PROPOSAL COVER PAGE

CAN 00-OES-01 - HPCC/ESS CAN

Proposal Title: A High-Performance Software Framework and Interoperable Applications for the Rapid Advancement of Earth System Science Part III: Data Assimilation Applications for the Earth System Modeling Framework (ESMF)

Date: January 8, 2002

Principal Investigator:

Dr. Arlindo da Silva
Data Assimilation Office
NASA Goddard Spcace Flight Center
Code 910.3
Greenbelt, MD 20771
Voice: (301) 614-6174, *FAX:* (301) 614-6297
E-mail: dasilva@gsfc.nasa.gov

Co-Investigators:

Stephen Lord	NOAA/NCEP, Stop 9910, Washington, DC 20233-9910, (301) 763-8000 x7202, slord@ncep.noaa.gov
John Derber	NOAA/NCEP, Stop 9910, Washington, DC 20233-9910, (301) 763-8000 x7230, jderber@ncep.noaa.gov
Ants Leetmaa	NOAA/GFDL, Princeton, NJ 08542, (609) 452-6500, ants.leetmaa@noaa.gov
Mark Iredell	NOAA/NCEP, Stop 9910, Washington, DC 20233-9910, (301) 763-8000 x7231, miredell@ncep.noaa.gov
Michele Rienecker	NASA/GFSC, Code 971, Greenbelt, MD 20771, (301) 614-5642, michele.rienecker@gsfc.nasa.gov
Philip Jones	Los Alamos National Laboratory, Los Alamos, NM 87545, (505) 667-6387, pwjones@lanl.gov
Christian Keppenne	NASA/GSFC, Code 971, Greenbelt, MD 20771, (301) 614-5874, clk@janus.gsfc.nasa.gov
Quentin Stout	University of Michigan, Ann Arbor, MI 48109-2122, (734) 763-1518, qstout@umich.edu
Byron Boville	NCAR CGD, Boulder, CO 80307, (303) 497-1337, boville@ucar.edu
Cecelia DeLuca	NCAR SCD, Boulder, CO 80307, (303) 497-1828, cdeluca@ucar.edu
Roberta Johnson	NCAR, Boulder, CO 80307, (303) 497-1828, cdeluca@ucar.edu
John Marshall	MIT EAPS, Cambridge, MA 02139, (617) 253-9615, marshall@gulf.mit.edu

NOTE: A number of Co-Investigators have changed since this proposal was first submitted in Nov. 2000

Authorizing Institutional Official:

Dr. Vincent V. Solomonson
Director of Earth Sciences
Goddard Space Flight Center
Greenbelt, MD 20771

Budget:

FY2002: $1,692,247
FY2003: $1,363,300
FY2004: $1,194,070
TOTAL: $4,249,617

(This page intentionally left blank).

Abstract

This is the third of three closely coordinated proposals written in response to the NASA HPCC Earth and Space Sciences (ESS) Project Cooperative Agreement Notice "Increasing Interoperability and Performance of Grand Challenge Applications in the Earth, Space, Life and Microgravity Sciences" (CAN-00-OES-01). These proposals are aimed at establishing a landmark collaboration between NASA, NOAA, NSF, DOE and the university community to develop a community Earth System Modeling Framework (ESMF). This Part IIII concentrates on the use of ESMF for Data Assimilation applications.

In Part I we propose the construction of an Earth System Modeling Framework (ESMF) to provide high-performance, extensible and interoperable codes for weather prediction, climate simulations, and data assimilation. The multi-institutional and cross-disciplinary interactions among the broad community involved in the proposed work, together with the greater ease in composing robust and exchangeable components that the framework will provide, will result in a new generation of diverse and performance portable earth science applications.

The framework to be developed in Part I will consist of component coupling services, a set of data constructs that support operations on a variety of data grids and decompositions, and a portable, optimized set of low-level utilities. The linked Part II proposal, addresses deploying community atmosphere, ocean, land and cryosphere models under the framework. In this Part III we will utilize framework services to build atmospheric and oceanic data assimilation applications, namely: 1) atmospheric 3D-VAR data assimilation systems involving components from NCEP, DAO and NCAR; 2) oceanic Ensemble Kalman Filter data assimilation system using NSIPP components; and 3) an infra-structure for building a 4D-VAR oceanic data assimilation system using MIT components. These systems will be built using the foundation classes developed in Part I, extending them when necessary to address specific needs of operational and research data assimilation algorithms.

Key words: Framework, parallel computing, earth system, data assimilation, weather prediction, climate modeling

Contents

(I) Science, Technology and Management

1 Introduction

This is the third of three linked proposals to construct an Earth System Modeling Framework (ESMF) that will allow models to interoperate and share components, improve their portability and ease of use, enable them to exploit rapidly evolving computer technology, and insulate them from changes in technology. The ESMF will support codes for weather prediction, climate simulations, and data assimilation. Part I deals with the core development work of the framework. Parts II deals with coupled model applications, while this Part III deals with the specifics of data assimilation applications.

Earth system models are increasingly large and complex endeavours for which new scientific demands have arisen, such as understanding the range of results among different Earth system models, and constraining these models with data. The current volatile high performance computing environment poses a distinct technical challenge for development of applications. The transition from low-parallelism, shared-memory, vector computers to massively-parallel, distributed-memory, cache-based computers is especially challenging for data assimilation given the intensive I/O nature of the application, the unstructured nature of the various data streams, and the global nature of the inverse/minimization problem.

The framework will consist of services for coupling and exchanging data between component models and model subcomponents, a hierarchy of data structures that support operations on a variety of data grids and decompositions, and a portable, optimized set of low-level utilities. A unique feature of our design is an integrated approach to data assimilation for earth systems modeling. General purpose utility routines will include a standard interface to a variety of memory management and communication protocols for scalable architectures, I/O formats and packages, performance profiling, time management, and error and signal handling. The framework will be designed in an object-oriented, layered manner to isolate machine dependencies, and offer an application programming interface natural for the Earth sciences. The data constructs and low-level utilities will be used by the coupling superstructure and may also be used separately to compose scientific applications. Thus our work will promote software reuse as well as interoperability.

The principal investigators on the three proposals will initially closely coordinate the definition of an *open standard* covering the application domains that the ESMF will target. This effort has already begun. The implementation of the standard under Part I will then get underway, accompanied by efforts on the part of the participants in Parts II and III to undertake rapid prototyping of key elements of the standard, perhaps using some existing frameworks that partially fulfill the requirements. When the Core ESMF itself is implemented, the participants on Parts II and III will undertake exhaustive testing and performance studies based on their own applications, and disseminate their results to the wider community.

For this Part III proposal, the leading national centers for operational atmospheric data assimiliation and two of the leading groups for ocean data assimilation for climate applications will collaborate to enhance development and interoperability of assimilation systems designed under the ESMF. The data assimilation applications we will implement in ESMF are atmospheric 3D-VAR assimilation systems, sequential ocean data assimilation systems, and the infrastructure to support 4D-VAR ocean data assimilation. The participants are NOAA's National Centers for Environmental Prediction (NCEP), focused on operational numerical weather prediction, NASA's Data Assimilation Office (DAO), focused on both real-time and re-analyses of atmospheric satellite observations, NASA's Seasonal-to-Interannual Prediction Project (NSIPP), focused on assimilation for ocean initialization in coupled ocean-atmosphere-land surface forecasts, and the Massachusetts Institute of Technology (MIT), a major participant in the "Estimation of the Circulation and Climate of the Ocean" (ECCO) Consortium, focused on dynamically consistent estimates of the ocean state for climate analyses.

The collaboration between NASA's DAO and NOAA's NCEP being proposed is consistent with the goals of the recent white paper "*A NASA and NOAA Plan to Maximize the Utilization of Satellite Data to Improve Weather Forecasts*" (Franco Einaudi, *personal communication*). This white paper outlines the required steps to enhance and solidify the synergism between the two agencies, with the goal of better utilizing satellite data for both operational and research purposes. In particular, this white paper calls for a common model and data assimilation infrastructure, as means to accelerate the use of data from advanced satellite sensors.

2 Scientific Objectives and Rationale

2.1 Rationale for the ESMF

The value of interoperable codes for Earth system modeling and data assimilation has become increasingly apparent. Forecast requirements for both weather and climate are becoming more stringent, and data assimilation is becoming a crucial component for both prediction and analysis of the climate record. Data assimilation has its own large and independent community, and its tools are not, in general, accessible across the earth science modeling community. Given the complexity of data assimilation systems, much is to be gained from interoperable code. Despite the advantages of code interoperability, very little progress has been made toward this goal. Groups share low-level forward models, e.g., radiative transfer codes such as OPTRAN, but fail to share more high level components such as on-line quality control systems and analysis equation solvers. Initial step towards the development of a framework for data assimilation has been taken by the *Project D'Assimilation par Logiciel Multi-methods* (PALM, http://www.cerfacs.fr/p̃alm).

The difficulty of the assimilation problem increases with the complexity of the model and the number and different types of observations to be assimilated. The more complex the systems the harder it will be to require interoperability of components. At the same time, it is precisely for the more complex systems — such as coupled climate models and data assimilation systems — that interoperability is most important.

To further complicate the task before us, the requirements for much more sophisticated assimilation methods have coincided with major changes in computer architecture. The need to utilize distributed memory architectures, or even more daunting, to mix in the same computer and the same application distributed and shared memory programming models, has placed a huge burden on the already strained software development efforts of most groups doing data assimilation. It has also led to the realization that the old approach, in which each center develops its own solutions, is not just preventing the interchange of scientific codes, but is simply becoming unaffordable. A common *reusable and interoperable* solution to these problems must be developed. This is the rationale for an ESMF.

2.2 Objectives of the ESMF

The problems and the needs are clear, but it is still important that we present a rigorous case that the development of an ESMF is the solution. A framework will involve political and administrative overheads and will, inescapably, require some compromises in performance and flexibility from all who comform to it. If the benefits do not outweigh these costs, the project is rightly doomed to fail. A particular challenge to the ESMF is that applications in the operational environment with its strict requirements on turnaround cannot afford a marked performance degradation merely to support interoperability. On the positive side, the distinct benefits to be gained from published standards in coding and common interfaces to access data streams should be a strong enabling factor in the transition of new developments from the external community to the operational groups.

The specific objectives of the ESMF are to provide the following benefits:

a) Facilitate the exchange of scientific codes (interoperability) so that researchers may more easily take advantage of the wealth of resources that are available in the US in smaller-scale, process modeling and to more easily share experience among diverse large-scale modeling and data assimilation efforts.

b) Promote the reuse of standard, non-scientific software, the development of which now accounts for a substantial fraction of the software development budgets of large groups. Any center developing or maintaining a large system for NWP, climate or seasonal prediction, data assimilation, or basic research will have to solve very similar software engineering and routine computational problems.

c) Focus community resources to deal with architectural changes and the lack of quality commodity middleware. The non-scientific parts of the codes that would be dealt with in a common framework are also the most sensitive to architectural changes and "middleware" quality.

d) Present the computer industry with a unified, well defined and well documented task for them to address in their software design. The scientific community's influence with the industry is much diminished, but it will be even smaller if it is exercised separately by five or six centers.

e) Share the overhead costs of the housekeeping aspects of software development: documentation and configuration management. These are the efforts that are most easily neglected when corners have to be cut.

f) Provide institutional continuity to model and data assimilation development efforts. Most US modeling and data assimilation efforts are necessarily tied to only a few individuals, and centers are hard-pressed to maintain continuity that transcends them. The competitive job market we are in will result in shorter tenure for programmers. At the same time, the increasing complexity of both our systems and the technology will produce more reliance on software specialists — and less on scientists — to maintain these aspects of the systems. Both of these factors will contribute to a more unstable workforce and much greater difficulty in maintaining "institutional continuity." A framework can help us do this by having a much larger institution to support it — the whole community. Also by having codes depend on a common, well-known framework, it will be easier to find and train new people to continue a line of development.

2.3 The ESMF Community and NASA's Role

The diversity of applications supported by the ESMF is what is required to support the national Earth science objectives and more specifically to support NASA's mission as outlined in the Earth Science Strategic Enterprise Plan 1998-2002 cited in the Announcement. NASA's mission in the Earth sciences focuses on data: its acquisition, dissimination, analysis, interpretation, and use to improve our knowledge of the Earth system and our ability to predict it behavior. An ESMF would make major contributions to NASA's activities in the last three of these.

Today, the analysis of geophysical data relies heavily on detailed and comprehensive numerical models of all components of the Earth system and on sophisticated data assimilation algorithms capable of synthesizing myriad of observations into a physically consistent picture. Further use of this information to predict the state of the system days or months in advance, or how it will respond to external perturbations years or decades from now also requires comprehensive models of the atmosphere, ocean and other components.

2.4 Scientific Benefits

In this section we describe ways in which the ESMF will increase scientific productivity and encourage new research in a range of Earth science domains. We focus here on Numerical Weather Prediction and Data Assimilation. The specific science benefits for for climate modeling and general circulation modeling are the focus of Part II and for the sake of conciseness will not be addressed here.

2.4.1 Numerical Weather Prediction

Operational forecasting centers have strict requirements for performance and robustness. By specifically optimizing parts of the ESMF, some of the hand-tuning optimization burden at operational centers can be shared with the rest of the modeling commmunity. Also, sharing utility code for error and signal handling, and having the whole community develop and test code, will contribute to the robustness of code overall.

The forecasting codes we will use as a testbed for ESMF are 1) the NCEP atmospheric global forecast code, and 2) DAO's Finite-volume forecasting code which is part of NCAR/DAO CCSM effort. NCEP's system is a key component of the Global Data Assimilation System (GDAS) that provides that backbone of all numerical weather prediction at NCEP. It also is used to make the 4 times per day 120- hour Aviation forecast, the daily 384-hour Medium Range forecast and the 22 per day 384-hour Ensemble forecasts at NCEP. The DAO system is run operationally in support of NASA/EOS intrument teams as well NASA missions and field campaigns.

2.4.2 Data Assimilation Applications

Data assimilation techniques have been developed and used for many years to provide initial conditions for numerical weather prediction. In more recent years, longer term, climate forecasts have become feasible and desirable, as have model-data syntheses of the myriad of observations of different types and accuracies into a physically consistent picture of the climate record. NASA maintains efforts in all of these areas. In addition, ongoing collaboration between NASA and NOAA are focussed toward development of next-generation operational components, to take advantage of work on algorithmic improvements and the incorporation of new satellite data types into operational systems. Technology transition is a significant challenge when it is at the software level, from research and development to the rigorous, stable standards of the operational environment. Thus, work at Goddard's Data Assimilation Office (DAO) and NASA's Seasonal-to-Interannual Prediction Project (NSIPP), and NOAA's National Centers for Environmental Prediction (NCEP), participants in this proposal, would benefit immediately from the development of the ESMF. Ocean products are only now emerging into the quasi-operational arena, so the inclusion of ocean state estimation tools within the same framework as atmospheric state estimation tools will be beneficial to speed development, portability, and interoperability of these systems.

An operational data system consists of three main components: 1) a forecasting model, 2) and analysis system which combines observations with a short term forecast (*first guess*), and 3) a quality control system which ensure that erroneous observations are not included in the assimilation. At least as complex as the atmospheric/ocean models used for forecasting, is the analysis software that comprises the core of data assimilation algorithms. As in the modelling community, these systems have been largely developed in isolation, making it difficult for groups to take advantage of development taking place outside their home institutions.

The data assimilation applications we will implement in ESMF are diverse and span a variety of dynamical models and assimilation algorithms. Two different global 3D-VAR atmospheric data assimilation systems used operationally at DAO and NCEP will be implemented under ESMF, affording a great degree of interoperability between the two centers. DAO's system will be based on the Physical-space Statistical Analysis System (PSAS, [4], [2]), while NCEP's system will be based on the Spectral Statistical Interpolation algorithm (SSI, [9]). Two other systems using NASA/GSFC-NSIPP component models, optimal interpolation and the Ensemble Kalman filter, focus on assimilation for ocean initialization in coupled ocean-atmosphere-land surface forecasts. A complementary approach is utilized for dynamically consistent estimates of the ocean state for climate analyses by the Massachusetts Institute of Technology, one of the major participants in the "Estimation of the Circulation and Climate of the Ocean" (ECCO) Consortium. This effort will develop the infrastructure under ESMF for 4D-VAR ocean data assimilation using the MITgcm and the tangent-linear and adjoint model compiler (TAMC, [3]).

The interoperability afforded by ESMF will enable the sharing of high level components such as on-line quality control systems, minimization algorithms and management of observational databases. For example, swapping DAO's PSAS with NCEP's SSI, while keeping the same forecasting model, will allow us to assess weaknesses and strengthes in each of these analysis schemes, leading to subsequent algorithm improvements. Similarly, the generality afforded by the observation space formulation of PSAS, combined with interoperability afforded by ESMF, will greatly facilitate the sharing of an analysis scheme between atmospheric and oceanic data assimilation systems.

2.5 Customers and Delivery of the ESMF

The initial customers of the ESMF are the application groups described in Parts II and III of this proposal set. These include the broad spectrum of university and government researchers who use codes such as the NCAR Community Climate System Model and the MITgcm, operational centers such as NCEP, and data assimilation groups such as NASA/DAO, NSIPP, and ECCO. Other near-term customers for the ESMF include research groups that have been active in the CMIWG effort, such as the Weather Research and Forecast model, and the International Research Institute for Climate Prediction (IRICP), from which letters of endorsement are attached.

In addition, the delivery and promotion of the ESMF will proceed on a number of fronts:

- A substantial number of applications will be entrained in using ESMF through direct relation to our individual Investigator's activities. This includes conversion of other codes at institutions participating in this proposal; codes from collaborators of who will wish to interoperate with our applications and will be exposed to the ESMF through major community events, such as the yearly CCSM workshop; and other groups active in the CMIWG, the impetus for our proposed collaboration. Some of these closely related codes may be introduced as plug-in applications. Both NSIPP and ECCO are participating in the international Global Ocean Data Assimilation Experiment (GODAE). The interactions and collaborations through GODAE will expose the ocean data assimilation community to the ESMF.

- Virtually anyone, internationally, with an interest in Earth system modeling will have access to the ESMF source code and documentation via the web. We will encourage moderated open source development.

- While we will strongly encourage community contributions to the framework, we feel that it is essential for one major center to promote, maintain and moderate continuing development of the ESMF. NCAR has committed to this task.

- A workshop is planned at the end of the proposed work to introduce the ESMF to a broad community. NCAR's ongoing support role will include the coordination of regular workshops.

- We plan to present our work at conferences and publish our experiences and results in technical and earth science journals

- As described in VII (*Education and Public Outreach Addendum*), we plan to develop a series of professional development workshops for middle and high school educators.

3 Technical Plan

The community collaborating on this proposal has already made significant progress in defining the structure of the ESMF. In June 2000, each participating group agreed to prepare a strawman document describing their application requirements, ESMF scope and architecture, and implementation strategies (see http://www.scd.ucar.edu/css/NASACAN.htm).

The initial degree of convergence was high, and a combined strawman was created. Collaborators met at NCAR in August 2000 to review and refine the combined strawman, and work on a more detailed design document for ESMF began. This work is underway. However, we have already converged upon a preliminary design, as outlined here.

The Core ESMF software will developed at NCAR under Part I of this suite of proposals. Under this Part III we will develop data assimilation applications compliant with ESMF. For the sake of completeness we outline in sections 3.1 and 3.2 the functionality and archicture of the ESMF as a whole, encompassing also those attributes which are specific of Parts I and II.

3.1 ESMF Functionality

3.1.1 Scope

The scope of the initial ESMF must be extensive enough to offer significant advantages to the Earth science community but modest enough to be completed with the resources offered through this Cooperative Agreement. There are two critical needs of the Earth science community: 1) robust, optimized, non-scientific *infrastructure* libraries with which to build data assimilation systems, models and model subcomponents, to promote code reuse and 2) a *superstructure* for coupling scientific components, to promote code interoperability.

In data assimilation applications, the coupling *superstructure* will perform regridding, interpolation and communication of gridded, distributed data, as well as manipulation of scattered observational data. The data may represent multiple fields or a single field or an observational data stream, may be in the same or different executables, may be in code segments executing serially or concurrently, and may be distributed among nodes and/or partitioned among multiple threads. The interfaces for components and couplers will embody an *open standard.*

The software necessary to support the above capabilities includes *infrastructure* for describing a wide variety of grids and decompositions, and for performing high-level manipulations of fields discretized on those grids, as well as unstructured observational data streams. The software for specifying decompositions will interface to a mechanism for performing dynamic load balancing. Operations on grids and fields must implement corresponding methods for the construction of tangent linear and adjoint models for use in data assimilation applications.

Both the coupling mechanism and application codes use common utility routines. This part of the ESMF *infrastructure* includes communication libraries, sychronization, optimized I/O, performance profiling, time management, and signal and error handling.

3.1.2 Requirements

In addition to achieving the ESMF overall goals (interoperable components, code reuse, simpler code maintenance, etc.) the ESMF code will conform to a set of functional requirements. These include:

Performance portability. Portability and computational efficiency over a wide range of platforms are essential. The framework should not significantly degrade performance of an existing code written without the framework. Optimized performance on scalable architectures for moderate numbers of processors (16-500) is the highest priority.

Flexible usage. The application writer should be able to choose how much or how little of the ESMF to use. For example, data assimilation applications may initially use ESMF only for the data structures and external shell of the inversion procedure and maintain the existing solver software as an external component.

Ease of use. A key principle of design is that every user is also a developer: the key "users" of the framework are "developers" of component models, and the ESMF must simplify, or at least not unduly complicate, their lives.

Extensive grid support. ESMF must be able to couple components that are discretized on: logically rectangular grids (which may be on a physically non-rectilinear metric and a variety of positional stencils, and including "exotic" but logically rectangular grids such the bipolar and cubed-sphere grids); reduced (cut-out or Kurihara) grids; unstructured grids; phase space grids (e.g., spectral, Fourier); nested and adaptive grids; and icosahedral grids. In addition we require support for describing masked regions and halo regions.

Observational data stream support. ESMF must be able to represent unstructured observational data, including methods for assessing BUFR observational databases used operationally at DAO and NCEP, as well as oceanic databases.

High performance, extensible, multi-format I/O. The ESMF utilities will support a generic interface for I/O of self-describing data in netCDF, binary, GRIB, BUFR and EOS HDF data formats. Other formats may be added later. There must be support for high-performance parallel I/O.

Multiple language bindings. ESMF utility and coupling software will be usable by applications written in C/C++, F77 and F90.

Other requirements. These include appropriate error and signal handling, runtime configurability, and an efficient, low maintenance implementation (e.g., auto-documentation from code).

3.2 Architecture

The architectural details of the ESMF are summarized briefly below.

3.2.1 Interoperability and reuse

Object-oriented design and design reuse - both standard framework techniques - are the methods we will use to achieve interoperability. By establishing a standard set of methods and coupling interactions, researchers across the country will be able to prepare codes that they can be confident will operate with a large set of others.

We will adopt an object-oriented, layered approach to the ESMF architecture. Object-oriented design enhances code reuse since a class structure encourages well-defined, general purpose, encapsulated code segments that can work in a variety of contexts.

A framework implies that overall design - how classes interoperate - is reused as well as code. In the ESMF design we propose, certain classes within an application will be provided by the framework and others will be supplied by the application developer. The classes provided by the developer must possess a core set of methods in order to interoperate within the framework. For example, an ocean model component of a data assimilation system might need to contain a method that returns a description of its data grid and distribution. This information would be used by a generic data assimilation mechanism to understand how data should be routed between that ocean component and the data inversion (analysis) component.

3.2.2 Performance portability and ease of use

We will achieve performance portability and ease of use by layering code, through the use of generic interfaces and by designing interfaces that are not burdensome for developers seeking to integrate their applications into ESMF.

Machine-dependent code will be isolated to the lowest level in the framework by wrapping it in a generic interface. Likewise, calls to higher-level communication functions such as transposes are isolated to a layer in the framework so that an application developer does not need to manage the details of distributed data transfers. This makes the application code easier to write and use.

3.2.3 Preliminary design

We have converged on a preliminary design for the architecture of ESMF that we expect will fulfil the requirements of the Earth systems modeling and data assimilation communities as laid out above. We describe here the layering strategy, with examples of the kinds of operations native to each layer.

At the bottom are **low-level utilities**, such as communication and memory management primitives, error and signal handling, timing, machine primitives, and basic I/O operations. These are often machine-dependent and may be coded procedurally for efficiency.

The second layer is a set of **parallel utility classes**. This layer includes a retrievable description of a *machine model* and a layout class that describes the portion of a machine over which a data object is distributed, and associates a simple topology with it.

Classes representing **distributed grids** are in the third layer. Distributed grids contain a layout and a grid specification that describes the grid coordinates and connectivity. A substantial portion of the ESMF resides in this layer. The methods here include *index-space* and *physical-space* methods, and are powerful enough to represent observational data streams.

Fields are in the fourth layer. Fields contain a distributed grid, and a field specification that describes attributes related to the physical field ("metadata"). Fields will also be able to describe observational data streams and associated metadata. We will also support a *field bundle* class, for fields that are discretized on the same distributed grid. High-level I/O operations reside in this layer. These will use the metadata information to create comprehensive header information for the self-describing data formats the ESMF will support, and the distributed grid information to create efficient, high-performance, parallel I/O in various modes, including single-threaded, multi-threaded, and distributed I/O.

The **coupling and components** layer includes large scale components, such as atmosphere models and data assimilation (analysis) components, and the classes used to simplify the transfer of data between them. This layer will provide *boundary state vector* objects that comprehensively describe the portion of a component model state that is necessary for coupling, and the operations thereon. This layer also establishes communicators between component models, and the highest-level control module for scheduling them, for serial or concurrent execution.

3.3 Application codes for Data Assimilation

The atmospheric data assimilation applications will utilize the full suite of real-time observations available from NCEP. The ocean data assimilation applications will assimilate subsurface temperature data from expendable bathythermographs (XBTs), which are available at unstructured, time-varying locations, from moored locations (i.e., fixed locations) and surface altimeter data (time-varying swath data) available from GODAE servers.

3.3.1 DAO Atmospheric Data Assimilation System

For integration into ESMF we will concentrated on DAO's next-generation Finite-volume Data Assimilation System (fvDAS) based on NASA-NCAR General Circulation Model and the Physical-space Statistical Analysis System. This system consists of the following main components:

General Circulation Model. The General Circulation Model used in fvDAS is the model jointly developed by the Data Assimilation Office (DAO) and the Climate and Global Dynamics Division (CGDD) at NCAR. This model is based on the *finite-volume dynamical core* developed at DAO ([7]) with physical parameterizations from the NCAR CCM3 [6]. During this project this model will be upgraded to the atmospheric component of NCAR's CCSM.
Quality Control. The Statistical Quality Control (SQC) System is used to screen observational data prior to assimilation. This QC system consists of simple check of the observations against a background field, followed by an adaptive buddy check which adjusts error bounds according to the *flow of the day.*
Analysis System. The Physical-space Statistical Analysis System (PSAS, [4], [2]) is used to combine a first guess from the NASA-NCAR GCM with observational data to provide an updated estimate of the state of the atmosphere.

The current implementation of fvDAS already includes object oriented concepts and will much benefit from the ESMF being proposed.

3.3.2 NCEP Atmospheric Data Assimilation System

The NCEP atmospheric global forecast code is a key component of the NCEP Global Data Assimilation System that provides the backbone of all numerical weather prediction at NCEP. This hydrostatic sigma coordinate model carries surface pressure, temperature, horizontal winds, moisture, ozone and other tracers as its prognostic variables. The code uses the spectral transform method to compute horizontal derivatives, to solve the semi- implicit Helmholz equation, and to apply subscale horizontal diffusion. The physical grid is used to compute single column physics, including clouds, solar radiation, longwave radiation, gravity wave drag, surface layer exchanges, planetary boundary layer vertical mixing, shallow convection, deep convection, large-scale condensation, and ozone chemistry. The code uses the transpose strategy to distribute data and work across processors. The entire model state is transposed several times every timestep, which for the operational T170 L42 resolution is 450 seconds. The code currently runs on an IBM SP using MPI for communications.

The NCEP atmospheric global analysis code is the second key component of the NCEP Global Data Assimilation System. This three-dimensional variational code solves a single minimization problem at a specific time given all the received observations within a 6-hour window and their assumed errors along with a forecast guess within the window and its background error covariance. The minimization problem also includes constraints that ensure that the final analysis will be dynamically balanced as well as fitting the observations and the guess. Many observations such as the satellite radiances require both a guess interpolation and a single column forward radiative transfer model to compute the guess radiances and hence the observation residuals. The code performs hundreds of iterations using the conjugate gradient method to perform the minimization. Each iteration contains spectral transforms and so the transpose strategy is again generally used to distribute data and work across processors, but the three-dimensional solver has a particularly high demand on the communications. The code currently runs on an IBM SP using MPI and IBM extensions for communications.

3.3.3 NSIPP Oceanic Data Assimilation System

NSIPP's current operational assimilation is Optimal Interpolation (OI) with simple prescribed error covariance functions. Under this proposal the OI software will be converted to use ESMF observation classes and the ESMF transformation services to calculate the innovations and to prepare the forecast and observational error covariance matrices prior to inversion.

The NSIPP OGCM is version 4 of the Poseidon quasi-isopycnal reduced gravity model developed by Paul Schopf (e.g., [14]). This version has been fully ported to generic parallel architectures using a Fortran90 object-oriented design and the GEMS communications framework developed by Max Suarez. The current Pacific configuration, at a resolution of approximately $1/3^\circ \times 1^\circ \times 20$ layers, runs one month of simulation on 64 PEs in approximately 20 minutes on a fully-loaded system.

A simple OI algorithm, implemented such that full-domain covariance matrices (using prescribed covariance functions) are assembled in observation space and inverted at analysis time, takes about 30 minutes per month, assimilating every 5 days on 64 PEs. A more sophisticated, multivariate OI algorithm (e.g., [12]) updates subsurface temperature, salinity and currents using satellite altimetric observations of sea-surface height. Because of memory limitations, the required multivariate forecast-error covariances are read from a file at analysis time, and a 1-month assimilation cycle takes almost 2 hours on 64 PEs.

NSIPP's plan is to transition this ODA systems to the Ensemble Kalman Filter (EnKF, e.g., [5]) which is a prognostic calculation of the multivariate forecast error covariances from an ensemble of ocean simulations running as asynchronous objects in a parallel design. The assembly of the covariance matrix will be converted to use the communications and transformation services of the ESMF. The development undertaken within ESMF will ensure the interchangeability of the EnKF with NSIPP's simpler covariance modeling using the OI methodology. Commonality with 3DVAR atmospheric assimilation will also be explored.

3.3.4 MIT Oceanic Data Assimilation System

The MIT General Circulation Model (MITgcm) is a versatile tool for simulating both small-scale and planetary scale ocean circulation. The system is designed to run efficiently on a broad range of parallel hardware. The model implementation supports automatic generation of tangent-linear and adjoint forms. This involves the use of an adjoint compiler (http://puddle.mit.edu/r̃alf/tamc/tamc.html). The MITgcm group will participate in the design, implementation and validation of adjoint forms of framework operators. The adjoint forms of the operators will be compatible with the automatic differentiation strategy employed by the ECCO consortium adjoint data assimilation project (http://www.ecco.ucsd.edu, http://ecco.jpl.nasa.gov/).

3.4 Integration and Test

3.4.1 Integration of the data assimilation applications into the ESMF

The specific data assimilation applications to be developed under this proposal will ensure that ESMF basic classes are powerful enough to represent observational data, and that the coupling mechanisms have the flexibility to handle forecast models and analysis equation solvers. Each of the prognostic models used in the data assimilation applications, as they exist at present, employ functions somewhat analogous to the ESMF elements. Over the course this project, as the core ESMF evolves, the individual modeling systems will transition to the ESMF. Similarly, data assimilation schemes such as PSAS and SSI will also transition, at least at the highest levels, to utilize ESMF services. However, given the complexity of these operational algorithms it is expected that a full conversion to use low-level ESMF services may not be feasible during the 3 year funding period of this proposal.

Our team will be actively involved in design and testing of ESMF Benchmark Suite being developed in Part I. This synthetic ESMF Benchmark Suite (EBench) will consist of a set of representative segments from the Earth system codes Part II and this Part III. EBench performance and functionality will be baselined using the ESS Testbed (**Milestone E**). The metrics will be time to solution, scalability, and functionality in several areas: mode of execution (SPMD/MPMD), number of distributed grid types supporting optimized operations, an integrated hybrid (distributed/multithreaded) programming model, and support for C++ and F90 components. These metrics were selected because they reflect both practical performance issues and an awareness of current trends (e.g., algorithm research, C++ components, more SMP-cluster architectures) **Milestone F** focuses on functionality while **Milestone G** focuses on performance and scalability; this is consistent with the standard software engineering practice of "get it right before you make it faster," see for example [13]. The interoperability of EBench with the application codes in Parts II and III will be demonstrated in **Milestones I and J.**

In addition, the testing of the data assimilation applications will be conducted in three phases. In the first phase, the assimilation software will be exercised with an offline driver where prognostic model states and observations are read from disk files and returned as ESMF objects; assimilation schemes interfaces and coupling mechanisms will be fully tested in this environment. For example, the calculation of innovation vectors (vector of the difference between the observations and the first guess) will be tested during this phase. In the second phase, model and assimilation schemes will be made ESMF compliant, and will be exclusively coupled using ESMF services. In phase three, the full-blown ESMF-compliant components, derived from the existing codes, will be employed in a series of interoperability "tests". The first tests will be simple experiments involving verifying computational interoperability – exchange of quality-control software, and 'coupling' between different assimilation software and different prognostic models, and perhaps even between assimilation software designed separately for atmospheric or for oceanic applications. It should be kept in mind that we are speaking of computational interoperability. The development of scientific validated interoperable applications is beyond the scope of this project.

3.4.2 Performance, portability and scalability experiments

A major challenge of framework based systems is avoiding heavy performance overheads. Our architecture allows numerical codes to remain in their native implementation language and to continue to use native code for most numerical operations. This should ensure strong performance will be preserved for codes that are already highly efficient. However, the framework will need to be scalable at both the driver level and at the support layer level. Performance analysis of both parts will be carried out in detail using both driver test codes and full components. These tests will be used to make sure that the implementation of these layers can give satisfactory performance.

3.5 Expertise in Scalable Grand Challenge Applications

All participants in this proposal are heavily involved in the development of Grand Challenge Earth science applications for massively parallel systems. Team members have been involved in this work for many years and have had experince with most MPP systems that have been produced over the last 15 years. Currently all participating institutions have highly optimized parallel implementations of at least some of their production codes. Figure 1 shows speed-up curves for a sampling of these applications. These were chosen to highlight the breadth of applications and platforms being used. Shown in the figure are computational engine of DAO's Physical-space Statistical Analysis System (PSAS), the NSIPP Oceanic GCM, DAO's Finite-volume dynamical Core and NSIPP's Oceanic Data Assimilation System (ODAS). We show results on platforms from two of the leading U.S. manufactures: SGI and Cray. Most of these models, however, run on multiple platforms. Additional results (e.g., NCEP's AGCM on IBM SP) appear in Parts I and II.

3.6 Vendor Support

Early input from hardware and software vendors on critical design decisions will help prevent major recoding efforts as we move from prototypes to optimized codes. Participation from vendors in code reviews and benchmarking, together with strong technical support contacts, will help us optimize the ESMF more effectively. IBM, Cray Research, Sun Microsystems, Api Networks, and High Performance Technologies have expressed interest in collaborate with us on the development of the ESMF.

4 Management Plan

4.1 Investigator Team

The Principal Investigator of the proposal is **Arlindo da Silva**, Meteorologist, Data Assimilation Office, NASA/GSFC. He has led the development of the DAO's Physical-space

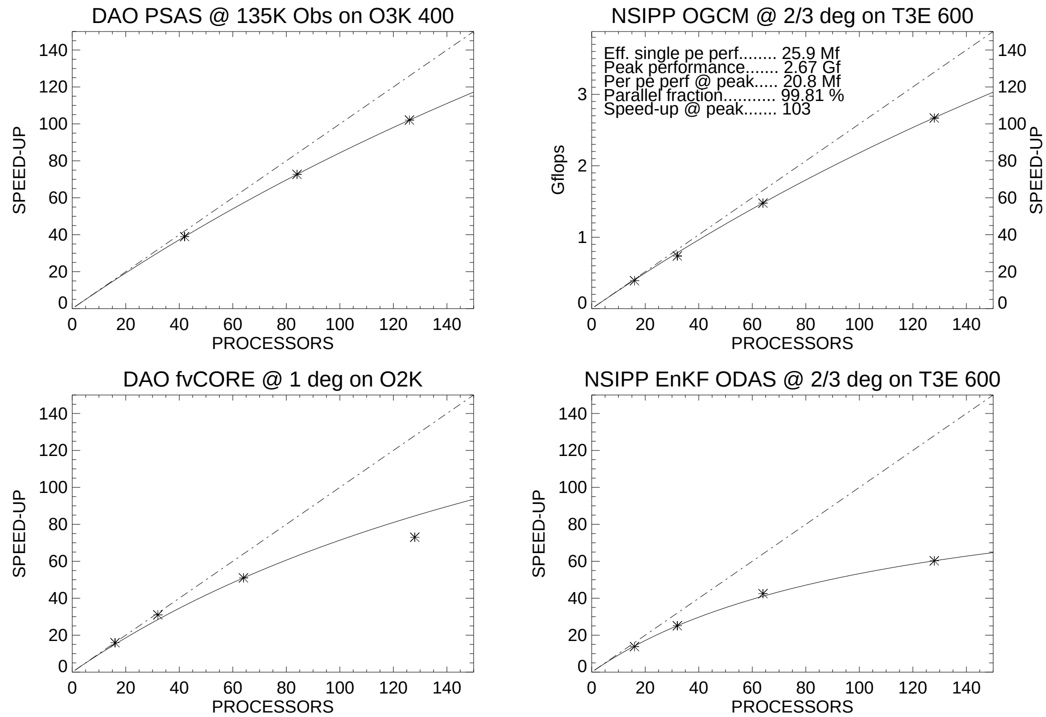

Figure 1: Speedup curves for sample modeling and data assimilation applications.

Statistical Analysis System (PSAS) and is currently leading the development of DAO's next-generation Finite-volume Data Assimilation System.

The Co-Investigators are: **Stephen Lord**, Director, Environmental Modeling Center, National Centers for Environmental Prediction; **Ants Leetmaa**, Director, Geophysical Fluid Dynamics Laboratory; **John Derber**, Meteorologist, Environmental Modeling Center, National Centers for Environmental Prediction has designed the NCEP operational atmospheric data assimilation and is one of the co-developers of the ocean data assimilation system used quasi-operationally for seasonal-to-interannual prediction; **Mark Iredell**, Meteorologist, Environmental Modeling Center, National Centers for Environmental Prediction; **Michele Rienecker**, Oceanographer, NASA Seasonal-to-Interannual Prediction Project; NASA/GSFC, heads NSIPP, and is a co-developer of the NSIPP parallel ocean data assimilation system; **Christian Keppenne**, Senior scientist, NASA Seasonal-to-Interannual Prediction Project; General Sciences Corporation, NASA/GSFC, leads the parallel implementation of the NSIPP ODAS and has extensive experience in parallel software design; **John Marshall**, Professor of Atmospheric and Oceanic Sciences, Massachusetts Institute of Technology, developed the mitGCM; **Philip Jones**, primary developer of the Parallel Ocean Program (POP) model at Los Alamos National Laboratory; **Quentin Stout**, Director of the Center for Parallel Computing at the University of Michigan; **Byron Boville**, Senior Scientist and Head, Climate Modeling Section, NCAR, was a founder and is a co-chair of the Community Climate System Model (CCSM) project. The CCSM is distributed to and supports hundreds of researchers internationally. **Cecelia DeLuca**, Software Engineer, Scientific Computing Division, NCAR, was a co-developer of the STAPL framework for real-time signal processing at MIT Lincoln Laboratory, a software tool currently used in multiple military radar systems; **Roberta Johnson**, Director of Education and Outreach, University Corporation for Atmospheric Research;

Our Investigator Team possesses the combination of skills and backgrounds necessary to help guide the design of the ESMF to support Earth system models and data assimilation, to develop the data assimilation applications under the ESMF, and establish it as a standard throughout the climate and weather communities. All of our investigators have experience constructing high-performance codes on parallel platforms.

The data assimilation applications will be based on the existing operational and quasi-operational systems at NCEP, DAO, NSIPP and on the system developed at MIT. The close proximity of the first three will facilitate the coordination of input to the design of the ESMF to support assimilation applications. It also eases coordination of the interoperability tests that are part of the demonstration of the utility of the ESMF. The team members will maintain contact through regular meetings and teleconferences.

4.2 Oversight Teams

As Principal Investigator, Dr. da Silva will serve as the primary contact and administrator of the proposed work. He will negotiate agreements with NASA and among Investigator Team members, and will arrange for disbursement of funds after payment. He will supervise the overall activities of the Investigator Team and help promote the ESMF project to the wider assimilation community. Dr. da Silva will also participate in appropriate *Oversight Teams* for the ESMF core software development at NCAR. These Teams will closely track and guide the design and implementation of the ESMF software.

4.3 Management summary

The key challenge in creating a management plan for the ESMF is to entrain broad expertise in the framework's development and implementation in several applications while ensuring that work can proceed efficiently, and that decisions can be made in an unambiguous manner. We believe that the Investigator Teams and management structure we propose across the three linked proposals will do exactly this.

Our management plan will accomplish the following:

- engage a broad spectrum of the Earth system modeling community in the specification of requirements and the overall design of the framework to maximize expert input and user buy-in;
- utilize groups with more focused interests to oversee the design and implementation of specific framework components in order to achieve timely, informed decisions;
- delegate much of the work of design drafts, prototyping and production coding to a closely integrated, central team of software engineers;
- resolve inconsistencies and differences of opinion throughout the project by allowing final software engineering decisions to be made by the software manager and oversight team leaders.

(II) Software Engineering Plan

In this section we present a software engineering plan for the ESMF.

1 Software Teams and Management

Core ESMF Development at NCAR. A *Core Implementation Team* will be established at NCAR consisting of five software engineers supervised by a *software engineering manager*. One of the software engineers will be an *integrator* (no relation to the NASA "ESMF Integrator") who is responsible for Team support functions such as configuration management and defect tracking. The Implementation Team will draft design specifications, prototype and implement ESMF components, test and validate the framework, and distribute releases. Software development will be guided by three partially overlapping *Oversight Teams* focused on different aspects of the framework: utilities, fields and grids, and coupling. The members of the Oversight Teams will include the CoInvestigators of this proposal, and will consist of appropriate mixes of software engineers, application scientists and computer scientists. Each Oversight Team will designate a lead. Responsibilities of the Oversight Teams will include reviewing software design and tracking implementation progress.

The software engineering manager and integrator will maintain a system view and ensure that development is coordinated. The Oversight Teams and Implementation Team will coordinate with staff at other institutions working on applications for Parts II and III.

Data Assimilation Application Development. Similarly, a *Data Assimilation (DA) Implementation Team* will be established at NASA/GSFC and NCEP for the development of DA applications for the ESMF. This DA implementation team will consist of five software engineers (two at DAO, two at NCEP and one at NSIPP), and will work in close contact with a postdoctoral scientist at MIT. One of these software engineers will be designated the *software engineer lead* who will be responsible for coordination among the participating institutions, as well as being the main point of contact with the Core Implementation Team at NCAR. This software development will be guided by a *DA Oversight Team* focused on the atmospheric and oceanic data assimilation applications. The members of the Oversight Teams will include the co-investigators of this proposal.

2 Software Process

The Core Implementation Team at NCAR will follow a structured software process commensurate with CMM Level II [10, 11]. The process will include many of the procedures recommended by the Software Best Practices Initiative [1]. Our documents and reviews will be simplified versions of those described in standard references [8, 15, 16]. We will aim for an effective process free of extraneous overhead.

The DA Implementation Team at NASA/GSFC and NCEP will follow a structured software process, but less rigorous than the one adopted at NCAR. In particular, we will adopt many of the procedures recommended by the Software Best Practices Initiative [1], such as software interface specification before implementation. Given the focus and size of the groups at each institution, documents and reviews will be simplified to a level necessary to ensure effective communication and eventual problem tracking.

Staged Software Development The major milestones described in Part V are the result of the coordinated completion of many smaller events, each of which has a "completion gate," such as approval of a document. Table 1 shows the progression of events in ESMF sofware development, and the product and gate associated with each event. The initial set of events, labelled "ESMF Definition" is focused on specifying the ESMF system and procedures as a whole. The second group of events, "Class Implementation" describes the development steps applied to individual software classes. As classes are completed they will be integrated into an evolving prototype of the ESMF. The final development stage, "Integration and Distribution", involves the integration of classes leading to a software release. The ESMF will have three major software releases corresponding to milestone I,J, and K; smaller releases and demonstrations will be scheduled to insure that the project is on track.

Source Availability and Distribution We plan to develop our code in an open source development environment such as SourceForge. After a prototype is released, we plan to engage the broader Earth science community in contributing to the ESMF. Community contributions will be integrated into new releases by the core team at NCAR.

Source code and documentation will be distributed via an ESMF website. The website may be an extension of that currently maintained by the CMIWG. We plan to hold a series of workshops to introduce the broader community to the ESMF.

Software Maintenance NCAR is committed to offering an ongoing program of user support, maintenance, promotion, and research into improved and extended capabilities for the ESMF. DAO, NCEP, NSIPP and MIT will equally maintain the DA applications produced. NCAR plans tp retain some or all members of the Core Implementation Team as permanent staff to carry out this work.

NASA ESS Team Role We welcome input from NASA technical personnel, and would anticipate detailing NASA involvement during the negotiation process.

3 Software Tools and Techniques

Configuration management We will likely use CVS for configuration management since it is mature, freely available, and the current community standard. We anticipate maintaining code at the following acceptance levels: **Active** (untested), **Unit tested**, and **Integrated** (code is part of a working ESMF prototype). These levels reflect the completion gates applied to code development shown in Table 1.

Software metrics The software engineering manager will track a simple set of software metrics throughout development to evaluate progress and predict schedules. These will include lines of source code and work hours metrics.

Defect tracking A tool such as Bugzilla will be used to maintain a database of defects and new feature requests.

Collaborative tools We plan to employ weekly teleconferences to keep Oversight Teams in close touch with the Implementation Team, as well as quarterly face-to-face meetings. We will continue to maintain project mailing lists and discussion forums.

Table 1 ESMF Software Event Progression

Event	*Product*	*Completion Gate*
	ESMF DEFINITION	
Requirements specification Outlines ESMF functional scope and requirements.	Requirements Document *Prepared by:* all collaborators	Document review *Reviewed by:* all collaborators
Architectural description Describes layering strategy, function and interaction of major components.	Architecture Document *Prepared by:* all collaborators	Document review *Reviewed by:* all collaborators
Software process definition Evolving documentation describing software procedures.	Developer's Guide Document *Prepared by:* integrator	Document review *Reviewed by:* software mgr.
Implementation study Assesses existing software, optimal language, threading strategy, more.	Implementation Report *Prepared by:* implementation team	Document review *Reviewed by:* all collaborators
Software implementation and test plan Plan for Implementation and testing based on class dependencies, milestones.	Software Impl. and Test Plan *Prepared by:* software mgr.	Plan review *Reviewed by:* all collaborators
	CLASS IMPLEMENTATION	
Class design Includes requirements, function, and interface specification.	Class Design Document *Prepared by:* class developer(s)	Design review *Reviewed by:* Oversight Team, software mgr.
Class implementation A class may be stubbed or partially implemented for a given release.	Prototype code *Prepared by:* class developer(s)	Code review *Reviewed by:* Oversight Team, software mgr.
Class unit test Class is tested stand-alone with a variety of inputs.	Unit test code *Prepared and tested by:* class developer(s)	Unit test *Verified by:* software mgr.
	INTEGRATION AND DISTRIBUTION	
Class integration Unit tested class is integrated into an evolving prototype of the ESMF.	ESMF system prototype *Prepared by:* class developer(s), integrator	System test and benchmarking *Verified by:* software mgr.
User documentation updated Class design documentation is updated and converted to user documentation.	User's Guide & Reference *Prepared by:* class developer(s), integrator	Review before software release *Reviewed by:* software mgr.
System release Code and documentation is released. Defects and requests for features are tracked and incorporated into future releases.	System test *Prepared by:* integrator, software mgr.	ESMF system release *Evaluated by:* ESS Project, user community

(III) References and List of Acronyms

References

[1] Airlie Software Council: Principal Best Practices. http://www.spmn.com/best_practices.html, 1994.

[2] Cohn, S., A. da Silva, J. Guo, M. Sienkiewicz, and D. Lamich. Assessing the effects of data selection with DAO PSAS. *Mon. Wea. Rev*, 126:2913–26, 1998.

[3] Giering, R. and T. Kaminsky. Recipes for adjoint code construction. *ACM Trans. Math. Software*, 24:437–474, 1997.

[4] Guo, J. and A. da Silva. Computational Aspects of PSAS. In *Numerical simulations in the environmental and earth sciences. Garcia et al, Eds.*, 1997.

[5] Keppenne, C. and M.M. Rienecker. Assimilation of temperature data into an ocean general circulation model with a parallel Ensemble Kalman Filter. In *Proceedings of 3rd WMO Symposium on Data Assimilation in Meteorology and Oceanography, Quebec*, 1999.

[6] Kiehl, J. T., J. J. Hack, G. B. Bonan, B. A. Boville, B. P. Briegleb, D. L. Williamson, and P. J. Rasch. Description of the NCAR Community Climate Model (CCM3). Technical Report NCAR/TN-420+STR, Boulder, CO, NCAR Technical Note, June 1996.

[7] Lin, S.-J. and R. B. Rood. A flux-form semi-Lagrangian general circulation model with a Lagrangian control-volume vertical coordinate. In *The Rossby-100 symposium, Stockholm, Sweden*, 1998.

[8] McConnell, S. *Code Complete.* Microsoft Press, 1993.

[9] Parrish, D.F. and J.C. Derber. The National Meteorological Center's Spectral statistical interpolation analysis system. *Mon. Wea. Rev.*, 120:1747–1763, 1992.

[10] Paulk, M.C., Curtis, W., Chrissis, M., and Weber, C.V. Capability Maturity Model for Software, Version 1.1. Technical Report CMU/SEI-93-TR-24, DTIC Number ADA263403, Software Engineering Institute, February 1993.

[11] Paulk, M.C., Weber, C.V., Garcia, S.M., Chrissis, M., and Bush, M.W. Key Practices of the Capability Maturity Model, Version 1.1. Technical Report CMU/SEI-93-TR-25, DTIC Number ADA263432, Software Engineering Institute, February 1993.

[12] Rienecker, M.M., A. Borovikov, C. Keppenne, and D. Adamec. Impact of multivariate assimilation on estimates of the state of the tropical pacific ocean. In *Proceedings of 3rd WMO Symposium on Data Assimilation in Meteorology and Oceanography, Quebec*, 1999.

[13] Royce, W. *Software Project Management.* Addison Wesley, 1998.

[14] Schopf, P.S. and A. Loughe. A reduced-gravity isopycnal ocean model: Hindcast of El Niño. *Mon. Wea. Rev.*, 123:2839–2863, 1995.

[15] Software Engineering Research Laboratory Series. Manager's Handbook for Software Development. Technical Report SEL-84-01, NASA Goddard Space Flight Center, November 1990.

[16] Software Engineering Research Laboratory Series. Recommended Approach to Software Development. Technical Report SEL-81-305, NASA Goddard Space Flight Center, June 1992.

List of Acronyms

BUFR	Binary Universal Form for the Representation of meteorological d ata
CAN	Cooperative Agreement Notice
CCSM	Community Climate System Model
CCM	Community Climate Model
CGDD	Climate and Glocal Dynamics Division
CMIWG	Common Modeling Infrastructure Working Group
CMM	Capability Maturity Model
CVS	Concurrent Versions System
DAO	Data Assimilation Office
DAS	Data Assimilation System
ECCO	Estimation of the Circulation and Climate of the Ocean
EnKF	Ensemble Kalman Filter
EOS	Earth Observing System
ESMF	Earth Systems Modeling Framework
ESS	Earth and Space Sciences
FMS	Flexible Modeling System
FVDAS	Finite-Volume Data Assimilation System
GCM	General Circulation Model
GEMS	Goddard Earth Modeling System
GFDL	Geophysical Fluid Dynamics Laboratory
GODAE	Global ocean Data Assimilation Experiment
GRIB	Gridded Binary
GSFC	Goddard Space Flight Center
HDF	Hierarchical Data Format
HPCC	High Performace Computing and Communications
IBM	International Business Machines
IRICP	International Research Institute for Climate Prediction
MPMD	Multiple Program, Multiple Data
MPI	Message Passing Interface
NCAR	National Center for Atmospheric Research
NCEP	National Centers for Environmental Prediction
NASA	National Aeronautics and Space Agency
NOAA	National Oceanographic and Atmospheric Administration
NSIPP	NASA Seasonal-to-Interannual Prediction Project
ODAS	Ocean Data Assimilation System
OGCM	Ocean General Circulation Model
OI	Optimal Interpolation
OPTRAN	Optical Path regression algorithm to compute TRANsmittances
QC	Quality Control
SGI	Silicon Graphics, Incorporated
SPMD	Single Program, Multiple Data
SQC	Statistical Quality Control
TAMC	Tangent Linear Model Compiler
XBT	Expendable bathythermographs

(IV) Biographical Sketches

ARLINDO da SILVA

Data Assimilation Office,
NASA Goddard Space Flight Center
Greenbelt, MD 20771

RESEARCH INTERESTS

My current research interests include techniques for global atmospheric data assimilation, physical-space analysis systems, error covariance modeling, bias estimation and correction, quality control, land-surface, precipitation and aerosol data assimilation, and efficient methods for assimilation of remotely sensed data. Other research interests not in the area of data assimilation include aerosol forcing of climate, hydrological cycle of the subtropics, estimatin of fluxes of heat, momentum and fresh water over the global oceans for observational studies and forcing ocean models.

EDUCATION

B.S.	1982	Physics, Catholic University of Rio de Janeiro, Brazil
M.S.	1984	Physics, Catholic University of Rio de Janeiro, Brazil
Ph.D.	1989	Meteorology, Massachusetts Institute of Technology

EMPLOYMENT

1994–Present	Meteorologist	Data Assimilation Office, NASA Goddard Space Flight Center
1990–1993	Assistant Professor	University of Wisconsin-Milwaukee
1989–1990	Visiting Scientist	Program in Atmospheric and Ocean Sciences Princeton University

RELATED PUBLICATIONS

1. Guo, J., and A. da Silva, 1997: Computational aspects of Goddard's Physical-space Statispical Analysis System (PSAS). In *Numerical simulations in the environmental and earth sciences.*, Garcia et al., Eds., ISBN 052158047, Cambridge University Press, 1997.

2. Dee, D. and A. da Silva, 1998: Data assimilation in the presence of forecast bias. *Q. J. R. Meteor. Soc.*, **124**, 269-295.

3. Joiner, J. and A. da Silva, 1998: Efficient Methods to Assimilate Satellite Retrievals Based on Information Content. *Q. J. R. Met. Soc.*, **124**, 1669–1694.

4. Cohn, S. E., A. da Silva, J. Guo, M. Sienkiewicz, D. Lamich. 1998: Assessing the Effects of Data Selection with the DAO Physical-space Statistical Analysis System. *Mon. Wea. Rev.* . **126**, 2913-26.

5. Dee, D., and A. da Silva, 1999: Maximum-likelihood estimation of forecast and observation error covariance parameters. Part I: Methodology. *Mon. Wea. Rev.*, **127**, 1822–1834.

6. Sawyer, W., L. Takacs, A. da Silva and P. Lyster, 1999: Parallel grid manipulations in earth science calculations. *Lecture Notes in Computer Science*, **1573**, 666-679. Springer, Berlin.

STEPHEN J. LORD

Environmental Modeling Center
National Centers for Environmental Prediction
NOAA Science Center, Rm. 207
Washington, DC 20233

RESEARCH INTERESTS

My interests are in managing and participating in all aspects of data assimilation and numerical model development for weather and seasonal climate forecasts. As Director of the Environmental Modeling Center, National Centers for Environmental Prediction, I oversee a staff of 90 who are dedicated to improving operational weather, ocean and climate modeling products to support the NWS mission. I have a strong background in data assimilation and tropical meteorology and have done original research on hurricane numerical modeling and data assimilation.

EDUCATION

B.S.	1969	Physics Yale University (cum laude)
M.S.	1975	Atmospheric Sciences University of California at Los Angeles
Ph.D.	1978	Atmospheric Sciences, University of California at Los Angeles

HONORS AND AWARDS

1997	AMS Fellow
1996	NOAA Dept. of Commerce Gold Medal for Implementation of the GFDL Hurricane Model
1993	NOAA Dept. of Commerce Bronze Medal for Applied research on hurricane track prediction

EMPLOYMENT

2000–Present	Director	Environmental Modeling Center, National Centers for Environmental Prediction
1993–2000	Acting Director/Deputy Director	Environmental Modeling Center, National Centers for Environmental Prediction
1989–1993	Meteorologist	National Meteorological Center
1980–1989	Meteorologist	Hurricane Research Division, Atlantic Oceanographic and Meteorological Laboratory

RELATED PUBLICATIONS

1. Pu, Zhao-Xia, S.J. Lord, and E. Kalnay, 1998: Forecast sensitivity with dropsonde data and targeted observations. In press (Tellus)

2. Surgi, N., H.L. Pan, and S.J. Lord, 1998: Improvement of the NCEP global model over the tropics: an evaluation of model performance during the 1995 hurricane season. Mon. Wea. Rev., 126, 1287-1305

3. Tuleya, R. E., and S. J. Lord, 1997: The impact of dropwindsonde data on GFDL hurricane model forecasts using global analyses. Wea. Forecasting, 12, 307-323.

4. Lord, S.J., and J. L. Franklin, 1990: The environment of Hurricane Debby (1982). Part II: Thermydynamic fields. Mon. Wea. Rev., 118, 1444-1459.

5. Lord, S. J., and J. M. Lord, 1998: Vertical velocity structures in an axisymmetric, nonhydrostatic tropical cyclone model. J. Atmos. Sci., 45, 1453-1461.

ANTS LEETMAA

NOAA/Geophysical Fluid Dynamics Laboratory,
P.O. Box 308,
Princeton, NJ 08542

PROFESSIONAL INTERESTS

As Director of NOAA/Geophysical Fluid Dynamics Laboratory, I am interested in fostering the development of improved predictive capabilities of climate change and variability on interseasonal to centennial timescales, including the development of comprehensive Earth system modeling capabilities.

EDUCATION

B.S.	1965	Physics, University of Chicago
Ph.D.	1969	Oceanography, Massachusetts Institute of Technology

EMPLOYMENT

2001–Present	Director	NOAA/Geophysical Fluid Dynamics Laboratory, Princeton, NJ
1997–2001	Director	NWS/Climate Prediction Center, Washington, D.C.
1995–1997	Senior Scientist	NWS/Office of the Director, Washington, D.C.
1990–1995	Chief	NWS/Coupled Model Project, National Meteorological Center, Washington, D.C.
1986–1990	Supervisory Oceanographer	Climate Analysis Center, Washington, D.C.
1972–1986	Research/ Supervisory Oceanographer	Environmental Research Laboratories, Miami, Florida

HONORS AND AWARDS

NOAA - Bronze Medal, 1996
NOAA - Gold Medal, 1998
AMS - Special Award

SELECTED PUBLICATIONS

1. Ji, Ming, and A. Leetmaa "Impact of Data Assimilation on Ocean Initialization and El Nino Prediction," *Mon. Weath. Rev.*, **125**, 724-753, 1997.

2. Latif, M., et al, A. Leetmaa, "A Review of Predictability and Prediction of ENSO, *Journal of Geophysical Research*, **103**, 14375-14393, 1998.

JOHN C. DERBER

Environmental Modeling Center
National Centers for Environmental Prediction
NOAA Science Center, Rm. 207
Washington, DC 20233

RESEARCH INTERESTS

My interests are in developing and implementing into operations, new data assimilation systems for weather and seasonal climate forecasts. As leader of the data assimilation team at the Environmental Modeling Center, National Centers for Environmental Prediction, I develop plans and advise the director of the Environmental Modeling Center on data assimilation matters for the ocean and the atmosphere. Also, I continue to perform my own original research on improved data assimilation techniques for the oceans and the atmosphere and on the incorporation of new observations in the data assimilation systems.

EDUCATION

B.S.	1979	Meteorology and Mathematics University of Wisconsin - Madison
M.S.	1981	Meteorology University of Wisconsin - Madison
Ph.D.	1985	Meteorology University of Wisconsin - Madison

EMPLOYMENT

2000–Present	Data Assimilation Team Leader	Environmental Modeling Center, National Centers for Environmental Prediction
1996–1997	Consultant,	European Centre for Medium-range Weather Forecasts
1989–2000	Meteorologist	National Meteorological Center (EMC, NCEP)
1987–1989	Meteorologist	Geophysical Fluid Dynamics Laboratory
1986–1987	Visiting Scientist	Geophysical Fluid Dynamics Program, Princeton University

RELATED PUBLICATIONS

1. Derber, J. C., D.F. Parrish and S. J. Lord, 1991: The new global operational analysis system at the National Meteorological Center. Wea. and Forecasting, 6, 538-547.

2. Derber, J. C. and W.-S. Wu, 1998: The use of TOVS cloud-cleared radiances in the NCEP SSI analysis system. Mon. Wea. Rev., 126, 2287 - 2299.

3. Derber, J. C. and F. Bouttier, 1999: A reformulation of the background error covariance in the ECMWF global data assimilation system. Tellus, 51A, 195-221.

4. Matsumura, T.,J.C. Derber, J.G. Yoe, F. Vandenberghe, X. Zou 1999: The inclusion of GPS limb sounding data into NCEP's global data assimilation system. NOAA/NWS/NCEP Office Note 426, Available from Environmental Modeling Center, W/NP2, Rm 207, WWBG, NOAA, 5200 Auth Rd, Camp Springs MD 20746-4304.

5. McNally, A.P., J.C. Derber, W.-S. Wu and B.B. Katz, 2000: The use of TOVS level-1B radiances in the NCEP SSI analysis system. Q.J.R.M.S., 126, 689-724.

6. Parrish, D. F. and J. C. Derber, 1992: The National Meteorological Center's spectral statistical interpolation analysis system. Mon. Wea. Rev., 120, 1747 - 1763.

MARK IREDELL

Environmental Modeling Center
National Centers for Environmental Prediction
NOAA Science Center, Rm. 207
Washington, DC 20233

RESEARCH INTERESTS

My interests are developing operational weather prediction systems. This includes atmospheric model development, related tools development, infrastructure coordination and collaboration with outside researchers.

EDUCATION

B.S.	1975	Mathematics University of Delaware
M.S.	1977	Mathematics The Johns Hopkins University
Ph.D.	1988	Atmospheric Sciences University of Washington

EMPLOYMENT

1990–Present	Meteorologist	Environmental Modeling Center, National Centers for Environmental Prediction
1988–1990	UCAR Postdoctoral	Environmental Modeling Center, National Centers for Environmental Prediction
1977–1982	Applications Programmer	Goddard Laboratory for Atmospheres, NASA Goddard Space Flight Center

MICHELE RIENECKER

Oceans and Ice Branch
NASA/Goddard Space Flight Center
Greenbelt, MD 20771

RESEARCH INTERESTS

My research interests include various aspects to understand the ocean's role in climate variability. Within NSIPP, I use numerical ocean models and develop ocean data assimilation methods for prediction of El Niño using coupled general circulation models. I am also interested in numerical analysis and in parallel computing algorithms.

EDUCATION

B.S.	1974	First Class Honours in Applied Mathematics Univeristy of Queensland, Australia
Ph.D.	1980	Applied Mathematics University of Adelaide, Australia

EMPLOYMENT

1991–Present	Oceanographer	NASA/Goddard Space Flight Center
1989–1991	Research Scientist	USRA, NASA/GSFC
1988–1989	NRC Senior RRA	NASA/GSFC
1986–1988	Scientist II	Institute for Naval Oceanography
1982–1986	Adjunct Research Professor	Naval Postgraduate School, Monterey

RELATED PUBLICATIONS

1. Borovikov, A., M.M. Rienecker and P.S. Schopf (2000), Surface heat balance in the equatorial ocean: climatology and the warming event of 1994-95, J. Clim, (in press).

2. Keppenne, C.L. and M.M. Rienecker (1999), Assimilation of temperature data into an ocean general circulation model with a parallel Ensemble Kalman Filter. *Proceedings of 3rd WMO Symposium on Data Assimilation in Meteorology and Oceanography*, Quebec.

3. Rienecker, M.M., A. Borovikov, C. Keppenne, and D. Adamec (1999), Impact of multivariate assimilation on estimates of the state of the tropical Pacific Ocean. *Proceedings of 3rd WMO Symposium on Data Assimilation in Meteorology and Oceanography*, Quebec.

4. Rienecker, M.M. and D. Adamec (1995), Assimilation of altimeter data into a quasigeostrophic model using optimal interpolation and eofs. J. Marine Systems, 6, 125–143.

CHRISTIAN L. KEPPENNE

Code 971
NASA/Goddard Space Flight Center
Greenbelt, Maryland 20771

RESEARCH INTERESTS

My research interests include scientific computing, parallel applications development, geophysical fluid dynamics, and data assimilation.

EDUCATION

B.S.	1986	Highest Honors, Civil Engineering Catholic University of Louvain, Belgium
M.S.	1987	Highest Honors, Applied Science Catholic University of Louvain, Belgium
Ph.D.	1989	Highest Honors, Applied Science Catholic University of Louvain, Belgium

EMPLOYMENT

1998–Present	Senior Research Scientist	General Sciences Corporation
1990–1997	Research Scientist	JPL, California Institute of Technology
1988–1990	Research Scientist	Dept. of Atmospheric Sciences, UCLA
1986–1988	Research Fellow	Belgian National Foundation for Scientific Research

RELATED PUBLICATIONS

1. Keppenne C.L. (2000), Data assimilation into a primitive equation model using a parallel ensemble Kalman filter, *Mon. Wea. Rev.*, 128, 1971-1981.

2. Keppenne C.L., S.L. Marcus, M. Kimoto, and M. Ghil (2000), Intraseasonal variability in a two-layer model and observations, *J. Atmos. Sci.*, 57, 1010-1028.

3. Keppenne C.L. and M.M.Rienecker (1999), Massively parallel sequential assimilation of temperature data from the TAO array into an ocean general circulation model, *Proceedings, 3rd World Meteorological Organization Symposium on the Assimilation of Observations in Meteorology and Oceanography*, held at Quebec, Canada, June 7-11, 1999.

4. Banfield D.J., A.P. Ingersoll, and C.L. Keppenne (1995), A steady-state Kalman filter for assimilating data from a single polar-orbiting satellite, *J. Atmos. Sci.*, 52, 738-753.

5. Keppenne C.L., and M. Ghil (1993), Adaptive filtering and prediction of noisy multi-variate signals: an application to atmospheric angular momentum, *Intl. J. Bifurcations and Chaos*, 3, 625-634.

6. Keppenne C.L., and M. Ghil (1992), Adaptive Spectral Analysis and Prediction of the Southern Oscillation Index, *J. Geophys. Res.*, 97, 20449-20554.

QUENTIN F. STOUT

Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, MI 48109-2122

PROFESSIONAL INTERESTS

My research interests include parallel computing and scientific computing, algorithms and data structures, adaptive statistical designs, and discrete mathematics.

EDUCATION

B.A.	1970	Mathematics, Centre College, Danville, Kentucky
Ph.D.	1977	Mathematics, Indiana University, Bloomington, Indiana

RECENT HONORS AND AWARDS

1999	College of Engineering Team Excellence Award, University of Michigan
1999	Partnership Award, IBM
1995	College of Engineering Service Award, University of Michigan

EMPLOYMENT

1997–Present	Director	Center for Parallel Computing, University of Michigan
1993–1998	Director	Software Systems Research Lab.
1992–Present	Professor	Department of Electrical Eng. and Computer Science University of Michigan

SELECTED PUBLICATIONS

1. R. Miller and Q.F. Stout, *Parallel Algorithms for Regular Architectures: Meshes and Pyramids*, MIT Press, 1996.

2. R. Miller and Q.F. Stout, "Seymour: a portable parallel programming language", *Structured Programming* **11** (1990), pp. 157–171.

3. Q.F. Stout, D.L. DeZeeuw, T.I. Gombosi, C.P.T. Groth, K.G. Powell, "Adaptive blocks: A high performance data structure," *Proc. SC*, 1997.

4. C.R. Clauer, T.I. Gombosi, D.DeZeeuw, A. Ridley, K. Powell, Q.F. Stout, B. Van Leer, T. Holzer, "High performance computer methods applied to predictive space weather simulations," *IEEE Trans. Plasma Science*, 2000, to appear.

5. R. Oehmke and Q.F. Stout, "Parallel adaptive blocks on a sphere", *SIAM Conf. Parallel Proc.*, 2001, to appear.

6. C.P.T. Groth, D.L. DeZeeuw, T.I. Gombosi, H.G. Marshall, K.G. Powell, Q.F. Stout, "Multiscale MHD simulations of a coronal mass ejection and its interactions with the magnetosphere-ionosphere system", submitted.

PHILIP W. JONES

Theoretical Fluid Dynamics (T-3),
Los Alamos National Laboratory,
Los Alamos, NM 87545

PROFESSIONAL INTERESTS

Current interests involve the use of massively parallel computers to study problems in geophysical and astrophysical fluid dynamics, including atmosphere, ocean and coupled climate modeling, middle atmosphere dynamics and fully-compressible thermal convection.

EDUCATION

B.S.	1985	Physics and Mathematics with distinction, Iowa State University
Ph.D.	1991	Astrophysical, Planetary, and Atmospheric Sciences, University of Colorado

EMPLOYMENT

1993–Present	Staff Member	Theoretical Fluid Dynamics (T-3), Los Alamos National Laboratory
1991–1993	Post-doctoral Research Associate	Geoanalysis Group (EES-5), Los Alamos National Laboratory
1986–1991	Research Assistant	Joint Institute for Laboratory Astrophysics and Center for Applied Parallel Processing, University of Colorado, Boulder

SELECTED PUBLICATIONS

1. Jones, P.W. 1999 "First- and Second-order Conservative Remapping Schemes for Grids in Spherical Coordinates," *Mon. Weath. Rev.*, **127**, 2204-2210.

2. Jones, P.W., Malone, R.C. and Lai, C.A. 1998 "The Los Alamos Coupled Model," *Proceeding of the Second International Workshop on Software Engineering and Code Design in Parallel Meteorological and Oceanographic Applications*, ed. M. O'Keefe and C. Kerr, NASA Publication GSFC/CP-1998-206860.

3. Jones, P.W. 1998 "The Los Alamos Parallel Ocean Program (POP) and Coupled Model on MPP and Clustered SMP Computers," *Making its Mark: Proceedings of the 7th ECMWF Workshop on the Use of Parallel Processors in Meteorology,* ed. G. R. Hoffmann and N. Kreitz (Singapore: World Scientific Publishing).

4. Jones, P.W., Hamilton, K.P. and Wilson, R.J. 1996 "A Very High-Resolution General Circulation Model Simulation of the Global Circulation in Austral Winter," *J. Atm. Sci.*, **54**, 1107-1116.

5. Jones, P.W., Kerr, C.L. and Hemler, R.S. 1995 "Practical Considerations in Development of a Parallel SKYHI General Circulation Model," *Parallel Computing,* **21**, 1677-1694.

BYRON A. BOVILLE

National Center for Atmospheric Research,
P.O. Box 3000, Boulder CO 80307

RESEARCH INTERESTS

My research has concentrated on developing and applying general circulation models of the lower and middle atmosphere for studies of atmospheric dynamics and climate. I have been one of the central figures in both the scientific and computational development of 4 generations of the NCAR atmospheric general circulation model. More recently, I have concentrated on coupled ocean-atmosphere modeling and was co-chair of the team which developed the NCAR Climate System Model (CSM). I am currently interested in the climate impact of solar variability and the role of the middle atmosphere in climate variability and climate change.

EDUCATION

B.S.	1975	1st Class Honours in Meteorology, McGill University, Montreal, Canada
Ph.D.	1979	Atmospheric Sciences, University of Washington, Seattle, Washington

EMPLOYMENT

1999–Present	Head	Climate Modeling Section, Climate and Global Dynamics Division National Center for Atmospheric Research
1992–Present	Senior Scientist	National Center for Atmospheric Research
1981–1992	Scientist I-III	National Center for Atmospheric Research
1979–1981	Postdoc	Advanced Study Program National Center for Atmospheric Research

RELATED PUBLICATIONS

1. Boville, B. A., J. T. Kiehl, P. J. Rasch, and F. O. Bryan, 2001: Improvements to the NCAR CSM-1 for transient climate simulations. J. Climate, 14, in press.

2. Boville, B. A., 2000: Toward a complete model of the climate system. In "Numerical modeling of the global atmosphere in the climate system", P. Mote and A. O'Neill, eds., Kluwer Academic Publishers, 419–442.

3. Kiehl, J. T., J. J. Hack, G. B. Bonan, B. A. Boville, D. L. Williamson, P. J. Rasch, 1998: The National Center for Atmospheric Research Community Climate Model: CCM3. J. Climate, 9, 1131–1149.

4. Boville, B. A., and P. R. Gent, 1998: The NCAR climate system model, version one. J. Climate, 11, 1115–1130.

5. Hack, J. J., J. M. Rosinski, D. L. Williamson, B. A. Boville, and J. E. Truesdale, 1995: Computational design of the NCAR community climate model. Parallel Computing, 21, 1545–1569.

CECELIA DeLUCA

National Center for Atmospheric Research,
P.O. Box 3000, Boulder CO 80307

RESEARCH INTERESTS

My interests include the design of large, high-performance software systems, particularly those relating to atmospheric science; parallel algorithms; real-time systems, and software engineering tools and processes. I was a design lead on the development of the Space-Time Adaptive Processing Library (STAPL) parallel framework for real-time radar applications. STAPL is an integral part of multiple operational next-generation radar systems and has been ported to several platforms. It extends the serial Vector Signal and Image Processing Library (VSIPL) standard to SMP-cluster architectures. Previous projects have included the development of parallel codes for the simulation of middle atmospheric dynamics, atmospheric chemistry, and remote sensing of atmospheric temperatures.

EDUCATION

A.L.B.	1990	Liberal Arts/Social Sciences, Harvard University, Cambridge, MA
M.S.	1994	General Engineering, Boston University, Boston, MA
M.S.	1996	Meteorology, Massachusetts Institute of Technology, Cambridge, MA

AWARDS

1994	Boston University College of Engineering Outstanding Achievement Award, first in graduating class

EMPLOYMENT

1999–Present	Software Engineer	Scientific Computing Division, National Center for Atmospheric Research
1998–1999	Lead Software Engineer	MIT Lincoln Laboratory
1996–1998	Software Engineer	MIT Lincoln Laboratory
1993–1994	Manager, Technical Support	Omnet Communications

RELATED PUBLICATIONS

1. Dickinson, R.E., S.E. Zebiak, J.L. Anderson, M. Blackmon, C. DeLuca, T. Hogan, M. Iredell, M. Ji, R. Rood, M. Suarez, K. Taylor, “Need for Infrastructure and Commonality in Climate and Weather Prediction Codes and Data,” submitted to *Bulletin of the American Meteorological Society*, 2000.

2. DeLuca, C., C. Heisey, R. Bond and J. Daly, “ A Portable, Object-Based Parallel Library and Layered Framework for Real-Time Radar Signal Processing,” In *Proceedings of Scientific Computing in Object-Oriented Parallel Environments*, ISCOPE 1997.

3. Heisey, C., C. Adamo, M. Arakawa, P. Baggeroer, J. Daly, C. DeLuca, W. Dale Hall, K. Pickard, and H. A. Spang, “Implementation of the STAP Library and Framework (STAPL) for Real-Time Matrix-Based Signal Processing,” In *Abstracts of High Performance Embedded Computing*, HPEC 1998.

ROBERTA JOHNSON

National Center for Atmospheric Research,
P.O. Box 3000, Boulder CO 80307

PROFESSIONAL INTERESTS

My research interests include modeling and analysis of aspects of the coupled magnetosphere/ionosphere/thermosphere system, paleoclimatology, isotope geochemistry and atmospheric chemistry, education, and multimedia.

EDUCATION

B.S	1980	Geophysics and Space Physics, University of California at Los Angeles
M.S.	1984	Geophysics and Space Physics, University of California at Los Angeles
M.S.	1987	Geophysics and Space Physics, University of California at Los Angeles

EMPLOYMENT

2000–Present	Director	Education and Outreach University Corporation for Atmospheric Research,
1993–2000	Associate Research Scientist	Space Physics Research Laboratory University of Michigan
1998–1999	Adjunct Associate Professor	Dept. Atmospheric, Oceanic and Space Science, University of Michigan
1989–1993	Assistant Research Scientist	Space Physics Research Laboratory, University of Michigan

SELECTED PUBLICATIONS

1. Johnson, R.M., C. Alexander, S. Barlett, M. Burek, T. Clarke, J. Durrance, J. Green, J. Kozyra, J. Linker, D. Mastie, P. Orselli, C. Rasmussen, R. Redding, and T. Weymouth, "Windows to the Universe," *WebNet '96*, Association for the Advancement of Computing in Education, ed. Hermann Maurer, 1996.

2. Azeem, S.M.I., T.L. Killeen, R.M. Johnson, Q. Wu, and D.A. Gell, "Space-time analysis of TIMED Doppler Interferometer (TIDI) measurements," submitted to *Geophysical Research Letters*, 1999.

3. Azeem, S.M.I., Johnson, R.M., "Lower-Thermospheric Neutral Winds at Sondre Stromfjord: A Seasonal Analysis," *JGR*, 102, 7379-7397, 1997.

4. Johnson, R.M. and J.G. Luhmann, "On the Dynamical Response of the High Latitude Mesopause to Solar Proton Events: Poker Flat MST Radar Observations and Results of a Simple Classical Tidal Model," *Journal of Atmos. and Terr. Physics*, **55**:9, 1203-1218, 1993.

JOHN MARSHALL

Department of Earth, Atmospheric and Planetary Sciences
Massachusetts Institute of Technology
Cambridge, MA 02139

RESEARCH INTERESTS

My research is directed at understanding key components of the general circulation of the atmosphere and ocean and the development of models to study them. I am interested in a variety of problems in geophysical fluid dynamics and their role in climate, ranging from rotating convection, the global circulation of the ocean and air-sea interaction. I use and develop numerical models of the atmosphere, ocean and climate.

EDUCATION

B.S.	1976	First Class Honors in Physics, Imperial College, London
Ph.D.	1980	Physics Imperial College

EMPLOYMENT

1992–Present	Professor	Massachusetts Institute of Technology
1992	Associate Professor	Massachusetts Institute of Technology
1991–1992	Reader in Physics	Imperial College
1984–1990	Lecturer in Physics	Imperial College
1982–1983	Post-doctoral fellow	University of Oxford

RELATED PUBLICATIONS

1. Marshall, J., C. Hill, L. Perelman, and A. Adcroft, (1997) Hydrostatic, quasi-hydrostatic, and nonhydrostatic ocean modeling, J. Geophysical Res., 102(C3), 5733-5752.

2. Marshall, J., A. Adcroft, C. Hill, L. Perelman, and C. Heisey, (1997) A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers, J. Geophysical Res., 102(C3), 5753-5766.

3. Adcroft, A.J., Hill, C.N. and J. Marshall, (1997) Representation of topography by shaved cells in a height coordinate ocean model Mon Wea Rev, vol 125, 2293-2315

4. Marshall, J., Jones, H. and C. Hill, (1998) Efficient ocean modeling using non-hydrostatic algorithms Journal of Marine Systems, 18, 115-134

5. Shaw; A. Arvind, Cho, K.-C., Hill, C., Johnson, R.P. and Marshall, J. (1998) A comparison of implicitly parallel multi-threaded and data-parallel implementations of an ocean model based on the Navier-Stokes equations. J. of Parallel and Distributed Computing, vol 48, No 1, 1-51