SANDIA REPORT

SAND2006-7626
Unlimited Release
Printed December 4, 2006

A Massively Parallel Scalable Atmospheric Model

William F. Spotz, Mark A. Taylor

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia is a multiprogram laboratory operated by Sandia Corporation,
a Lockheed Martin Company, for the United States Department of Energy's
National Nuclear Security Administration under Contract DE-AC04-94-AL85000.

Approved for public release; further dissemination unlimited.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

NOTICE: This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone:	(865) 576-8401
Facsimile:	(865) 576-5728
E-Mail:	reports@adonis.osti.gov
Online ordering:	http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone:	(800) 553-6847
Facsimile:	(703) 605-6900
E-Mail:	orders@ntis.fedworld.gov
Online ordering:	http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2006-7626
Unlimited Release
Printed December 4, 2006

A Massively Parallel Scalable Atmospheric Model

William F. Spotz
Mark A. Taylor
Exploratory Simulation Technologies Department
Sandia National Laboratories
P.O. Box 5800
Albuquerque, NM 87185

Abstract

This is the Final Report for an LDRD project whose purpose was to demonstrate the parallel scalability of the spectral element method applied to a model of the atmosphere, and to mature the model for inclusion into a production-level climate system model. This atmospheric model is indeed highly scalable, achieving over 5 TFLOPS on Red Storm (pre-upgrade) running on almost 10,000 processors. Much of the software engineering to integrate it with the Community Climate System Model (CCSM) has been completed. The CCSM is a nation-wide climate modeling effort sponsored by The National Science Foundation, the Department of Energy and the National Aeronautics and Space Administration. Once this integration is completed, it will make the spectral element model one of four atmospheric model options from which researchers can choose when running CCSM.

Acknowledgment

We would like to thank the Sandia National Laboratories Lab Directed Research and Development program for their support.

We would also like to thank our NCAR collaborators: Rich Loft, Steve Thomas, Henry Tufo, Amik St-Cyr, John Dennis, Aimé Fournier and Jim Edwards; and their funders: the National Science Foundation and the Office of Biological and Environmental Research, Office of Science, US Department of Energy.

Finally, we would like to thank Frank Giraldo of the Naval Research Laboratory (now at the Naval Postgraduate School), Lorenzo M. Polvani of Columbia University and R. Saravanan of Texas A&M University for their collaboration and informative discussions.

Contents

Preface **9**

Summary **10**

Nomenclature **11**

1 Introduction **15**

Background: The Spectral Transform Method 15

The Community Climate System Model 16

The Community Atmosphere Model 18

Spectral Element Dynamical Cores 19

2 Parallel Scalability **23**

Cache Blocking 24

Domain Decomposition 24

HOMME on ASCI Red 25

HOMME on IBM BlueGene/L 26

HOMME on Red Storm 27

Semi-implicit Results 27

3 The Polar Vortex Problem **33**

Initial Condition 33

Time-Dependent Forcing at the Bottom 35

Sponge Layer at the Top 37

Results 37

4 Triangular Spectral Elements **43**

Spatial Discretization 45

Interpolation in the Reference Element 45

Cubature in the Reference Element 46

Computing Cubature Points 47

Evaluating the Nodal Basis Functions 48

Filtering the High-Frequency Waves 49

Integration and Local Element-wise Operators 51

Formulation of the Method 52

Poisson Equation on the Plane 53

Advection Equation on the Sphere 53

Shallow Water Equations on the Sphere 53

5 Conclusions **55**

References **57**

List of Figures

1.1 An example spectral element grid for the Earth. 21

2.1 Performance of HOMME on ASCI Red 25

2.2 Performance of HOMME on BlueGene/L 26

2.3 Performance of HOMME on Red Storm 28

2.4 Semi-implicit Integration Rates 30

2.5 Semi-implicit Accelerations 31

3.1 Polar vortex problem initial state profiles 36

3.2 Isosurfaces of scaled potential vorticity $\Pi = 1$ for the 36km run 38

3.3 Stereographic projection of constant Π surface contours 39

3.4 Constant Π contours 41

3.5 Horizontal vortex-area 42

List of Tables

2.1 HOMME dynamical core (dry dynamics) benchmark runs on ASCI-Red 26

2.2 HOMME dynamical core benchmark runs on BlueGene/L 27

2.3 HOMME dynamical core benchmark runs on Red Storm 28

2.4 Semi-implicit Experiment Parameters 29

4.1 Properties of the cubature points and the points used in the Fekete-Gauss method 48

4.2 Number of cubature points and highest unaffected modes 51

Preface

This SAND Report is the Final Report for an LDRD project entitled "A Massively Parallel Scalable Atmosphere Model," funded from FY04 through FY06. Its content comes from published papers and presentations developed under the auspices of the project, primarily

> J. Dennis, A. Fournier, W. Spotz, A. St-Cyr, M. Taylor, S. Thomas and H. Tufo, "High-resolution convergence properties and parallel efficiency of a spectral element atmospheric dynamical core," *The International Journal of High Performance Computing Applications*, V. 19, No. 3, Summer 2005, pp. 225–235.
>
> W. Spotz, "Massively Parallel Performance of HOMME," presented at *2006 Workshop for the Solution of PDEs on the Sphere*, July, 2006, Monterey, CA.
>
> F. Giraldo and M. Taylor, "A diagonal mass matrix triangular spectral element method based on cubature points," *Journal of Engineering Mathematics*, December, 2006.

Summary

The purpose of this project is to demonstrate scalable performance of a global atmospheric model on massively parallel computer architectures. Reliable climate modeling is an increasingly important capability for social, economic and security policy planning. It is also one of the most computationally intensive applications in science. Of the different components that comprise a climate system model (atmosphere, ocean, sea-ice and land, plus the coupler that ties them together), the atmosphere consumes the greatest computational resources. Traditional atmosphere models based on the spectral transform method have always vectorized well, albeit with a somewhat limited granularity. On more general and less expensive distributed-memory systems, this type of atmosphere model does not parallelize well, for several reasons: expensive all-to-all transposes of data; Legendre transforms which do not scale linearly; and the aforementioned granularity issue.

This project approached the scalability issue by collaborating with the National Center for Atmospheric Research (NCAR) on a spectral element atmospheric model. In this method, the surface of the earth is tiled with quadrilaterals, each of which is a high-order finite element (typically 8 by 8). These elements define locally-supported basis functions which eliminate the need for expensive data transposes, replacing them with efficient edge communications. The Legendre transforms are replaced with scalable, piecewise polynomial computations, and the granularity limitation is greatly reduced because the problem can be decomposed down to one element per processor.

The scalability properties of the spectral element method were known before the start of the project. The purposes of the project were to demonstrate the model's capabilities at higher resolutions and at larger processor counts than had been achieved before, and to help mature the model so that it could be adopted by the climate modeling community as a component within the Community Climate System Model (CCSM), a production-quality climate model maintained at NCAR. This will allow the advantages of the spectral element dynamical core to be coupled to existing atmospheric physics parameterizations, ocean dynamical cores and physics parameterizations, and sea-ice and land models, so that climatologists can better take advantage of the terascale computers of today and the petascale computers of tomorrow.

Nomenclature

CAM Community Atmosphere Model. The atmospheric component of the CCSM. It supports a variety of dynamical cores that a user may choose from: two spectral transform dynamical cores, one Eulerian and one semi-Lagrangian; and a finite volume dynamical core. Work is underway to add the spectral element dynamical core, which was the focus of this project.

CAM Physics A collection of sub-grid parameterizations within CAM that is shared by all of the dynamical cores supported by CAM. These parameterizations include cloud physics, aerosols, chemistry, and radiation, to name but a few.

CCSM Community Climate System Model. This is the premier climate system modeling effort in the U.S., sponsored by the National Science Foundation, the Department of Energy and the National Aeronautical and Space Administration, and is administered at the National Center for Atmospheric Research. It consists of four major components: an atmosphere model (CAM); an ocean model (POP); a sea-ice model (CSIM); and a land model. These four components are coupled together at their boundaries via flux conditions by a fifth component known as the flux coupler.

Dycore Common abbreviation for "dynamical core." A Dycore refers to the algorithm responsible for solving the equations of motion in an atmosphere or ocean model. Specifically, it refers to the Euler equations in spherical coordinates with a central-pointing gravitational field and a Coriolis term. An atmosphere or ocean model is comprised of its dycore and its physics modules.

Embarrassingly parallel is a term used to describe algorithms that can run on multiple processors with little or no communication. This lack of communication overhead results in a high degree of parallel efficiency and scalability.

Explicit algorithms march the solution for transient equations forward in time by relating data at the new time step exclusively to data at previous time steps. Explicit algorithms are characterized by highly efficient computations, but strict stability restrictions, typically resulting in small time steps.

HOMME High-Order Multi-scale Modeling Environment. This is a follow-on to SEAM developed in the Scientific Computing Division at NCAR. Special attention was paid to efficient memory usage and parallel scalability in the development of this code, and it is here where most new development is implemented. HOMME is the code that was the focus of much of the collaboration for this project and is in the process of being integrated with CAM.

Implicit algorithms march the solution for transient equations forward in time by solving the all data at the new time step simultaneously, via either linear or nonlinear systems, depending on the problem. Implicit algorithms are characterized by expensive computations, but usually have no stability restrictions, typically resulting in large time steps that can be chosen purely on the basis on accuracy.

Latitude describes the angular distance in spherical coordinates north or south of the equator, which is 0 by definition. The north pole represents a latitude of $\pi/2$ or $90°$, and the south pole is $-\pi/2$ or $-90°$.

Longitude describes east-west angular distance in spherical coordinates. By convention, the Greenwich meridian is defined to be zero, and eastward distances are positive, while westward distances are negative.

NCAR National Center for Atmospheric Research. Located in Boulder, Colorado, and primarily funded by the NSF, one of NCAR's many responsibilities is administering the CCSM, which is a collaborative effort with many other government agencies, laboratories, institutions and universities.

Pole problem refers to a collection of problems that numerical methods have in spherical coordinates at or near the poles. Terms like the curl are bounded, but they can be the difference of terms that become unbounded as they approach the poles. The form of spherical derivative operators can pollute the accuracy of approximation techniques. Latitude-longitude grids are highly refined near the poles, which can affect stability conditions and introduce high-frequency oscillations. Methods based on the spherical harmonics address the pole problem "automatically" because the underlying basis functions are isotropic on the sphere. Other methods must deal with the pole problem more explicitly, often with alternate grids or filters that reduce the spectrum of the data near the poles.

SEAM Spectral Element Atmosphere Model. The very first implementation of a spectral element dynamical core for the global atmosphere was called SEAM. It is still in use by certain researchers in the Climate and Global Dynamics Division at NCAR for non-production studies and research.

Semi-implicit algorithms march the solution for transient equations forward in time by treating linear terms implicitly and nonlinear terms explicitly, resulting in a linear system to solve. Semi-implicit algorithms are characterized by computations that are less expensive than fully implicit, but with stability conditions less restrictive than explicit, typically resulting in larger time steps than explicit.

STM Spectral Transform Method. In the STM, spherical coordinate data is represented by spherical harmonic coefficients. Linear operators can be applied directly to this data. Nonlinear operators are applied by transforming the data

to a point-wise representation on a latitude-longitude grid, performing the operation in "physical" space, and then transforming the data back to "spectral" space. Aliasing errors from quadratic terms can be eliminated by truncating the top third of the spectrum.

Chapter 1

Introduction

Climate modeling has been a key application since the earliest days of supercomputing. The very first Cray-1A supercomputer was installed at NCAR in 1974. At the time, its computing resources were strained by a relatively low resolution, global atmosphere model [Oliger(1970)]. Today, much larger and faster computers are devoted to climate modeling, but they are still strained by relatively modest resolution system models that simulate the atmosphere, ocean, sea ice and much, much more. To obtain resolutions that can give scientists meaningful predictions of regional climate change, faster and larger capacity computational hardware will be required. Also needed will be algorithms that utilize these machines more efficiently, by scaling to tens, if not hundreds of thousands of processors. This scalability is the focus of this LDRD project.

Background: The Spectral Transform Method

The atmosphere model developed for early vector processor architectures was the spectral transform method [Machenhauer(1972)]. Numerical fields in such a model are stored as multiple layers of two-dimensional, triangular truncations of spherical harmonic coefficients. To perform nonlinear operations on such fields, they are transformed to "physical" space, i.e. point-wise data on a latitude-longitude grid. The results of these nonlinear operations are then transformed back to "spectral" space. Aliasing errors are eliminated by use of a "two-thirds rule," where the top third of the spectral coefficients is truncated.

This approach is extremely efficient on a vector machine. Indeed, the Japanese Earth Simulator, which sat atop the Top 500 list of the world's fastest computers for an unprecedented two and a half years, was designed in part to take advantage of the highly vectorizable Fourier and Legendre transforms required by the spectral transform atmosphere model [Shingu(2002)]. However, massively parallel, distributed-memory architectures are far more economic and applicable to a wider range of applications. The next planned Japanese supercomputer is reported to be a more general commodity-based machine.

The spectral transform method is far less efficient on distributed memory computers. Conceptually, a transform, from say physical space to spectral space, requires a fast Fourier transform (FFT) in the longitude direction, followed by a Legendre transform in the latitude direction. It is possible to decompose the data in one dimension so that one of the transforms is completely on-processor, but then the other transform would have to be distributed. A 2D decomposition, which is more scalable, would require that both transforms be distributed. Alternatively, we can employ an all-to-all data transposition between the transformation algorithms, so that all transforms are computed on-processor. It is this last approach that is most efficient in almost all computing regimes [Worley(1992)] [Walker(1992)].

Nevertheless, the data transpose communication operation represents a bottleneck in the computations. It is very time-consuming compared to "local methods" that only have to communicate edge data across processor domain boundaries. In addition to the transpose problem, the Legendre transform takes up more and more of the compute time as the resolution increases, because Legendre transforms do not scale linearly.

Despite these drawbacks, the spectral transform method is still used today. One of the reasons is that climate models are still run operationally at relatively modest resolutions, and so these inefficiencies are tolerable. Another reason is that the spectral transform method can be run using larger time steps than other methods. This is because the semi-implicit formulation, which treats all linear terms implicitly in the time-stepping algorithm, has a less restrictive CFL stability condition. The semi-implicit formulation requires the solution of a Helmholtz problem, which is trivial for the spectral transform method because the spherical harmonics are the eigenfunctions in spherical coordinates of the underlying Laplace operator. The matrix that needs to be inverted is therefore diagonal and hence trivial. No other numerical method has this property for semi-implicit time stepping.

The Community Climate System Model

Climate models have evolved considerably over the last several decades. Early climate models were purely atmospheric models, with simple moisture, cloud and radiation sub-grid parameterizations added. Various components and improvements have been added over the years to this starting foundation. Development efforts have branched apart and recombined. The problem of modeling the entire climate system is so large, and requires such a breadth of expertise, that there is far more cooperation in the field of climate modeling than there is competition.

In the U.S., this cooperation has resulted in the Community Climate System Model (CCSM) [Kiehl(2004)], which is funded by the National Science Foundation (NSF), the Department of Energy (DOE), and the National Aeronautics and Space Admin-

istration (NASA). The CCSM is administered by NCAR, an NSF-sponsored research center.

The CCSM is comprised of five components: an atmosphere model, an ocean model, a sea-ice model, a land surface model, and a flux coupler that transfers boundary condition data between the other components. These components are described briefly below:

Atmosphere. The CCSM atmosphere component is called the Community Atmosphere Model (CAM) [Collins(2004)], whose development is led by NCAR. CAM is conceptually divided into two parts. The first is the dynamical core (dycore), which is essentially the Euler equations in spherical coordinates with a central gravitational field and a Coriolis term. The second part is the physics parameterizations, which include just about everything else: solar radiation, moisture, clouds, aerosols, and chemistry, to name just a few. Originally, CAM supported only an Eulerian spectral transform dycore on a latitude-longitude grid. A second, optional dycore was added that supported the semi-Lagrangian spectral transform paradigm on a lat-lon grid. Prior to this project, a third optional dycore has also been added, a finite volume dycore, also on a lat-lon grid.

Ocean. The CCSM ocean model is the Parallel Ocean Program (POP) [Smith(2002)], developed at Los Alamos National Laboratory. POP supports a single finite volume dycore on a "warped" lat-lon grid whose "north pole" has been transferred to Greenland. It also supports a more limited set of parameterizations that model such things as salinity and biomass.

Sea Ice. The CCSM sea-ice model is known as the Community Sea-Ice Model (CSIM) [Briegleb(2004)], developed primarily at NCAR. Its purpose is to track the growth and recession of ice in the polar regions of the earth's oceans. Los Alamos also develops a sea-ice model named CICE [Hunke(2006)], which is compatible with POP and CCSM.

Land Surface. The CCSM land surface model is the Community Land Model (CLM) [Oleson(2004)], which is a collaboration among many groups within NCAR. It simulates the physical, chemical, and biological processes by which terrestrial ecosystems interact with the climate across a variety of spatial and temporal scales.

Flux Coupler. The CCSM flux coupler [Kauffman(2004)] is developed at Argonne National Laboratory. It computes boundary fluxes from state data from a given component on a given grid distributed over a given processor partition, interpolates the data to a different grid and communicates it to a different component on a different processor partition. In this manner, sea surface heat flux is communicated from the ocean to the atmosphere, precipitation is communicated from the atmosphere to the land surface, etc.

The Community Atmosphere Model

Of all the CCSM components, the atmosphere consumes the greatest amount of computing time. The ocean is similar in terms of memory usage, but the atmosphere dominates the computations because the stability condition is more strict for the atmosphere and the time step is much smaller. So, while all components must execute in parallel, it is the efficiency of the atmosphere that is of prime importance to the overall efficiency of the climate system model. Within CAM, the physics modules are "embarrassingly parallel" and so scale without issue. The dycore, however, can have severe scaling limitations, and it is the atmospheric dycore that is the focus of this project.

Eulerian STM. The Eulerian STM [Machenhauer(1972)] was the first dycore implemented in CAM, and is still used today. Atmospheric fields are represented by spherical harmonic coefficients. Linear operators are applied directly to the coefficients. Nonlinear operators are applied by transforming the data to a latitude-longitude grid, performing the operation, and transforming the results back to spectral space. Time-stepping is comprised of applying the appropriate linear and nonlinear operators of the Eulerian form of the equations of motion to obtain the fields at the new time step.

Semi-Lagrangian STM. The semi-Lagrangian STM [Robert(1982)] achieves time-stepping by following atmospheric "particles" over their trajectories through space. Unlike full Lagrangian, the semi-Lagrangian method follows those particles that end their trajectory at a given time step at the vertices of a lat-lon grid. A semi-Lagrange solver must therefore "back out" the starting location of these particles, and their properties are interpolated from the lat-lon data at the beginning of the time step. This algorithm, combined with a semi-implicit formulation, results is a much longer stable time step. It also introduces interpolation error to obtain the initial conditions for the trajectories.

Finite Volume. The finite volume (FV) dycore [Lin(1997)], also know as Lin-Rood, is a relatively recent addition. CAM had to be redesigned to become more modular in order to accept such a different numerical method. The FV (currently) uses a lat-lon grid, and time-stepping is achieved by relating flux conditions across cells. Special filters have to be applied to address the collection of issues known as the "pole problem." While less accurate than a spectral method, the FV method does have the advantage being locally conservative. The FV method allows for the use of flux limiters, which suppress artificial oscillations in the solution variables. These are important properties for tracers such as moisture, aerosols and chemistry, as well as cloud physics.

Spectral Element Dynamical Cores

The spectral element method is a finite element method in which a high degree spectral method is used within each element. The method provides spectral accuracy while retaining both parallel efficiency and the geometric flexibility of unstructured finite elements grids. The method has proven accurate and efficient for a wide variety of geophysical problems, including global atmospheric circulation modeling [Taylor(1997)] [Taylor(1998)] [Giraldo(2001)] [Thomas(2002)] [Giraldo(2004)] [Fournier(2004)] [Thomas(2005)] ocean modeling [Haidvogel(1997)] [Iskandarani(2002)] [Molcard(2002)] [Iskandarani(2003)] and planetary scale seismology [Komatitsch(1999)] [Komatitsch(2002)]. The method has unsurpassed parallel performance; it was used for earthquake modeling by the 2003 Gordon Bell Best Performance winner, running on 1944 CPUs of the Earth Simulator [Komatitsch(2003)] and for climate modeling by a 2002 Gordon Bell Award honorable mention, running on 2048 processors of an IBM SP [Loft(2001)].

Our spectral element dynamical core solves the primitive equations using a hybrid η pressure vertical coordinate system [Simmons(1981)], using the continuous Galerkin spectral element discretization for the horizontal directions (the surface of the sphere), and second order finite differences in the vertical direction [Fournier(2004)] [Thomas(2005)]. The spectral element method relies on quadrilateral elements. We use a subdivided inscribed cube to generate quasi-isotropic tilings of the sphere with such elements. An example mesh is shown in Fig. 1.1. To characterize the horizontal resolution of these meshes, let M be the total number of elements and N be the number of polynomials in each direction used within each element. In the figure, $M = 384$. The spectral transforms performed within each element rely on an $N \times N$ grid and thus the total horizontal resolution is specified by $N \times N \times M$. The number of nodes along the equator is given by $4(M/6)^{1/2}N$, and the average equatorial grid spacing in kilometers is given by $2.45 \times 10^4 M^{-1/2} N^{-1}$.

Using a spectral element discretization on the sphere has several advantages for global climate modeling. First of all, handling the spherical geometry presents no problem since the sphere can be tiled with quadrilateral elements of approximately the same size, thus avoiding clustering points at the poles. Secondly, by using a local coordinate system within each element, the singularities associated with spherical coordinates can also be avoided. Additionally, for climate applications, the method can obtain the accuracy of traditional spherical harmonics based models with only a slight increase in the number of degrees of freedom [Taylor(1997)] [Thomas(2002)].

In what follows we present results from two models, SEAM and HOMME. SEAM is a spectral element atmospheric model research code [Fournier(2004)]. Most of the algorithms in SEAM have been re-implemented in NCAR's High Order Multi-scale Modeling Environment (HOMME). HOMME adds modern software engineering practices in addition to many new features and algorithms, including advanced time stepping, discontinuous Galerkin, adaptive mesh refinement and several domain decomposition

strategies [Thomas(2005)] [St-Cyr(2004)] [St-Cyr(2005)] [Nair(2004)]. Here we only use the simplest form of HOMME: explicit time stepping, a continuous Galerkin treatment of the prognostic variables and quasi-isotropic conforming element meshes. In this form, the numerical results of HOMME and SEAM are indistinguishable.

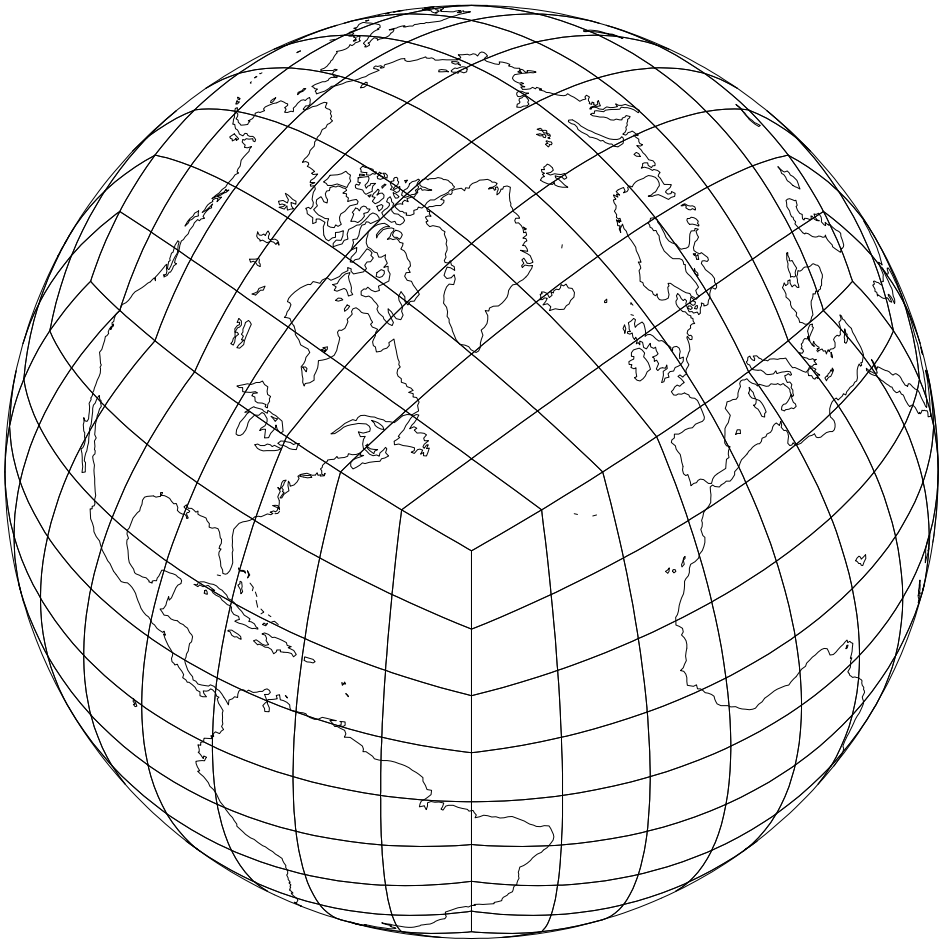

Figure 1.1. An example spectral element grid for the Earth. Continental outlines are shown for reference. The grid is generated by projecting an inscribed cube onto the surface, and then subdividing each of the 6 faces of the cube into an 8×8 grid of elements. Typically each element uses a degree 7 polynomial expansion. Finite differences are used in the vertical direction (not shown).

Chapter 2

Parallel Scalability

Spectral element methods are well suited to modern cache-based parallel computers for several reasons. First, the basic data structure in the method, the spectral element, is naturally cache blocked. Secondly, due to the $O(N^3)$ cost of the spectral transforms, the method has a very low ratio of communication to computation. Finally, the spectral element discretization allows for efficient two-dimensional domain decomposition strategies.

We demonstrate this performance by presenting results for HOMME running on the IBM BlueGene/L (up to 7776 processors) and ASCI Red (up to 8938 processors). We use vertical resolutions from 20 to 100 levels, and horizontal resolutions from 156km down to 10km. All cases use $N = 8$. For the benchmark problem, we use timings from HOMME configured to run the Held-Suarez test [Held(1994)]. The Held-Suarez tests were designed to allow for the inter-comparison of the climate produced by different atmospheric dynamical cores. Results from SEAM and HOMME have been previously reported [Taylor(1998)] [Thomas(2005)].

The Held-Suarez tests are also useful for benchmarking since they represent the entire dynamical core of an atmospheric model such as CAM. Only the physics components (sub-grid scale models of physical processes) in CAM are not represented. CAM physics is column based, meaning the calculations are performed using only data from the vertical column associated with each element. Using an element based domain decomposition, all the column physics will be performed on processor and require no additional communication. Thus the addition of physics can only improve the parallel scalability of the dynamical core; however, the single processor performance will be affected by the single processor performance of the column physics models. One other difference between CAM and the Held-Suarez tests which can effect parallel performance is that CAM contains additional constituents which must be advected by the dynamical core. These prognostic variables will have an identical parallel behavior as the prognostic variable for temperature, and thus they are represented in the Held-Suarez test but going from one such variable to many such variables will result in a proportional increase in the required CPU time.

These factors show that the time-to-solution from a Held-Suarez test problem can provide a reasonable estimate of the time-to-solution for a full climate simulation. For a given resolution, adding column physics and additional constituents may ac-

tually increase the parallel scalability of the model and the number of floating point operations per second. The time-to-solution will increase, but in the worst case only proportional to the complexity of the column physics and number of additional constituents. In typical climate simulations, this increase is a factor in the range of 1.2 to 2.

Cache Blocking

The elements in the spectral element method provide a natural cache blocking. By way of example, consider the 64 bit data storage requirement for a typical 8×8 element with 20 vertical levels requires only 10 Kbytes of cache per variable. One hundred such 3-D variables will fit into a typical 1 Mbyte data cache. If the data is stored in memory so as to avoid cache conflicts, then all the computations performed within an element can be done entirely in cache. This blocking is independent of resolution since we can increase the resolution by simply using more elements. This situation is unlike codes optimized for vector architecture, where the natural block size and data access patterns grow with resolution. Thus a spectral element model maintains good performance even at high resolutions.

Domain Decomposition

The most natural way to parallelize the spectral element method is to simply assign several elements to each processor. Each element only needs information from adjacent elements, so the domain decomposition reduces to a standard graph partitioning problem. To solve this problem, we use an algorithm based on space filling curves [Dennis(2003)].

The resulting communication patterns are similar to finite difference/finite element methods which parallelize with the same type of domain decomposition. Thus the parallel efficiency of this method would be similar to these other methods, except for the fact that the spectral element method involves computing high order derivatives using spectral transforms instead of low order stencils. These computationally intensive transforms are performed within each element and are localized to each processor thus requiring no communication. The result is that for a wide range of problem sizes, number of processors, and computer architectures, the parallel efficiency remains above 80%, for parallel decompositions as fine as two elements per processor, and reasonable efficiency is obtained at the finest decomposition of one element per processor.

HOMME on ASCI Red

We first present results from the ASCI-Red computer at Sandia National Laboratories. ASCI-Red was the first of DOE's Advanced Strategic Computing Initiative (ASCI) machines. It has 4510 nodes in a mesh interconnect, each with two 333 MHz Pentium II processors. In Fig. 2.1, we present the total MFLOPS obtained per processor at several resolutions and processor counts. Each curve represents a fixed resolution, so that the the total amount of work was kept constant while the processor count was increased. Data for two resolutions includes parallel decompositions as fine as one element per processor. As can be seen in the figure, all resolutions achieve good scalability - better than 80% in all cases. In this range the performance fluctuates between 30 and 40 MFLOPS per processor. The best performance obtained for each resolution, along with the integration rate is given in Table 2.1.

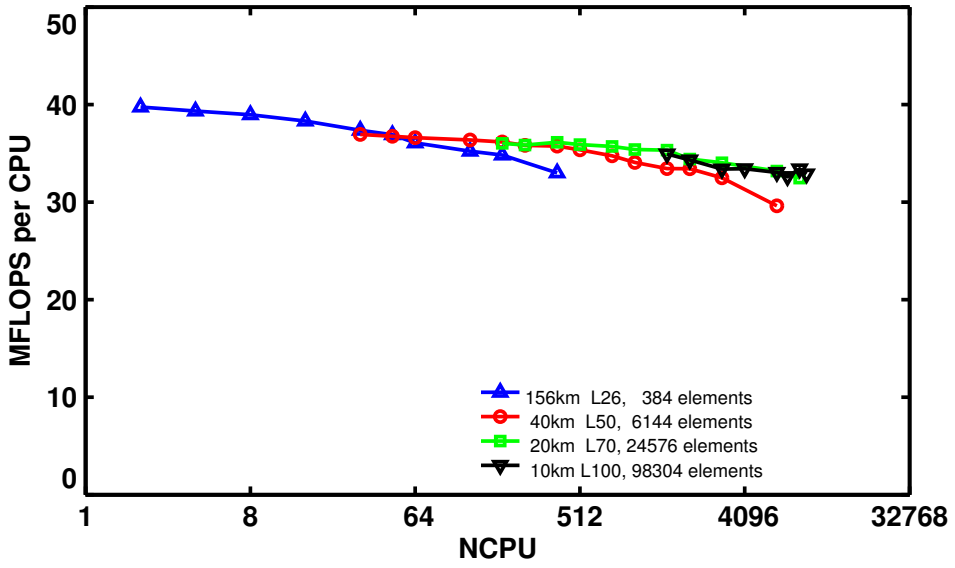

Figure 2.1. Performance of the HOMME dynamical core (dry dynamics) benchmark runs on ASCI Red. Each curve shows the performance for a fixed resolution as the number of processors is increased.

ASCI Red was of interest at the time of this study because of Cray's upcoming Red Storm architecture. Red Storm was designed to have the same balance between processor performance and inter-processor communications as ASCI Red, and thus we expected to obtain very similar scalability results on Red Storm.

Resolution	Processors	GFLOPS	Simulated days/day
156km/26L	384	13	1,600
40km/50L	6,144	181	154
20km/70L	8,192	265	18
10km/100L	8,936	294	1.5

Table 2.1. HOMME dynamical core (dry dynamics) benchmark runs on ASCI-Red

HOMME on IBM BlueGene/L

Our benchmark runs for the IBM BlueGene/L system were run on systems at IBM Watson and Rochester, on up to 7776 processors using one processor per node and only using one of the floating-point pipelines. BlueGene/L uses a toroidal interconnect and PowerPC processors. Sustained MFLOPS per processor are shown in Fig. 2.2, for several resolutions and processor counts. Each curve represents a fixed resolution, so that the the total amount of work was kept constant while the processor count was increased. Data for two resolutions includes parallel decompositions as fine as one element per processor. The best performance obtained for each resolution, along with the integration rate is given in Table 2.2.

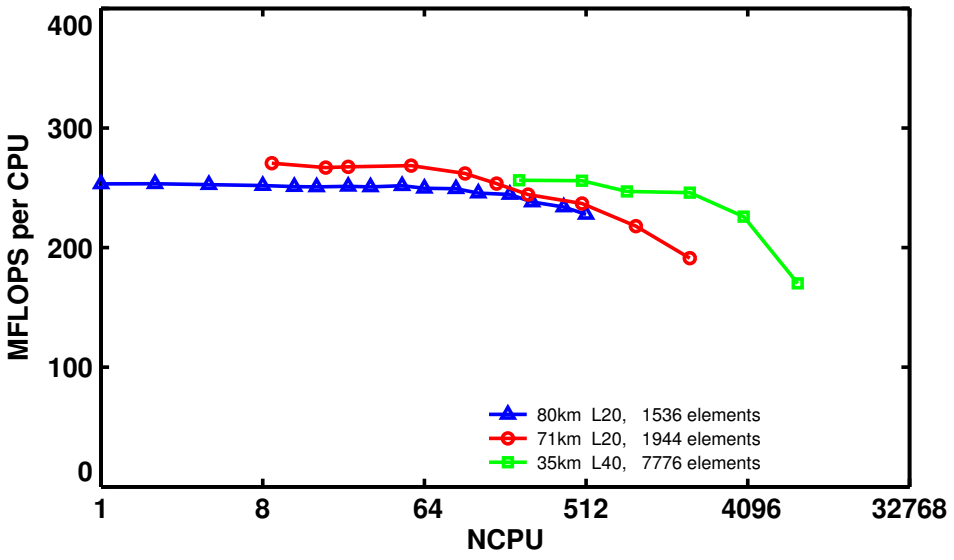

Figure 2.2. Performance of the HOMME dynamical core (dry and moist dynamics) benchmark runs on BlueGene/L. Each curve shows the performance for a fixed resolution as the number of processors is increased.

Resolution	Processors	GFLOPS	Simulated days/day
80km/20L	512	117	7,200
71km/20L	1,944	372	6,500†
35km/40L	7,776	1,320	1,100

Table 2.2. HOMME dynamical core benchmark runs on BlueGene/L. The run denoted by a † included one prognostic moisture variable.

BlueGene/L is of interest since IBM intends to build some very large systems, with over 64,000 processors. Noting that the performance on the dry dynamics problem is around 250 MFLOPS per processor even when running with as few as 2 elements per processor, we can estimate a performance at 9.8km/100L on 49152 processors (2 elements per processor) to be 9.8 TFLOPS, thus allowing an integration rate of 50 simulated days per day.

HOMME on Red Storm

Red Storm, of course, is Cray's Opteron-based supercomputer at Sandia National Laboratories. Our benchmark runs were performed prior to the recent upgrade, when Red Storm had 10,368 single-core 2.0 GHz AMD Opteron compute nodes and a 3D mesh interconnect for a peak performance of 41.5 TFlops. Fig. 2.3 shows the total MFLOPS at several resolutions and processor counts. For the most part, efficiencies are kept above 80%, with the exception of one- or two-element per processor decompositions. Single processor performance is roughly 725 MFLOPS.

Table 2.3 shows the maximum GFLOPS rates and simulated days/day for each resolution. At the maximum resolution on the maximum processor count, Red Storm is achieving 36 days/day. By way of comparison, the Japanese Earth Simulator running the significantly more expensive spherical harmonic model, requires 24TF to achieve an integration rate of 57 model days per day, with 84% of the computations devoted to the dynamical core.

Semi-implicit Results

The results shown thus far are for explicit time-stepping. HOMME also supports semi-implicit time-stepping, in which the linear terms of the governing equations are treated implicitly. This results in an increase in the maximum stable timestep by a factor of 3-8. The drawback is that the each time step requires the solution of a

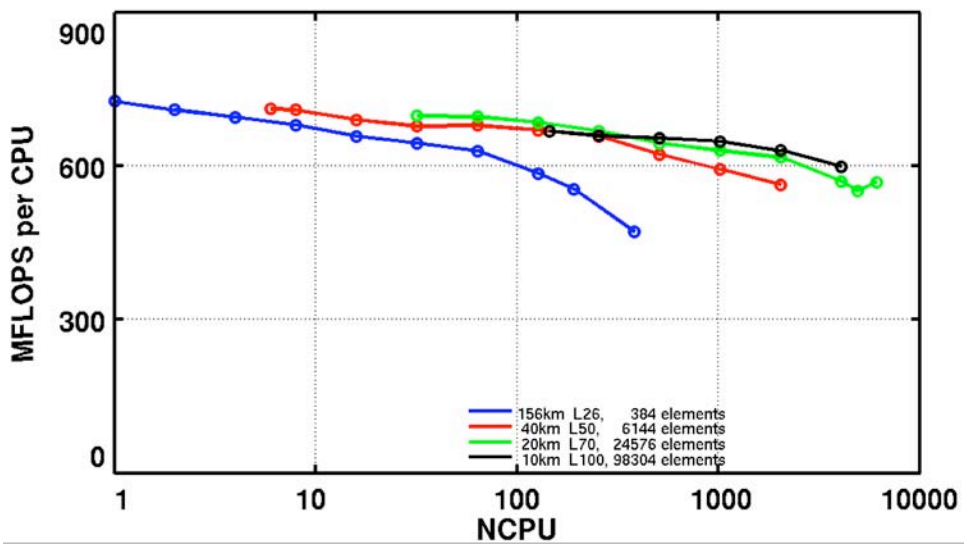

Figure 2.3. Performance of the HOMME dynamical core (dry dynamics) benchmark runs on Red Storm. Each curve shows the performance for a fixed resolution as the number of processors is increased.

Resolution	Processors	GFLOPS	Simulated days/day
156km/L26	384	181	20,293
40km/L50	6,144	2,946	2,684
20km/L70	8,192	4,006	325
10km/L100	9,831	5,159	36

Table 2.3. HOMME dynamical core benchmark runs on Red Storm.

Helmholtz system, which is not as efficient as explicit calculations. Thus, the semi-implicit approach is faster than explicit only when the solver is efficient enough to exploit the gain in time-step.

Typically, as you increase the resolution of a Helmholtz problem, you increase the number of iterations to converge by a factor of $O(h^{-2}$. This does not scale. Fortunately, however, the diagonal term for the atmosphere equations is proportional to $\Delta t^2/h^2$ and because $\Delta t/h$ is held roughly constant in order to satisfy the CFL condition, we find that the number of iterations required to converge is essentially independent of resolution. Theoretically, the semi-implicit approach can scale to high resolutions.

The method implemented in HOMME is the split semi-implicit scheme, analogous

to what is used in ocean models. For an atmosphere with L levels, an $L \times L$ eigensystem is solved as a preprocessing step. This eigensolution is used to represent decoupled vertical modes. These modes can be solved independently, and most of them converge very quickly, in say 3 iterations or fewer. To reduce communication overhead, processor boundary data for all unconverged modes is packed together at each iteration. Converged modes are then masked from the solver algorithm. The longest-converging zeroth mode typically converges in less than 20 iterations. The savings of this approach comes from all of the early-converging modes. The iterative method employed is preconditioned conjugate gradient, specifically stabilized bi-conjugate gradient with an optimized additive Schwarz preconditioners the details can be found in [St-Cyr(2006)].

In practice, however, we see a failure-to-converge rate of about 4%. This phenomenon is still being studied, but its effect is that the approach is still considered highly experimental. The experiments that follow use statistics from experimental runs to determine the average number of iterations each mode takes to converge and then hard-codes those iteration counts to obtain performance estimates assuming the algorithm is one day stabilized to prevent the failures-to-converge.

Table 2.4 shows the various parameters used in our experiments to study semi-implicit time-stepping. We ran Held-Suarez, and typically saw an increase by a factor of 5 for the stable time step. These maximum stable time-steps were modified to durations that divide evenly into 24 hours, which is a common practice in climate and weather modeling. Thus, the ratio of semi-implicit to explicit times steps is not precisely constant.

N_e	h (km)	Levels	Exp Δt (sec)	SI Δt (sec)
8	156	26	75	400
16	80	26	35	200
32	40	50	20	100
64	20	70	10	50
128	10	100	5	30

Table 2.4. Semi-implicit Experiment Parameters. N_e is the number of elements along an edge of the cubed sphere; h is the average element length; "Exp" and "SI" refer to the explicit and semi-implicit time steps, respectively.

Fig. 2.4 shows the timing results of our semi-implicit experiment, run on Sandia's Thunderbird computer. The plotted quantity is simulated days per day (or integration rate) and the results show a promising linear speedup with respect to number of processors. The only real drop off is seen in the coarsest resolution when the decomposition is reduced to two elements per processor.

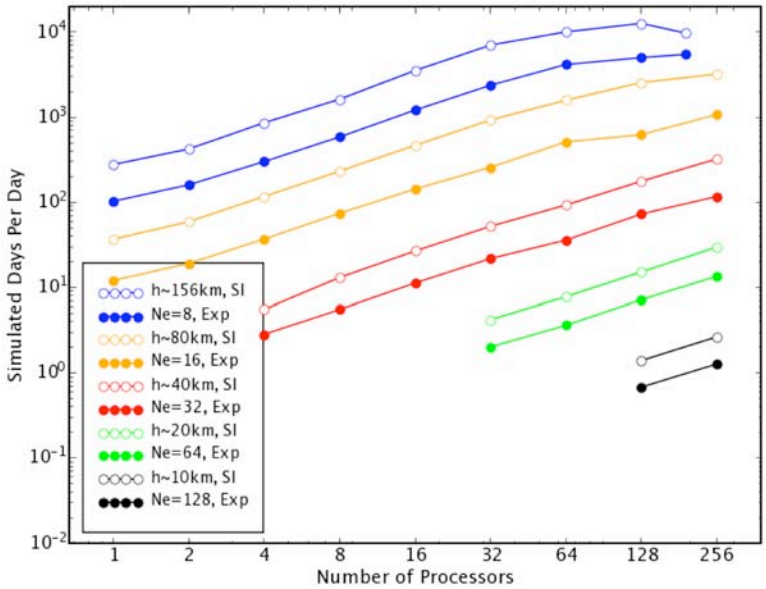

Figure 2.4. Semi-implicit Integration Rates. Each color represents a constant resolution. Open circles represent semi-implicit integration rates and closed circles represent explicit integration rates.

Fig. 2.5 shows acceleration rates, defined as the semi-implicit integration rate divided by the explicit integration rate. If this quantity ever falls below 1, then the semi-implicit scheme is actually costing us efficiency. For the range of experiments conducted here, the acceleration never drops below unity. We do see some volatility in the accelerations, which relate to small deviations from linear in the previous Figure. Nevertheless, the accelerations are consistently above 2 and often above 3. If the method can be made more stable for practical simulations, it will be a benefit to the spectral element atmosphere model.

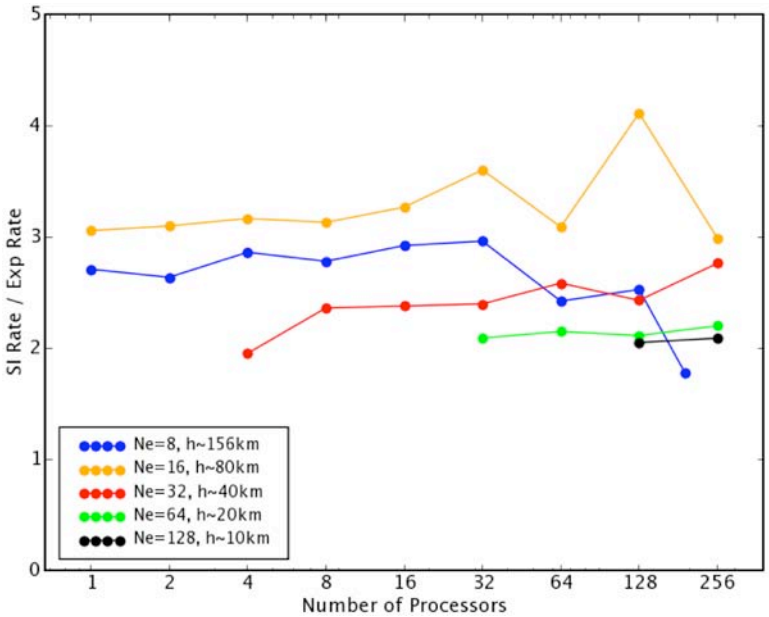

Figure 2.5. Semi-implicit Accelerations. Each color represents a constant resolution. Acceleration is defined as the semi-implicit integration rate divided by the explicit integration rate.

Chapter 3

The Polar Vortex Problem

The stratospheric polar vortex is bounded by strong potential-vorticity gradients which isolate polar air from lower latitudes. A phenomenon of major atmospheric-research importance is the eventual mixing of these air masses, when the vortex is distorted by breaking planetary-scale Rossby waves. This was the focus of [Polvani(2000)]. In addition to that study and those referenced therein, the importance of this phenomenon is underscored by a number of recent studies of its various aspects, in various journals. [Thompson(2000)] studied the positive feedbacks involving ozone destruction and other phenomena that may delay polar-vortex breakdown and enhance spring ozone loss. [James(2000)] provided objective diagnostics of characteristic winter polar-vortex displacements towards Europe or Canada, linked to transient dynamically induced ozone mini-holes. Planetary Rossby-wave refraction was found by [Limpasuvan(2000)] to link stratospheric polar-vortex strength to tropospheric annular-mode variability. [Koh(2000)] evaluated the crucial choice of boundaries for computing material transport across the polar-vortex edge. Finally, [Randall(2005)] found that an "unusually strong upper stratospheric vortex . . . caused enhancements" of NO_x that led to "reductions in O_3 of more than 60% in some cases".

All these authors also cite numerous other examples of this topic's importance. Here we document a minimally complex model for the numerical simulation of polar-vortex breakdown based on that presented in [Polvani(2000)], and then present results from SEAM for several different resolutions up to 36km with 200 levels.

Initial Condition

Following [Polvani(2000)], our initial model state consists of a gradient-flow balanced axisymmetric-vortex: at time $t = 0$, as a function of longitude λ, colatitude φ and pressure p, let the horizontal wind vector

$$\mathbf{u}(\lambda, \varphi, p) = \mathbf{u}^0(\varphi, p) \equiv a_0 \Omega_0 u^0(\varphi, p)\mathbf{i},$$

the isobaric velocity $\omega(\lambda, \varphi, p) = 0$, the surface pressure $p_{\text{sfc}}(\lambda, \varphi) = 10^4$Pa, the temperature

$$T(\lambda, \varphi, p) = T^0(\varphi, p) \equiv T_0 - \frac{a_0^2 \Omega_0^2}{R_0} \frac{\partial \Phi^0}{\partial \ln p}$$

and the geopotential

$$\Phi(\lambda, \varphi, p) = g_0 z_p + a_0^2 \Omega_0^2 \Phi^0(\varphi, p),$$

where a_0, Ω_0 and g_0 are the earth's radius, angular frequency and gravity, $\mathbf{i}$ is the longitudinal unit vector, $T_0 = g_0 H_0 / R_0 = 239.14$K is the isotherm for a scale height $H_0 = 7$km, R_0 is the dry-air gas constant and $z_p \equiv H_0 \ln p_{\rm sfc}/p$ is log-pressure height. The nonlinear gradient-flow balance is

$$\frac{\partial \Phi^0}{\partial \varphi} = -\left(2\mu + u^0 \tan\varphi\right) u^0,$$

where $\mu = \sin\varphi$, so the complete initial state only depends on the dimensionless velocity component u^0, assuming $\int_{-1}^{1} \Phi^0(\varphi, p) d\mu = 0$.

Let us now describe the initial profile as a function of (φ, p). Observing equator-ward from $\varphi = \pi/2$, we assume the absolute vorticity

$$\zeta \equiv 2\mu - \frac{\partial u \cos\varphi}{\partial \mu} \tag{3.1}$$

is initially positive, almost constant until the vortex-edge latitude $\varphi_{\rm v} \equiv 60°$, where it decreases rapidly over a zone of width $\Delta\varphi \equiv 6°$, then again nearly constant until the *surf-zone* edge $\varphi_{\rm s} \equiv 35°$, followed by cubic-sine decrease to zero at the equator, and solid-body rotation for $\varphi \leq 0$. We impose $u_{\rm z}^0 = 0$, where subscript z denotes evaluation at $\varphi = \varphi_{\rm z} \equiv 37°$ and any p. The equator-ward progression just described is controlled by four vertical-structure coefficients $c_n(p)$ to be determined:

$$\zeta^0(\varphi, p) \equiv \begin{cases} r'(\mu) c_1 + c_2, & \mu_{\rm s} \leq \mu, \\ 2\mu c_3 - 4\mu^3 c_4, & 0 \leq \mu \leq \mu_{\rm s}, \\ 2 M_{\rm e}^0 \mu, & \mu \leq 0, \end{cases} \tag{3.2}$$

where the step-like function

$$r'(\mu) \equiv \frac{1}{2} + \frac{1}{2} \tanh \frac{\mu - \mu_{\rm v}}{\Delta\mu}, \tag{3.3}$$

the specific absolute angular momentum from (3.1)

$$M(\mu, p) \equiv 1 - \mu^2 + u \cos\varphi = \int_{\mu}^{1} \zeta(\mu') d\mu', \tag{3.4}$$

where subscript e denotes evaluation at the equator $\varphi = 0$, subscripts s and v denote evaluation at $\varphi_{\rm s}$ and $\varphi_{\rm v}$, and $\Delta\mu \equiv \Delta\varphi \cos\varphi_{\rm v}$. The coefficients $c_n(p)$ are determined as follows. From (3.2,3.4) one has

$$M^0(\mu, p) = \begin{cases} -r(\mu) c_1 + (1-\mu) c_2, & \mu_{\rm s} \leq \mu, \\ M_{\rm e}^0 - \mu^2 c_3 + \mu^4 c_4, & 0 \leq \mu \leq \mu_{\rm s}, \\ (1-\mu^2) M_{\rm e}^0, & \mu \leq 0, \end{cases} \tag{3.5}$$

where by (3.3) the hyperbolic ramp-like function

$$r(\mu) \equiv \int_1^\mu r'(\mu')d\mu' = \frac{\mu-1}{2} + \frac{\Delta\mu}{2}\left(\ln\cosh\frac{\mu-\mu_\mathrm{v}}{\Delta\mu} - \ln\cosh\frac{1-\mu_\mathrm{v}}{\Delta\mu}\right).$$

Since $\varphi_\mathrm{v} > \varphi_\mathrm{z} > \varphi_\mathrm{s}$ we may immediately solve a linear system to obtain

$$\begin{aligned} c_2 &= \frac{M_\mathrm{v}^0 r_\mathrm{z} - M_\mathrm{z}^0 r_\mathrm{v}}{(1-\mu_\mathrm{v})r_\mathrm{z} - (1-\mu_\mathrm{z})r_\mathrm{v}}, \\ c_1 &= [(1-\mu_\mathrm{v})c_2 - M_\mathrm{v}^0]r_\mathrm{v}^{-1}. \end{aligned}$$

Ensuring μ-continuity of (3.2,3.5) leads to another linear system with solution

$$\begin{aligned} c_4 &= \mu_\mathrm{s}^{-4}(M_\mathrm{e}^0 - M_\mathrm{s}^0) - 2^{-1}\mu_\mathrm{s}^{-3}\zeta_\mathrm{s}^0, \\ c_3 &= (2\mu_\mathrm{s})^{-1}\zeta_\mathrm{s}^0 + 2\mu_\mathrm{s}^2 c_4. \end{aligned}$$

Observing that

$$M_\mathrm{z}^0 = 1 - \mu_\mathrm{z}^2$$

and

$$M_\mathrm{s}^0 = -r_\mathrm{s}c_1 + (1-\mu_\mathrm{s})c_2,$$

it only remains to define $a_0\Omega_0 u_\mathrm{e}^0 \equiv -20\mathrm{m\,s^{-1}}$ in order to get M_e^0 from (3.4), and then

$$u_\mathrm{v}^0 \equiv (1 + z_p/z_\mathrm{top})u_\mathrm{e}^0$$

in order to get M_v^0 from (3.4), where z_top is the model height. Finally, (3.4,3.5) yields $u^0(\varphi, p)$, as shown in Fig. 3.1a.

The initial scaled potential-vorticity (Π) profile is shown in Fig. 3.1b. $\Pi \equiv P(T,\zeta)/P(T_0, 2\Omega_0)$, where

$$P(T,\zeta) \equiv \left(\sec\varphi\frac{\partial v}{\partial p}\frac{\partial}{\partial\lambda} - \frac{\partial u}{\partial p}\frac{\partial}{\partial\varphi} - \zeta\frac{\partial}{\partial p}\right)\theta(p,T)$$

approximates the Ertel potential vorticity in isobaric coordinates, $\theta(p,T) \equiv (p_0/p)^{R_0/c_p}T$ is the potential temperature and $p_0 = 10^5$Pa. In fact Π thus defined is identically the *SPV* of [Polvani(2000), section 3].

Time-Dependent Forcing at the Bottom

The model is forced by setting Φ at the surface for all time equal to

$$\Phi_\mathrm{sfc}(\lambda,\varphi,t) = \begin{cases} \Phi_\mathrm{max}\cos\lambda\left(\sin\frac{9}{2}\varphi\right)^2\left(1-e^{-t/\tau}\right), & 40^\circ < \varphi < 80^\circ, \\ 0, & \text{otherwise}, \end{cases}$$

with decay time $\tau = 3$d and amplitude $\Phi_\mathrm{max} = g_0 \times 800$m. The forcing peaks at $\varphi = 60^\circ = \varphi_\mathrm{v}$, and increases monotonically in time. The effect is to instigate upward-propagating Rossby waves.

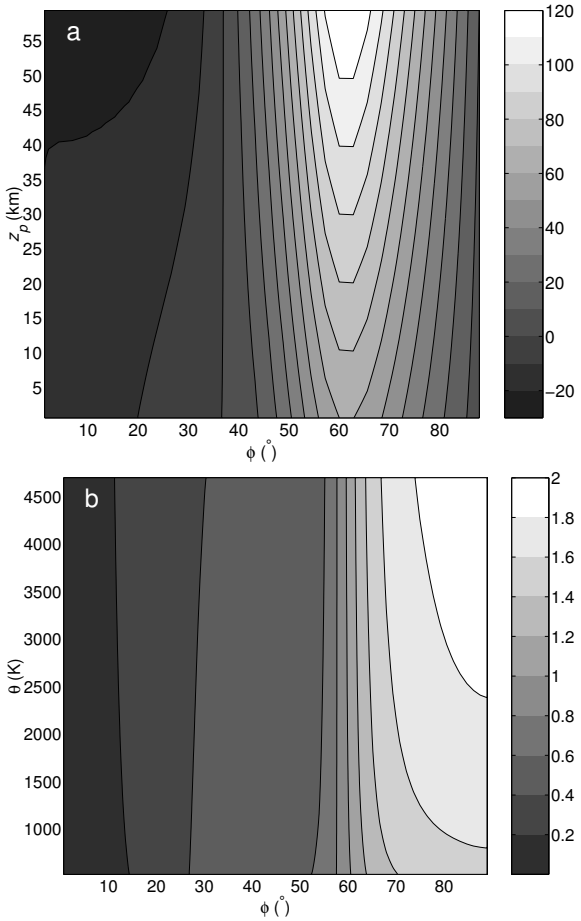

Figure 3.1. Initial state profiles vs latitude φ (°, abscissa). (a) Zonal-wind isotachs (m s^{-1}) vs log-pressure height z_p (km, ordinate). Contour interval is 10 m s^{-1} and $u(\varphi = 37°, z_p) = 0$. (b) Isopleths of scaled potential vorticity vs potential temperature θ (K, ordinate). Contour interval is 0.2 and $\Pi(\varphi = 0, \theta) = 0$.

Sponge Layer at the Top

Depending on z_{top}, spurious reflection may be prevented by adding the following sponge-layer forcing near the model top, which amounts to Rayleigh damping of $\mathbf{u}$ and Newton damping of T:

$$\left.\frac{\partial(\mathbf{u},T)}{\partial t}\right|_{\text{sl}} = \nu \cdot (\mathbf{u}^0 - \mathbf{u}, T^0 - T), \tag{3.6}$$

where the damping coefficient

$$\nu \equiv \begin{cases} \frac{1}{2} + \frac{1}{2}\tanh\frac{z_p - \bar{z}_{\text{sl}}}{\Delta z}, & z_{\text{sl}} < z_{\text{top}}, \\ 0, & z_{\text{sl}} \geq z_{\text{top}}, \end{cases} \tag{3.7}$$

$\bar{z}_{\text{sl}} \equiv 2^{-1}(z_{\text{sl}} + z_{\text{top}})$, $\Delta z \equiv \frac{3}{10}(z_{\text{top}} - z_{\text{sl}})$ and $z_{\text{sl}} = 45\text{km}$.

Results

Using SEAM, we were able to carry out the polar-vortex simulation at several resolutions, up to 36km with 200 levels. The highest resolution required $\approx 10^5$ time steps on $\approx 10^8$ collocation points in ≈ 2 wall-clock days on 256 IBM SP RS/6000 processors. The Π isosurfaces from that run are shown in in Fig. 3.2. They indicate the complex dynamical features of the polar vortex: a primary Π "tongue," succumbing to a secondary instability, leading to a roll-up into a ring of five or six smaller sub-vortices.

We look at the convergence question systematically, as shown in Fig. 3.3. It appears that at moderate horizontal resolution of 156km, some additional Π structure is resolved in increasing from 50 to 100 levels (Fig. 3.3a-b), but not as much, from 100 to 200 levels (Fig. 3.3b-c). At the higher horizontal resolution of 70km, the sub-vortices are much better resolved for 100 levels, with little change at 200 levels (Fig. 3.3d-f). Increasing the horizontal resolution to 36km at 200 levels (Fig. 3.3g) produces mainly small and qualitative changes, and only in the details, not in the overall simulation. The problem appears to have converged w.r.t. horizontal and vertical resolution.

The precise dissipation and filtering were different for all runs, but for each run were empirically minimized while sufficient to stabilize the simulation. That is, we attempt to investigate the inviscid limit. This differs from other convergence studies in which dissipation and filtering are fixed, and only resolution is varied.

To investigate convergence under mesh refinement more systematically, we consider an the isentropic Π-tongue-tip position diagnostic. This may be quantified as the point $\mathbf{x}_{\max}(\theta, t)$ of maximum angular change $\beta(\mathbf{x})$ of the $\Pi = 1$ contour in stereographic projection coordinates

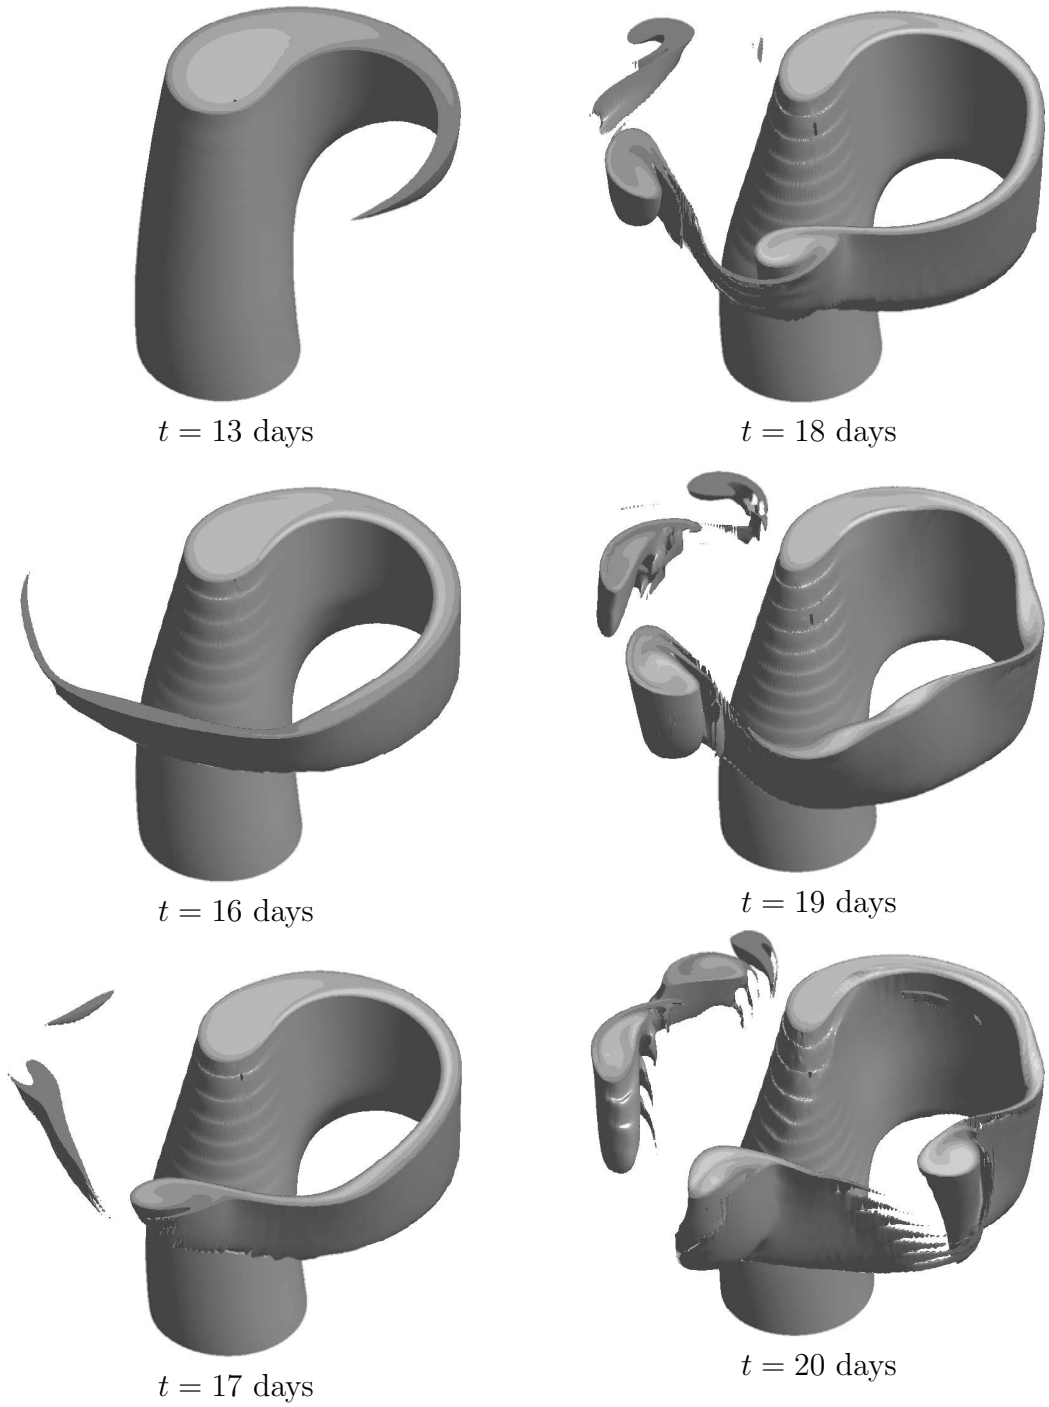

Figure 3.2. Isosurfaces of scaled potential vorticity $\Pi = 1$ for the 36km run. Horizontal coordinates have been stereographically projected (3.8), and there are 22 vertical levels $\theta_l = 10^{2/7} T_0 \exp(2.5 \frac{l+3/2}{49})$. View is towards stereographic origin and $\theta_1 = 524$K, from $\lambda = 110°$W, elevation 40°. Top plane shows contours of $\Pi > 1$ on $\theta_{22} = 1530$K.

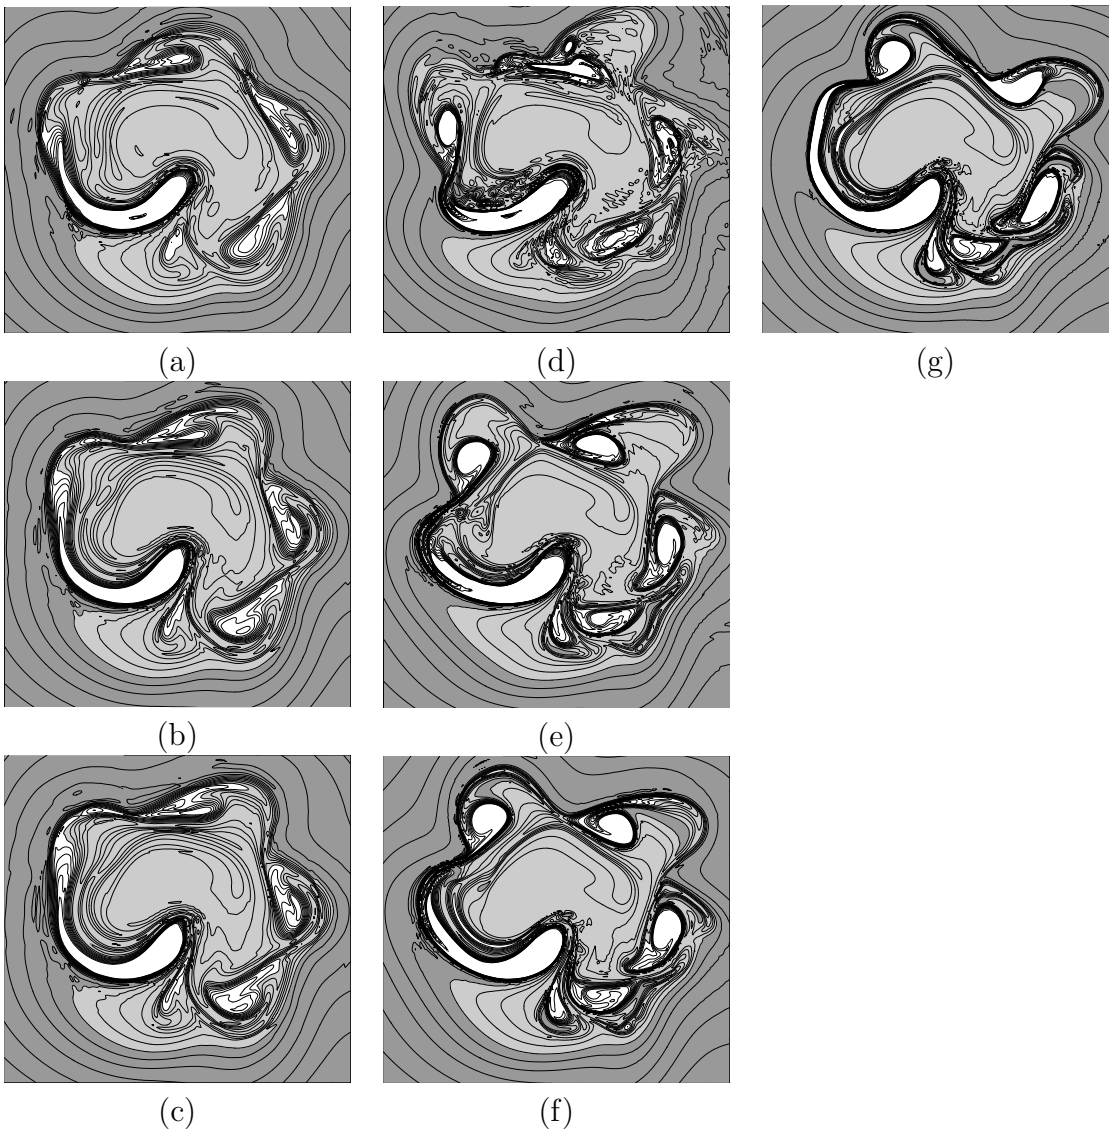

Figure 3.3. Stereographic projection in the the $\theta = 1500$K surface of $\Pi(\lambda, \varphi)$ contours from -0.3 to 1.3 by 0.1, at $t = 20$d. Values below 0.2 are dark gray and above 0.8 are white. Resolutions (a) 156km/50L, (b) 156km/100L, (c) 156km/200L; (d-f) as in (a-c) but for 70km; (g) as in (f) but for 36km.

$$\mathbf{x} \equiv \tan(\pi/4 - \varphi/2) \begin{bmatrix} \cos\lambda \\ \sin\lambda \end{bmatrix}. \tag{3.8}$$

The diagnostic involves computing the L tangent vectors $\mathbf{t}_j$ along the L points of each contour:

$$\mathbf{t}_j \equiv \Delta \begin{cases} \mathbf{x}_j, & j = 1, \\ (\mathbf{x}_j + \mathbf{x}_{j-1})/2, & j = 2, \cdots L \\ \mathbf{x}_{j-1}, & j = L+1, \end{cases}$$

where $\Delta s_j \equiv s_{j+1} - s_j$. From the $\mathbf{t}_j$ one computes the angular change

$$\beta(\mathbf{x}_j) \equiv \arccos \hat{\mathbf{t}}_j \cdot \hat{\mathbf{t}}_{j+1} \leq \beta(\mathbf{x}_{\max}(\theta, t)), \tag{3.9}$$

where $\hat{\mathbf{a}} \equiv \mathbf{a}/|\mathbf{a}|$. This is illustrated in Fig. 3.4.

The area $A(\theta, t)$ enclosed by values $\Pi \geq 1$ evolves with a qualitatively similar pattern as resolution is increased, as shown in Fig. 3.5. This is a good measure of vortex erosion, as discussed in further detail by [Polvani(2000)]. We estimated $A(\theta, t)$ by summing over northern-hemisphere points with $\Pi(\lambda, \varphi, \theta, t) > 1$:

$$A(\theta, t) = \frac{2\pi}{N_\lambda} \sum_{m=1}^{N_\varphi} \Delta \sin \varphi_m^+ \sum_{l=0}^{N_\lambda - 1} \begin{cases} 1, & \Pi(\lambda_l, \varphi_m, \theta, t) > 1, \\ 0, & \Pi(\lambda_l, \varphi_m, \theta, t) \leq 1, \end{cases} \tag{3.10}$$

where $\lambda_l \equiv (2N_\lambda^{-1} l - 1)\pi$, $\varphi_m \equiv (2N_\varphi)^{-1} m\pi$ and $\varphi_m^+ \equiv \varphi_m$ $(m = 1, \cdots N_\varphi)$, $2\varphi_m - \varphi_{m-1}$ $(m = N_\varphi + 1)$.

As an indication of convergence as resolution is increased, the ratio

$$\frac{s_{36\mathrm{km}}/s_{70\mathrm{km}} - 1}{s_{70\mathrm{km}}/s_{156\mathrm{km}} - 1}$$

for $s = A(\theta, t)$ stayed below 0.35 for all θ in Fig. 3.5, averaged over $t \in [0, 12]$d, and was usually much smaller.

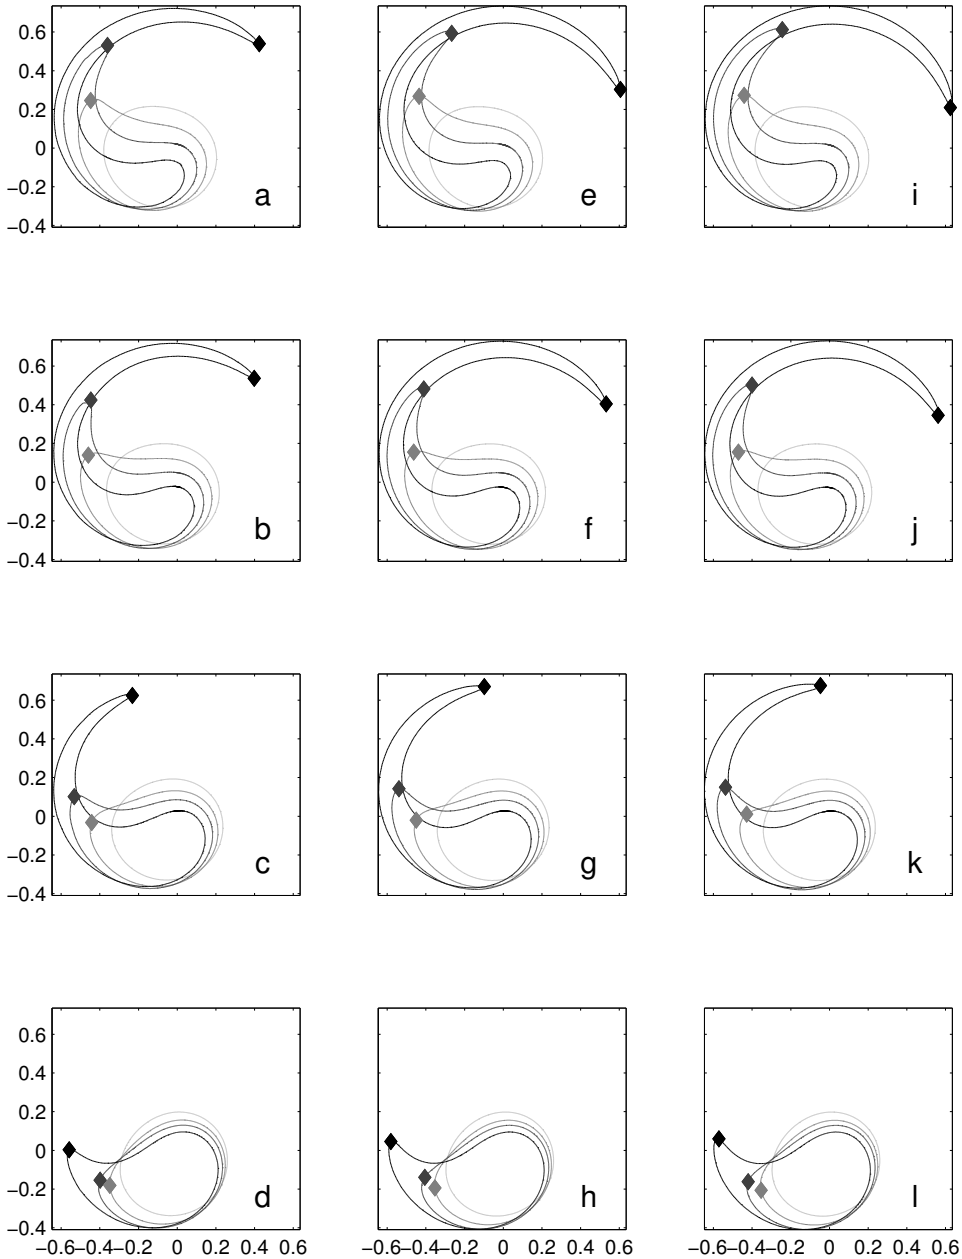

Figure 3.4. $\Pi = 1$ contours, for $t \in \{5, 7, 9, 11\}$d (light gray to black). Diamonds indicate $\mathbf{x}_{\max}(\theta, t)$ (Eq. 3.9). (a-d) Descending levels $\theta = 3000, 2500, 2000, 1500$K for 156km resolution. (e-h) As (a-d) but for 70km. (i-l) As (a-d) but for 36km.

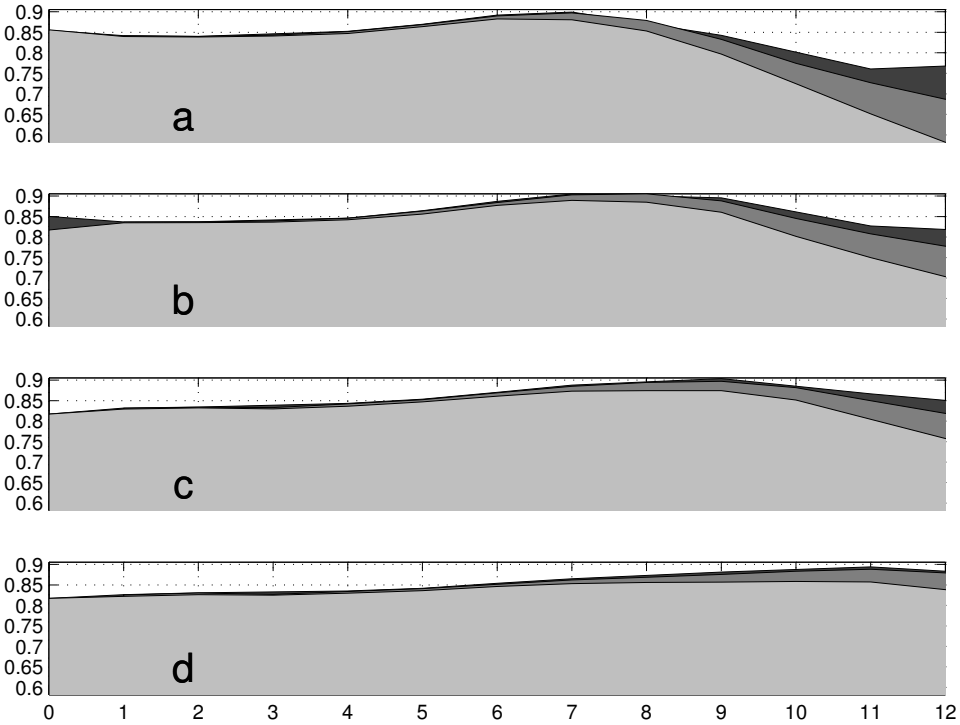

Figure 3.5. Horizontal vortex area $A(\theta, t)$ (3.10) vs t (d, abscissa) for (a) $\theta = 3000$, (b) 2500, (c) 2000, (d) 1500K. Resolutions are 156km/50L (light gray), 70km/100L (gray) and 36km/200L (dark gray).

Chapter 4

Triangular Spectral Elements

The spectral element method applied to the atmosphere greatly enhances our ability to run high-resolution simulations. But it also provides the advantage of unstructured grids, which in turn allow for statically or adaptively refined grids. In theory, these refinements could be in regions of interest, around phenomena of interest such as hurricanes, or for higher fidelity in the vicinity of sea coasts or mountain ranges. Some of these applications are best addressed with triangular elements. However, the classic spectral element method was designed for, and took advantage of, quadrilateral elements. This chapter discusses work done under this LDRD project that extends spectral element theory to triangular grids. which would be an enabling technology for refined grid applications.

Finite element methods were initially applied to self-adjoint operators (e.g., elliptic equations) but eventually found widespread use in non-self-adjoint operators such as those arising from hyperbolic equations. For many applications especially those having smooth solutions (i.e., infinitely differentiable) it is far more efficient to use high-order methods instead of low-order ones. On quadrilateral elements, the high-order accuracy is obtained by using the nodal polynomial basis generated from a tensor product of the Legendre-Gauss-Lobatto (LGL) points; these points have both good polynomial interpolation and integration (cubature) properties. This approach was introduced by Patera [Patera(1984)] and dubbed the *spectral element* method.

The importance of the LGL points to the diagonal mass matrix spectral element method (DMM SE) cannot be understated. In the square, a $(N+1)\times(N+1)$ tensor product of LGL points has a near optimal Lebesgue constant for the polynomial space $\mathcal{Q}_N = \text{span}\{\xi^n\eta^m, m, n \leq N\}$. This small Lebesgue constant means the LGL points generate a well conditioned nodal basis for $\mathcal{Q}_N$. By nodal basis, we mean the Lagrange interpolating polynomials associated with the LGL points; these basis functions are also known as cardinal functions. This nodal basis naturally separates into vertex, edge and interior modes, and can be used in a standard high p finite element method [Solin(2003)]. In addition, the LGL points which interpolate $\mathcal{Q}_N$ have a quadrature formula (cubature in more than one dimension) which will exactly integrate all polynomials in $\mathcal{Q}_{2N-1}$. The inner products which appear in the mass matrix of such a formulation will be polynomials of up to degree $2N$. Thus a high-quality approximation of these inner products (although not exact) can be obtained by evaluating the

integrals using the LGL cubature. Combining this cubature approximation with a nodal basis yields an accurate method with a diagonal mass matrix.

Unfortunately, points analogous to the LGL points do not appear to exist for the triangle. Thus spectral element methods in triangles have focused on two different approaches. In the first approach, a more traditional basis of vertex, edge and interior modes is used to construct C_0 test functions, and cubature formulas for the triangle are used to exactly evaluate the resultant inner products (see [Sherwin(1995)]); this approach results in a *modal* triangular spectral element method. In the second approach, a nodal basis is constructed using nodal sets in the triangle with a small Lebesgue constant. These points must be found numerically (see [Hesthaven(1998)] [Taylor(2000)]). They can then be coupled with exact cubature formulas, resulting in a *nodal space* approximation where two different sets of points are used for interpolation (Fekete points) and integration (Gauss points) [Warburton(2000)] [Giraldo(2005)]); this approach results in a *nodal* triangular spectral element method. Both the *modal* and *nodal* high-order triangular finite element approaches yield exponential (or spectral) convergence and for this reason are known as either *spectral elements* or *spectral/hp elements*. The difficulty that these methods face is that they both require the inversion of a sparse global mass matrix because the interpolation and integration points are not co-located as in the quadrilateral case. Thus these two triangle-based approaches cannot quite compete in terms of efficiency with quadrilateral-based spectral elements.

For this reason some attempts have been made to construct triangular spectral element (TSE) methods with diagonal mass matrices (DMM) [Cohen(2001)] [Taylor(2000b)] [Helenbrook(2004)]. The reason for developing triangular high-order methods is due to the triangle (2-simplex) being much more geometrically flexible than quadrilaterals for constructing grids especially for complex domains. In the past, to make triangular finite element methods more efficient mass lumping had been used which, as pointed out by Cohen et al. [Cohen(2001)], is related to seeking co-located interpolation and integration points. However, such points for the triangle are not very easy to derive. Cohen et al. [Cohen(1995)] obtained points for degree $N = 2$ and $N = 3$ in the triangle by enriching the polynomial space with additional interior modes that vanish at the edges and vertices of the elements and increase the cubature accuracy. Following similar ideas Mulder [Mulder(2001)] obtained points of degree $N = 4$ and $N = 5$. Because these points are constructed with integration accuracy in mind we refer to them as *cubature points*; we generalize the approach for deriving these cubature points to re-derive the sets $N = 1, ..., 5$ and derive two new sets $N = 6$ and $N = 7$ keeping in mind that higher degree sets can be achieved with this approach. The main point of this paper, however, is to show how to build numerical models for various partial differential equations using these cubature points.

Spatial Discretization

In this section we describe the spatial discretization of the equations by the DMM TSE method including: the derivation of the cubature points, the choice of basis functions, and the Vandermonde matrix used for filtering. However, before discussing these issues let us first review a few relevant points concerning interpolation and cubature.

Interpolation in the Reference Element

We start by giving some definitions and our notation in the reference element T. Let $z = (\xi, \eta)$ be a point in $\mathbf{R}^2$ and let T be the right triangle given by

$$T = \{(\xi, \eta) \mid -1 \leq \xi, \eta \leq 1;\ \xi + \eta \leq 0\}.$$

Let $\mathcal{P}_N$ denote the traditional space of all polynomials of degree $\leq N$,

$$\mathcal{P}_N = \text{span}\{\xi^n \eta^m, m + n \leq N\}.$$

Here, following [Cohen(2001)] we will also be working in an augmented space of polynomials denoted by $\mathcal{P}_{N,M}$. To construct this space, we first let $\mathcal{P}_N^0$ denote the space spanned by the interior modes,

$$\mathcal{P}_N^0 = \{f \in \mathcal{P}_N \mid f(z) = 0\ \forall z \in \partial T\}$$

where ∂T is the boundary of T. The augmented space $\mathcal{P}_{N,M}$ is then given by adding interior modes to the space $\mathcal{P}_N$,

$$\mathcal{P}_{N,M} = \mathcal{P}_N \cup \mathcal{P}_{(N+M)}^0$$

for $M \geq 0$. Thus $\mathcal{P}_{N,M}$ is the space of polynomials up to at most degree N along the boundary and up to degree $N + M$ in the interior. Note that $\mathcal{P}_N = \mathcal{P}_{N,0}$ and

$$\mathcal{P}_N \subset \mathcal{P}_{N,M} \subset \mathcal{P}_{(N+M)}.$$

Also, $\dim\mathcal{P}_N = (N+1)(N+2)/2$ and, for $N \geq 3$, $\dim\mathcal{P}_{N,M} = \dim\mathcal{P}_N + M(N-3) + (M+1)/2$.

Now consider a set of K points in the triangle T, $\{z_i = (\xi_i, \eta_i), i = 1, \ldots, K\}$, with $K = \dim\mathcal{P}_{N,M}$. If the points are non-degenerate, the nodal basis for $\mathcal{P}_{N,M}$ can be defined uniquely as the cardinal functions in $\mathcal{P}_{N,M}$ which satisfy

$$\psi_i(z) = \begin{array}{ll} 1, & \text{if} \mathrm{z} = \mathrm{z_i} \\ 0, & \text{if} \mathrm{z} = \mathrm{z_j}, \mathrm{j} \neq \mathrm{i}. \end{array} \tag{4.1}$$

The nodal basis is directly related to interpolation, since for an arbitrary function q the interpolant $I(q) \in \mathcal{P}_{N,M}$ is given by

$$I(q) = \sum_{i=1}^{K} q(z_i)\psi_i(z)$$

where $I(q)(z_i) = q(z_i)$. The quality of the interpolation operator (and thus the nodal basis) for a given space and set of points is usually measured by the Lebesgue constant $\|I\|$ (the L^∞-norm of the interpolation operator).

Cubature in the Reference Element

Now we turn to the cubature properties of the points z_i. A cubature rule for these points and weights w_i is of strength d if and only if

$$\int_T g = \sum_{i=1}^{K} w_i g(z_i) \quad \forall g \in \mathcal{P}_d. \tag{4.2}$$

Any set of $K = \dim\mathcal{P}_{N,M}$ non-degenerate points $\{z_i\}$ will have a cubature rule of at least strength N by using the generalized Newton-Cotes weights,

$$w_i = \int_T \psi_i.$$

Because of Eq. (4.1), this choice gives a quadrature rule that exactly integrates the cardinal functions ψ_i. Since these functions are a basis for $\mathcal{P}_{N,M}$ the cubature rule is exact for all $g \in \mathcal{P}_{N,M}$ and thus for all $g \in \mathcal{P}_N$ and so the cubature rule is at least of strength N. However, for some point sets, the Newton-Cotes quadrature formula will integrate a larger space and the cubature rule can be of strength $2N$ or higher.

As mentioned in the introduction, for the 1-simplex and its tensor products, the LGL points of degree N are both optimal interpolation and cubature points. Thus a highly accurate method can be developed which results in a diagonal mass matrix since the interpolation and integration points are co-located; having a diagonal mass matrix is important for achieving efficiency. On the 2-simplex such points have not yet been found and thus far one must be content to choose either good interpolation or integration but not both.

It should be noted that in previous works the Fekete points have been used as integration points which then results in a diagonal mass matrix [Taylor(2000b)] [Komatitsch(1998)]; however, the cubature rule for Fekete points of degree N is only of strength N, and thus using these cubature points is a poor approximation to the degree $2N$ inner products which appear in the integral formulation of the equations. For some problems, cubature of strength N is insufficient to achieve exponential convergence. The fact that the interpolation and integration points are co-located means that the mass matrix will be diagonal and thereby will be trivial to invert; a non-diagonal mass matrix has been a thorn in the side of triangular spectral element methods and is the main reason why triangles have not been able to compete with quadrilaterals in terms of efficiency.

From this discussion it seems logical to assume that to get around the current dilemma requires the construction of a new set of points for the triangle that have cubature

strength of close to $2N$ while having reasonably low Lebesgue constants. In the next section we describe the approach used for constructing points that satisfy both of these two criteria. These points offer a similar accuracy to the Fekete-Gauss method while having the attractive property of the Fekete points, that is, a diagonal mass matrix.

Computing Cubature Points

Finding points with optimal cubature properties has been extensively studied independently of spectral element applications and has a long history of both theoretical and numerical development. For a recent review, see [Cools(1993)] [Lyness(1994)] [Cools(1997)] and [Cools(1999)]. An on-line database containing many of the best known quadrature formulas is described in [Cools(2003)]. One successful approach for numerically finding quadrature formulas dates to [Lyness(1975)]. A generalized version was used recently in [Wandzura(2003)]. Newton's method is used to solve the nonlinear system of algebraic equations for the quadrature weights and locations of the points. Symmetry is used to reduce the complexity of the problem. The complexity can be further reduced with a cardinal function based algorithm [Taylor(2006)].

If one consults the database and the newest numerical results in the above references, it appears that cubature points of degree N with strength $2N$ or even $2N-1$ do not exist in T. Furthermore, the best cubature formulas for T do not have sufficient points (if any) on the boundary of T. If a set of cubature points will also be used to construct the nodal basis in the space $\mathcal{P}_{N,M}$ there must be $N+1$ points along each edge of T. This is because the nodal basis must generate vertex, edge and interior modes, and the interior modes must be uniformly zero on ∂T (this is critical for the construction of C^0 approximations across element boundaries). If a cardinal function $\psi_i \in \mathcal{P}_{N,M}$ is zero on $N+1$ points on an edge of T then $\psi_i(z) = 0$ for any point z on that edge. This is because the restriction of any function in $\mathcal{P}_{N,M}$ to that edge is a polynomial of degree N. Thus any ψ_i for z_i in the interior of T is automatically an interior mode.

Since the goal of achieving a cubature formula using points which interpolate $\mathcal{P}_N$ and are of strength $2N-1$ is not achievable, we thus follow the ideas first put forth by Cohen et al. [Cohen(1995)] [Cohen(2001)] and Mulder [Mulder(2001)]. Instead of working in $\mathcal{P}_N$, we use the enriched space $\mathcal{P}_{N,M}$ with $M > 0$. A point set which can interpolate $\mathcal{P}_{N,M}$ necessarily contains more points than those for $\mathcal{P}_N$. Thus there are more degrees of freedom which can be used to satisfy the system of cubature equations given in Eq. (4.2). The resulting points, which we call $\mathcal{P}_{N,M}$ cubature points, are tabulated and compared against the two point sets used in the Fekete-Gauss method in Table 4.1. The cubature points for $N = 2$ and $N = 3$ are due to Cohen et al. [Cohen(1995)] [Cohen(2001)] and degrees $N = 4$ and $N = 5$ are due to Mulder [Mulder(2001)].

To extend these results, we used the cardinal function cubature algorithm of [Taylor(2006)]. Straightforward modifications were required, first to work with interpolation points for $\mathcal{P}_{N,M}$ instead of $\mathcal{P}_N$, and second to impose that there be $N+1$ points along each edge of T (for a total of $3N$ points along the perimeter). We were able to reproduce the previously computed results for degrees $N \leq 5$ along with the two new degrees $N = 6$ and $N = 7$. In fact, higher degrees of N for either triangles or tetrahedra are achievable with this procedure.

Fekete-Gauss in $\mathcal{P}_N$					Cubature in $\mathcal{P}_{N,M}$				
N	K	K'	d	$\parallel I \parallel$	N	M	$K = K'$	d	$\parallel I \parallel$
1	3	3	2	1.00	1	0	3	1	1.00
2	6	6	4	1.67	2	1	7	3	1.45
3	10	12	6	2.11	3	1	12	5	2.21
4	15	16	8	2.58	4	1	18	7	3.75
5	21	25	10	3.19	5	2	30	10	5.23
6	28	36	12	4.08	6	3	46	12	7.40
7	36	46	14	4.78	7	3	51	14	7.50

Table 4.1. Properties of the cubature points and the points used in the Fekete-Gauss method. N and M determine the polynomial space $\mathcal{P}_{N,M}$, K is the number of interpolation points, K' the number of integration points, d the strength of the integration points and $\parallel I \parallel$ is the Lebesgue constant of the interpolation points.

Table 4.1 shows the number of interpolation points, K, integration points, K', and the Lebesgue constant, $\parallel I \parallel$, for both the Fekete-Gauss and cubature points as a function of the polynomial degree, N. The first thing to notice about these two sets of points is that the Fekete-Gauss points have superior Lebesgue constants than the cubature points. However, this should not be surprising since the Fekete points are constructed in order to minimize interpolation error and the Lebesgue constant is a measure of the quality of the points to achieve good interpolation (a lower Lebesgue constant implies better interpolation).

Evaluating the Nodal Basis Functions

To define the local operators which shall be used to construct the global approximation of the solution we begin by decomposing the domain Ω into N_{e} non-overlapping triangular elements Ω_e such that

$$\Omega = \bigcup_{\mathrm{e}=1}^{N_{\mathrm{e}}} \Omega_{\mathrm{e}}.$$

We then further map the arbitrary triangles Ω_{e} into the reference right triangle T. To perform differentiation and integration operations, we introduce the nonsingular mapping $\boldsymbol{x} = \Psi(\boldsymbol{\xi})$ which defines a transformation from the physical Cartesian coordinate system, $\boldsymbol{x}$, within each triangle Ω_e to a local reference coordinate system, $\boldsymbol{\xi}$, in the reference right triangle T (see [Giraldo(2005)] for details on this mapping for curved elements in $\mathbf{R}^3$).

Let us now represent the local element-wise solution $\boldsymbol{q}$ by its nodal expansion in $\mathcal{P}_{N,M}$ as

$$\boldsymbol{q}(\boldsymbol{\xi}) = \sum_{k=1}^{K} \boldsymbol{q}(\boldsymbol{\xi}_k)\psi_k(\boldsymbol{\xi})$$

where $K = \mathrm{dim}\mathcal{P}_{N,M}$ and the nodal basis functions ψ_k are defined as in Eq. (4.1). For the points (ξ_i, η_j) we choose the newly derived cubature points which, while trivial to evaluate nodal basis functions, complicate any other operation such as the evaluation of their derivatives. Here we follow [Taylor(2000)] and represent the nodal basis functions in terms of the Proriol polynomials [Proriol(1957)] [Dubiner(1991)] which in T are an orthogonal basis for $\mathcal{P}_{N,0}$. Stable recurrence relations which can be used to evaluate these polynomials and their derivatives are given in [Karniadakis(1999)]. Note that these polynomials are traditionally denoted with a double index (m, n) representing the top degree in ξ and η. But here we use a single index and denote them by $\varphi_i, \quad i = 1, \ldots, \mathrm{dim}\mathcal{P}_N$. To generate an easily computable basis for $\mathcal{P}_{N,M}$ we need to add additional basis functions for the interior modes in $\mathcal{P}_{N,M}$ which are not in $\mathcal{P}_N$. For these polynomials we use the formulas given in [Karniadakis(1999)] [Solin(2003)], and we denote them by $\varphi_i, \quad i = (\mathrm{dim}\mathcal{P}_N + 1), \ldots, K$.

Filtering the High-Frequency Waves

As with most high-order methods, when solving non-linear problems some filtering is needed to control the accumulation of aliasing errors. The ability to selectively filter only the highest wave numbers is an advantage of the spectral element method. However it does require that we use an expansion in only orthogonal Proriol polynomials, since a nodal expansion or an expansion that involves interior modes will not be orthogonal and thus not isolate the high frequency content to only the high wave number modes.

Thus to implement filters, we need to compute the expansion of the local element-wise solution $\boldsymbol{q}$ in terms of only orthogonal Proriol polynomials. In order to prevent confusion with the augmented basis for $\mathcal{P}_{N,M}$, here we denote the basis for $\mathcal{P}_{(N+M)}$ by $g_i, i = 1, \ldots, K_1$, where $K_1 = \mathrm{dim}\mathcal{P}_{(N+M)}$.

Since $q \in \mathcal{P}_{N,M} \subset \mathcal{P}_{(N+M)}$, q has a unique expansion in terms of g. Denoting this

expansion by

$$q(\xi_i, \eta_i) = \sum_{k=1}^{K} \tilde{q}_k R_{i,k} \tag{4.3}$$

where $\tilde{q}_k$ are the Proriol coefficients of q and R is the rectangular Vandermonde matrix for the basis g, allows us to define the $K \times K_1$ rectangular Vandermonde matrix R by

$$R_{i,k} = g_k(\xi_i, \eta_i).$$

Note that the appropriate right inverse of R for filtering is constructed as follows: given the set of point values $q(\xi_i, \eta_i)$, the expansion in terms of φ (the basis for $\mathcal{P}_{N,M}$) can be computed by applying V^{-1}, the inverse of the Vandermonde matrix. This polynomial is in $\mathcal{P}_{(N+M)}$ and thus has a unique expansion in terms of g (the basis for $\mathcal{P}_{(N+M)}$), and so we have a mapping from the set of point values $q(\xi_i, \eta_i)$ to the expansion coefficients $\tilde{q}$; let us denote this map by R^r. Applying Eq. (4.3) to expansion coefficients computed with R^r must recover the original grid point values, and thus $RR^r = I$ where I is the identity matrix. In matrix form, we now write Eq. (4.3) as

$$\mathsf{q} = \mathsf{R}\ \tilde{\mathsf{q}} \qquad \tilde{\mathsf{q}} = \mathsf{R}^r\ \mathsf{q}. \tag{4.4}$$

We can now apply filters to q directly to its Proriol coefficients in $\mathcal{P}_{(N+M)}$. There are many possible filters, but here, based on past experience [Taylor(1997)] [Giraldo(2005)], we choose the Boyd-Vandeven transfer function [Boyd(1996)] which we denote by $\mathsf{\Lambda}$. Applying the filter to the amplitudes and then transforming to nodal (physical) space is achieved in the following matrix-vector multiply operation

$$\mathsf{q}^F = \mathsf{F}\ \mathsf{q}$$

where

$$\mathsf{F} = \mathsf{R}\ \mathsf{\Lambda}\ \mathsf{R}^r \tag{4.5}$$

is the $K \times K$ filter matrix and is applied every time-step at full strength. However, in the original Boyd-Vandeven filter only the highest modes (polynomials of degree $N+M$) are completely annihilated and this may be too severe for either quadrilateral spectral elements or exact integration triangular spectral elements. Thus the issue with the new cubature points is how to view the approximation space which these points span. For example, should one take the degree to be the degree of the edge modes N, or the degree of the interior modes, $N + M$? Let us denote the order of the space which the filter acts on as N_F. In Table 4.2 we show the values used for filtering the cubature points. As an example, for $\mathcal{P}_{7,3}$ the interpolation functions of the cubature points contain some degree 10 (N+M) modes. The filter then acts on all the degree 9 and 10 modes. Unfortunately, we currently have no theory to support our filter choices; the values listed in Table 4.2 were found experimentally but they appear to work for a variety of applications.

Filtering for $\mathcal{P}_{N,M}$

N	M	K	N_F
1	0	3	2
2	1	7	3
3	1	12	4
4	1	18	5
5	2	30	6
6	3	46	8
7	3	51	8

Table 4.2. The indices N, M for the polynomial space, the number of cubature points K, and the highest mode unaffected by the filter, N_F.

Integration and Local Element-wise Operators

In order to complete the discussion of the local element-wise operations required to construct discrete spectral element operators we must lastly describe the integration procedure required by the weak formulation of all Galerkin methods. For any two functions f and g the integration $\mathcal{I}$ proceeds as follows

$$\mathcal{I}[f,g] = \int_{\Omega_e} f(\boldsymbol{x})\, g(\boldsymbol{x}) d\boldsymbol{x} = \sum_{i=1}^{K} w_i \mid J(\boldsymbol{\xi}_i) \mid f(\boldsymbol{\xi_i})\, g(\boldsymbol{\xi_i}) \tag{4.6}$$

where w are the cubature weights and J the Jacobians of the transformation from physical space to the local space of the reference element. Note that for straight-edged triangles J and all other metric terms are constant but not so for curved elements.

To simplify the description of the numerical algorithm, let us define the following local element operators: let

$$M^e_{ij} = \int_{\Omega_e} \psi_i(\boldsymbol{x})\psi_j(\boldsymbol{x}) d\boldsymbol{x}, \tag{4.7}$$

$$L^e_{ij} = \int_{\Omega_e} \operatorname{grad} \psi_i(\boldsymbol{x}) \cdot \operatorname{grad} \psi_j(\boldsymbol{x}) d\boldsymbol{x}, \tag{4.8}$$

$$\boldsymbol{D}^e_{ij} = \int_{\Omega_e} \psi_i(\boldsymbol{x}) \boldsymbol{\nabla} \psi_j(\boldsymbol{x}) d\boldsymbol{x}, \tag{4.9}$$

represent the mass, Laplacian, and differentiation matrices where $i, j = 1, \ldots, K$.

Using Eq. (4.6) allows us to rewrite Eqs. (4.7), (4.8), and (4.9) as

$$M^e_{ij} = w_i \mid J(\boldsymbol{\xi}_i) \mid \delta_{ij}, \tag{4.10}$$

$$L_{ij}^e = \sum_{i=1}^{K} w_i \mid J(\boldsymbol{\xi}_i) \mid \operatorname{grad} \psi_i(\boldsymbol{\xi}_i) \cdot \operatorname{grad} \psi_j(\boldsymbol{\xi}_i), \tag{4.11}$$

$$\boldsymbol{D}_{ij}^e = w_i \mid J(\boldsymbol{\xi}_i) \mid \boldsymbol{\nabla} \psi_j(\boldsymbol{\xi}_i), \tag{4.12}$$

where δ_{ij} is the Kronecker delta function. As Eq. (4.10) shows, having the interpolation and integration points, $\boldsymbol{\xi}_i$, co-located results in the mass matrix, M^e, being diagonal which significantly simplifies the construction of the global matrix problem and its solution.

Formulation of the Method

We follow the generalized diagonal mass matrix spectral element formulation outlined in [Taylor(2000b)]. This method is based directly on the original DMM spectral element method proposed by [Maday(1987)] except that the domain is decomposed into triangular instead of quadrilateral elements and we are working in the space $\mathcal{P}_{N,M}$. The method proceeds as follows: by piecing together appropriate nodal basis functions in neighboring elements, a set of global functions can be constructed that are C_0 and piecewise polynomial. These global functions are used as the test functions in the weak form of the equations of interest, and the unknowns are expanded in terms of these global functions. The resultant integral equations are decomposed as a sum of integrals over each element. Finally, the quadrature rule associated with the points used to construct the cardinal function basis is used to evaluate each integral. Because of the nodal nature of the global functions, this results in much simplification, including giving a diagonal mass matrix.

The result is a method which can satisfy the equations globally by simply summing the local element matrices, Eqs. (4.10), (4.11), and (4.12), to form their global representation [Maday(1987)]. This summation procedure is known as the global assembly or direct stiffness summation. Let us represent this direct stiffness summation (DSS) procedure by the summation operator

$$\bigwedge_{\mathrm{e}=1}^{N_\mathrm{e}}$$

with the mapping $(i, e) \longrightarrow (I)$ where $i = 1, \ldots, K$ are the local element grid points, $e = 1, \ldots, N_\mathrm{e}$ are the spectral elements covering the global domain, and $I = 1, \ldots, N_\mathrm{p}$ are the global grid points. Applying the DSS operator to the local element matrices results in the following global matrices:

$$M = \bigwedge_{\mathrm{e}=1}^{N_\mathrm{e}} M^e, \quad L = \bigwedge_{\mathrm{e}=1}^{N_\mathrm{e}} L^e, \quad \boldsymbol{D} = \bigwedge_{\mathrm{e}=1}^{N_\mathrm{e}} \boldsymbol{D}^e$$

where M, L, and D are matrices of dimension $N_p \times N_p$ and M is diagonal and thereby trivial to invert.

Poisson Equation on the Plane

For the Poisson equation

$$\nabla^2 q = f \tag{4.13}$$

we define its variational statement as: find $q \in H_0^1(\Omega)\ \forall\ \psi \in H^1$ such that

$$-Lq = f \tag{4.14}$$

where $H_0^1(\Omega)$ is the space of all functions (with zero Dirichlet boundary conditions) with functions and first derivatives belonging to $L^2(\Omega)$ - the space of all functions that are square integrable over Ω. The domain used for this test is $\boldsymbol{x} \in [0,1] \times [0,1]$.

Advection Equation on the Sphere

Similarly, for the advection equation on the sphere

$$\frac{\partial q}{\partial t} + \boldsymbol{u} \cdot \operatorname{grad} q = 0 \tag{4.15}$$

we define the variational statement as: find $q \in H^1(\Omega)\ \forall\ \psi \in H^1$ such that

$$\frac{\partial q}{\partial t} = -M^{-1}\boldsymbol{u}^T \boldsymbol{D} q. \tag{4.16}$$

On the sphere, however, no additional boundary conditions are required other than periodicity which is satisfied by the connectivity of the grid.

Shallow Water Equations on the Sphere

The shallow water equations on the sphere are

$$\frac{\partial \boldsymbol{q}}{\partial t} = S(\boldsymbol{q}) \tag{4.17}$$

$$S(\boldsymbol{q}) = -\begin{pmatrix} \operatorname{grad} \cdot (\phi \boldsymbol{u}) \\ \boldsymbol{u} \cdot \operatorname{grad} \boldsymbol{u} + f\,(\boldsymbol{x} \times \boldsymbol{u}) + \operatorname{grad} \phi + \mu \boldsymbol{x} \end{pmatrix} \tag{4.18}$$

where $\boldsymbol{q} = (\phi, \boldsymbol{u}^T)^T$, the nabla operator is defined as $\boldsymbol{\nabla} = (\partial_x, \partial_y, \partial_z)^T$, ϕ is the geopotential height ($\phi = gh$ where g is the gravitational constant and h is the vertical height of the fluid), $\boldsymbol{u} = (u, v, w)^T$ is the Cartesian wind velocity vector, $f = \frac{2\omega z}{a}$ is the Coriolis parameter and (ω, a) represent the rotation of the earth and its radius, respectively.

The term $\mu \boldsymbol{x}$, where $\boldsymbol{x} = (x, y, z)^T$ is the position vector of the grid points, is a fictitious force introduced to constrain the fluid particles to remain on the surface of the sphere (see [Giraldo(2005)] for details). Note that this equation set represents

an initial value problem with no boundary conditions; the only condition required is that of periodicity which is imposed by the geometry of the spherical domain.

The variational statement of the problem is: find $(\phi, \boldsymbol{u}^T)^T \in H^1(\Omega)\ \forall\ \psi \in H^1$ such that

$$\frac{\partial \phi}{\partial t} = -M^{-1}\boldsymbol{D}^T\left(\phi \boldsymbol{u}\right) \tag{4.19}$$

$$\frac{\partial \boldsymbol{u}}{\partial t} = -M^{-1}\boldsymbol{u}^T\boldsymbol{D}\boldsymbol{u} - f\left(\boldsymbol{x} \times \boldsymbol{u}\right) - M^{-1}\boldsymbol{D}\phi - \mu\boldsymbol{x} \tag{4.20}$$

where for ϕ and $\boldsymbol{u}$ we choose the polynomial space $\mathcal{P}_{N,M}$-$\mathcal{P}_{N,M}$.

Chapter 5

Conclusions

The spectral element method is known to perform well on parallel computers. This performance is maintained for a spectral element atmospheric model on very large, previously unattained, processor counts. The parallel efficiency remains between 70–80%, for parallel decompositions as fine as one element per processor and for all problem sizes and processor counts. These results give evidence that the massively parallel systems being built now (IBM's BlueGene/L and Cray's Red Storm) will be able to sustain between 4-10 TFLOPS of performance, allowing one to run 10km global atmospheric simulations at a rate of 22–50 simulated days per day. This establishes a performance level for atmospheric modeling competitive with that obtained using the specialized vector supercomputer architectures of the Japanese Earth Simulator. On that machine, the AFES (Atmospheric model for the Earth Simulator) obtained 27 TFLOPS of performance and an integration rate of 57 simulated days per day [Shingu(2002)]. AFES uses a global spectral model for its dynamical core, and thus requires more flops to achieve similar integration rates.

The semi-implicit time-stepping approach is promising, as we have evidence that it results in faster simulated days per day integration rates, with typical speedups in the range of 2-3. However, in practice we see a failure to converge approximately 4% of the time, therefore this approach remains experimental, and the subject of continued research.

The parallel performance of the spectral element method allowed us to conduct a mesh convergence study using a polar vortex model problem. This problem is the focus of much recent research. It has a complex unstable evolution dominated by strong potential-vorticity gradients. Only at high resolution (36km and 200 vertical levels) does evidence for mesh convergence of large scale features start to become apparent.

We have also presented a promising new set of cubature points on the triangle which can be used both for interpolation and integration. The fact that the interpolation and integration are co-located means that the points yield a diagonal mass matrix (DMM). We presented implementation strategies of this DMM triangular spectral element (TSE) method for elliptic and hyperbolic equations on planar and spherical surfaces. We compared the cubature points to the usual strategy of employing one set of points for interpolation (Fekete) and another set for integration (Gauss) which

we call Fekete-Gauss TSE; this method, however, does not yield a diagonal mass matrix. The cubature points compared quite favorably with the Fekete-Gauss points especially for the system of nonlinear hyperbolic equations. In addition, these good solutions can be obtained quite cheaply with the cubature method because they yield a diagonal mass matrix which allows for simple and efficient time-integration strategies; this may now allow the triangular SE method to compete in terms of both accuracy and efficiency with the quadrilateral SE method while allowing far more flexibility in the choice of adaptive unstructured grids.

References

[Briegleb(2004)] B. P. Briegleb, C. M. Bitz, E. C. Hunke, W. H. Lipscomb, M. M. Holland, J. L. Schramm and R. E. Moritz (2004). Scientific Description of the Sea Ice Component in the Community Climate System Model, Version Three. *NCAR Technical Note TN-463+STR*, 78pp.

[Boyd(1996)] J. Boyd (1996). The erfc-log filter and the asymptotics of the Euler and Vandeven sequence accelerations. in A. V. Ilin and L. R. Scott (eds), *Proceedings of the Third International Conference on Spectral and High Order Methods*, Houston Journal of Mathematics, Houston, Texas, 267–276.

[Cohen(1995)] G. Cohen, P. Joly, and N. Tordjman (1995). Higher order triangular finite elements with mass lumping for the wave equation, in *Proc. 3rd Int. Conf. on Mathematical and Numerical Aspects of Wave Propagation*. Eds. G. Cohen, E. Becache, P. Joly, and J.E. Roberts, SIAM, Philadelphia.

[Cohen(2001)] G. Cohen, P. Joly, J.E. Roberts, and N. Tordjman (2001). Higher order triangular finite elements with mass lumping for the wave equation. *SIAM Journal of Numerical Analysis* **38** 2047–2078.

[Collins(2004)] Collins, W. D., Rasch, P. J., Boville, B. A., Hack, J. J., McCaa, J. R., Williamson, D. L., Kiehl, J. T., Briegleb, B., Bitz, C., Lin, S.-J., Zhang, M., and Dai, Y. (2004). Description of the ncar community atmosphere model (cam 3.0). Technical Report TN–464+STR, NCAR.

[Cools(1993)] R. Cools and P. Rabinowitz (1993). Monomial cubature rules since *Stroud*: a compilation. *J. Comp. and Appl. Math.* **48** 309–326.

[Cools(1997)] R. Cools (1997). Constructing cubature formulae: the science behind the art. Acta Numerica **6** 1–54.

[Cools(1999)] R. Cools (1999). Monomial cubature rules since *Stroud*: a compilation—part 2. *J. Comp. and Appl. Math.* **112** 21–27.

[Cools(2003)] R. Cools (2003). An encyclopaedia of cubature formulas. *Journal of Complexity* **19** 445–453. Online database located at http://www.cs.kuleuven.ac.be/~nines/research/ecf/ecf.html.

[Dennis(2003)] Dennis, J. M. (2003). Partitioning with space-filling curves on the cubed-sphere. In *Proceedings of the Workshop on Massively Parallel Processing at IPDPS'03*. Nice, France.

[Dubiner(1991)] M. Dubiner (1991). Spectral methods on triangles and other domains. *Journal of Scientific Computing* **6** 345–390.

[Fournier(2004)] Fournier, A., Taylor, M. A., and Tribbia, J. (2004). The spectral element atmosphere model (SEAM): High-resolution parallel computation and localized resolution of regional dynamics. *Mon. Wea. Rev.*, **132**, 726–748.

[Giraldo(2001)] Giraldo, F. X. (2001). A spectral element shallow water model on spherical geodesic grids. *International Journal for Numerical Methods in Fluids*, **35**, 869–901.

[Giraldo(2004)] Giraldo, F. X. and Rosmond, T. E. (2004). A scalable spectral element eulerian atmospheric model (see-am) for nwp: Dynamical core tests. *Mon. Wea. Rev.*, **132**, 133–153.

[Giraldo(2005)] F.X. Giraldo, and T. Warburton (2005). A nodal triangle-based spectral element method for the shallow water equations on the sphere. *Journal of Computational Physics* 207 129–150.

[Haidvogel(1997)] Haidvogel, D., Curchitser, E. N., Iskandarani, M., Hughes, R., and Taylor, M. A. (1997). Global modeling of the ocean and atmosphere using the spectral element method. *Ocean-Atmosphere*, **35**.

[Held(1994)] Held, I. M. and Suarez, M. J. (1994). A proposal for the intercomparison of the dynamical cores of atmospheric general circulation models. *Bull. Amer. Meteor. Soc.*, **75**, 1825–1830.

[Helenbrook(2004)] B. Helenbrook (2004). Polynomial Bases Suitable for Mass Lumping on Triangles. *International Conference on Spectral and High-Order Methods (ICOSAHOM).*

[Hesthaven(1998)] J.S. Hesthaven (1998). From electrostatics to almost optimal nodal sets for polynomial interpolation in a simplex. *SIAM Journal on Numerical Analysis* **35** 655–676.

[Hunke(2006)] Elizabeth Hunke and William H. Lipscomb *CICE: the Los Alamos Sea Ice Model Documentation and Software User's Manual*, Draft available at http://climate.lanl.gov/Models/CICE.

[Iskandarani(2002)] Iskandarani, M., Haidvogel, D., Levin, J., Curchitser, E. N., and Edwards, C. A. (2002). Multiscale geophysical modeling using the spectral element method. *Computing in Science and Engineering*, **4**.

[Iskandarani(2003)] Iskandarani, M., Haidvogel, D., and Levin, J. (2003). A three-dimensional spectral element model for the solution of the hydrostatic primitive equations. *Journal of Computational Physics*, **186**.

[James(2000)] James, P. M., Peters, D., and Waugh, D. W. (2000). Very low ozone episodes due to polar vortex displacement. *Tellus*, **52B**(4), 1123–1137.

[Karniadakis(1999)] G.E. Karniadakis and S. J. Sherwin (1999). *Spectral/hp Element Methods for CFD*, Oxford Univ Press.

[Kauffman(2004)] B. G. Kauffman, R. Jacob, T. Craig and W. G. Large (2004). *The CCSM Coupler, Version 6.0—User's Guide, Source Code Reference, and Scientific Description*, Draft available at http://www.ccsm.ucar.edu/models/ccsm3.0/cpl6.

[Kiehl(2004)] Kiehl, J. T. and Gent, P. R. (2004). The community climate system model, version two. *J. Clim.*, **17**, 3666–3682.

[Koh(2000)] Koh, T.-Y. and Plumb, R. A. (2000). Lobe dynamics applied to barotropic Rossby-wave breaking. *Phys. Fluids*, **12**(6), 1518–1528.

[Komatitsch(1998)] D. Komatitsch, R. Martin, J. Tromp, M.A. Taylor, B.A. Wingate (2001). Wave propagation in 2-D elastic media using a spectral element method with triangles and quadrangles. *J. Comp. Acoustics* **9** 703–718.

[Komatitsch(1999)] Komatitsch, D. and Tromp, J. (1999). Introduction to the spectral-element method for 3-d seismic wave propagation. *Geophys. J. Int.*, **139**, 806–822.

[Komatitsch(2002)] Komatitsch, D. and Tromp, J. (2002). Spectral-element simulations of global seismic wave propagation - i. validation. *Geophys. J. Int.*, **149**, 390–412.

[Komatitsch(2003)] Komatitsch, D., Tsuboi, S., Ji, C., and Tromp, J. (2003). A 14.6 billion degrees of freedom, 5 teraflops, 2.5 terabyte earthquake simulation on the earth simulator. In *Proceedings of the ACM / IEEE Supercomputing SC'2003 conference*.

[Limpasuvan(2000)] Limpasuvan, V. and Hartmann, D. L. (2000). Wave-maintained annular modes of climate variability. *J. Climate*, **13**, 4414–4429.

[Lin(1997)] S.-J. Lin and R. B. Rood (1997). An explicit flux-form semi-Lagrangian shallow water model on the sphere. *Q. J. R. Meteorol. Soc.*, **123**, 2531–2533.

[Loft(2001)] Loft, R., Thomas, S., and Dennis, J. (2001). Terascale spectral element dynamical core for atmospheric general circulation models. In *Proceedings of the ACM / IEEE Supercomputing SC'2001 conference*.

[Lyness(1975)] J. Lyness and D. Jespersen (1975). Moderate degree symmetric quadrature rules for the triangle. *J. Inst. Math. Appl.* **15** 19–32.

[Lyness(1994)] J. Lyness and R. Cools (1994). A survey of numerical cubature over triangles. *Applied Mathematics* **48** 127–150.

[Machenhauer(1972)] B. Machenhauer and R. Daley (1972). A Baroclinic Primitive Equation Model with Spectral Representation in Three Dimensions. Inistitute for Theoretical Meteorology, Copenhagen University, **4**.

[Maday(1987)] Y. Maday, and A.T. Patera (1987). Spectral element methods for the incompressible Navier Stokes equations. in: A. K. Noor, J. T. Oden (Eds.), *State of the Art Surveys on Computational Mechanics*, ASME, New York, 71–143.

[Malone(2004)] Malone, R. C., Drake, J. B., Jones, P. W., and Rotman, D. A. (2004). Computing the climate. In *A Science-based Case for Large-scale Simulation, Volume II.* to appear in the SIAM series on Computational Science and Engineering. Draft available at http://www.pnl.gov/scales.

[Molcard(2002)] Molcard, A., Pinardi, N., Iskandarani, M., and Haidvogel, D. (2002). Wind driven circulation of the mediterranean sea simulated with a spectral element ocean model. *Dynamics of Atmospheres and Oceans*, **35**.

[Mulder(2001)] W.A. Mulder (2000). Higher-order mass-lumped finite elements for the wave equation. *Journal of Computational Acoustics* **9** 671–680.

[Nair(2004)] Nair, R. D., Thomas, S., and Loft, R. (2004). A discontinuous galerkin transport scheme on the cubed-sphere. *under review, Mon. Wea. Rev.*

[Oleson(2004)] Keith W. Oleson, *et.al* (2004). Technical Description of the Community Land Model (CLM). *NCAR Technical Note TN-461+STR*, 174pp.

[Oliger(1970)] J.E. Oliger, R.E. Wellck, A. Kasahara, and W.M. Washington (1970). Description of NCAR Global Circulation Model. *NCAR Technical Note TN-56+STR*, 94 pp.

[Patera(1984)] A.T. Patera (1984). A spectral element method for fluid dynamics - laminar flow in a channel expansion. *Journal of Computational Physics* **54** 468–488.

[Polvani(2000)] Polvani, L. M. and Saravanan, R. (2000). The three-dimensional structure of breaking Rossby waves in the polar wintertime stratosphere. J. Atmos Sci., **57**, 3663–3685. [Available on-line at www.columbia.edu/~lmp/lmp_pubs.html.].

[Proriol(1957)] J. Proriol (1957). Sur une famille de polynomes á deux variables orthogonaux dans un triangle. *C.R. Acad. Sci Paris* **257** 2459.

[Randall(2005)] Randall, C., Harvey, V., Manney, G., Orsolini, Y., Codrescu, M., Sioris, C., Brohede, S., Haley, C., Gordley, L., Zawodny, J., and Russell, J. (2005). Stratospheric effects of energetic particle precipitation in 2003–2004. *Geophysical Research Letters*, **32**.

[Robert(1982)] André Robert (1982). A semi-Lagrangian and semi-implicit numerical integration scheme for the primitive meteorological equations *Journal of the Meteorological Society of Japan*, **60**, 1, 319–324.

[Sherwin(1995)] S.J. Sherwin and G.E. Karniadakis (1995). A triangular spectral element method; applications to the incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering* **123** 189–229.

[Shingu(2002)] Shingu, S., Takahara, H., Fuchigami, H., Yamada, M., Tsuda, Y., Ohfuchi, W., Sasaki, Y., Kobayashi, K., Hagiwara, T., Habata, S., Yokokawa, M., Itoh, H., and Otsuka, K. (2002). A 26.58 tflops global atmospheric simulation with the spectral transform method on the earth simulator. In *Proceedings of the ACM / IEEE Supercomputing SC'2002 conference*.

[Simmons(1981)] Simmons, A. J. and Burridge, B. M. (1981). An energy and angular momentum conserving vertical finite-difference scheme and hybrid vertical coordinates. *Mon. Wea. Rev.*, **109**, 758–766.

[Smith(2002)] R. Smith and P. Gent (2002). Reference Manual for the Parallel Ocean Program (POP). Los Alamos National Laboratory LAUR-02-2484, revised 2004, 75pp.

[Solin(2003)] P. Solin, K. Segeth, and I. Dolezel, Higher-Order Finite Element Methods. *Chapman and Hall/CRC Press* (2003).

[St-Cyr(2004)] St-Cyr, A., Dennis, J., Thomas, S., and Tufo, H. (2004). An adaptive nonconforming spectral element atmospheric model. *under review, Journal of Scientific Computing*.

[St-Cyr(2005)] St-Cyr, A. and Thomas, S. (2005). Nonlinear operator integration factor spliting for the shallow water equations. *Applied Numerical Mathematics*, **52**(4), 429–448.

[St-Cyr(2006)] A. St-Cyr, M. J. Gander and S. J. Thomas (2006). Optimized multiplicative, additive and restricted additive Schwarz. Submitted to the *SIAM Journal on Scientific Computing*.

[Taylor(1997)] M. Taylor, J. Tribbia, and M. Iskandarani (1997). The spectral element method for the shallow water equations on the sphere. *J. Comput. Phys.* **130** 92–108.

[Taylor(1998)] Taylor, M., Loft, R., and Tribbia, J. (1998). Performance of a spectral element atmospheric model (SEAM) on the HP Exemplar SPP2000. Technical Report TN–439+EDD, NCAR.

[Taylor(2000)] M.A. Taylor, B.A. Wingate, and R.E. Vincent (2000). An algorithm for computing Fekete points in the triangle. *SIAM Journal on Numerical Analysis* **38** 1707–1720.

[Taylor(2000b)] M.A. Taylor, B.A. Wingate (2000). A generalized diagonal mass matrix spectral element method for non-quadrilateral elements. *Applied Numerical Mathematics* **33** 259–265.

[Taylor(2006)] M.A. Taylor, B.A. Wingate, and L.P. Bos (2006). A cardinal function algorithm for computing multivariate quadrature points. *SIAM Journal on Numerical Analysis*, to appear.

[Thomas(2002)] Thomas, S. and Loft, R. (2002). Semi-implicit spectral element atmosphere model. *Journal of Scientific Computing*, **17**, 339–350.

[Thomas(2005)] Thomas, S. and Loft, R. (2005). The NCAR spectral element climate dynamical core: Semi-implicit eulerian formulation. *J. Sci. Comput.* to appear.

[Thompson(2000)] Thompson, D. W. J., Wallace, J. M., and Hegerl, G. C. (2000). Annular modes in the extratropical circulation. Part II: Trends. *J. Climate*, **13**, 1018–1036.

[Walker(1992)] D. W. Walker, P. H. Worley, and J. B. Drake (1992). Parallelizing the spectral transform method. Part II. *Concurrency: Practice and Experience*, **4**, 7, pp. 509–531.

[Wandzura(2003)] S. Wandzura and H. Xiao (2003). Symmetric quadrature rules on a triangle. *Computers and Mathematics with Applications*, **45** 1829–1840.

[Warburton(2000)] T. Warburton, L.F. Pavarino, and J.S. Hesthaven (2000). A pseudo-spectral scheme for the incompressible Navier-Stokes equations using unstructured nodal elements. *Journal of Computational Physics* **164** 1–21.

[Worley(1992)] P. H. Worley and J. B. Drake (1992). Parallelizing the spectral transform method. *Concurrency: Practice and Experience*, **4**, 4, pp. 269-291.

DISTRIBUTION:

1	MS 0370	W. Spotz, 1433
1	MS 0370	M. Taylor, 1433
1	MS 0370	J. Strickland, 1433
2	MS 9018	Central Technical Files, 8944
2	MS 0899	Technical Library, 4536
1	MS 0123	D. Chavez, LDRD Office, 1011

Sandia National Laboratories