the modelling domain to place the finer subdomain (Constantinescu *et al*. 2008). For instance, at Norwegian Meteorological Institute (DNMI), the fifth-generation Mesoscale Model (MM5) is nested with High Resolution Limited Area Model (HIRLAM). A subdomain with 3 km resolution was statically set up for the Oslo region for nesting with MM5 (Baklanov *et al*. 2002).

For a real-time dust-storm prediction system, users are not aware of where the dust storm will occur. Therefore, the traditionally static tightly coupled nested approach is not suitable for the sporadic event forecasting, and adaptively adjusting the high-resolution subdomains should be required for dust storm simulation. Some studies have been conducted to perform adaptive atmospheric modelling where different subdomains are able to be computed with different spatial resolutions dynamically during the simulations (Hart *et al*. 1998, Borthwick *et al*. 2000, Constantinescu *et al*. 2008). Those adaptive approaches use refinement criteria to adaptively control the high-resolution subdomain placement. However, the success of those approaches depends on the design of the refinement criteria which can highly influence the model's accuracy (Constantinescu *et al*. 2008). Also, there is an overhead associated with the management of the non-uniform subdomains by interpolating the solution between different resolution levels (Constantinescu *et al*. 2008).

This research will investigate how to integrate multi-resolution models to tackle the computing challenges of complex problems in an 'adaptive loosely coupled nested' manner. Within this approach, the high-resolution model will adaptively perform the forecasting over the areas that have potential dust events and are identified by the coarse model results. This approach can resolve the computing demands of large scale problems efficiently without or with only slight modification to the original models. In addition, this approach would greatly promote the research of communication and integration of different models, be easily applied to other research, e.g. air quality, and contribute to flexible and extensible global modelling framework.

2.2. High-performance computing for dust storm forecasting

2.2.1. Parallelization of dust storm models

Dust storm models are developed by adding dust solvers into the regional atmospheric models (Nickovic *et al*. 1997, Shao *et al*. 2007). Dust storm models parallelization is to parallelize the kernel atmospheric modules. The atmosphere is modelled by dividing the study area into three-dimensional spatial cells and atmospheric modelling is the process to solve a system of coupled non-linear partial differential equations on each cell with appropriate boundary conditions (Purohit *et al*. 1999). The calculations of the equations on each cell are repeated with a time-step to model phenomena evolution. Therefore, the computing cost of an atmospheric model is a function of the number of cells in the domain and the number of time-steps (Baillie *et al*. 1997).

Parallel computing architectures are greatly used as an instrumental mechanism for the execution of atmospheric phenomena simulation (Jin *et al*. 2003), such as Eta Model (Henderson *et al*. 1994), Rapid Update Cycle (RUC) (Rodriguez *et al*. 1995), Quasi Non-Hydrostatic meteorological model (QNH) (Baillie *et al*. 1995), HIRLAM (Wolters *et al*. 1995), MM5 (Davis *et al*. 1999), Advanced Regional Prediction System (ARPS) (Xue *et al*. 2003), and Regional Atmospheric Modeling System (RAMS) (Cotton *et al*. 2003). HPC and computing advancements made by faster Central Processing Units (CPUs) allow modern Numerical Weather Prediction (NWP) models to reach a very high resolution (e.g. Davis *et al*. 1999). For example, MM5 has been used for real-time weather prediction at a resolution of 1 km by the US Army Test and Evaluation Command (Davis *et al*. 1999).

An atmospheric model performs essentially the same set of computations in each simulation cell within the domain. Normally, a Single Program Multiple Data stream (SPMD) domain decomposition approach and nearest neighbour communication in the physical subdomains is required (Nanjundiah 1998). Data dependencies between neighbouring cells in the vertical direction are much larger than that in the horizontal directions. Horizontal decomposition in the model is usually used to minimize communication overhead (Wolters *et al*. 1995). The communication includes halo exchanges, periodic boundary updates, parallel transposes, and others. The halo region is the part of the local memory allocated around a cell for exchanging information with neighbouring cells using message passing. The process of data decomposition should define the halo regions assigned to each processor and also define a virtual array of processors used to execute these subdomains and create neighbourhood relations between subdomains.

2.2.2. Performance improvement

Performance improvement of parallel systems has been studied by improving the data structure, algorithm, libraries for parallelization, and compiler for code compilation

(Rodriguez *et al*. 1995, 1996). Rodriguez *et al*. (1995) studied and compared the performance issues in the parallelization of weather prediction models using two different parallelization libraries. He also discussed optimization strategies for parallel systems by using redundant computations to minimize the data exchanges (Rodriguez *et al*. 1996). Baer and Zhang (1998) proposed to reconstruct the prediction equations in a format that will allow a longer time-step without loss of accuracy.

In addition to modifying, constraining, and reengineering the application architecture and algorithms, an HPC supporting geospatial sciences should leverage spatiotemporal principles and constraints to better optimize and utilize HPC in a spatiotemporal fashion (Calstroka and Waston 2010, Yang *et al*. 2011a). Earlier investigations found that not only can HPC be used to support geospatial sciences, but it can also be optimized with spatiotemporal principles to best utilize available distributed computing resources (Yang *et al*. 2011b). Yang *et al*. (2011a) demonstrated that spatiotemporal principles can be used to better parallelize models, arrange the computing resources based on the spatiotemporal scales and spatial resolution, and select the computing resources based on the network connection and topology. In this article, more spatiotemporal patterns and examples are explored and employed to improve the HPC performance for enabling the computability of dust storm forecasting.

3. Methodology

3.1. Adaptive loosely coupled strategy

Several numerical dust models have been proposed and developed. The Dust Regional Atmospheric Model (DREAM) (Nickovic *et al*. 2001), designed to simulate dust entrainment and transport on a regional scale, is widely used for dust cycle modelling systems. DREAM can be easily configured and incorporated into other atmospheric models. For instance, it has been successfully coupled with National Centers for Environmental Prediction (NCEP)/Eta as both the Eta-4bin (four particle-size classes) and the Eta-8bin (eight particle-size classes) dust forecast models to simulate the dust cycle in the atmosphere. The performance of the system has been tested for a variety of dust storm episodes in a variety of locations and resolutions.

The Eta-8bin model has shown considerable potential in forecasting severe storms. However, the Eta-8bin model runs in sequence rather than in parallel (Xie *et al*. 2010), and the model has a coarse spatial resolution of 1/3 degree that cannot be used for many applications. With current required horizontal resolutions, for example, 3 km, the Eta model used with NWP would go beyond its accuracy limit. The Eta model was replaced in the US National Weather Service (NWS) operations by a non-hydrostatic mesoscale model (NMM), which has a higher resolution and greater computational efficiency (Janjic 2003). The coupling of the DREAM dust forecasting algorithm and the NMM meteorological module (NMM-dust) forms a much higher resolution model, which enables an increased spatial resolution to the postcode level, or about 3 km $\times$ 3 km. NMM-dust model can produce higher-resolution results and is executable in parallel mode on distributed systems. Parallel processing is supported through the message passing interface (MPI) programming model.

It would be ideal if the NMM-dust model could be run for a large forecasting domain (i.e. the southwest of the continental United States, 37 $\times$ 20 degree in total). However, the NMM-dust model itself and computing capacity cannot support such a large domain size running at a 3 km resolution due to high computing and memory consumption (Yang *et al*. 2011a). To support the runs, we should either (1) redesign the existing algorithms, codes, and data structures or (2) increase the speed of the CPU and network connection (Yang *et al*. 2011a). Even though the model, after re-engineering the code, can support such a large domain, the forecasting cannot successfully complete within reasonable time for 1-day forecasting (Figure 1).

Therefore, it is not feasible to run the high-resolution NMM-dust model for the entire southwest United States. Instead, the coarse-resolution Eta-8bin model could perform a quick forecasting for a large domain, whereas the NMM-dust model could run at a higher resolution for subdomains with potential dust storm events. Specifically, the adaptive loosely coupled strategy would (1) first run the low-resolution model, Eta-8bin; (2) identify subdomains of high predicted dust concentrations by analysing the results of Eta-8bin model; and (3) run the higher-resolution NMM-dust model for only those subdomains with much smaller area in parallel by assimilating the output of the Eta-8bin model. In this approach, high-resolution model results for specific subdomains of interest could be obtained more rapidly than an execution of a high-resolution model over the entire domain.

The Eta-8bin model output can be automatically taken by the NMM-dust models as an input to reduce the communication time between the two different models. Modifications to the Eta-8bin model pre- and post-processing components are required to enable this model interoperability. Standard input and output file formats are implemented and utilized for both the Eta-8bin and the NMM-dust models so that the outputs of one could be used as the input for another. Specifically, we decided to use the Network Common Data Form (NetCDF) as the file format for the outputs of both models, allowing standard meteorological data processing tools to access and process these products. Both models, after modification of pre- and post-processing components, are able to read NetCDF files for model initialization. The use of a common, well-supported data format for both model initialization and output significantly improves the communication and streamlines the process of developing multi-model workflows. In this way, the output of one model can be used to initialize another, either for a model run for a subsequent time-step, or for the execution of a higher-resolution model for the same time period when the low-resolution model has already been run.

3.2. Performance improvement through spatiotemporal patterns

3.2.1. Parallelization of NMM-dust model

Similar to other atmospheric model parallelization, an SPMD data decomposition approach is used for parallelization by decomposing the domain into multiple subdomains and distributing the computing load of each subdomain onto one CPU core as a process. Figure 2 shows parallelizing a (4.5 × 7.1 degree) domain with spatial resolution 0.02083 degree to 24 subdomains (4 × 6 decomposition) for one vertical layer. Based on the boundary size and spatial resolution, the grid cells of one vertical layer would be 215 × 345. The decomposition algorithm will try to make all subdomains the same dimension but the border subdomains on the last column and row will be allocated with the remaining grid cells. Therefore, there would be generally 54 × 57 grid cells for each subdomain except for subdomains on the border. The process processing the subdomains will need to communicate with their neighbouring processes for local computation and synchronization. The process responsible for processing the internal subdomains, such as subdomain 5 in blue, will require communication among four neighbour processes. The subdomains at the edge of the domain, such as subdomain 4 in green, will exchange data among three neighbours while the rest, such as subdomain 0 in red, will only require data exchange

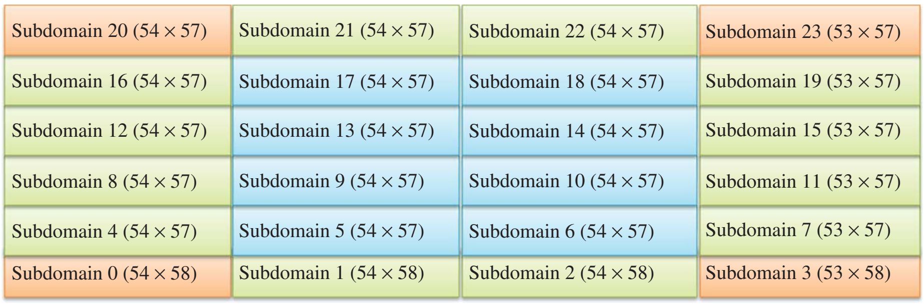

Figure 2. Parallelizing a 4.56 × 7.12 degree domain to 24 subdomains (4 × 6) with 0.02083 degree spatial resolution (about 3 km, 215 × 343 grid cells in total).

among two neighbours. During the computation, the intermediate data results of each subdomain are produced in the local memory of a computing node. Other computing nodes need to access the data through file transfer across the computer network. The cost of data transfer due to communication among neighbour subdomains is a key efficiency issue because it adds significant overhead (Baillie *et al*. 1997). Figure 3 and Table 1 illustrate all the NMM-dust kernel subroutines in computing sequence for each subdomain.

3.2.2. Computability strategies

The spatiotemporal principles enlighten the direction of improving HPC performance (Yang *et al*. 2011a, 2011b) by (1) parallelizing the model and (2) selecting, arranging, and scheduling the computing resources. In this article, more strategies are explored to improve the performance of HPC by considering the spatiotemporal patterns of computing resources.

3.2.2.1. Subdomain and computing node mapping. Different subdomain and computing node mapping methods result in different communication overheads. Figures 4a and b show two mapping methods for dispatching 12 subdomains to two computing nodes A

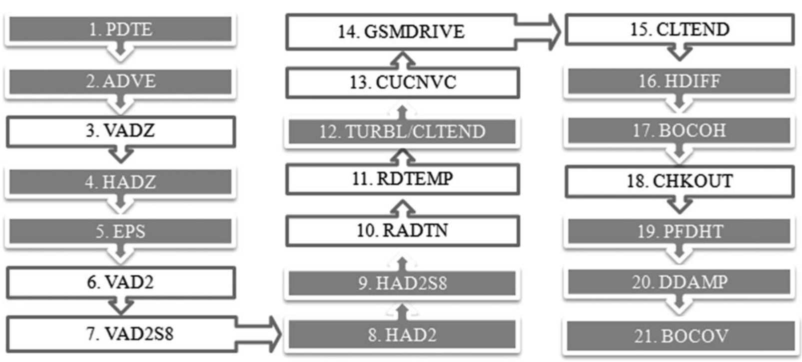

Figure 3. Computing subroutines and communication and synchronization for NMM-dust model (the subroutines in grey require communication and synchronization).

Table 1. The subroutines of dust storm model.

Routine number	Routine name	Communication	Description
1	PDTE	Yes (communication and synchronization of the hydrostatic surface pressure PD (Pa))	Integrate mass flux divergence, compute vertical velocity, and update the pressure field
2	ADVE	Yes (The communication among U, V, and geopotential height z is required for the subroutine ADVE)	Horizontal advection of the variables of temperature T(K), u wind components U (m/s), v wind components V (m/s), and coriolis effect and curvature terms are applied
3	VADZ	No (happens in the vertical layers.)	The process of vertical advection of geopotentional height
4	HADZ	Yes (the synchronization of W is required before the subroutine)	Used for the horizontal advection of height and the vertical wind speed W (dz/dt) is updated
5	EPS	Yes	Used for both vertical and horizontal advection of dz/dt and the vertical wave treatment is added in the subroutine
6	VAD2	No	Vertical advection of the variables of specific humidity q, total cloud water condensate (CWM), turbulent kinetic energy (m/s)
7	VAD2S8	No	Eight classes of dust particle load ($s8$)
8	HAD2	Yes	The horizontal advection of the variables of turbulent kinetic energy Q, total cloud water condensate CWM(kg/kg), $2\times$ turbulent kinetic energy $Q2$ (m^2/s^2)
9	HAD2S8	Yes	Eight classes of dust particle load ($s8$), local halo data of q, CWM, $Q2$, and $s8$ are communicated with neighbouring processes
10	RADTN	No	For radiation
11	RDTEMP	No	For applying temperature tendency due to radiation
12	TURBL/CLTEND	Yes (the communication of pd, T, q, CWM, s8, and the vertical component of the initial velocityUZ0 is conducted)	Used to perform vertical turbulence and store original temperature array (CLTEND)

(*Continued*)

Table 1. (*Continued*).

Routine number	Routine name	Communication	Description
13	CUCNVC	No	Used for convective precipitation
14	GSMDRIVE/ CLTEND	No	To grid scale microphysics and then store the original temperature array (CLTEND)
15	CLTEND	No	To update temperature tendency due to cloud processes
16	HDIFF	Yes(the communication of T, q, U, V, Q2, and S8 needs to be exchanged)	Used for horizontal diffusion
17	BOCOH	Yes(the exchange of q, CWM, Q2, S8, pd, and T are required)	To update boundary conditions for subdomains
18	CHKOUT	No	Every three hours in this case, the post profile data subroutine CHKOUT will be performed
19	PFDHT	Yes (the variables pd, T, U, V, q, CWM (kg/kg), dw/dt, and non-hydrostatic pressure PINT (Pa) are exchanged)	The subroutine to calculate pressure gradient force (PGF), update winds due to PGF, and compute divergence
20	DDAMP	Yes (the variable div would be exchanged)	Used for divergence damping
21	BOCOV	Yes (the communication of U and V are required)	Used to update boundary conditions at the wind points

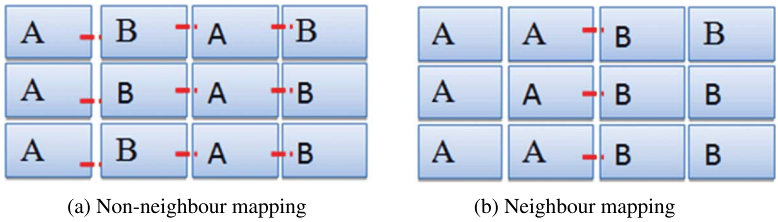

Figure 4. Two mapping methods for dispatching 12 subdomains to two computing nodes A and B.

and B. Figure 4b requires only three grid cell communications over two different computing nodes while Figure 4a requires nine adjacent boundaries to exchange data over two computing nodes. Obviously, the method in Figure 4b can reduce the communication overhead by making more communication occur within the same node rather than over computer networks.

By default, the middleware Message Passing Interface Chameleon version 2 (MPICH2) will dispatch the subdomains to the computing nodes sequentially. For instance, if we have two computing nodes and six subdomains, then the first, third, and fifth subdomains will be dispatched to the first computing node and the second, fourth, and sixth subdomains will be dispatched to the second computing node. Therefore, MPICH2 uses the typical non-neighbour mapping method. In the subdomain and computing nodes experiment, two computing nodes are utilized and half continuous subdomains are dispatched on the first computing node and the rest are dispatched on the other computing node, in a neighbour mapping fashion.

3.2.2.2. Storage. During the simulation, each process will produce the intermediate results representing the subdomain for neighbouring subdomains to integrate after simulation. In the HPC architecture, each computing node is usually designed to access the same remote data storage to execute tasks in parallel. A network file system (NFS) or other methods are used to share the storage and to ensure the synchronization of data access.

However, the performance of the parallel system would be compromised if the network connection and topology of the remote storage and computing nodes were not properly configured (Huang and Yang 2010). In addition, different storage models could have different impacts on the performance for I/O and data-intensive applications. In this article, we will test the impact of the most popular storage models, NFS and Parallel Virtual File System version 2 (PVFS2; Latham *et al*. 2010) on the performance of dust storm forecasting, compared with the local storage strategy, where each computing node has the same copy of data on its local storage.

3.3. Experiment design

To enable the computability of dust storm forecasting (Koh *et al*. 2005) for higher resolutions and larger geographic domain, we designed three sets of experiments to utilize and understand various aspects of the computing strategies and spatiotemporal patterns of the dust storm simulation: (1) *Adaptive loosely coupled model* is to test the efficiency of our loosely coupled nesting strategy; (2) *Subdomain and computing node mapping* is used to test the performance difference of different methods for dispatching subdomains to different computing nodes; and (3) *File system* is to analyse the impact of different file systems on performance.

To investigate the performance improvement, we calculated the performance improvement factor s:

$$s = \frac{\Delta t}{T} \tag{1}$$

where Δt is the decrease in the computing time of dust storm simulation and T is the original computing time before parallel systems are optimized.

4. Experiments and result analyses

4.1. Experiment environment

Facility A includes 25 computing nodes and all nodes are connected through local area networks (LANs of 1Gbps). Each node has 16GB memory and two quad-core processors (eight physical cores) with a clock frequency of 2.33 GHz, a peak performance of 7.6 Gflops/core, and a sustained performance of 1 Gflops/core.

4.2. Adaptive loosely coupled models

This experiment first executes a lower-resolution (22 km) Eta-dust for a large geographic domain to identify areas of interests (AOIs) with high dust concentration. Spatiotemporal correlation analysis is performed in near real-time fashion for 2-day, low-resolution results of Eta-dust to identify AOI regions by utilizing the dust spatial distribution and evolution patterns. High resolution but much smaller regions will be simulated by NMM-dust model with 3 km as resolution in parallel. The dust event on 1 July 2007 was used to test the feasibility of the adaptive loosely coupled nested dust storm framework. Figure 5 shows the entire domain and 18 AOIs identified.

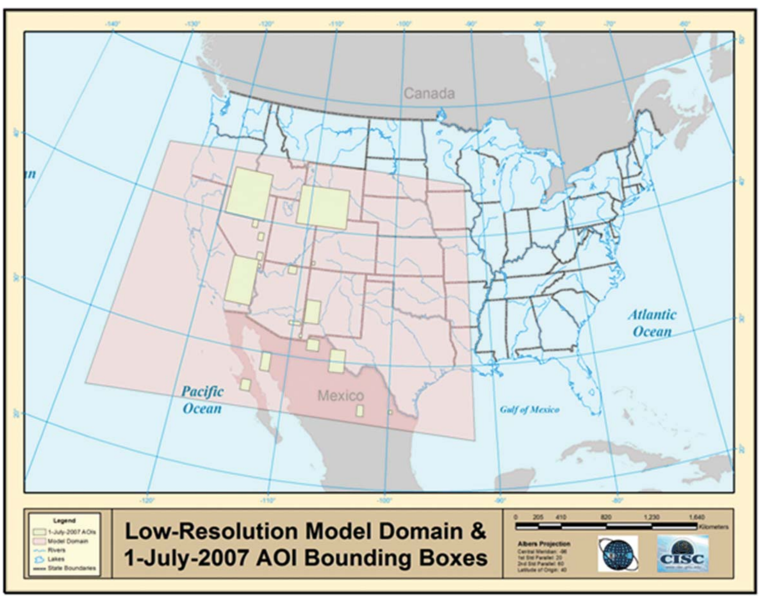

Figure 5. Low-resolution model domain and sub-domains (AOIs) identified for high-resolution model execution.

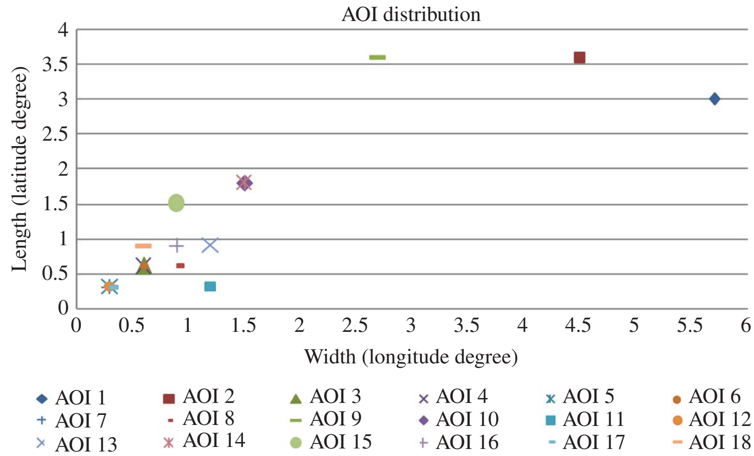

Figure 6. AOIs width and length distribution.

Figure 6 shows the AOI width and length distribution where most AOIs are within 2 × 2 degrees geographic coverage that can be well supported by higher-resolution NMM-dust model. Figure 6 shows only 12 AOIs because some AOIs have similar domain size and only one domain of each different size is illustrated. Figure 7 shows the execution time required for different AOIs when HPC handles all AOIs in parallel. Therefore, it is expected to finish the entire AOIs within 2.7 hours if all the AOIs are simulated by the NMM-dust model in parallel. This does not satisfy the 2-hour constraint for 1-day forecasting. Therefore, facility A should be optimized through better parallelization and scheduling strategies to enable the forecasting within 2 hours.

4.3. *Performance improvement*

4.3.1. *Subdomain and computing nodes mapping*

Figure 8 shows the computing time for NMM-dust model using different process numbers with neighbour mapping and non-neighbour mapping methods. The experiment results of different mapping methods support that if we map the neighbouring subdomains to the neighbouring processors of computing nodes, much higher performance can be obtained.

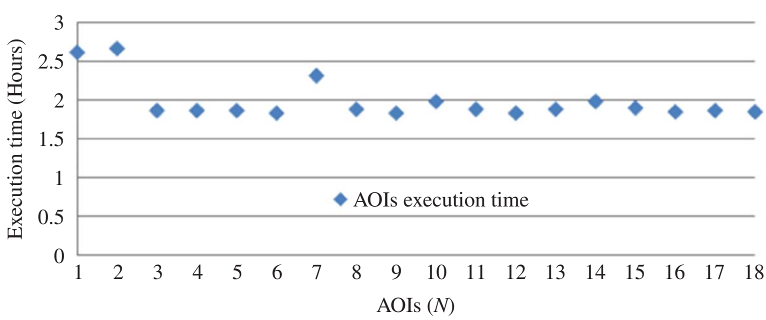

Figure 7. NMM-dust model execution time for each AOI on facility A in parallel.

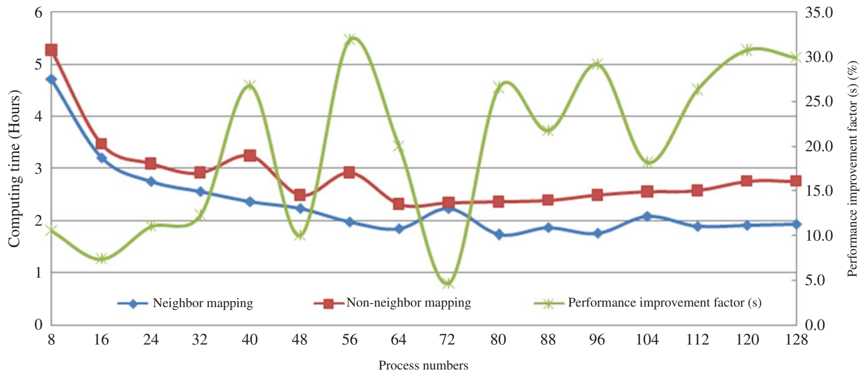

Figure 8. Neighbour mapping can improve the performance of NMM-dust storm modelling by a factor of 20% than that of non-neighbour-mapping methods.

It is observed that around 20% performance improvement factor (*s*) could be achieved on an average if we use the neighbour mapping method. Therefore, this result indicates that we should dispatch neighbour subdomains to the same computing node as much as possible to reduce the communication over computer networks.

4.3.2. File system model

During the simulation, each process will produce temporary files for its subdomains to be integrated with that of other neighbouring subdomains/processes after simulation. The experiment results (Figure 9) demonstrate that it is possible to get a similar performance when using local storage to store the temporary files and then transfer results to the master node after finishing the simulation and using NFS to share remote storage. When more process numbers are used, the local storage strategy gets a little better performance. This result indicates that the I/O performance is not the key bottleneck. However, PVFS2 would be

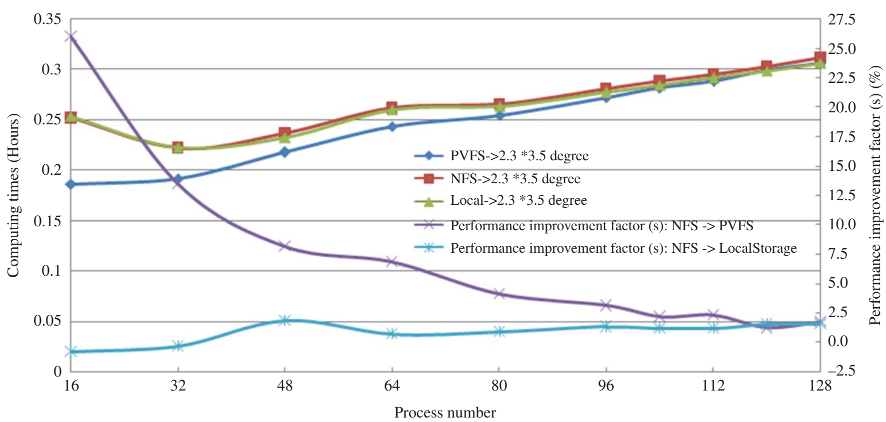

Figure 9. Using different file systems to run the NMM-dust storm model for domains 2.3 × 3.5 degree in the southwest United States with 3 km resolution, for 3-hour predictions. The PVFS2 achieved a 20% performance improvement (s) than that of NFS.

helpful to improve the performance by reducing the communication overhead. In addition, PVFS would obtain better performance when using less than 32 processes.

4.4. Experimental result analysis and integration

4.4.1. Result analysis

The three sets of experiment results verified the feasibility of using loosely coupled nesting models, and reveal spatiotemporal patterns and guidance for configuring and scheduling the HPC facilities for better performance. Table 2 shows how to integrate the experiment results into practical applications: (1) experiment 1 indicates that the spatiotemporal pattern of a phenomenon can be utilized to enable the computability of simulating a large geographic domain; (2) experiment 2 supports that if we map the neighbouring subdomains to the neighbouring computing nodes, better performance can be obtained. The neighbour subdomains assigned to the same machine would reduce the communication overhead by making the communication occur at the same machine without going through the network; and (3) experiment 3 suggests that a good file system model would greatly improve the performance.

4.4.2. Integration

Finally, we integrated adaptive loosely coupled models to divide the large area into multiple small subdomains by performing a fast, low-resolution model simulation, and then using neighbour mapping and local storage strategies to improve the performance of the parallel system for high-resolution model executions.

Figure 10 shows the improvements of computing time and performance for each AOI execution on the 24 computing nodes for 48 process numbers before and after optimization using the NMM-dust model. It is observed that the performance could be improved by more than 36% in average for each AOI execution after using neighbour mapping and local storage strategies. In addition, it indicates facility A can successfully complete the forecasting within 1.8 hours. These results show that the approaches proposed in this article address the computing demands of a large-area and high-resolution forecasting within the time constraints.

5. Conclusion

This research reports strategies to enable the computability of dust storm forecasting. Two dust-storm simulation models, Eta-8bin and NMM-dust, were used as examples to illustrate the adaptive loosely coupled strategy. The strategy divides the computation for a large domain into multiple small subdomains that require much less computing. We presented a solution based on the parallelization of the high-resolution model, NMM-dust, and a set of strategies, including subdomain and computing nodes mapping and storage model that can be used to better leverage HPC by considering the spatiotemporal patterns of computing resources. Three sets of experiments are conducted to demonstrate the feasibility of adaptive loosely coupled model executions and the improvements by using performance optimization strategies. The integration results show that the approaches proposed in this article can successfully perform high-resolution forecasting over a large geographic domain within a reasonable time constraint.

Table 2. Experimental results for better leveraging HPC with spatiotemporal patterns.

Experiment	How to leverage distributed HPC systems	Reason
1. Loosely coupled nesting models	Divide and run hotspots area, which is identified based on the spatial pattern and temporal evolution of dust storm in parallel	Dividing a large geographic scope into multiple small geographic coverage areas based on the spatiotemporal pattern of dust storm produces small areas whose simulation takes much less time to complete
2. Subdomain and processor mapping	Dispatching subdomains to computing resources based on the spatial relationships of subdomains as well as computing resources	Much higher performance can be obtained through mapping the neighbour subdomains to the neighbour processors
3. File system model	Data-intensive applications should use optimized file system.	PVFS2 helps in reducing the I/O and network bottlenecks

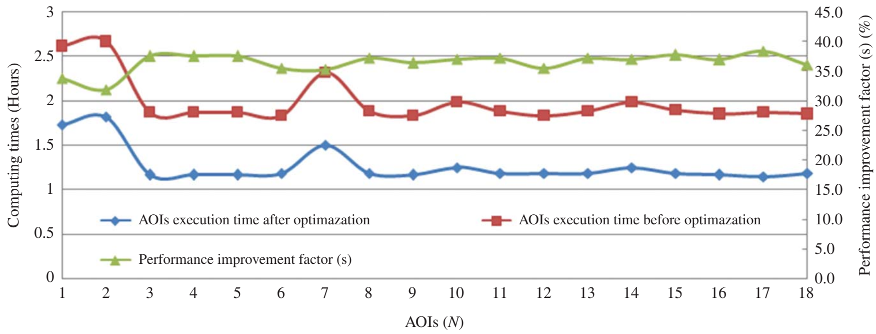

Figure 10. After optimization, 18 subdomains run on facility A using NMM-dust model can be completed within 2 hours.

The research results demonstrate a great potential to solve computing problems that require conducting clustered high-resolution phenomena prediction for a large geographic domain. More research and advancements in relevant fields would greatly enhance the applicability of the methodologies introduced in the article in the next decade including the following research aspects:

Cost model: The computing and communication time is based on many factors, such as the amount of computing nodes, network bandwidth, CPU speed, memory, parallelization degree, parallelization methods, mapping of the subdomain to computing nodes, spatial domain and spatial resolution of the model, and the geospatial nature of the scientific problem. The factors will determine the total computing time and communication required for each computing subroutine. A cost model, considering all these factors, would be very helpful to design scheduling algorithms and allocate suitable computing resources for similar forecasting problems.

Integration and validation: The validation of the adaptive loosely coupled nesting models to perform real-time dust storm forecasting should be performed. The integration of the performance improvement technologies and validation of model results by comparing to real-time observations should be conducted.

Model combination: The success of the loosely coupled nested model depends on two assumptions: (1) the coarse model can perform a quick and accurate forecasting and the results are accurate enough to identify the high dust concentration areas and (2) the finer model can perform a more accurate and detailed forecasting on subdomains after ingesting the results of the coarse model. In this article, Eta-8bin model is used as the coarse model and NMM-dust model is used as the finer model. In the future, the two models will be put in operational environments.

Spatiotemporal computing: In practice, the simulation of dust storms is very dynamic in spatiotemporal scales; therefore, it requires the dynamic allocation of computing resources. The spatiotemporal patterns are also of great importance to address other aspects of geographic sciences, such as visualizing spatiotemporal intensive datasets from simulation models. Therefore, conducting research using HPC techniques in combination with 3D and 4D visualization techniques would help understand the results of complex simulations to enable the modelling of complex problems in a convenient fashion (Li *et al*. 2011). A middleware considering these spatiotemporal patterns would optimize and enable the allocation and use of computing resources for geospatial applications effectively and efficiently (Huang and Yang 2010). In the future, it would be necessary to develop a middleware that can schedule the tasks in a way that improves the scalability and performance of networked computing nodes by fully considering the spatiotemporal patterns. Such an effort would also help to construct a better geospatial cyber infrastructure (Yang *et al*. 2011a).

Spatial cloud computing (SCC): SCC is expected to be the next generation platform to support geospatial science applications (Yang *et al*. 2011b). Adaptive loosely coupled models require a large computing pool to run various AOIs identified by a coarse model in parallel to achieve the best performance. In addition, real-time dust storm simulation is a big data application with computing and concurrent access spikes when a dust storm approaches. Different amounts of computing resources are required at different times to run low- or high-resolution models for a large geographic region. This causes computing spike requirements that can be best handled by elastic and on demand SCC.

Acknowledgements

This work was supported by NSF (IIP-1160979 and CNS-1117300), NASA (NNX12AF89G), and the Microsoft Research's Earth, Energy, and Environment Program.

References

Anthes, R., 1983. Regional models of the atmosphere in middle latitudes. *Monthly Weather Review*, 111, 1306–1335.

Armstrong, M.P., Cowles, M., and Wang, S., 2005. Using a computational grid for geographic information analysis. *Professional Geographer*, 57 (3), 365–375.

Baer, F. and Zhang, B., 1998. Optimizing computations in weather and climate prediction models. *Meteorology and Atmospheric Physics*, 67 (1–4), 153–168.

Baillie, C.F., MacDonald, A.E., and Sun, S., 1995. QNH: a portable, massively parallel multi-scale meteorological model. *In*: *Proceedings of the fourth international conference on the applications of high performance computers in engineering*, 19–21 June, Milan, Italy.

Baillie, C., Michalakes, J., and Skilin, R., 1997. Regional weather modeling on parallel computers. *Parallel Computing*, 23 (14), 2135–2142.

Baklanov, A., Rasmussen, A., Fay, B., Berge, E., and Finardi, S., 2002. Potential and shortcomings of numerical weather prediction models in providing meteorological data for urban air pollution forecasting. *Water, Air, & Soil Pollution: Focus*, 2 (5–6), 43–60.

Borthwick, A., Marchant, R., and Copeland, G., 2000. Adaptive hierarchical grid model of water-borne pollutant dispersion. *Advances in Water Resources*, 23 (8), 849–865.

Calstroka, J. and Watson, P., 2010. Automatic software deployment in the Azure cloud, distributed applications and interoperable systems. *In*: Eliassen, F. and Kapitza, R., Eds. *Proceedings of 10th IFIP WG 6.1 International Conference, DAIS 2010*, Amsterdam, Netherlands, June 7–9, 2010. Lecture Notes in Computer Science: New York; 155–168.

Constantinescu, M., Sandu, A., and Carmichael, G.R., 2008. Modeling atmospheric chemistry and transport with dynamic adaptive resolution. *Computational Geosciences*, 12 (2), 133–151.

Cotton, W.R., *et al*., 2003. RAMS 2001: current status and future directions. *Meteorology and Atmospheric Physics*, 82 (1–4), 5–29.

Davis, C., *et al*., 1999. Development and application of an operational, relocatable, meso-gammascale weather analysis and forecasting system. *Tellus*, 51A, 710–727.

Drake, J. and Foster, I., 1995. Introduction to the special issue on parallel computing in climate and weather modeling. *Parallel Computing*, 21 (10), 1539–1544.

Dulac, F., Moulin, C, and Lambert, CE., 1996. Quantitative remote sensing of African dust transport to the Mediterranean. *In*: S. Guerzoni, R. Chester, eds. *The impact of desert dust across the Mediterranean*. Berlin: Springer Verlag, 25–49.

Fennessy, M.J., Shukla, J., 2000. Seasonal prediction over north America with a regional model nested in a global model. *J. Climate*, 13, 2605–2627.

Giorgi, F., Mearns, L.O., Shields, C., and McDaniel, L., 1996. Regional nested model simulations of present day and 2 × CO2 climate over the central plains of the U.S. *Climatic Change*, 40 (3–4), 457–493.

Gong, S.L., *et al*., 2003. Characterization of soil dust aerosol in China and its transport and distribution during 2001 ACE-Asia: 2. Model simulation and validation. *Journal of Geophysical Research*, 108 (D9), 4262. doi:10.1029/2002JD002633.

Goudie, A.S. and Middleton, N.J., 1992. The changing frequency of dust storms through time. *Earth and Environmental Science*, 20 (3), 197–225.

Han, Z., *et al*., 2004. Model study on particle size segregation and deposition during Asian dust events in March 2002. *Journal of Geophysical Research*, 109, D19205. doi:10.1029/2004JD004920.

Hart, G., *et al*., 1998. Multi-scale atmospheric dispersion modeling by use of adaptive gridding techniques. *Environmental Monitoring and Assessment*, 52 (1–2), 225–238.

Henderson, T., *et al*., 1994. Parallelizing the Eta weather forecast model: initial results. *In*: *Proceedings of High Performance Computing '94. Society for Computer Simulation*, La Joya, CA, 10–14 April, 76–81.

Huang, Q., 2011. *Utilizing Model interoperability and spatial cloud computing to support dust storm forecasting*. Thesis (PhD), George Mason University, Fairfax, VA, USA.

Huang, Q. and Yang, C., 2010. Optimizing grid computing configuration and scheduling for geospatial analysis–An example with interpolating DEM. *Computers and Geosciences*, 37 (2), 165–176.

Janjic, Z.I., 2003. A nonhydrostatic model based on a new approach. *Meteorology and Atmospheric Physics*, 82, 271–285.

Jasper, K., Gurtz, J., Lang, H., 2002. Advanced flood forecasting in Alpine watersheds by coupling meteorological observations and forecasts with a distributed hydrological model. *Journal of Hydrology*, 267 (1–2), 40–52.

Jin, H., *et al*., 2003. Automatic multilevel parallelization using OpenMP. *Scientific Programming*, 11 (2), 177–190.

Koh, M., Peng, L., and See, S., 2005. Integration of parallel MM5 with distributed resource manager and performance evaluation. *In*: *Proceedings of the eighth international conference on high-performance computing in Asia-Pacific Region (HPCASIA'05)*. 30 November–3 December, Beijing, China. Washington, DC: IEEE Computer Society, 289–298.

Kuligowski, R.J., and Barros, A.P., 1999. High-resolution short-term quantitative precipitation forecasting in mountainous regions using a nested model. *J. Geophys. Res.*, 104 (D24):31, 553–31, 564, doi:10.1029/1999JD900938.

Latham, R., *et al*., 2010. A next-generation parallel file system for Linux cluster. *LinuxWorld Magazine*, 2 (1), 54–57.

Li, J., *et al*., 2011. Visualizing dynamic geosciences phenomena using an octree-based view-dependent LOD strategy within virtual globes. *Computers and Geosciences*, 37 (9), 1295–1302.

McGregor, J.L., 1997. Regional climate modelling. *Meteorol. Atmos. Phys*, 63, 105–117.

Michalakes, J., Dudhia, J., Gill, D., Klemp, J., and Skamarock, W., 1998. Design of a next-generation regional weather research and forecast model. *In*: W. Zwieflhofer, N. Kreitz, eds. *Towards Teracomputing*, World Scientific: River Edge, New Jersey, 117–124.

Nanjundiah, R.S., 1998. Strategies for parallel implementation of a global spectral atmospheric general circulation model. *In*: *High Performance Computing, 1998. HIPC '98. 5th International Conference*, 17–20 December, Chennai, India. Washington, DC: IEEE Computer Society, 452–458.

Neilson, R.P., Pitelka, L.F., Solomon, A.M., Nathan, R., Midgley, G.F., Fragoso, J.M.V., Lischke, H., and Thompson, K., 2005. Forecasting regional to global plant migration in response to climate change. *American Institute of Biological Sciences*. 55 (9), 749–759.

Nickling, W.G. and Gillies, J.A., 1993. Dust emission and transport in Mali, West Africa. *Sedimentology*, 40 (5), 859–868.

Nickovic, S., *et al*., 1997. Aerosol production/transport/deposition processes in the Eta model: desert dust cycle simulations. *In*: G. Kallos, V. Kotroni, and K. Lagouvardos, eds. *Proceedings of the Symposium on Regional Weather Prediction on Parallel Computer Environments*, University of Athens, Greece, 109–122.

Nickovic, S., *et al*., 2001. A model for prediction of desert dust cycle in the atmosphere: quantifying the radiative impacts of mineral dust (DUST). *Journal of geophysical research*, 106 (D16), 18113–18129.

Pleim, J.E., Chang, J.S., and Zhang, K., 1991. A nested grid mesoscale atmospheric chemistry model. *J. Geophys. Res*., 96 (D2), 3065–3084.

Purohit, S., *et al*., 1999. Development of parallel climate/forecast models on 100 GFlops PARAM computing systems. *In*: *Proceedings of the eight ECMWF workshop on the use of parallel processors in meteorology*, 16–20 November, Reading, UK. Reading: World Scientific.

Ramón, D.E., Laprise, R., and Denis, B., 2002. Forecasting skill limits of nested, limited-area models: A perfect-model approach. *Mon. Wea. Rev*., 130, 2006–2023.

Rodriguez, B., Hart, L., and Henderson, T., 1995. Comparing scalable programming techniques for weather prediction. *In*: *Proceedings of the conference on Programming Models for Massively Parallel Computers (PMMP '95)* 9–12 October, Berlin, Germany. Washington, DC: IEEE Computer Society, 111–120.

Rodriguez, B., Hart, L., and Henderson, T., 1996. Parallelizing operational weather forecast models for portable and fast execution. *Journal of Parallel and Distributed Computing*, 37 (2), 159–170.

Sela, J.G., 1980. Spectral modeling at the National Meteorological Center. *Mon. Wea. Rev*., 108, 1279–1292.

Shao, Y. and Dong, C.H., 2006. A review on East Asian dust storm climate, modeling and monitoring. *Global and Planetary Change*, 52 (1–4), 1–22.

Shao, Y., *et al*., 2007. Numerical simulation of the October 2002 dust event in Australia. *Journal of Geophysical Research*, 112, D08207. doi:10.1029/2006JD007767.

Sokolik, I.N. and Toon, O.B., 1996. Direct radiative forcing by anthropogenic airborne mineral aerosols. *Nature*, 381, 681–683.

Sterling, T., *et al*., 1995. Communication overhead for space science applications on the Beowulf parallel workstation. *In*: *Fourth IEEE international symposium on high performance distributed computing (HPDC-4 '95),* 2–4 August 1995, Pentagon City, VA. Washington, DC: IEEE Computer Society, 23.

Westphal, D.L., Toon, O.B., and Carlson, T.N., 1988. A case study of mobilization and transport of Saharan dust. *Journal of the Atmospheric Sciences*, 45 (15), 2145–2175.

Wolters, L., Cats, G., and Gustafsson, N., 1995. Data-parallel numerical weather forecasting. *Scientific Programming*, 4 (3), 141–153.

Xie, J., *et al*., 2010. High performance computing for the simulation of dust storms. *Computers, Environment, and Urban Systems*, 34 (4), 278–290.

Xue, M., *et al*., 2003. The Advanced Regional Prediction System (ARPS), storm-scale numerical weather prediction and data assimilation. *Meteorology and Atmospheric Physics*, 82 (1–4), 139–170.

Yang, C., *et al*., 2008. Distributed geospatial information processing: sharing Earth science information to support Digital Earth. *International Journal of Digital Earth*, 1 (3), 259–278.

Yang, C., *et al*., 2011a. Using spatial principles to optimize distributed computing for enabling the physical science discoveries. *Proceeding of National Academy of Sciences of USA*, 108 (14), 5498–5503.

Yang, C., *et al*., 2011b. Spatial Cloud Computing: how could geospatial sciences use and help to shape cloud computing. *International Journal on Digital Earth*, 4 (4), 305–329.