


Advancing Cross-Cutting Ideas for Computational Climate Science



Workshop Report

Sponsored by: U.S. Department of Energy Biological and Environmental Research Program
and Advanced Scientific Computing Research Program

September 12–13, 2016
Rockville, MD

Cover

Clockwise, beginning with the top of the cube:

Description: Mesh representation of the coastlines of Greenland within the CISM-Albany model using a new algebraic multigrid solver, an ice sheet component option within ACME.

Courtesy of Ray Tuminaro, Mauro Perego, Irina Tezaur, Andrew Salinger, Sandia National Laboratories, and Stephen Price, Los Alamos National Laboratory.

Description: Monthly averaged Sea Surface Temperature from years 71-100 of a test simulation of the beta v1 ACME model with active ocean, sea ice, atmosphere, and land surface components.

Courtesy of Milena Veneziani, Los Alamos National Laboratory, and the DOE ACME ocean model development team.

Description: Snapshot of instantaneous integrated water vapor from a development version of the ACME model atmosphere component.

Courtesy of Kate Evans, Oak Ridge National Laboratory, and the DOE ACME atmosphere model development team.

Availability of This Report

This report is available, at no cost, at <http://www.osti.gov/bridge>. It is also available on paper to the U.S. Department of Energy and its contractors, for a processing fee, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone (865) 576-8401
fax (865) 576-5728
reports@adonis.osti.gov

Please cite this report as follows: E. Ng, K. J. Evans, et al. (2017) Advancing Cross-Cutting Ideas for Computational Climate Science Workshop Report, Rockville, MD. Sept. 11–12, 2016. DOI:10.2172/1341564

Disclaimer

This document was prepared as an account of a workshop sponsored by the U.S. Department of Energy. Neither the United States Government nor any agency thereof, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents 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 or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof and shall not be used for advertising or product endorsement purposes.

Advancing Cross-Cutting Ideas for Computational Climate Science

Workshop Report

Sponsored by
U.S. Department of Energy
Office of Biological and Environmental Research
and
Office of Advanced Scientific Computing Research

*September 12–13, 2016
Rockville, MD*

Workshop Chairs

Esmond G. Ng (Lawrence Berkeley National Laboratory)
Katherine Evans (Oak Ridge National Laboratory)

Program Committee

Peter Caldwell (Lawrence Livermore National Laboratory)
Forrest M. Hoffman (Oak Ridge National Laboratory)
Charles Jackson (University of Texas at Austin)
Kerstin Kleese Van Dam (Brookhaven National Laboratory)
Lai-Yung (Ruby) Leung (Pacific Northwest National Laboratory)
Dan Martin (Lawrence Berkeley National Laboratory)
George Ostrouchov (Oak Ridge National Laboratory)
Raymond Tuminaro (Sandia National Laboratories)
Paul Ullrich (University of California, Davis)
Stefan M. Wild (Argonne National Laboratory)
Samuel Williams (Lawrence Berkeley National Laboratory)

DOE/ASCR Point of Contact

Randall Laviolette

DOE/BER Point of Contact

Dorothy Koch

Contents

Table of Contents	i
Executive Summary	iii
1 Introduction	1
1.1 Purpose	1
1.2 Workshop overview	1
1.3 Goals and structure of the report	2
2 Big Ideas	2
2.1 Process-resolving models	2
2.1.1 Grand challenge	2
2.1.2 Opportunities and potential solutions	2
2.1.3 Example: Global cloud feedbacks	3
2.1.4 Future directions	4
2.2 Integrated model credibility	4
2.2.1 Grand challenge	4
2.2.2 Opportunities and potential solutions	5
2.2.3 Example: Uncertainty in water vapor feedbacks	7
2.2.4 Future directions	7
2.3 Understanding and managing climate system complexity	8
2.3.1 Grand challenge	8
2.3.2 Opportunities and potential solutions	9
2.3.3 Example: Ocean biogeochemical modeling	9
2.3.4 Future directions	10
2.4 Continuum model framework	10
2.4.1 Grand challenge	10
2.4.2 Opportunities and potential solutions	11
2.4.3 Example: Pore-to-cloud continuum	11
2.4.4 Future directions	11
3 Cross-Cutting Issues	13
3.1 Addressing model (software) complexity	13
3.1.1 Grand challenge	13
3.1.2 Opportunities and potential solutions	13
3.1.3 Future directions	14
3.2 New hardware	15
3.2.1 Grand challenge	15
3.2.2 Opportunities and potential solutions	16
3.2.3 Future directions	16
3.3 Model performance and time to solution	18
3.3.1 Grand challenge	18
3.3.2 Opportunities and potential solutions	18
3.3.3 Future directions	20
3.4 Data management and analysis	20
3.4.1 Grand challenge	20
3.4.2 Opportunities and potential solutions	21

3.4.3	Future directions	22
4	Structure of Collaboration	22
4.1	Existing tools and collaborations	22
4.2	Facilitation of communication across disciplines	23
4.3	Outreach and learning as potential solutions	24
4.4	Possible follow-on activities	25
5	Conclusions	25
	Acknowledgments	27
	References	28
	Appendices	32
	List of Acronyms	32
	Workshop Participants	33
	Workshop Agenda	35
	Contributed Ideas Papers	39

Executive Summary

This report presents results from the DOE-sponsored workshop entitled “Advancing X-Cutting Ideas for Computational Climate Science Workshop,” known as AXICCS, held on September 12–13, 2016 in Rockville, MD. The workshop brought together experts in climate science, computational climate science, computer science, and mathematics to discuss interesting but unsolved science questions regarding climate modeling and simulation, promote collaboration among the diverse scientists in attendance, and brainstorm about possible tools and capabilities that could be developed to help address identified computational climate challenges.

Several research opportunities emerged from discussions at the workshop that participants believed could significantly advance climate science. These include (1) process-resolving models to provide insight into important processes and features of interest and inform the development of more advanced physical parameterizations, (2) a community effort to develop and provide integrated model credibility, (3) incorporating, organizing, and managing increasingly connected model components that improve model fidelity and potentially complexity, and (4) treating Earth system models as one interconnected organism without numerical- or data-based boundaries that limit interactions.

Participants also identified several cross-cutting advances in mathematics, computer science, and computational science that would be needed to enable one or more of these big ideas. It is critical to address the need for organized, verified, and optimized software, which enables the models to grow and continue to provide solutions in which the community can have confidence. Effectively utilizing the newest computer hardware enables simulation efficiency and the ability to handle output from increasingly complex and detailed models. This will be accomplished through hierarchical multiscale algorithms in tandem with new strategies for data handling, analysis, and storage.

These big ideas and cross-cutting technologies for enabling breakthrough climate simulation advancements also need the “glue” of outreach and learning across relevant scientific domains to be successful. The workshop identified several strategies to allow productive, continuous engagement with those who have a broad knowledge of the various parts of the problem. Specific ideas to foster education and tools to make material progress were discussed. Examples include follow-on cross-cutting meetings that enable unstructured discussions of the types this workshop fostered. A concerted effort to recruit undergraduate and graduate students from all relevant domains and provide them experience, training, and networking across their immediate expertise is needed. This will broaden and expand their exposure to the future needs and solutions and provide a pipeline of scientists with a diversity of knowledge and know-how. Providing real-world experience with subject matter experts from multiple angles may also motivate the students to attack these problems and even come up with the missing solutions.

1 Introduction

1.1 Purpose

Beyond the certainty of globally increasing atmospheric temperatures, there are many remaining unknowns regarding the specifics of how our Earth’s climate will be affected over the next several years and into the next century. Examples include lengthening droughts, stronger and more quickly developing hurricanes, more frequent severe winter storms in the northeastern United States, and rising sea levels due to collapsing ice sheets in Antarctica and Greenland. These more localized changes have important scientific, economic, and societal impacts. As the scientific community explores these critical aspects of climate change, with the aim of informing stakeholders over the next 10+ years, there is a growing recognition of the expanding requirements for multiscale, global, coupled Earth System Models (ESMs) as a tool for climate scientists. Such ESMs are expected to provide much more detail and fidelity through the addition of new physics, chemistry, and biology and the use of high spatial and temporal resolutions. They are also expected to provide improved accuracy and automated diagnostics highlighting the model credibility and degree of uncertainties in their predictions. As a result, there are significant challenges to overcome in terms of their construction, algorithmic developments, and computational requirements. Also, as computation capability transitions from petascale to exascale over the next 5+ years, high-performance computer systems are expected to become larger and more complex. This poses additional challenges in terms of the ability to execute and process information from ESMs robustly and efficiently on current and future high-performance computer systems.

Advances in applied mathematics and computer science are crucial to overcoming many of the algorithmic and computational challenges that climate scientists will be facing. A multidisciplinary approach, involving climate scientists, applied mathematicians, and computer scientists, is required to tackle the issues properly and make the necessary breakthroughs in climate science. While many successes have been demonstrated through existing programmatic investments, we propose that longer term and broader efforts are also necessary to realize the promise of new scientific discoveries.

1.2 Workshop overview

Recognizing the need for much closer collaborations across multiple domains to meet the challenges outlined above, a grassroots effort to generate fresh thinking was initiated by a group of climate scientists, applied mathematicians, and computer scientists. The goal was to discover and chart the optimal directions for climate modeling and motivate the latest and as yet to be uncovered developments in mathematics and computer science that will be needed to address new scientific and computational requirements. The first step forward was a workshop, which was convened September 12–13, 2016, in Rockville, MD, with the objective of discussing bold new computational ideas to address longer term science needs for climate modeling. An open invitation was extended to researchers involved with any combination of climate science, applied mathematics, and computer science to submit “ideas” papers on any topic that addressed the future needs of climate science and/or how mathematics and computer science advances could be leveraged to help. Lead authors or their designees were invited to attend the workshop. All authors were encouraged to read and consider the papers before the event to begin conversations and generate ideas and solutions. A total of 59 ideas papers were accepted, and each of them is included in the Appendix of this report.

The workshop was structured to provide ample opportunities for computational climate scientists to present the state of the science and critical bottlenecks to research progress and for applied mathematicians and computer scientists to offer potential solutions. Several rounds of follow-up discussions, in which both groups discussed issues and solutions, provided a venue for the further

maturation of ideas. To generate vigorous thought and discussion, several plenary talks and a poster session were held to exchange ideas from the papers. This report summarizes the challenges, strategies to address them, and possible avenues for execution that were identified by the participants at the workshop.

1.3 Goals and structure of the report

The purpose of this report is to summarize the findings at the workshop and identify potential opportunities to advance climate research through collaborations among climate scientists, applied mathematicians, and computer scientists. In Section 2, several grand challenges in climate science are presented, together with examples and descriptions of future directions. In Section 3, a number of areas in applied mathematics and computer science that were identified at the workshop as relevant to computational climate modeling are presented. Issues related to the structure of collaborations and outreach were also discussed at the workshop and are summarized in Section 4. The report ends with some concluding remarks in Section 5.

2 Big Ideas

2.1 Process-resolving models

2.1.1 Grand challenge

Key processes, particularly convection and cloud processes [Papers 19, 23], cryosphere feedbacks [Papers 3, 7, 22, 29, 47], biogeochemistry [Paper 58], and human-environment interactions [Paper 13]—along with their emergent features such as organized convection and surface heterogeneity—are poorly represented in modern global ESMs [28]. This is largely because the native scales of those processes are too small to be explicitly simulated, so they are mostly parameterized in these models. With intrinsic assumptions about the subgrid processes and their interactions with the explicitly resolved larger-scale environment, subgrid parameterizations have large structural and parametric uncertainties that are not well constrained by observations. In addition, many parameterizations used in current ESMs do not demonstrate proper convergence with increasing temporal or spatial resolution [Papers 44, 53]. In order to improve the representation of fine-scale processes and their upscaled effects in global modeling systems, efforts have emerged to develop specialized models that are built with a minimal set of assumptions. These *process-resolving models* provide insight into the emergent properties that occur in association with the processes and features of interest and can inform the development of advanced physical parameterizations for global climate modeling systems.

2.1.2 Opportunities and potential solutions

Several capabilities have been proposed to advance the development of cutting-edge process-resolving models:

- Advances in the development of lightweight and massively scalable numerical methods that can handle multiple scales of behavior are needed to ensure rapid throughput at extremely high model resolutions. This research aims to construct global process-resolving models with the capability to better utilize computational resources to maximize years of simulation per wall-clock day.

- Improvements in the capability to place computation “where it is needed.” For instance, boundary layer clouds are currently poorly represented in Global Climate Models (GCMs) due to models’ relatively coarse vertical resolution, leading to significant uncertainty in computations of energy balance. In the horizontal direction, the most obvious improvements in model performance occur due to better representation of topography; using the variable-resolution capabilities, where available. Grid resolution can be enhanced over rough topography but kept relatively coarse elsewhere. In this sense, optimal configurations of the model can be found that produce the best representation of the climate at minimal computational cost.
- If the former point identifies moving regions of needed attention, research could be focused on the ability to adapt computational meshes during run-time to reduce model error and improve the representation of fine-scale features [Papers 39, 42, 49].
- Load-balancing strategies capable of distributing a heterogeneous workload among processors could potentially do so over heterogeneous architectures as well [Paper 31].
- Development of methods for rapid calibration of free parameters with climate observations. Traditional approaches that rely on long-term climate simulations to tune the top-of-the-atmosphere radiative balance will no longer be feasible, so alternative techniques are needed [Paper 25].
- Development of techniques for handling of output data associated with these simulations. Storing all prognostic variables on disk is likely infeasible given the high temporal and spatial resolutions required by these models.
- A co-design effort to define domain-specific languages relevant for translating models of physical processes to heterogeneous architectures [Paper 16].

Many of these capabilities can only be attained by bringing together Earth system researchers and scientists, applied mathematicians, and computer scientists. A guiding objective of this work would be to optimize the scientific value produced for each unit of computational effort, including a focus on computational methods that are capable of reaching peak performance.

2.1.3 Example: Global cloud feedbacks

Among the processes described above, global cloud feedbacks are perhaps the largest contributor to climate model errors, and one of the largest uncertainties in future projections of global climate. A precise and accurate representation of underlying feedback processes requires extremely high model resolution that is unattainable with present general-purpose modeling frameworks such as the U.S. Department of Energy (DOE)’s Accelerated Climate Model for Energy (ACME) or the National Science Foundation’s Community Earth System Model (CESM). It is difficult to adapt present frameworks to modeling controlling mechanisms due to the traditional separation between physics and dynamics that is assumed in coarse-resolution atmosphere models. Consequently, efforts have focused on developing extremely high-resolution Large Eddy Simulation (LES) and cloud-resolving models (CRMs) at local and regional scales [22, 33]. Further, the recent development of Superparameterizations and Ultraparameterizations [21, 29] has highlighted a potential avenue by which cloud-resolving models embedded within global climate models as a parameterization can directly improve simulation skill. The success of this methodology for improving emergent atmospheric features such as the Madden-Julian Oscillation (MJO) has suggested the benefits of process-resolving parameterizations [41].

The aforementioned modeling approaches have charted a path toward global cloud-resolution model development efforts now under way at several modeling centers [3, 34]. Global CRMs are potentially desirable future replacements for the atmospheric component of ESMs, as their high spatial resolution means that convective and macrophysical parameterizations are no longer needed, in turn leading to reduced model uncertainties. Transformative advances in computing have already availed substantial supercomputing power, with further rapid advancements in leadership-class facilities (LCFs) expected in the near future. Nonetheless, present GCMs leverage only a fraction of available computing power (see Section 3.3).

2.1.4 Future directions

Many outstanding questions remain on the feasibility of scaling process-resolving models to the global domain. In particular, it is unclear at present what the needed basic assumptions are to constrain the “search space” of possible models. For example, for CRMs: Do we require low-order or high-order methods? Finite volume or finite element dynamics? Column-wise or 3D radiation? What choice of microphysics parameterization? Further, future models will likely need alternative representations of topography; the improved representation of rapidly varying orography at high resolution may make terrain-following coordinates untenable. In its place, should we employ immersed-boundary or cut-cell technologies [Paper 59]?

Additional investigation will be needed in order to determine how to best tackle these questions, particularly in light of future hardware and software infrastructure. A larger challenge is to resolve processes that connect multiple components (atmosphere, ocean, land, cryosphere) in the coupled Earth system. For example, resolving continental scale land surface hydrology processes requires model developments in multiple model components. One example of this at continental scales has provided important insights on how groundwater flow influences evapotranspiration [24]. However, coupled, global ESMs that include process-resolving models for all components will exceed the capacity of exascale computing, and pose additional challenges for model coupling that will require advances in applied mathematics, computational science, and software engineering.

2.2 Integrated model credibility

2.2.1 Grand challenge

Providing model credibility for ESMs has always been a priority; it is necessary for building confidence in model output and analysis. However, the complexity of fully coupled global ESM has prevented a comprehensive and integrated approach to all aspects of credibility. A good fraction of ideas papers submitted to AXICCS concerned the many facets of creating a more integrated approach to model credibility including, for example, strategies emphasizing verification, validation, and uncertainty quantification [Papers 14, 20, 25, 26, 34, 38, 51], new ideas for exploiting scientific or computational insights to allow greater efficiencies in sampling sources of uncertainty [Papers 4, 8, 9, 36, 37, 38, 46, 55], and ideas for leveraging code design to more easily synthesize models and observations [Papers 2, 3, 15, 35]. The topic that most undermines efforts to assess model credibility is the challenge to quantify the effects of model biases on model predictions, particularly when extrapolating into regimes for which we do not have observational data [27]. Here we emphasize structural uncertainty to call attention to the important roles that scientists have within a domain long dominated by applied mathematics and statistics. However, the ideas that were presented could apply to many aspects of the challenges that are faced to assess model credibility given the often ad hoc nature of building climate models.

The current paradigm for addressing structural uncertainties at the global scale is to compare climate predictions from simulations across the Coupled Model Intercomparison Project (CMIP) ensemble. Predictions that are shared among multiple independently developed models are a powerful indicator of robustness insofar as many models differ numerically and exploit different notions for how sub-gridscale processes are represented (known unknowns). However, these models are not entirely independent and do not identify processes for which we are not aware (unknown unknowns). The most important additional ingredients for assessing model credibility are our scientific understanding of the processes that are responsible for reliable predictions, robust comparisons to observations within a historical context, and our assessment of the skill of individual models to capture these processes. Some growth in understanding may be elicited from the analysis of multi-model ensembles. However, pinpointing cause and effect within a multi-model ensemble can be very difficult when one does not have the ability to perform more controlled experiments. By embracing the challenge posed by structural uncertainty to include information about data, models, and sensitivities within a single model framework, one can better align the scientific, math, and computational solutions that are outlined below.

At present, it takes the concerted effort and computational resources of an entire community to develop and test an advanced climate system model. This expense is a major factor limiting efforts to explore alternate, credible solutions to simulating climate phenomena. The sampling process is also stymied by system complexity and the dimensionality of sources of uncertainty. Strategies for meeting limitations must be addressed if we are to make any practical progress on even small steps in applying advances in the applied math and statistics communities. Moreover, there is not a lot of shared experience between communities working on uncertainty quantification [Papers 17, 18, 24] and climate system models and their components. Any practical solutions will need to exploit our understanding of the system and provide more theoretical expertise across disciplines so that solutions make the most of limited resources.

2.2.2 Opportunities and potential solutions

A set of directions that would put us in a much stronger position to assess model credibility is listed below in no particular order. The challenge of accounting for structural uncertainty may be viewed as something that we can have in mind, even as we make solid progress in developing more tools and gaining experience in how to apply more formal strategies to climate model development and testing.

- Embedded error modeling addresses the question of how model errors affect predictions [Paper 17]. The approach treats model error by “embedding” it in some of the key model parameters, which means those parameters are treated as random variables. The distribution of these random variables is calibrated against data on the model outputs. Consequently, a model ensemble for sampled values of these embedded model parameters represents the effect of model error or structural uncertainty on the model outcomes so it can be studied. This approach is sometimes contrasted with Gaussian Process modeling [19], which uses a discrepancy term added to model output within assessments of model likelihood as a kind of discount when testing models against observations. The embedded error model is intended to represent the effects of errors closer to their origin, before those errors propagate, and as such, can capture the effect of model error on all model outcomes, even the ones for which no calibration data are available. This approach is well suited to complex systems where it may not be clear how errors develop. Although embedded error modeling has not been applied to phenomena as multifaceted as the climate system, it is perhaps the most direct approach

to quantifying structural errors on predictions. It also depends on the close collaboration between applied mathematicians and climate scientists for its success.

- Data assimilation (DA) encompasses a broad range of mathematical methods to combine data and models. These methods may differ widely in their degree of sophistication and application purposes. Common goals of these model–observation syntheses are to allow models to reside in a state close to what is being observed and to do so by taking advantage of known sources of uncertainty in observations and models. This formal approach has several scientific and computational advantages [Papers 22, 29, 35, 38]. First, it allows a model developer to test predictions in hindcast-mode of newly introduced processes. This is what is being done already through the strategy of developing climate system component models independently using observations (often in the form of climatologies) as boundary conditions. This approach is simplistic because the approach neither offers a formal way to calibrate uncertain parameters, nor does it provide a formal mechanism to account for observational or boundary condition uncertainties when these data are prescribed. Second, from a computational point of view, DA allows us to initialize models for prediction without an expensive spin-up. To the extent that the time evolution of the model to the inferred initial state carries information that is important for predictability, we will need DA that can follow nonsteady behavior and time-resolved observation that constrains transients. We note that while dynamic equilibrium provides a useful simplifying concept for inferring time-mean behavior, it may be of limited value for predicting changes in a real-world context that is governed by nonsteady dynamics and implied time scales. Third, the use of formal DA methods allows for a more quantitative and targeted approach to selecting phenomena and identifying observations that would be the most informative for testing hypotheses and estimating uncertainties in model development. A powerful, computationally efficient DA tool being explored is the adjoint or Lagrange multiplier method to solve gradient-based optimization problems. In contrast to this intrusive method, nonintrusive approaches that operate on large complex systems are also being developed. They may have advantages in circumstances of high nonlinear dynamics. In many cases, it helps to have the application and mathematics communities working in tandem to tailor the approach to take advantage of application-specific knowledge of the physics and observations.
- Multi-fidelity methods: The computational expense of initialization and running experiments is a major factor limiting our ability to explore sources of uncertainty in model predictions. Several papers centered around ideas for improving computational efficiencies beyond numerics. One of the strategies that was mentioned repeatedly in the papers and discussions is the use of multi-fidelity methods [Papers 3, 8, 17, 22, 25, 26]. The idea is to take advantage of what can be learned from cheaper, possibly simplified versions of high-fidelity models to help anticipate what would be important to learn from high-fidelity experiments. Climate scientists often like making use of simplified versions of models for gaining scientific insights into phenomena. The potential exists for both developing a modeling hierarchy and exploiting them within a uncertainty quantification framework using multi-fidelity methods. The challenging aspect of this effort stems from the sometimes tenuous relationship that exists between phenomena at different levels of the hierarchy. Model versions with increasing resolution are perhaps the most straightforward examples of this as is done in Multilevel Monte Carlo (MLMC) [Paper 22]; however, one could envision different kinds of hierarchies based on levels of interactivity/feedbacks permitted. More experience is needed with both the science and mathematics of models with varying degrees of complexity to make this approach

feasible.

- **Emergent constraints:** One approach for reducing uncertainties using contemporary observations is to identify relationships between contemporary variability and future trends from a suite of model results. When such a relationship is found and contemporary variability can be bounded with observations, future trends are thereby constrained. This strategy was employed by Hall and Qu [14], who evaluated the strength of the springtime snow albedo feedback ($\Delta\alpha_S/\Delta T_S$) from 17 models used for the Intergovernmental Panel on Climate Change (IPCC) Fourth Assessment Report and compared them with the observed springtime snow albedo feedback from the International Satellite Cloud Climatology Project (ISCCP) and ERA-40 reanalysis data. This approach was recently applied to the carbon cycle by Cox et al. [7] and to constrain carbon cycle feedbacks on future atmospheric CO₂ levels by [15]. While this method provides useful insights into the influence of model biases on future projections, it may not directly identify specific improvements to model structures that would result in reduced projection uncertainties. Methods for representing and sampling structural uncertainty within a single modeling framework could be used to determine and test metrics that may be more robust to structural uncertainty.

2.2.3 Example: Uncertainty in water vapor feedbacks

Water vapor feedbacks have twice the impact of greenhouse gases but are not a major source of uncertainty. The process that scientists took to gain confidence in this result provides an interesting case study for our discussion on model credibility. We know of no law that regulates size of the water vapor feedback. However, virtually none of the spread in multi-model simulations of 20th century climate originated from differences in the water vapor feedback [4], maintaining a virtual lock on relative humidity even as global climates evolve. For global warming, the region that is most critical to water vapor feedbacks is where the concentrations are maintained in the upper troposphere. Soden et al. [37] looked at interannual variations in upper troposphere moisture within satellite observations and compared them to climate model simulations with and without the water vapor feedback. The model was only able to reproduce the observed variations with water vapor feedbacks enabled. This gave confidence that the models capture these feedbacks in the right amplitude and location. The argument made use of multi-models to determine the relationship between an observable (water vapor) and the amplitude of global warming. Theory and experiments identified the upper troposphere as being critical to predictions, and an experiment was set up with a modified version of the model to test whether models were capturing this feedback for the right reasons. This example underscores the importance of developing methodologies for identifying sources of uncertainty and identifying the observables that can be used to test model skill.

2.2.4 Future directions

While there is general enthusiasm for these topics, the scientific, mathematical, and computer science communities need more experience in understanding the science questions, models, and methods for model–data synthesis. For this reason, it may be prudent to start with easier targets such as inverse and forward modeling of sets of parameters deemed to be important to a tractable science question (e.g., cloud feedbacks affecting climate sensitivity, uncertainties in ice flow initialization on rates of sea level rise). For instance, one could focus on the processes by which observations are used to configure the components of a climate system model to generate realistic modes of variability such as the El Niño Southern Oscillation (ENSO). Once these systems are in

place, one could begin to explore how the formalized process could be improved through modeling hierarchies and embedded error modeling. The development of component model adjoints and the communities of individuals who are able to make use of such capabilities should be fostered. The availability of model adjoints will be especially important for systems that have long-term memories such as the ocean circulation, ice sheet states, and terrestrial and marine biogeochemistry, where the costs associated with spin-up are prohibitively expensive and time-consuming. Clearly, any improvements that result in greater computational efficiencies or throughput of simulating coupled Earth system processes also result in a greater capacity to assess model credibility.

2.3 Understanding and managing climate system complexity

2.3.1 Grand challenge

Earth's climate system consists of a balance between many complex and interrelated processes. Omitting or misrepresenting any of these processes or their interactions in Earth system models will lead to unrealistic predictions. Over the years, new processes have been added to climate models in an attempt to reduce model bias. Several papers from this workshop describe additional processes that should be included in comprehensive climate models. Developing a comprehensive list of processes critical to climate prediction is an ongoing and important task.

It is necessary to include all relevant processes to improve climate prediction, but getting them to work well together is challenging. Teasing out process deficiencies from observations (in which all processes are always operating and interacting) is tricky but critically important. Further work on techniques for decomposing observations into useful process-level information for model development and validation would be very useful. Sufficient testing is a constant deficiency in climate models because there are so many features and interactions to track.

Another difficulty in developing comprehensive climate models is the need for tuning to create the best possible model out of a collection of processes. In particular, imperfections in one part of the model will induce errors in other related pieces of the model. For example, even perfect models for interactive vegetation or ice sheets will give poor results if provided with inaccurate rainfall information. To avoid this, calibration is often performed on combinations of parameterized processes (rather than focusing on one parameterization at a time). In this context, improving one parameterization can disturb the balance of processes achieved through calibration that contains cancelling errors. As a result, improved representation of the Earth system often results in decreased model skill [16]. More effort to develop modularity, with tests that exercise various processes and also highlight tuning knobs that may not be linked to Earth system behavior, is needed.

Another issue is that the processes included in climate models are often not well understood. Interactions between multiple processes are particularly critical for capturing climate feedbacks. One approach to understanding model uncertainty is to perform a sensitivity analysis, whereby ensembles of simulations are performed with different values for uncertain parameters. Even if ensemble sensitivity analysis was exhaustive in its span of parameter space (which is itself untenable computationally), this approach is not sufficient for quantifying uncertainty because the structure of the parameterizations (as embodied in their equations and closure assumptions) may also be wrong [Papers 17, 36]. In addition, parametric uncertainty cannot capture the effect of processes that are not yet included in models. Parameter sensitivity studies may not capture data uncertainties; the initial and forcing conditions that drive the model are derived from information that is imperfect to an undefined degree [Paper 35] and requires methods such as data assimilation [Papers 18, 22]. Developing a methodology that addresses all these uncertainties is an important challenge for the climate and mathematics communities.

Capturing the important timescales in the climate system (including extreme weather events, ENSO, biogeochemical cycling, and deep ocean circulation) requires ensembles of long runs with high spatial resolution and full system complexity. These requirements are impossible to meet with current computational capabilities. Thus improving model efficiency and scalability are required for solving the climate complexity challenge.

2.3.2 Opportunities and potential solutions

This workshop identified several important processes currently missing from climate models, including, in particular, better boundary treatments [Paper 59], biogeochemical cycling [Papers 37, 43], subglacial hydrology [Paper 29], turbulence [Papers 8, 57], and human and urban systems [Papers 13, 57]. Better integration of multiple processes was also suggested, such as combining models of the hydrological and ecosystems cycles [Paper 58] and coupling iceberg calving with global ice sheet flow [Paper 7]. There are other highlighted areas where better processes are needed, but they have not yet been developed to sufficient maturity [Paper 48]. In some cases, processes need to be treated with better numerics, so that they are more accurate representations of the climate system [Papers 11, 44].

The computational challenges surrounding model complexity, while daunting, can be tackled. Section 3.3 describes many opportunities to improve algorithmic efficiency. It is also worth noting that additional processes or more sophisticated process representations can be added to existing models *without increasing wall-clock time* by running processes in parallel [Papers 4, 28, 33]. Implementing parallel physics is, however, a challenging research question because processes running in parallel don't know what each other is doing, which can lead to potential conservation errors [Paper 54]. Instead of executing individual processes in parallel, performing many simulations in parallel provides a quick way to create large ensembles. This approach to parametric sensitivity analysis is particularly useful when the timescale of interest is relatively short [Papers 34, 46]. Automation is useful for efficiently finding parameter settings that optimize model skill in these ensembles [Paper 24]. The need for process complexity can be balanced with computational constraints by using detailed offline calculations to inform online models with minimal expense [Paper 37]. On a similar note, models with very fine resolution and at high numerical accuracy can be used as benchmarks for global models to ensure that the impact of smaller scales are captured in large-scale models [Paper 23].

There are also several steps we could take to improve our ability to make sense of output from complex simulations [Paper 37]. When simulations have some processes replaced by prescribed observed values, one can isolate the process interactions responsible for comprehensive model behavior of interest. The plenary talk given by Chris Bretherton emphasizes the use of idealizations to distill the essence of complex systems into understandable pieces. Breaking climate models into individual components is also useful for avoiding compensating errors while tuning. An easy-to-use hierarchy of increasingly idealized climate-model configurations would make isolating the source of model behavior and avoiding compensating errors easier. Comprehensive model evaluation packages will also be important for ensuring that comprehensive models capture all aspects of the climate for which they are responsible [Papers 26, 51].

2.3.3 Example: Ocean biogeochemical modeling

Ocean biogeochemical models [Papers 37, 43] provide an example of the challenges and opportunities that come with complexity in Earth system modeling. The goal of such models is to quantify CO₂ uptake and to predict responses of marine productivity to climate change. Such predictions

potentially involve 100s of species that respond to varying levels of different kinds of nutrients and sunlight. As an alternative to this traditional approach, one paper [43] suggested taking a new genome-based approach to predicting the biologic response to changing environments by using bioinformatics to predict the biochemical pathways and estimate fluxes through ecosystem-scale biochemical networks.

An additional challenge is that biogeochemical models require decades to centuries to equilibrate, making the development and testing of new ideas expensive and time-consuming. Thus all the strategies that were mentioned concerning model credibility (Section 2.2) also apply here. A range of options exists for making the scientific and computational problem more tractable. First is that the local problem is not as complex as it seems since there are usually only a few factors that dominate the evolution of the local system. Which factors are important can vary in time and space and may not always be obvious. Some ideas to manage this complexity include (1) building statistical models to predict the behavior of unrepresented species, (2) adaptively learning the minimum sufficient representation needed to resolve a behavior of interest, (3) making use of process-resolving column models to inform parameterizations (in an analogy to the cloud-resolving paradigm mentioned above), and/or (4) applying new advances in stochastic parameterization to better represent scale-dependent parameters within a coarser grid cell. Since observations of the primary variables (species and nutrients) are limited, computational approaches [Paper 38] will need to be developed that optimally make use of available observations while identifying what new observations would help constrain uncertainties in model predictions.

2.3.4 Future directions

Climate model complexity has reached a critical state that demands we take a new approach. More formal approaches to model reduction, initialization, calibration, and testing will be needed to keep track of future model versions. Maintaining a process-level understanding of model behavior will be challenging in the face of this formalization. Many of the ideas proposed during this meeting involve exploiting a process-level understanding to simplify the bigger challenge of understanding a complex model. Much work needs to be done to bring together these strategies into a broader work flow that synthesizes the learning that takes place within simpler contexts so that the level of complexity that remains for the fuller system is reduced. At present, we use observational data to accept or reject a parameterization based on how it performs within a coupled climate system model. It would be preferable if we had developed a set of observation-based challenges that test and possibly tune/calibrate individual model processes. Metrics and parameters could be discovered that associate the level of accuracy required to achieve a desired level of skill. The problems listed above are examples of problems that can help the community develop the capacity, knowledge, and tools for driving new advances in climate system modeling.

2.4 Continuum model framework

2.4.1 Grand challenge

As Earth system models become increasingly complex, coupling states and fluxes across component models with strongly interacting processes (e.g., land to atmosphere), finer discretizations, and growing degrees of freedom pose a significant opportunity to improve the methodology by which climate models connect processes across existing boundaries. However, stronger and more integrated couplings of separate Earth system components also create a larger and more complex computational challenge and require a mathematically rigorous coupling strategy to handle the data transfers [Paper 1]. Reduced-order models employing a continuum methodology could

have significantly fewer degrees of freedom, provide tight coupling for critical processes spanning traditional component models, eliminate errors associated with coupling (e.g., temporal averaging, remapping, converting units, and computing derived quantities), and reduce communications, computation, and latency incurred from maintaining component model structures. In addition, correlated information can be treated consistently and numerical instabilities minimized since the number of different solvers can be reduced to a few or even one for some process representations. In a continuum model framework, the role of the coupler component to provide mediated data exchanges can be altered to drive other features such as solver execution or performance tracking, which could improve scalability and overall computational efficiency.

2.4.2 Opportunities and potential solutions

By taking a holistic top-down view of the entire Earth system, one can envision the design of integrated multiscale, multiphysics models of critical energy, water, and carbon and nutrient processes that span ocean, land, sea and land ice, and atmosphere realms and instead separate processes by physical coupling rather than “political” or traditional domain science boundaries. Focused on complete algorithmic descriptions of all processes, couplings would be synchronized based on their degree of strong-versus-weak and tight-versus-loose connections as defined by the computational science community [20].

Finer discretizations over continuous domains could be solved with functional approaches, which represent an infinite discretization and provide a flexible and rigorous platform for component model coupling. This encompasses data assimilation and UQ [Paper 22], calibration [Paper 17], and data analysis [Paper 55]. Functional approaches replace computational complexity due to finer discretization with mathematical complexity of handling flexible functions instead of data points. These approaches already generate very substantial activity in mathematics and statistics research and are only beginning to be noticed in climate research [Papers 22 and 55].

It is unnecessary to build a single model spanning all realms and process timescales, so a key strategy will focus on the determination of how each Earth system process is related to others and performing a prioritization based on available computing resources and algorithmic techniques. This will require additional scientific research and development into process understanding, algorithmic tools, and computational impacts of tight coupling of processes when developing parallelization techniques. Evaluating and constructing process parameterizations that should be solved simultaneously to a converged state represents a significant climate and computational science research challenge.

2.4.3 Example: Pore-to-cloud continuum

One example to illustrate the opportunity is the connection of soil moisture reduction through evapotranspiration to cloud condensation, precipitation, and infiltration (the pore-to-cloud continuum). These tightly coupled processes could be treated together and could then be coupled to the rest of the Earth system across interfaces determined to be more loosely connected (see Figure 1).

2.4.4 Future directions

Implementing continuum process representations to construct a new Earth system model is a high-risk endeavor with significant implications for current climate science. A complete reorganization of current Earth system model structures would be required, and implementation of a continuum strategy will necessitate new techniques for balancing computation and communication, with




Figure 1: Directly connecting water cycle processes across land, ocean, and atmosphere realms, following a continuum modeling methodology, poses a significant climate and computational science challenge, but this approach could offer a more consistent treatment, eliminate coupler-mediated data exchanges, and improve scalability and computational efficiency. Water storage and flux units are 10^3 km^3 and $10^3 \text{ km}^3 \text{ year}^{-1}$, respectively.

potentially large payoffs in model performance. Thus, climate scientists must work closely with computational scientists and engineers if such a pioneering effort is to be undertaken.

Component models for land, ocean, sea ice, land ice, and the atmosphere have proved useful for a wide variety of climate science questions, and maintaining the ability to conduct research with component models is important for investigating domain-specific mechanisms. Since continuum modeling is designed to solve equations of state under conditions of mass and energy conservation throughout all parts of the Earth system, this methodology conflicts with the traditional component model structure. However, strategies for modular design of key parameterizations could enable dynamic “plug and play” creation of traditional component models, built from the structures underpinning the continuum process representations. The degree to which code can be reused and the ability to dynamically build numerically stable and accurate models of sub-components are open questions.

One path forward for initiating a research effort to investigate continuum modeling strategies for the design of a new Earth system model would be to select one or two continuum processes to implement. The pore-to-cloud example described above and illustrated in Figure 1 is a prime candidate of relevance to ongoing DOE research. Having strong connections to the mathematics community (e.g., algorithms, solution techniques, and libraries, Section 3.3) and the statistics community (e.g., functional data analysis, Section 3.4.2) may enable faster development and proper implementation of the correct methodologies. Testing methods for coupling the energy cycle and vegetation processes into this continuum model are key initial research tasks, requiring an integrated multidisciplinary team.

3 Cross-Cutting Issues

3.1 Addressing model (software) complexity

3.1.1 Grand challenge

The simulation of a full Earth system (atmosphere, ocean, land ice, sea ice, and land surface components) requires that a multitude of physical processes be accurately represented over a wide range of temporal and spatial scales. As full simulations are built across several models developed by different domain scientists, managing complexity is a nontrivial but essential task. Tackling this is further complicated by the fact that climate models include a large number of parameters associated with approximated quantities, which are often impossible to determine precisely. Ultimately, climate modeling is intended to inform decision makers about possible future climate scenarios, and this in turn relies on verification and validation. Verification centers on assessing that the computer implementation accurately represents a conceptual description of the model and its solution, which includes partial differential equations (PDEs) as well as other types of models and assumptions. Yet, verification is linked to complexity management.

Verification in the presence of complexity, uncertainty, and simulation parameters is discussed in different contexts within many papers including Papers 10, 14, 26. Numerical accuracy and its relationship with parameter calibration [Papers 44, 53] were highlighted in the presentation associated with Paper 53 using an example based on an existing climate capability. In this case, multiple inaccuracies cancel and give the appearance of a validated solution. This can easily occur when models are erroneously calibrated with defective sub-components, as was called out in Section 2.3. That is, the calibration might falsely correct for deficiencies. Not only does this lead to inaccuracies, but it skews future assessments of potentially improved sub-components (e.g., with better accuracy/stability properties) if the assessment employs poorly calibrated parameters (determined with a previously errant sub-component).

3.1.2 Opportunities and potential solutions

Properly managing simulation complexity provides great benefits to domain scientists, to computational mathematicians, and to computer scientists. Many papers touched on complexity and the needed to improve simulation confidence (e.g., Papers 17, 26, 51). Ideas from computer science and computational mathematics can help manage complexity. This includes ideas from verification, software ideas associated with modularity, software engineering ideas and practices, and software ideas associated with programming environments for cross-platform performance on sophisticated computer hardware.

The presentation associated with Paper 53 highlighted opportunities to improve numerical accuracy/stability or, more generally, enhance verification at the most fundamental level. To address current concerns, rigorous benchmarking and testing of important algorithm kernels must be promoted. These tests must include examples that employ representative grids over realistic geometries (using typical elements types and aspect ratios that are characteristic of how a sub-component is utilized in larger simulations). More generally, tests should include increasing levels of complexity and be made available for those exploring potential sub-component improvements. Such tests might include mesh convergence studies perhaps employing the method of manufactured solutions. In addition to improving simulation confidence, benchmarks and published tests provide entry points for computer scientists and mathematicians to explore new concepts, algorithms, and implementations and to make meaningful comparisons with published results. Most likely, the computer science and applied mathematics communities would develop new benchmarks and publish additional results.

In another direction, uncertainty quantification techniques can also assist in investigating and analyzing the sources of different types of errors. For example, [Paper 17](#) emphasizes the need to understand both structural errors as well as spatial and temporal discretization errors and how certain novel UQ techniques (based on carefully embedding statistical model error terms in model parameters) can possibly account for structural errors.

An important consideration for improving complexity management is a closer association between mathematical concepts and software components. That is, the software framework should be designed to mirror mathematical relations (e.g., software abstractions based on vectors, fields, matrices, vector spaces, function evaluators, etc.). Mathematics is a universal precise language. Software written in this way is generally accessible to a wide community. For example, the discussion associated with [Papers 50](#) and [54](#) noted that experimenting with time integrators can be cumbersome if it is not easy to associate any piece of code with the discretized form of the time level update to the ordinary differential equation. [Papers 2, 3, and 4](#) touch on a related theme, that is, tools that help increase modularity and the building of sophisticated software by combining already developed agile components. This idea was adopted in the PISCEES project using the Albany and Chombo frameworks to build the FELIX and BISICLES ice sheet models [[Papers 39, 45](#)]. By leveraging preexisting development, newly developed dynamical cores can be built rapidly with advanced features (e.g., AMR or UQ). While these examples are mathematical, there are also software/performance incentives. A number of developments are under way associated with languages that facilitate asynchronous many-task programming as well as capabilities that facilitate performance optimization across multiple platforms (e.g., Intel Phi and NVIDIA GPUs). These software projects and their abstractions allow developers to obtain high performance over a set of different complex hardware platforms at a reduced effort [[Papers 12, 28, 31](#)]. It should be noted that existing refactorization efforts that leverage next generation platforms provide an opportunity to introduce much greater modularity and improve the management of complexity overall.

3.1.3 Future directions

The future is likely to bring increased complexities to the science of climate modeling. These increases will be due to the inclusion/combination of more physically realistic models, the greater availability of climate data, and the use of more sophisticated hardware platforms.

At this point, it is unclear exactly how these changes will impact climate model design decisions. However, it is clear that greater modularity and an increased emphasis on sub-component flexibility, testing, and performance will be necessary. This will allow climate scientists to manage the growth in complexity while leveraging contributions and foster collaborations from computational scientists in the applied mathematics and computer science fields. As noted above, modularity should consider a closer association between mathematical concepts and software components. Additionally, future modularity efforts should consider the separation of software implementation from the underlying hardware design while at the same time exposing fine-grain parallelism and allow for the possibility of task-based programming models that can address the heterogeneous nature of climate simulations.

While this section has focused on modularity and the importance of verifying individual modules as a means to addressing complexities, model interactions are extremely important and will grow in importance over the next decade. [Paper 1](#) discusses a framework for addressing important aspects of heterogeneous numerical methods in a mathematically rigorous way. An important idea here revolves around using optimization and control ideas to provide a solid mathematical coupling strategy so that dissimilar numerical methods can function within a unified simulation. [Paper 51](#) focuses on understanding the interactions between models that have been developed independently.

Models that do not directly interact with each other (noncontiguous model components) might have indirect effects on each other, so pragmatically understanding this propagation can be challenging. Despite the challenges of component *fractionation* (where each model has its own data, assumptions, and quality measures), the climate community will need to develop approaches for full-system model evaluation. This in turn requires computer science tasks such as building an infrastructure that can efficiently and flexibly transfer data between a large number of models as well as mathematical tasks such as developing appropriate full system metrics that might involve several model ensembles.

3.2 New hardware

3.2.1 Grand challenge

Historically, processor performance has generally followed Moore’s law, the observation that the computing power that can be cost-effectively integrated on an integrated circuit doubles every 24 months [26]. Unfortunately, the performance of climate simulations has not kept pace with the improvements in processor performance because all improvements have been more than consumed by increased degrees of freedom and model complexity. In addition, performance is likely affected by several fundamental changes in processor scaling and architecture.

Despite the end of Dennard scaling (constant power density scaling) [9], processor performance has continued to improve according to Moore’s law without exponential increases in power. This has been realized primarily through design changes such as energy-efficient circuits, core architectures designed for energy efficiency at lower frequencies, massive increases in parallelism (multicore, manycore, wide vector), and more software-controlled functionality (e.g., scratch pad memories). Lower frequencies tend to increase latencies and overheads (bottlenecks in strong-scaled climate simulations), while the massive increase in hardware parallelism has demanded a commensurate increase in parallelism expressed in the software. Until recently, many compilers for commodity hardware would fail to vectorize even the simplest stencil codes (let alone iterative point-wise equations or complex microphysics routines). Additionally, users’ attempts to thread their simulations often ran afoul of bottlenecks in the maximum theoretical parallelization. Programming models and hardware that once virtualized data movement now increasingly demand programmers micro-manage data movement through the memory hierarchy, including exploitation of any data locality. Failure to do so can result in substantial performance degradation.

Over the past decade, while peak processor performance continued to scale well, main memory bandwidth scaled much more slowly [30]. As a result, many processors and accelerators available today have machine balances of less than 0.1 Bytes of memory bandwidth per Flop, which is far below what is needed in many algorithms to achieve good parallel efficiency in general. The future trends in semiconductor manufacturing will only exacerbate and compound these challenges. In the strictest sense, Moore’s law is dead (exponential performance progresses at a slower pace than doubling every 2 years). By 2016, we were already observing the implications on vendor processor performance and roadmaps. For example, Intel’s “tick-tock” cadence, which alternated between process shrinks and novel architectural changes on a regular 2-year basis for the past decade, has been abandoned with the introduction of Kaby Lake between the 14nm SkyLake architectural “tick” and the 10nm Cannon Lake “tock”. Similarly, it required 4 years for Intel’s Knights Landing (2016) processor to double the performance of the Knights Corner processor (2012).

Whether the climate science community realizes it or not, they are in a race to the bottom with the semiconductor industry. Whereas the semiconductor industry is driving towards atom-scaled devices with exponentially increasing performance (perhaps no longer every 2 years), the climate community is driving towards ever-finer resolutions (e.g., global cloud-resolving models) with expo-

nentially increasing computational demands. With perhaps only four more planar complimentary metal oxide semiconductor (CMOS) lithographic nodes, it may not be possible for the current trajectory of vendor processor offerings alone to enable more than a 2–4× increase in resolution. Any further increase in resolution will require a corresponding increase in power and system cost (e.g., 2× the resolution requires 8–16× the power and system cost). This suggests that it may not be cost-effective to increase the resolution of climate simulations beyond that attained in the 2025–2030 time frame without radical changes to algorithm, implementation, architecture, and semiconductor technology.

3.2.2 Opportunities and potential solutions

Several approaches have been proposed to bridge the potential performance impediments to simulation and analysis.

One obvious strategy to bridge the gap is to improve methods and algorithms. Whereas changes to (or adoption of) a linear solver may reduce computational demands by a constant factor, changes to discretization in space or time (e.g., higher order) can result in an exponential reduction in the computational requirements and mitigate increasing supercomputing costs [Papers 27, 41]. Unfortunately, ensuring a global benefit is an challenging prospect.

Due to the rapid size increase in both observational data and model output, a co-design effort to develop future computational resources to meet the unique needs of large-scale climate data analytics is needed [Paper 56]. Such systems typically require less computational capacity (e.g., fewer processor cores) but would benefit most from large, fast memory systems and high bandwidth input/output (I/O). To meet the growing demands of climate analysis and model benchmarking, a balance must be struck between high computational capacity resources and high-throughput resources at major computing facilities.

Vendors have begun to embrace the concept of benchmarks and proxy applications as a means of specializing their offerings for the computational demands of the various computing centers. Doing so provides them an advantage over their competitors since, with the same process technology and power constraints, they can deliver superior performance. Unfortunately, as vendors are profit motivated to create general solutions that balance the cost-performance trade-off across many computational domains, any solution for climate is likely to be suboptimal. Nevertheless, it has become increasingly cost-effective to design general-purpose, domain-optimized supercomputers specifically for their computational needs.

To mitigate the effects of evermore novel and complex vendor processor offerings, several discussions at the workshop focused around domain-specific languages (DSLs) [Paper 16]. The value of a DSL is its ability to describe the functionality of a method in the context of the domain (e.g., PDEs on structured grids) without prescribing a solution or approach to execution (e.g., Fortran/C with MPI+OpenMP). This separation of concerns is desirable as it amortizes the development cost required to port to each new generation of architecture. Unfortunately, substantial investigation is required to determine what the DSL should look like, and substantial effort is required to develop and maintain a DSL compiler.

3.2.3 Future directions

At the workshop, a number of discussions examined alternatives to the current and emerging computational roadblocks. These alternatives included a number of algorithmic [Papers 5, 6, 30, 40, 41] and model [Papers 13, 33] changes, performance optimization techniques [Papers 4, 28], and advances in semiconductor technology (optical interconnects, carbon nanotubes, etc.) as well as

cost-effective semi-custom supercomputers tailored for the needs of climate science [Paper 19].

The past decade has seen the proliferation of IP-based (intellectual property) hardware building blocks for system-on-a-chip processor designs in everything from network routers to cell phones. These building blocks range from memory and network controllers to full, superscalar out-of-order cores (e.g., ARM, Tensilica). More recently, we have witnessed the emergence of open-source hardware, including high-speed I/O circuits and memory controllers as well as full processors (e.g., Berkeley’s RISC V core [2]). With a fabrication-less design flow, it is possible not only to tailor the core architecture (out-of-order, cache sizes, vector widths, frequencies, etc.) but also the entire node (including memory) and system (including network) designs to match the computational needs of a scientific problem [11, 23, 36, 40]. Although one may discount this approach, several systems have successfully adopted it, including D.E. Shaw’s ANTON and ANTON2 supercomputers designed for extremely fast and efficient molecular dynamics simulations [36]. Sunway’s TaihuLight, the world’s fastest supercomputer, followed a similar approach and produced a system arguably optimized for solving dense linear systems while still being capable of solving problems from other domains [10, 13].

The cost of such systems should not be underestimated, and several researchers have extensively tabulated the costs for such machines [8], dividing the cost into the nonrecurring engineering costs associated with designing and implementing the system and fabrication and integration costs. As the scale of the procurement increases, the speedup required to reach cost-performance parity with a traditional supercomputer quickly diminishes (e.g., to break even, the custom design needs only be 25% faster—a modest factor for domain-customized hardware—on a \$100M procurement). As a result, for large procurements, there is a clear potential cost savings by leveraging the emerging customizable hardware market to tailor the computational capabilities of the supercomputer to the computational requirements of a domain such as climate science.

In the climate community, one of the biggest impediments to exploit emerging manycore and accelerated hardware is the substantial and required optimization and porting to new programming models (CUDA, OpenACC, etc.). By contrast, the basic building blocks of these semi-custom designs are traditional RISC cores, often the same as those found in commodity devices like cell phones. As such, existing C/Fortran code written using standardized programming models is immediately portable. To ensure one may leverage semi-custom instructions without unduly burdening the programmer with intrinsics, an optimizing compiler (cognizant of the new instructions) is auto-generated along with a synthesizable (e.g., Verilog) description of the core.

Substantial preliminary investigation is required to determine the efficacy of a domain-optimized supercomputer for climate. Architecting a domain-optimized supercomputer requires benchmarks to guide optimization. As the design space expands (combinatoric explosion of architectural parameters) and the evaluation infrastructure slows (from proxy hardware down to cycle-accurate simulators), the cost of evaluating configurations becomes expensive. As such, it is incumbent on the climate community to deliver compact, configurable benchmarks to the computer science and applied mathematics communities. These benchmarks must be future looking and designed to proxy the computational methods intended to be used in the 2025–35 time frame. Concurrently, the climate community needs to track the evolution of vendor solutions to quantify the emerging performance gap and to determine whether the architected solution can bridge the gap (constrained by expected procurement funds). Failure to be proactive in this manner will result in hardware and/or software solutions that are poorly matched to the computational challenges with climate science — a continuation of the status quo.

3.3 Model performance and time to solution

3.3.1 Grand challenge

The underlying dynamics driving changes in climate include a wide spectrum of complex interacting physical processes with many different types of nonlinear phenomena that span a range of spatial and temporal scales. These complexities often require high spatial resolution and drive increased model uncertainty, which is addressed with large model ensembles. The push for high-resolution simulations and increased ensemble size has exposed a pressing need for faster time-to-solution. These challenges require a strong collaborative research environment that promotes interaction between computational climate scientists, applied mathematicians, and computer scientists. Domain knowledge about the important physics of the problem is needed to inform appropriate choices for algorithms that best suit the problem. While hardware improvements provide a bedrock for faster completion, progress in the algorithms themselves will be essential in attaining significant reductions in the computational time needed to simulate different features of the climate system. A reduction in solution time allows for a combination of quicker solutions, longer simulation times to consider longer timescales, and/or higher fidelity models that consider more complexity for a given model execution.

3.3.2 Opportunities and potential solutions

To improve solution times, it is expected that climate simulations must leverage and drive advances in time advancement, solvers (nonlinear and linear), adaptivity (spatial and temporal), and where appropriate, the use of ensembles to gather statistics and sensitivities.

Ultimately, understanding the changes in our climate relies on accurately computing the spatial and temporal dynamics of a full Earth system model. This in turn requires improvements in the efficiency and sophistication of the discretization, for example, in finite-element strategies [Paper 6]. These schemes must generally be adapted to specific problem characteristics and must be appropriate for increasingly complex models. As many of the underlying sub-systems involve a wide range of spatial and temporal scales, improvements in model efficiency may come from targeting spatial and temporal resolution at specific processes and regions as needed.

Time-to-solution is also highly dependent on the inter-node communication requirements of spatial discretizations (latency, in particular). New spatial discretizations have the potential to target this cost by increasing the maximum stable time-step size [Paper 5]. Also, many schemes that support arbitrary order-of-accuracy can be tuned to enable improved performance on specific architectures. For instance, L1/L2 cache sizes and vector widths are hardware parameters that can provide dramatically improved performance depending on the order-of-accuracy of the numerical method.

Many advanced numerical techniques used in climate science rely on nonlinear and/or linear solvers. These include implicit time integrators, numerical optimization methods, and some uncertainty quantification/sensitivity analysis approaches. Preconditioner developments are often tied to specific problem classes. This is because an effective preconditioner must roughly approximate the inverse of a linear system, which may include interactions across a range of scales. This underscores the importance of the need for close collaboration between climate scientists and applied mathematicians in developing solvers/preconditioners that are effective for the linear systems arising in climate modeling.

In particular, due to the wide range of temporal scales, efficient time advancement techniques for stiff systems are essential. A number of recent advances in solver technologies have been made to improve solver convergence and robustness. These advances should be further investigated, ex-

tended, and adapted to climate situations. This includes implicit [Papers 30, 40], implicit-explicit (IMEX) [Papers 50, 54], operator split, and fully explicit [Paper 5] approaches, although it was widely agreed through ideas papers and discussions that algorithms with large time-step sizes that minimize data transfer across nodes and provide acceptable accuracy are the goal. Future efforts into nonlinear techniques should consider strategies to move beyond time integration advances such as traditional Newton-Krylov (e.g., [17]) to address the increase in model complexity. Two highlighted examples include nonlinear composite combination and nonlinear preconditioning [Paper 30] and Anderson accelerated fixed point methods [Paper 50]. Efforts to provide additional gains are suggested by using Newton-Krylov as a coarse operator for parallel-in-time and spectral-in-time methods [Paper 41]. Additionally, including optimized error constants for specific problems or methods that yield better conditioned linear sub-systems to improve linear solution times is of interest, especially within nonlinear solvers.

For linear problems, there has been significant progress in multigrid, domain decomposition, and physics-based solvers. Linear solver advances include convergence/robustness enhancements for certain problem classes or PDE operators (e.g., incompressible Navier-Stokes operators and/or Maxwell's equations), for problems that contain particular features (e.g., discontinuities or singularities), for certain discretizations (e.g., high-order discontinuous Galerkin methods, mimetic approaches, mixed finite elements), and to address meshes that might be thin with bad aspect ratios or might include boundary layers. There have also been advances in more general solver ideas such as multigrid methods based on K-cycles, compatible relaxation, and bootstrap- or energy-minimization-based multigrid schemes. Additionally, there have been advances in using auxiliary operators to precondition difficult linear systems. The general idea is that the auxiliary operator is more amenable to solver techniques and can still be used to precondition the original operator (e.g., shifted Laplacians for Helmholtz operators). For extremely difficult linear systems, direct solvers may be needed even though they typically have much higher memory and computational requirements. In addition to robustness/convergence enhancements for difficult problems, there continues to be ongoing research to improve solver scalability and performance on advanced/emerging architectures that might include hundreds of thousands of computing units. This includes methods such as communication-avoiding techniques, solvers with increased parallelism (e.g., concurrent processing of levels within a multigrid hierarchy), and solvers that have increased locality (e.g., nonlinear domain decomposition or recursive domain decomposition approaches where algorithm choices take into account the hierarchical nature of specific compute architectures). To leverage recent solver advances within the climate sciences, further research is vital to adapt and enhance solver ideas to the types of realistic complex situations that arise in the modeling of Earth systems. Ultimately, promising solver directions are often driven by a combination of domain-specific knowledge that comes from a combination of climate science and preconditioner expertise.

Many climate processes span a wide range of dynamic scales where fine spatial or temporal resolution is needed in a small subset of the domain in order to correctly model key dynamic processes, while large portions of the domain exhibit much coarser-scale dynamics. Modeling the entire system at the finest spatial and temporal resolutions needed is often both impractical and enormously inefficient. Adaptive mesh refinement (AMR) is an emerging technology that dynamically refines the grid during run-time so as to provide localized improvement in error norms and direct computational cost where it is needed [Papers 39, 42, 49]. Spatial mesh refinement refers to the addition of spatial degrees of freedom to specific areas of the mesh based on a *refinement criterion*. Temporal adaptivity provides the additional capacity to locally modify the time-step size based on the local resolution of the mesh, so as to ensure maximal performance while still satisfying the local Courant-Friedrichs-Lewy (CFL) condition. AMR offers the potential to better represent extreme weather events, such as tropical cyclones or atmospheric rivers, which are

similarly associated with fine-scale features that require high model resolution. In ice sheet models, variable-resolution and AMR models have already shown their value due to high-resolution requirements to resolve the dynamics of grounding lines and ice streams, which only occupy very small portions of the continental-scale ice sheets but exert a controlling influence on the dynamics of the entire ice sheets. Research on AMR for global climate models is ongoing, with recent work highlighting its potential [12, 25, 38].

3.3.3 Future directions

Other concerns regarding time to solution consider methods that reclaim computational time due to latency via context switching. One option is to simultaneously compute two ensemble simulations and swap between ensembles when a barrier is reached. Another option is to employ parallel-in-time physics parameterizations and perform physics computations while the barrier is in place. Also, the need for better controls on solution accuracy in coupled systems was highlighted [Paper 53]. In several examples, the output of coupled models is no longer convergent with increasing spatial and temporal resolution. This is an issue that needs to be continually addressed as coupled system complexity increases, since a large part of the increased time to solution is due to the combination of increased model complexity (more sub-models coupled together) and increased model resolution. If the coupled system is not convergent, then it's unclear that the increased cost is actually improving the model results.

Currently much effort is spent tuning physical parameterizations in various models, which is time-consuming and potentially brittle in the context of dynamic changes in climate systems. At some of the breakouts, the idea of replacing physical parameterizations with black-box machine learning or data tools was discussed as a promising alternative that would allow models to tune themselves, using observational data or through deterministic or stochastic [Papers 17, 18, 24] parameterization/superparameterization/ultraparameterization output. This would both improve time to solution for these models and enhance confidence in model output [Papers 25, 52].

3.4 Data management and analysis

3.4.1 Grand challenge

Climate is unique in terms of the scale of production, retention, and throughput of simulation output. This is due to the size and diversity of the international community that consumes the output and the breadth of questions spanning broad time and space scales for which the data is tasked to answer. Operationally, the expense of each simulation is such that all data that might possibly be desired or required for post-processing is kept for insurance. As a result, data reduction decisions are particularly difficult. Yet as we approach exascale, hardware constraints will force some novel solutions including relative data reduction via in situ analysis, more extensive utilization of low-power profile co-processors, and checking for potential data corruption. The increasing complexity of models, while a challenge for model throughput, unlocks many opportunities for new avenues for analyzing model output in ways that will provide understanding. However, the expense and complexity of methods to target data from these models present their own challenges of performance.

Uncertainty quantification and model ensembles bring further challenges to analytics that exacerbate the need for scalability and require the development of new techniques. They also bring experimental design issues for optimal exploration of model parameter space so that minimal-size ensembles can deliver information needed for understanding of the parameter space and uncertainty.

3.4.2 Opportunities and potential solutions

Functional data analysis The development of functional data analysis (FDA) in statistical science [Paper 55] presents an opportunity to facilitate a gradual transition to continuum process representations in climate science while addressing data reduction. At a high level, FDA begins with data, converts it to nonparametric functions (splines, wavelets, Fourier, etc.), and remains as functions for further reduction, storage, analysis, and visualization. Addressing the transition to continuum challenge is not unlike the transition from mathematical tables to mathematical functions that took place in the late part of the past century. We trade storage for mathematical complexity. The same can be said about the transition from dense matrix to sparse matrix methods. This has several consequences:

- Reduced storage and data movement from simulation to storage.
- No need to store tables of numbers. Instead, specifications of functions are stored.
- Data analysis is performed in functional space with a smaller memory footprint than traditional methods. This is where the mathematical complexity comes in, but FDA tools already exist for common analytics such as variability attribution, principal components, canonical correlation, clustering, and many other multivariate methods as well as new techniques that are unique to functional data.
- Enables principled model component coupling at any resolution.
- Provides rigorous interpolation schemes for visualization.
- Allows a fallback to traditional methods because data can be reconstructed at any resolution.

Standard decompositions, such as principal components analysis (or empirical orthogonal functions), can be extended to tensor decompositions (also known as higher-order principal components analysis). Algorithms for these are already available [1] even in the functional space, but their application to climate data remains open. Function space sampling algorithms [Paper 22] take the functional approach to reduce computational complexity of Monte Carlo for inference and UQ by sampling functions instead of points [6].

A wealth of methodology for statistical analysis and wrangling of data is encapsulated in the R Environment for Statistical Computing [32] and its thousands of packages. This includes a variety of multivariate methods with many already available to operate in functional space. Recent developments enable R scalability on large systems [5, 35] with various in situ options, including sharing a communicator and data staging. These developments enable novel data analysis algorithm scripting powered by the same scalable mathematical and communication libraries that currently power many petascale simulation codes.

Additional strategies Other ideas were suggested to improve data processing and work flow, which would enhance efficiency outside of the time-to-solution envelope as well as model credibility, through reproducibility.

Clustering and empirical orthogonal functions for variability attribution are commonly used in analysis of climate data. Their ubiquity and their potential data reduction properties make them good candidates for in situ processing. Scalable algorithms for exascale architectures are still lacking, although some have been developed in the past [Paper 56] and more recently in [5, 35], the latter being able to utilize pluggable dense linear algebra co-processor libraries.

In situ analysis as well as fast surrogates [Paper 21] provide an external observer the opportunity to check for potential data corruption. Such data reductions rely on smoothness properties in time or in space to detect anomalies, some of which may be corruptions. As data analysis is based on an explicit or implicit model, it provides an opportunity for building data-based reduced models as smart proxies. This is discussed in [Paper 9] and is also a potential result of the models in all the other approaches discussed in this section. An overarching strategy to automate large portions of the climate model work flow system could shift important but expensive and mundane steps in the model development, analysis, and publication process [Paper 32].

3.4.3 Future directions

Implementing scalable in situ and post-processing analysis algorithms requires common frameworks for data transport and data sharing. Options for in situ analysis include running on the same resources as the simulation, sharing a communicator, or as a service on separate resources through in-memory or burst buffer staging. Establishing middleware standards that enable these choices with a common interface (such as ADIOS [31]) will enable the use of the same analysis software whether in situ or in post-processing.

Analytics challenges brought by ensembles and uncertainty quantification can be addressed by stronger connections to the statistics community, where uncertainty plays a fundamental role. Techniques that include statistical design of experiments can lead to optimal information content in ensembles and simplify analysis. Fundamental mathematics in analytics, such as dense and sparse linear algebra, is available in scalable libraries and continues to be developed by the applied mathematics community. Engaging this as infrastructure in easily reconfigurable ways is needed for scalable analytics. This has already started in the R language [35], but a lot more needs to be done to provide diverse scalable techniques from modern statistical science and machine learning for climate.

4 Structure of Collaboration

4.1 Existing tools and collaborations

Due to the computational demands of modeling the climate system with high fidelity, there is already a rich history of collaboration and communication between applied mathematicians, computer scientists, statisticians, and climate scientists. In fact, climate models have led the way in organizing successful collaborations between these disciplines in the service of improving the state of climate science. One example of such a collaboration was through the DOE-supported ISICLES effort. The IPCC AR5 report [39] called attention to the large uncertainties in predictions of sea level rise resulting from the inadequacy of then-current ice sheet models to the task of modeling the response of the Greenland and Antarctic ice sheets to projected climate forcing. As a result, DOE ASCR, as an expansion of an existing partnership with DOE BER, issued a unique call for ice sheet modeling efforts that would bridge the gap between computational expertise in ASCR and the near-term needs for improved ice sheet modeling efforts within BER. The result was six funded projects that covered a range of ice sheet modeling strategies ranging from Lagrangian particle models to spatial meshes that better capture ice fracture. While some of the projects were more successful than others, the net result was an infusion of applied mathematics and computer science expertise into the relatively young field of ice sheet modeling. This helped to propel the DOE portion of that effort forward to the point where the DOE is now a world leader in ice sheet model sophistication and maturity.

The joint ASCR and BER partnership in the Scientific Discovery through Advanced Computing (SciDAC) Program is a long-term investment from the DOE Office of Science that has propelled climate modeling to petascale and has a new call towards the exascale level of computing. It varies with each call but generally consists of centralized mathematics and computer science “Institutes” that support a number of “Applications Partnerships” in each of the sub-offices in the Office of Science. This model has proved an effective way to organize the deployment of advances into the various application fields. Climate modeling has benefited from this collaboration since the first SciDAC program, and the key driver for success has been the joint participation from applied mathematicians, computer scientists, and a variety of science domain experts with expertise in climate science.

More recently, the ACME Project has stood up to create a global coupled Earth system model that is designed to answer key science questions of interest to the DOE mission. It is achieving its science goals through effective use of DOE leadership-class computing via a significant focus on advanced performance, data management and workflow, and software engineering. Other international climate centers, including the work at the United Kingdom’s European Centre for Medium-Range Weather Forecasts (ECMWF) presented in a plenary talk by George Mozdzynski at the workshop, also recognize the importance of and invest in research to best utilize the largest available computing resources to achieve the best modeling results for weather and climate prediction.

4.2 Facilitation of communication across disciplines

A primary issue is communication across the various disciplines that all feed into a successful climate modeling enterprise. While it is tempting to think of collaboration in terms of tools, as in “which tools facilitate collaboration” or “how can we adapt tools to be of use to climate scientists,” in the breakout session discussions of effective communication, climate scientists and mathematics and computer science specialists invariably dwell on processes rather than concrete tools.

Embedding applied mathematics, statistics, and computer science experts in the domain science side has proven through the successes of SciDAC and in the instantiation of ACME to be a useful strategy for intellectual and technology transfer, particularly from the applied mathematics domain to the climate and computer science domains. For example, applied mathematics researchers in the ISICLES and SciDAC-funded PISCEES ice sheet projects have regularly contributed both to applied mathematics and glaciology-related conferences and journals. Embedding works because it creates a hybrid person who is conversant in both the specific demands of the application space along with the mathematical background to either apply directly or to be able to communicate back to others in the mathematical and computer science sides. The key challenge is for these hybrid scientists to be recognized and motivated in their efforts. Recognizing statistics as distinct from applied mathematics puts more emphasis on the increasing need for modern data analysis and data assimilation required for analysis going forward in addition to current and future simulation science’s use of PDE solution methodologies to advance model execution speed and accuracy.

The importance of meaningful (but tractable) benchmarks came up multiple times during AXI-CCS discussions. Besides the significant utility in building model confidence, benchmark problems are the way we communicate between fields, making it much easier to transfer information about model attributes and issues rather than simply “kicking something over the fence” back to the applied mathematics and computer science communities. A useful set of benchmark problems also provides a well-defined path to entry into a science domain for the embedded interloper from the mathematics side.

Discussion about collocation of applied mathematicians, computer scientists and climate sci-

entists indicated a likely benefit, although the logistics of such an arrangement is complicated, especially when climate scientists themselves are also not colocated, as in the case of ACME. The biggest challenge is how to ensure continuous engagement among the climate science, mathematics and computer science communities—how to continue to find and encourage people who are effectively trained in climate science, mathematics, and computer science. Usually, job advertisements targeting the full range of postdoc to senior scientist career stages that prioritize personnel with combined mathematics and/or computer science along with Earth science knowledge do not generate many candidates.

4.3 Outreach and learning as potential solutions

Direct marketing to undergraduate and graduate mathematics and computer science students can be useful—the perception that climate science is one of the grand challenge scientific problems of our time is a strong motivator in many instances to get non-climate scientists involved. Workforce training efforts like the DOE-managed Computational and Stewardship Science Graduate Fellowships (CSGF, SSGF) and Office of Science Graduate Student Research (SCGSR) programs provide a useful mechanism for channeling graduate students in the right directions to be able to contribute meaningfully as early-career scientists. For example, the CSGF actively steers the courses of study of its fellows, ensuring solid backgrounds in mathematics and computer science while maintaining a focus on producing true computational scientists who are productive in their respective scientific fields. The newly formed Science Graduate Student Research (SCGSR) program within DOE involves graduate students being hosted at a national laboratory for a 3-month appointment. While this is not much time to develop skills, it does start the process for future development and interaction. Additional venues to train multidisciplinary scientists include topical workshops, webinars, and/or summer schools. A team composed of ACME scientists, named team “HACME,” have attended the last two Oak Ridge Leadership Computing “Hackathons.” These events connect domain scientists, mathematicians, computer scientists, and compiler vendors to learn how to optimize the performance of their code on the Oak Ridge Leadership Computing Facility (OLCF) machines.

One of the side benefits of the AXICCS workshop was a large number of side conversations between members of the different communities who otherwise would not have met. It is likely that at least some of these conversations will lead to productive lines of work. In the workshop, these informal gatherings became known as the “hot tub” idea—the more chances for people from the different fields to interact, the more chances for meaningful and substantive interaction to take place. While it may be hard to justify in a project-driven and end-result-driven funding ecosystem, such opportunities for cross-cutting contact and discussions stand to enormously benefit the state of climate science as a whole, which will then feed back into the applied mathematics and computer science disciplines. In the past, the annual SciDAC project meetings have provided such opportunities, although by necessity they were limited to those projects already in SciDAC collaborations and covered many subject areas beyond climate as the target area within which to form connections.

Networking opportunities for graduate students and postdoctoral scholars have also been developed through dual-purpose (education and workshop) events such as the Dynamical Core Model Intercomparison Project (DCMIP) summer school and workshop [18]. Every 4 years since 2008 this program has brought together students, postdoctoral researchers, application scientists, model developers, and other experts with the objective of (a) supporting a massive learning initiative directed at dynamical core research and (b) leveraging an enthusiastic cohort of participants to undertake the model intercomparison. The outcome of this intercomparison has been praised by both participants and modeling groups and featured wide support from major federal agencies. The

event in 2016 was sponsored by the National Center for Atmospheric Research’s (NCAR) Computational Information Systems Laboratory (CISL), NOAA, NASA, DOE, the National Science Foundation (NSF), and the Office of Naval Research (ONR). Similar efforts could likewise be developed for other model intercomparisons that require mathematics and computer science experts to bring together a cohort of experts and students, for instance, one specifically focused on coupling of model components. However, workshops such as these rely on the target model component(s) to have a sufficiently diverse and broad community to make intercomparison worthwhile. Further, such an effort also relies on community agreement on a “standardized” suite of benchmarks to be employed. Nonetheless, projects such as these provide an excellent bridge between laboratories and academia, with positive outcomes for both education and productivity.

4.4 Possible follow-on activities

The discussions at the workshop have generated a collection of big ideas and cross-cutting issues that climate scientists, mathematicians, and computer scientists may want to pursue collaboratively in order to advance climate modeling and simulation. It is important to keep the momentum going so that the big ideas can be realized. In addition, it will be beneficial to engage additional researchers from the three communities beyond this workshop group. This allows the big ideas and cross-cutting issues identified in this report to be fleshed out by bigger groups, and possibly motivating additional ideas and issues. For example, follow-up workshops can be organized to focus on each of the big ideas. Another possibility is to have special journal articles and/or issues calling for research in such areas.

One step this group has taken is to organize a minisymposium at the upcoming SIAM Conference on Computational Science and Engineering, which will be held at the end of February 2017. Several members of the program committee of this workshop will present the findings in this report at the minisymposium. As alluded to earlier, this is not meant to signal the end of the workshop. Rather, it should be considered the beginning of a larger and broader effort.

5 Conclusions

The Advances in Mathematical and Computational Climate Modeling (AXICCS) workshop was different from many workshops that focus on scientific discussions and brainstorming of new research directions. It was unique because the format was designed so that the participants focused on identifying future areas of research that would materially impact climate simulation without restrictions of time, funding, or expertise. A key attribute was to discuss, in tandem, the key mathematical and computer science advances that would be required to pursue new ideas we identified. By allowing discussions across many scientific domains and among colleagues at different stages of their career, the goal was to provide the opportunity for many voices and ideas to be heard and discussed.

Although many new science ideas were considered, we developed several major findings about how to improve Earth system models into more formal strategies within Section 2. These spanned improvements in managing model complexity, strategies to maintain and improve confidence in these complex and interacting models, and the coupling of processes to enable a better understanding of our rich and diverse Earth system and its changes. We also presented a novel path to provide a simulation of the Earth without arbitrary boundaries fostered by community tradition.

We prioritized key tools from the mathematics and computer science communities we thought would be most beneficial to bring these ideas to a more mature state of evaluation and implementation. These included computer science methods such as language and compiler tools to address

upcoming and dramatic hardware changes, mathematical tools for model development that will target the new hardware effectively to maximize space and timescale limitations, strategies to treat and minimize software complexity so that climate scientists and non-climate scientists can develop and evaluate code use to produce these large and interacting models, and data analysis and management, which address statistical tools to unlock the new insights within simulation output as well as new methods for data storage and workflow to enable the more complex models to be managed. These have been outlined in Section 3.

Finally, in Section 4, we discussed the possible path forward. Additional opportunities, through more interactions with these and additional climate scientists, would flesh out these ideas and also motivate additional ideas. Next steps would include workshops to generate expansions and specifics of each idea, additional and incentivized interactions with colleagues from the climate, mathematics, and computer science communities beyond the limited attendance of this workshop, and more communication of these groups through traditional and novel methods, for example a special journal article and issue calling for research in these areas.

Acknowledgments

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Offices of Biological and Environmental Research (BER) and Advanced Scientific Computing and Research (ASCR). The workshop organizers thank Drs. Dorothy Koch from BER and Randall Lavolette from ASCR for their encouragement and support. Logistics were handled with excellent care by Andrew Flatness and Tracey Vieser. The organizers are grateful to the participants of the workshop, especially the program committee and those who served as plenary speakers, poster and oral presenters, note-takers, and engaged participants in the discussion. The organizers are also grateful to Deborah Stevens, Kathy Jones, and Madeleine Evans for editing the report.

References

- [1] G. Allen. Sparse higher-order principal components analysis. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22, page 27:36, 2012.
- [2] K. Asanovi and D. A. Patterson. Instruction sets should be free: The case for RISC-V. Technical Report UCB/EECS-2014-146, EECS Department, University of California, Berkeley, Aug 2014. URL <http://www.eecs.berkeley.edu/Pubs/TechRpts/2014/EECS-2014-146.html>.
- [3] C. S. Bretherton. Frontiers in multiscale and global simulation of boundary layer clouds, 2016. Presentation at the Workshop on Advancing X-cutting Ideas for Computational Climate Science (AXICCS), Monday, September 12, 2016.
- [4] R. Cess, G. L. Potter, J. P. Blanche, G. J. Boer, S. J. Ghan, J. T. Kiehl, H. L. Treut, Z.-X. Li, X.-Z. Liang, J. Mitchell, J.-J. Morcrette, D. A. Randall, M. R. Riches, E. Roeckner, U. Schlese, A. Slingo, K. E. Taylor, W. M. Washington, R. T. Wetherald, and I. Yagai. Interpretation of cloud–climate feedback as produced by 14 atmospheric general circulation models. *Science*, 245(4917):513–516, 1989.
- [5] W.-C. Chen, G. Ostrouchov, D. Pugmire, Prabhat, and M. Wehner. A parallel EM algorithm for model-based clustering applied to the exploration of large spatio-temporal data. *Technometrics*, 55(4):513–523, 2013. doi:[10.1080/00401706.2013.826146](https://doi.org/10.1080/00401706.2013.826146).
- [6] S. L. Cotter, G. O. Roberts, A. M. Stuart, D. White, et al. MCMC methods for functions: Modifying old algorithms to make them faster. *Statistical Science*, 28(3):424–446, 2013.
- [7] P. M. Cox, D. Pearson, B. B. Booth, P. Friedlingstein, C. Huntingford, C. D. Jones, and C. M. Luke. Sensitivity of tropical carbon to climate change constrained by carbon dioxide variability. *Nature*, 494(7437):341–344, Feb. 2013. doi:[10.1038/nature11882](https://doi.org/10.1038/nature11882).
- [8] M. Deneroff. Building an SOC: How to do it? What will it cost? In *Workshop on System-on-a-chip design for HPC*, 2014.
- [9] R. Dennard, F. Gaensslen, H. Yu, V. Rideout, E. Bassous, and A. Leblanc. Design of ion-implanted MOSFETs with very small physical dimensions. *IEEE Journal of Solid-State Circuits*, 9(5):256–268, 1974.
- [10] J. Dongarra. Report on the Sunway TaihuLight system. Technical Report UT-EECS-16-742, University of Tennessee, 2016.
- [11] D. Donofrio, L. Oliker, J. Shalf, M. F. Wehner, C. Rowen, J. Krueger, S. Kamil, and M. Mohiyuddin. Energy-efficient computing for extreme-scale science. *IEEE Computer*, 42(11):62–71, 2009. doi:[10.1109/MC.2009.353](https://doi.org/10.1109/MC.2009.353).
- [12] J. O. Ferguson, C. Jablonowski, H. Johansen, P. McCorquodale, P. Colella, and P. A. Ullrich. Analyzing the Adaptive Mesh Refinement (AMR) characteristics of a high-order 2D cubed-sphere shallow-water model. *Monthly Weather Review*, 2016. In press.
- [13] H. Fu, J. Liao, J. Yang, L. Wang, Z. Song, X. Huang, C. Yang, W. Xue, F. Liu, F. Qiao, W. Zhao, X. Yin, C. Hou, C. Zhang, W. Ge, J. Zhang, Y. Wang, C. Zhou, and G. Yang. The Sunway TaihuLight supercomputer: System and applications. *Science China Information Sciences*, 59(7):072001, 2016. ISSN 1869-1919. doi:[10.1007/s11432-016-5588-7](https://doi.org/10.1007/s11432-016-5588-7). URL <http://dx.doi.org/10.1007/s11432-016-5588-7>.

- [14] A. Hall and X. Qu. Using the current seasonal cycle to constrain snow albedo feedback in future climate change. *Geophys. Res. Lett.*, 33(3):L03502, Feb. 2006. doi:[10.1029/2005GL025127](https://doi.org/10.1029/2005GL025127).
- [15] F. M. Hoffman, J. T. Randerson, V. K. Arora, Q. Bao, P. Cadule, D. Ji, C. D. Jones, M. Kawamiya, S. Khatiwala, K. Lindsay, A. Obata, E. Shevliakova, K. D. Six, J. F. Tjiputra, E. M. Volodin, and T. Wu. Causes and implications of persistent atmospheric carbon dioxide biases in Earth system models. *J. Geophys. Res. Biogeosci.*, 119(2):141–162, Feb. 2014. doi:[10.1002/2013JG002381](https://doi.org/10.1002/2013JG002381).
- [16] F. Hourdin, T. Mauritsen, A. Gettelman, J.-C. Golaz, V. Balaji, Q. Duan, D. Folini, D. Ji, D. Klocke, Y. Qian, F. Rauser, C. Rio, L. Tomassini, M. Watanabe, and D. Williamson. The art and science of climate model tuning. *Bull. Amer. Meteor. Soc.*, 2017. doi:[10.1175/BAMS-D-15-00135.1](https://doi.org/10.1175/BAMS-D-15-00135.1). URL <http://dx.doi.org/10.1175/BAMS-D-15-00135.1>.
- [17] J.-F. Lemieux, S. F. Price, K. J. Evans, D. A. Knoll, A. G. Salinger, and T. Payne. Implementation of the Jacobian-free Newton-Krylov method for solving the first-order ice sheet momentum balance. *J. Comp. Phys.*, 230:6531–6545, 2011. doi:[10.1016/j.jcp.2011.04.037](https://doi.org/10.1016/j.jcp.2011.04.037).
- [18] C. Jablonowski, P. A. Ullrich, J. Kent, K. A. Reed, M. A. Taylor, P. H. Lauritzen, and R. D. Nair. The 2012 dynamical core model intercomparison project (DCMIP). In *AGU Fall Meeting Abstracts, A53C-0160*, 2012. URL http://fallmeeting.agu.org/2012/files/2012/12/AGU_poster_Jablonowski_Dec_2012.pdf. (Accessed 3 July 2016).
- [19] M. Kennedy and A. O’Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 63(3):425–464, 2001. doi:[10.1111/1467-9868.00294](https://doi.org/10.1111/1467-9868.00294).
- [20] D. E. Keyes, L. C. McInnes, C. Woodward, W. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. P. Randles, D. Reynolds, B. Rivire, U. Rde, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth. Multiphysics simulations: Challenges and opportunities. *Int. J. High Perf. Comp. Apps.*, 27(1):4–83, 2013. doi:[10.1177/1094342012468181](https://doi.org/10.1177/1094342012468181).
- [21] M. Khairoutdinov, D. Randall, and C. DeMott. Simulations of the atmospheric general circulation using a cloud-resolving model as a superparameterization of physical processes. *Journal of the Atmospheric Sciences*, 62(7):2136–2154, 2005.
- [22] M. F. Khairoutdinov and D. A. Randall. Cloud resolving modeling of the arm summer 1997 iop: Model formulation, results, uncertainties, and sensitivities. *Journal of the Atmospheric Sciences*, 60(4):607–625, 2003.
- [23] J. Krueger, D. Donofrio, J. Shalf, M. Mohiyuddin, S. Williams, L. Oliker, and F.-J. Pfreund. Hardware/software co-design for energy-efficient seismic modeling. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC ’11*, pages 73:1–73:12, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0771-0. doi:[10.1145/2063384.2063482](https://doi.org/10.1145/2063384.2063482). URL <http://doi.acm.org/10.1145/2063384.2063482>.
- [24] R. M. Maxwell and L. E. Condon. Connections between groundwater flow and transpiration partitioning. *Science*, 353(6297):377–380, 2016.

- [25] P. McCorquodale, P. Ullrich, H. Johansen, and P. Colella. An adaptive multiblock high-order finite-volume method for solving the shallow-water equations on the sphere. *Communications in Applied Mathematics and Computational Science*, 10(2):121–162, 2015.
- [26] G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8), April 1965.
- [27] National Research Council. *Assessing the Reliability of Complex Models: Mathematical and Statistical Foundations of Verification, Validation, and Uncertainty Quantification*. National Academies Press, May 2012. doi:10.17226/13395. URL <https://www.nap.edu/download/13395>.
- [28] NRC-Committee. *A National Strategy for Advancing Climate Modeling*. The National Academies Press, 2012. ISBN 9780309259774. URL http://www.nap.edu/openbook.php?record_id=13430. Committee on a National Strategy for Advancing Climate Modeling; Board on Atmospheric Studies and Climate; Division on Earth and Life Studies.
- [29] H. Parishani. Explicit representation of low clouds in a global atmospheric model: initial results from Ultra-Parameterized CAM5. In *2015 AGU Fall Meeting*. AGU, 2015.
- [30] D. A. Patterson. Latency lags bandwidth. *Commun. ACM*, 47(10):71–75, Oct. 2004. ISSN 0001-0782. doi:10.1145/1022594.1022596. URL <http://doi.acm.org/10.1145/1022594.1022596>.
- [31] N. Podhorszki, Q. Liu, J. Logan, J. Mu, H. Abbasi, J.-Y. Choi, and S. A. Klasky. Adios 1.9 user’s manual. Technical report, ORNL/TM-2009/100, 2015.
- [32] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016. URL <https://www.R-project.org>.
- [33] D. M. Romps. The dry-entropy budget of a moist atmosphere. *Journal of the Atmospheric Sciences*, 65(12):3779–3799, 2008.
- [34] M. Satoh, T. Matsuno, H. Tomita, H. Miura, T. Nasuno, and S.-I. Iga. Nonhydrostatic icosahedral atmospheric model (nicam) for global cloud resolving simulations. *Journal of Computational Physics*, 227(7):3486–3514, 2008.
- [35] D. Schmidt, W.-C. Chen, M. A. Matheson, and G. Ostrouchov. Programming with BIG data in R: Scaling analytics from one to thousands of nodes. *Big Data Research*, 2016. URL <http://dx.doi.org/10.1016/j.bdr.2016.10.002>.
- [36] D. E. Shaw, M. M. Deneroff, R. O. Dror, J. S. Kuskin, R. H. Larson, J. K. Salmon, C. Young, B. Batson, K. J. Bowers, J. C. Chao, M. P. Eastwood, J. Gagliardo, J. P. Grossman, C. R. Ho, D. J. Ierardi, I. Kolossváry, J. L. Klepeis, T. Layman, C. McLeavey, M. A. Moraes, R. Mueller, E. C. Priest, Y. Shan, J. Spengler, M. Theobald, B. Towles, and S. C. Wang. Anton, a special-purpose machine for molecular dynamics simulation. *Commun. ACM*, 51(7): 91–97, July 2008. ISSN 0001-0782. doi:10.1145/1364782.1364802. URL <http://doi.acm.org/10.1145/1364782.1364802>.
- [37] B. J. Soden, D. L. Jackson, V. Ramaswamy, M. D. Schwarzkopf, and X. Huang. The radiative signature of upper tropospheric moistening. *Science*, 310(5749):841–844, 2005. doi:10.1126/science.1115602.

- [38] A. St-Cyr, C. Jablonowski, J. M. Dennis, H. M. Tufo, and S. J. Thomas. A comparison of two shallow-water models with nonconforming adaptive grids. *Mon. Wea. Rev.*, 136(6):1898–1922, 2008.
- [39] T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. Midgley, editors. *Summary for Policymakers*, book section SPM, page 130. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2013. ISBN ISBN 978-1-107-66182-0. doi:[10.1017/CBO9781107415324.004](https://doi.org/10.1017/CBO9781107415324.004). URL www.climatechange2013.org.
- [40] M. F. Wehner, L. Oliker, J. Shalf, D. Donofrio, L. A. Drummond, R. Heikes, S. Kamil, C. Kono, N. Miller, H. Miura, M. Mohiyuddin, D. Randall, and W.-S. Yang. Hardware/software co-design of global cloud system resolving models. *Journal of Advances in Modeling Earth Systems*, 3(4), 2011. ISSN 1942-2466. doi:[10.1029/2011MS000073](https://doi.org/10.1029/2011MS000073). URL <http://dx.doi.org/10.1029/2011MS000073>. M10003.
- [41] H. Zhu, H. Hendon, and C. Jakob. Convection in a parameterized and superparameterized model and its role in the representation of the MJO. *Journal of the Atmospheric Sciences*, 66(9):2796–2811, 2009.

Acronyms used in the report

ACME: Accelerated Climate Model for Energy
AMR: Adaptive Mesh Refinement
ASCR: Advanced Scientific Computing Research
BER: Biological and Environmental Research
CS: Computer Science
CESM: Community Earth System Model
CISL: Computational Information Systems Laboratory
CMIP: Coupled Model Intercomparison Project
CMOS: Complimentary metal oxide semiconductor
CRM: Cloud-resolving Model
CSGF: Computational Science Graduate Fellowship
DA: Data Assimilation
DOE: U.S. Department of Energy
DSL: Domain-Specific Language
ECMWF: European Centre for Medium-Range Weather Forecasts
ENSO: El Niño Southern Oscillation
ESM: Earth System Model
FDA: Functional Data Analysis
GCM: Global Climate Model
IMEX: Implicit-Explicit
IPCC: Intergovernmental Panel on Climate Change
ISCCP: International Satellite Cloud Climatology Project
ISICLES: Ice Sheet Initiative for CLimate ExtremeS
LCF: Leadership-Class Facility
LES: Large Eddy Simulation
MJO: Madden-Julian Oscillation
MLMC: Multilevel Monte Carlo
NCAR: National Center for Atmospheric Research
NASA: National Aeronautics and Space Administration
NOAA: National Oceanographic and Atmosphere Administration
NSF: National Science Foundation
ONR: Office of Naval Research
PDE: Partial Differential Equation
PISCEES: Predicting Ice Sheet and Climate Evolution at Extreme Scales
SciDAC: Scientific Discovery through Advanced Computing
SCGSR: Office of Science Graduate Student Research
SSGF: Stewardship Science Graduate Fellowship
UQ: Uncertainty Quantification

Workshop Organizers

Katherine J. Evans, Oak Ridge National Laboratory
Esmond Ng, Lawrence Berkeley National Laboratory

Program Committee

Peter Caldwell, Lawrence Livermore National Laboratory
Forrest M. Hoffman, Oak Ridge National Laboratory
Charles Jackson, University of Texas, Austin
Kerstin Klees Van Daam, Brookhaven National Laboratory
L. Ruby Leung, Pacific Northwest National Laboratory
Dan Martin, Lawrence Berkeley National Laboratory
George Ostrouchov, Oak Ridge National Laboratory
Ray Tuminaro, Sandia National Laboratories
Paul Ullrich, University of California, Davis
Stefan Wild, Argonne National Laboratory
Samuel Williams, Lawrence Berkeley National Laboratory

Plenary Speakers

Chris Bretherton, University of Washington
William Collins, Lawrence Berkeley National Laboratory
Petros Koumoutsakos, ETH Zurich
George Mozdzyński, European Center for Medium Range Forecasts

Workshop Participants

Melissa Allen, Oak Ridge National Laboratory
Pavel Bochev, Sandia National Laboratories
Jed Brown, University of Colorado, Boulder
Peter Caldwell, Lawrence Livermore National Laboratory
Donna Calhoun, Boise State University
Nathan Collier, Oak Ridge National Laboratory
Emil Constantinescu, Argonne National Laboratory
Bert Debusschere, Sandia National Laboratories
George Delic, HiPERiSM Consulting, LLC
Pedro Diniz, University of Southern California
David Donofrio, Lawrence Berkeley National Laboratory
David Gardner, Lawrence Livermore National Laboratory
David Hall, University of Colorado Boulder
Matthew Hoffman, Los Alamos National Laboratory
Christiane Jablonowski, University of Michigan
Robert Jacob, Argonne National Laboratory
Philip Jones, Los Alamos National Laboratory

Joseph Kennedy, Oak Ridge National Laboratory
Andreas Kreienbuehl, Lawrence Berkeley National Laboratory
Kody Law ,Oak Ridge National Laboratory
Li Li, Penn State University
Yangang Liu , Brookhaven National Laboratory
Fei Lu, UC Berkeley/Lawrence Berkeley National Laboratory
Vijay Mahadevan, Argonne National Laboratory
Salil Mahajan, Oak Ridge National Laboratory
Benjamin Mayer, Oak Ridge National Laboratory
Richard T. Mills, Intel
Sri Hari Krishna Narayanan, Argonne National Laboratory
Matthew Norman, Oak Ridge National Laboratory
George Ostrouchov, Oak Ridge National Laboratory
Maruro Perego, Sandia National Laboratories
Tom Peterka, Argonne National Laboratory
Stan Posey, Nvidia
Stephen Price, Los Alamos National Laboratory
David Randall, Colorado State University
Jaideep Ray, Sandia National Laboratories
Andy Salinger, Sandia National Laboratories
Bill Spatz, Sandia National Laboratories
Georg Stadler, Courant Institute, New York University
Irina Tezaur, Sandia National Laboratories
Raymond Tuminaro, Sandia National Laboratories
Paul Ullrich, University of California, Davis
Nathan Urban, Los Alamos National Laboratory
Haim Waisman, Columbia University
Hui Wan, Pacific Northwest National Laboratory
Dali Wang, Oak Ridge National Laboratory
Shanlin Wang, Los Alamos National Laboratory
Carol Woodward, Lawrence Livermore National Laboratory

Workshop Agenda

Workshop on Advancing X-cutting Ideas for Computational Climate Science (AXICCS)

September 12-14, 2016

Hilton Rockville, 1750 Rockville Pike, Rockville, MD 20852

Agenda

Monday, September 12, 2016

<i>Start Time</i>	<i>End Time</i>	<i>Topic</i>
9:00AM	9:30AM	Welcome
9:30AM	10:30AM	Plenary: Bill Collins , Lawrence Berkeley National Laboratory <i>Climate Simulation at Impactful Scales: Charge for a New Physics Paradigm</i>
10:30AM	11:00AM	Coffee Break
11:00AM	12:30PM	Breakout #1

Breakout Topic 1A: Climate Science Problems in Coupling

*Moderators: **Peter Caldwell and Forrest Hoffman***

Speakers: 40 minutes total

Discussion: 50 minutes

L. Li, Y. Shi, C. Duffy. *Building Computational Bridges Across the Water, Ecosystem, and Soil Biogeochemistry Disciplines*

R. Mills and F. Hoffman. *Machine-learning guided, multi-resolution approaches to high-fidelity representation of global hydrology in ESMs*

H. Waisman, J. Bassis, S. Price, R. Tuminaro and I. Tezaur. *A physics based iceberg calving model coupled with a global ice-sheet flow model for accurate assessment of sea level rise*

M. Hoffman, L. Bertagna, M. Gunzburger, M. Perego and Stephen Price. *Realistic Subglacial Hydrology For Improved Ice Sheet-Climate Coupling and Sea Level Prediction*

Breakout Topic 1B: Climate Model Complexity and Scaling

*Moderators: **Ruby Leung and Paul Ullrich***

Speakers: 45 minutes total

Discussion: 45 minutes

M. Allen, M. Branstetter, O. Omitaomu. *Embedded Urban Framework for ACME Regions of Refined Resolution*

P. Bochev, K. Evans, M. Gunzburger and K. Peterson. *Optimization-Based Heterogeneous Numerical Methods: an Abstraction for Mathematically Rigorous Coupling of Earth System Models*

C. Jablonowski, J. Ferguson, H. Johansen and P. Colella. *Transforming Climate Modeling via Scale-Adaptive Computational Techniques*

W. Maslowski, A. Roberts, E. Hunke, F. Giraldo and M. Kopera. *Sea Ice Modeling Across Scales at Exascale and Beyond*

Breakout Topic 1C: Climate Model Ensembles and Uncertainty Quantification

*Moderators: **Charles Jackson and Michael Prather***

Speakers: 40 minutes total

Discussion: 50 minutes

Workshop Agenda

A. Salinger, E. Phipps and J. Fyke. *Embedded Ensembles*

S. Mahajan, K. Evans and M. Norman. *Expanding the Utility of High-Resolution Global Climate Models via Short Ensembles*

S. Price, M. Perego and G. Stadler. *Optimization and Uncertainty Quantification of Ice Sheet Models*

S. Wang, N. Urban, M. Maltrud and Alexandra Jonko. *Automation of parameterization and structure selection of ocean biogeochemical models*

12:30PM	2:00PM	Lunch
2:00PM	2:30PM	Outbriefs from Breakout #1 (all)
2:30PM	3:30PM	Plenary: Christopher S. Bretherton , University of Washington <i>Frontiers in Multiscale and Global Simulation of Boundary Layer Clouds and Their Interactions with Climate</i>
3:30PM	4:00PM	Coffee Break
4:00PM	5:30PM	Breakout #2: Math and Computer Science Advances

Breakout Topic 2A: Coupling, PDEs, and linear algebra

*Moderators: **Ray Tuminaro and Dan Martin***

Speakers: 45 minutes total

Discussion: 45 minutes

M. Perego, S. Price and A. Salinger. *Next generation implicit solvers and analysis algorithms for ice sheet modeling*

J. Brown. *Higher Standards on the Control of Numerical Accuracy*

M. Norman. *New Temporal and Spatial Algorithms for Atmospheric Climate Models*

Breakout Topic 2B: Optimization and Statistics

*Moderators: **Stefan Wild and George Ostrouchov***

Speakers: 40 minutes total

Discussion: 50 minutes

O. Ghattas and **G. Stadler**. *From Data through Inference to Optimization under Uncertainty: Towards End-to-End Climate Model-Based Decision-Making*

J. Ray, L. Swiler, G. Pau, G. Bisht, F. Hoffman, M. Huang, Z. Hou and X. Chen. *Improving predictive capability of land surface models through robust statistical calibration techniques*

N. Urban. *Climate Model Uncertainty Quantification*

Breakout Topic 2C: Computational Performance and Data Management

*Moderators: **Sam Williams and Kerstin Kleese Van Dam***

Speakers: 40 minutes total

Discussion: 50 minutes

P. Ullrich, G. Jost, B. Lelbach and H. Johansen. *Exascale-Ready Programming Models for Climate*

D. Wang, O. Hernandez, G. Lopez and F. Winkler. *Compiler-based software analysis toolkit for climate model development*

5:30PM	7:00PM	Poster Session
--------	--------	----------------

Workshop Agenda

Workshop on Advancing X-cutting Ideas for Computational Climate Science (AXICCS)

September 12-14, 2016

Hilton Rockville, 1750 Rockville Pike, Rockville, MD 20852

Agenda

Tuesday, September 13, 2016

<i>Start Time</i>	<i>End Time</i>	<i>Topic</i>
9:00AM	9:30AM	Outbriefs from Breakout #2 (all)
9:30AM	10:30AM	Plenary: Petros Koumoutsakos , ETH Zurich, Switzerland <i>The Art of Computational Science: Closing Gaps, Forming Alloys</i>
10:30AM	11:00AM	Coffee Break
11:00AM	12:30PM	Breakout #3: Climate Response to Math and CS Ideas Breakout Topic 3A: Same as 1A Breakout Topic 3B: Same as 1B Breakout Topic 3C: Same as 1C
12:30PM	2:00PM	Lunch
2:00PM	2:30PM	Outbriefs from Breakout #3 (all)
2:30PM	3:30PM	Plenary: George Mozdzynski , European Centre for Medium-Range Weather Forecasts, UK <i>Addressing Future Scalability and Power Challenges at the European Centre for Medium-Range Weather Forecasts (ECMWF)</i>
3:30PM	4:00PM	Coffee Break
4:00PM	4:30PM	Wrap-up

Workshop Agenda

Workshop on Advancing X-cutting Ideas for Computational Climate Science (AXICCS)

September 12-14, 2016

Hilton Rockville, 1750 Rockville Pike, Rockville, MD 20852

Agenda

Wednesday, September 14, 2016

<i>Start Time</i>	<i>End Time</i>	<i>Topic</i>
8:30AM	1:00PM	PC Only Report Writing

Contributed ideas paper 1

Optimization-Based Heterogeneous Numerical Methods: an Abstraction for Mathematically Rigorous Coupling of Earth System Models November 20, 2015

Pavel Bochev

Center for Computing Research, Sandia National Laboratories¹, Albuquerque, NM 87185

Katherine Evans (ORNL), Max Gunzburger (FSU), Kara Peterson (SNL)

Abstract

Earth System Models (ESM) are an instance of *Heterogeneous numerical models* (HNM), which incorporate diverse mathematical and computational parts to describe complex, multidisciplinary systems. Predictive simulation of such systems requires a mathematically rigorous HNM coupling strategy to enable these parts to function as one whole. We posit that a coupling strategy, which couches HNM formulation and data transfers between its parts into a constrained optimization problem has the potential to deliver stable, accurate and property-preserving ESMs.

Background/Research to Date

Earth system models (ESM), such as the DOE funded Accelerated Climate Modeling for Energy (ACME) project, comprise multiple components; see Fig. 1(L), admitting a range of options for mathematical and numerical models. For example, atmosphere, ocean, land surface, sea ice and ice sheet models can use equation hierarchies wherein each level is a reduced dynamic of the level above. Variations in the complexities among these levels engenders a corresponding diversity of the numerics resulting in ad hoc tradeoffs between preservation of key physical properties and performance. The ESM components generally stand alone, connecting to each other via a flux coupler component. They resemble cities built without public planning.

A successful ESM should support complex interactions between heterogeneous models and methods, but the resulting systems are not fully understood, making it difficult to evaluate their impact. This scenario replicates itself in individual ESM components, which may include additional heterogeneous couplings; Fig. 1(R).

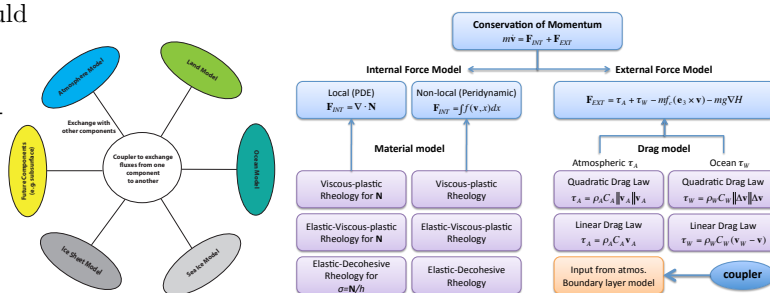


Figure 1: Left: ESMs comprise a number of components, which operate independently but periodically exchange boundary data through a flux coupler component. Right: an ESM component such as sea ice can involve further model and method hierarchies.

Proposed Direction of work

Formulation of predictive ESM faces a number of crosscutting challenges, such as

1. physics-compatible representation of the flow of information between sub-models and scales;
2. development of stable, high-order accurate, physical property preserving, couplings; and
3. efficient mapping of the numerical models onto emerging architectures.

¹Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. Sandia Document Tracking Number 364849.

Contributed ideas paper 1

Addressing these challenges requires a shift from the dominant solution paradigm for multi-physics, multiscale problems that for the past 30 years has been based on first-order-accurate operator splitting, semi-implicit and explicit time integration methods and decoupled nonlinear solution methods. The DOE Town Hall report [4, §3.1] points out that this strategy is rapidly approaching a point of diminishing returns because it lacks the stability properties needed to carry out simulations over dynamic scales of interest, relies on heuristics to control the splitting errors, does not keep up with the increases in the accuracy of models for the individual components of the climate system model, and it is prone to nonintuitive instabilities and fragile solutions.

We assert that optimization and control ideas provide a powerful foundation for a rigorous, mathematically sound coupling strategy, which enables a collection of dissimilar numerical methods from multiple scientific disciplines to function together as a unified simulation tool (an HNM), and which is capable of successfully addressing the above crosscutting challenges.

Connections to Math, Comp Sci & and Climate Science

Recent research efforts by the applied mathematics community have explored optimization ideas in several settings, relevant to the crosscutting challenges mentioned above.

Efficient optimization-based methods for property-preserving data transfers have been developed and studied numerically in [7]. Theoretical foundations for optimization-based multidisciplinary simulations are established in [12, 13], while [6, 11] demonstrate the potential of these ideas for decomposition and synthesis of multiphysics problems. The papers [9, 10] demonstrate virtual control formulations of HNMs that are stable, accurate and free of spurious coupling artifacts. Finally, [8, 5] apply optimization to obtain property-preserving methods for passive tracer transport. The scheme in [5] has been implemented in the High Order Modeling Environment (HOMME), which is the default atmospheric dynamical core in ACME.

These papers establish theoretically the viability of the optimization approach, while numerical studies confirm the potential of the approach for ESM. As a result, optimization ideas are ripe for consideration in ESM. Their success will largely depend on a focussed effort to develop, analyze, implement, test, and apply the necessary scalable optimization algorithms. These tasks may be facilitated by careful design of the optimization formulation, leading to computationally efficient structures of the optimization problem [7].

Potential Impact on the field

ESMs must provide accurate estimates of climate variability over the instrumental record, and project future variations over decadal to centennial periods. The coupling of more components with more physics at increased resolutions creates ESMs with daunting levels of sophistication. This complexity can make quantitative estimation difficult and limits the ability to provide credible uncertainty estimates for derived climate indices. The proposed work has the potential to reduce and/or eliminate many of the uncertainties associated with the coupling of ESM components by ensuring physically consistent, stable and accurate flow of information between sub-models.

An optimization-based coupling strategy for ESMs should also facilitate numerical algorithms that maximize local computation for a given level of communication [1, p.56], [3] and whose computational parts can operate more asynchronously to increase the computation to communication ratio. The latter will allow simulations to move away from bulk-synchronization, which has been identified as an important aspect of exascale-appropriate mathematical problem formulations [2, p.20].

Contributed ideas paper 1

References

- [1] S. Ashby, P. Beckman, J. Chen, P. Colella, B. Collins, D. Crawford, J. Dongarra, D. Kothe, R. Lusk, P. Messina, T. Mezzacappa, P. Moin, M. Norman, R. Rosner, V. Sarkar, A. Siegel, F. Streitz, A. White, and M. Wright, *Report on exascale computing*, Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee, U.S. Department of Energy, Advanced Scientific Computing Research (ASCR) program, 2010.
- [2] *Scientific Grand Challenges. Crosscutting technologies for computing at the exascale.*, Report from the Workshop Held February 2-4, 2010, U.S. Department of Energy, Advanced Scientific Computing Research (ASCR) program, May 2010.
- [3] F. Alexander, M. Anitescu, J. Bell, D. Brown, M. Ferris, M. Luskin, S. Mehrotra, B. Moser, A. Pinar, A. Tartakovsky, K. Willcox, and S. Wright, *A multifaceted mathematical approach for complex systems. report of the doe workshop on mathematics for the report of the doe workshop on mathematics for the analysis, simulation, and optimization of complex systems*, tech. rep., U.S. Department of Energy, 2011.
- [4] H. Simon, T. Zacharia, and R. Stevens, *Modeling and simulation at the exascale for energy and the environment.*, Town hall meetings report, U.S. Department of Energy, Advanced Scientific Computing Research (ASCR) program, 2007.
- [5] P. B. Bochev, S. A. Moe, K. J. Peterson, and D. Ridzal. A conservative, optimization-based semi-Lagrangian spectral element method for passive tracer transport. In B. Schrefler, E. Onate, and M. Papadrakakis, editors, COUPLED PROBLEMS 2015, VI Int. Conf. on Comp. Meth. for Coupled Problems in Science and Engineering, pages 23–34, April 2015.
- [6] P. Bochev and D. Ridzal. *An optimization-based approach for the design of PDE solution algorithms*. SIAM Journal on Numerical Analysis, 47(5):3938–3955, 2009.
- [7] P. Bochev, D. Ridzal, and M. Shashkov. *Fast optimization-based conservative remap of scalar fields through aggregate mass transfer*. Journal of Computational Physics, 246(0):37 – 57, 2013.
- [8] K. Peterson, P. Bochev, and D. Ridzal. *Optimization-based conservative transport on the cubed-sphere grid*. In I. Lirkov, S. Margenov, and J. Wasniewski, editors, Proceedings of LSSC 2013, Springer Lecture Notes in Computer Science, Berlin, Heidelberg, 2014.
- [9] M. D’Elia, M. Perego, P. Bochev, D. Littlewood *A coupling strategy for nonlocal and local diffusion models with mixed volume constraints and boundary conditions*. To appear in Comp. Math. with Applic. 2016.
- [10] P. Gervasio, J.-L. Lions, and A. Quarteroni, *Heterogeneous coupling by virtual control methods*, Numerische Mathematik, 90 (2001), pp. 241–264. 10.1007/s002110100303.
- [11] J. Lions, *Virtual and effective control for distributed systems and decomposition of everything*, Journal d’Analyse Mathématique, 80 (2000), pp. 257–297. 10.1007/BF02791538.
- [12] Q. Du and M. D. Gunzburger, *A gradient method approach to optimization-based multidisciplinary simulations and nonoverlapping domain decomposition algorithms*, SIAM Journal on Numerical Analysis, 37 (2000), pp. pp. 1513–1541.
- [13] M. D. Gunzburger, M. Heinkenschloss, and H. K. Lee, *Solution of elliptic partial differential equations by an optimization-based domain decomposition method*, Applied Mathematics and Computation, 113 (2000), pp. 111 – 139.

Contributed ideas paper 2

Component-Based Application Code Development, Part 1: The Agile Components Strategy and Albany Code

Andy Salinger

Sandia National Laboratories

Irina Tezaur [SNL], Mauro Perego [SNL], and the Albany and Trilinos development teams

Abstract

In this submission, we present a code development strategy and collaboration model that is designed to maximize the leverage and long-term impact of algorithmic investments on application code projects such as earth system models.

Our computational science organization has been following a strategy -- dubbed “Agile Components” -- for code development that puts the emphasis on generating and leveraging reusable software components [1]. Software components are broadly defined as libraries, interfaces, software quality tools, and demonstration applications. We have found there to be numerous benefits of this approach in the writing of new application codes in terms of development time, costs, and algorithmic capabilities. Furthermore, we have found that the Agile Components approach provides a framework where scientists and algorithm experts can collaborate on a code project and each can be productive and successful in their own realm, and leveraging of expertise in funding is maximized. In this way, it maps well to the vision and funding mechanism of DOE's SciDAC program.

In this submission (Part I), the Agile Components strategy and collaboration model is presented, along with its manifestation in the Albany code framework. In Part II (submitted by Irina Tezaur), the experiences of using the Agile Components approach in the development of the Albany/FELIX Ice Sheet application code under the PISCEES SciDAC project will be presented. We believe this example is instructional for creating a model for future successful Climate Science – Computational Science collaborations.

Background: The Agile Components Strategy

The Agile Components strategy for computational science organizations involves requiring projects to **both leverage and contribute to** a comprehensive set of software components, consisting of libraries, interfaces, software quality tools, and demonstration applications. Just as compilers, BLAS, Lapack, and MPI have long been standard external dependencies for application codes -- and it is also common to depend on external linear solvers and meshing tools -- the Agile Components strategy extrapolates this trend to include dozens of algorithmic capabilities that can be generalized into reusable libraries. This includes such diverse capabilities as automatic differentiation, discretization tools, and mesh databases. The strategy also extends to a common and mature set of software engineering tools, so that each application project doesn't have to become experts in this arena as well. Figure 1 conveys the large amount of leverage from other development efforts that align with the development of the Albany/FELIX ice sheet code.

Many benefits and challenges to this strategy have been enumerated in a technical report [1]. Benefits include the ability to leverage several funding sources to establish an algorithmic expertise, the ability to leverage that existing expertise on new applications, the amortization of verification, porting, and maintenance costs across projects, and the modular design of codes to allow for project agility. The main drawback is the dependence on code written by others.

Research To Date: Albany and Trilinos

Contributed ideas paper 2

At Sandia, we have been actively pursuing and refining this strategy for a number of years. The common set of reusable libraries and interfaces is deployed primarily through the Trilinos suite of computational math algorithms [2]. The Albany code project [3] is the chief incarnation of the strategy. Albany is a finite element code base that hosts several funded application projects, is the testbed for several algorithm research projects, and provides a template for a science applications code's software engineering tools and processes.

In Albany, we have successfully developed application codes for quantum device modeling [4], computational mechanics [5], and an ice sheet velocity solver [6], and are making progress on an atmosphere dycore. These codes are all born with a quality software engineering infrastructure, scalable algorithms, and embedded analysis and UQ algorithms, and are on a path to performance portability across all architectures. For many of those involved, the productivity in developing new applications with sophisticated capabilities in Albany using Trilinos libraries has been striking. While this assertion is anecdotal, not easily quantifiable, the success of these projects is supporting evidence, as is the fact that Agile Components is central to Sandia's execution strategy for its ATDM effort, a significant new code development effort under NNSA's ASC program.

In the companion Part II idea paper submission by Irina Tezaur, the experience of using this strategy for the development of the Albany/FELIX Ice Sheet code [6] will be detailed.

Proposed Direction: Agile Components for Climate Science Impact

The strategy described here is designed to maximize the long-term productivity in developing coupled earth system models, in terms of the speed and cost of development, the breadth and sophistication of the algorithmic capabilities, and the trustedness of the simulation results.

With those goals in mind, the path forward for impacting Climate Science **requires the climate application codes to embrace the use of software components developed by computational scientists and in developing long-term collaborations.** The benefits are numerous, including:

1. Inclusion of the algorithmic capability into the climate code.
2. Access to the algorithmic expert for continued engagement, who can work in his/her own familiar code base. From this foundation, fractional funding streams and even consultations can have large impact.
3. Automatic, free upgrades in capability as the library development continues through other funding sources and using developments motivated by other applications.
4. A decreased code base that must be supported, and ported to new architectures.
5. Encourages more modular code design, for improved agility and maintainability.
6. Broadening of the network of people thinking about your application.
7. Ability to leverage algorithm and inter-disciplinary funding streams to impact climate science application development.

In contrast, an alternate collaboration model is for the application scientist to learn about a promising algorithm and implement it directly in the application code. This shares benefit #1 above, and has the added benefit of the application scientist having complete control and familiarity of the code base, but sacrifices the remaining benefits from the list above.

A separate phase of the strategy can involve spinning-off algorithmic capabilities that were originally part of climate applications into reusable libraries to be deployed for other applications. This has many benefits, including the expansion of impact of the initial investment, the maturation of the algorithms in other settings, the access to other funding sources to continue the work in a more general setting, and career growth for the developer.

Contributed ideas paper 2

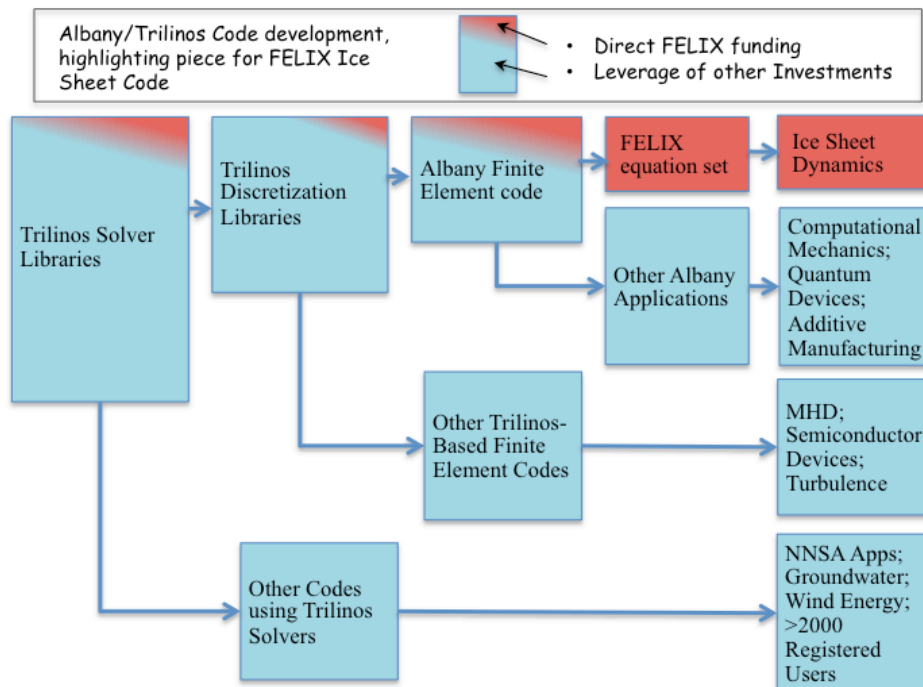


Figure 1: A schematic showing the leverage of the Agile Components code strategy, highlighting the slice of the code base written specifically for the Albany/FELIX ice sheet application. It can be seen how large parts of the code base are also used by other applications -- several of which are listed in the right column -- which all contribute to the development costs, research drivers, and creation of expertise.

References

1. A.G. Salinger, "Component-based scientific application development," SAND2012-9339, 2012.
2. M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, A. Williams, and K. Stanley, "An Overview of the Trilinos Project," *ACM TOMS*, 31(3), 397-423, 2005.
3. A.G. Salinger, R.A. Bartlett, Q. Chen, X. Gao, G.A. Hansen, I. Kalashnikova, A. Mota, R.P. Muller, E. Nielsen, J.T. Ostien, R.P. Pawlowski, E.T. Phipps and W. Sun, "Albany: A Component-Based Partial Differential Equation Code Built on Trilinos", Sandia National Laboratories Tech Report SAND2013-8430J, 2013.
4. X. Gao, E. Nielsen, R. P. Muller, R.W. Young, A. G. Salinger, M. S. Carroll, "The QCAD Framework for Quantum Device Modeling," *15th international Workshop on Computational Electronics*, 1-4, 2012.
5. W. Sun, J.T. Ostien, A.G. Salinger, "A stabilized assumed deformation gradient finite element formulation for strongly coupled poromechanical simulations at finite strain," *International Journal for Numerical and Analytical Methods in Geomechanics*, 37(16), 2755-2788, 2013.
6. I. Tezaur, M. Perego, A. G. Salinger, R.S. Tuminaro, and S.F. Price, "Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis," *Geosci. Model Dev.*, 8, 1197-1220, 2015.

Contributed ideas paper 3

Component-Based Application Code Development, Part 2: Demonstration on a Land-Ice Model and Proposed Extension to Other Climate Components November 23, 2015

Irina Tezaur

Sandia National Laboratories

Andrew Salinger, Mauro Perego, Ray Tuminaro (Sandia National Laboratories)

Abstract

This paper illustrates the success of the components-based code development strategy for developing a next-generation world-class climate code in the specific case of the *Albany/FELIX* land-ice solver created as a part of the PISCEES SciDAC3 project. The proposed idea is to find opportunities to apply this approach in other climate areas (e.g., atmosphere, sea-ice, ocean) to make these models more scalable, robust and portable to emerging architectures, to endow these models with improved analysis capabilities (e.g., adjoint-based optimization, embedded uncertainty quantification), all towards a more integrated climate modeling framework that shares software, data standards and tools, and model components. This submission is a companion paper to a submission by A. Salinger on the Agile Components Strategy and *Albany* code.

Background/Research to Date

According to a 2012 report by the National Research Council [1], there is a critical need for a next generation of advanced climate models. More specifically, the report calls for climate models to take a more integrated path and use a common software infrastructure while adding regional detail, new simulation capabilities, and new approaches for collaborating with their user community. Although climate models have improved in recent years, much work is needed to make these models reliable and efficient on continental scales, to quantify uncertainties in the models' outputs, and to port the models to a new generation of high performance computer (HPC) architectures. Many legacy climate models lack advanced analysis capabilities, like sensitivity and adjoint calculations. Moreover, legacy codes need to be rewritten substantially in order to run accurately and efficiently on new architecture machines (e.g., GPUs), as they are based on algorithms optimized for existing architectures (e.g., CPUs).

A promising approach for developing next-generation performance-portable solvers with advanced analysis capabilities is the so-called components-based code development strategy¹ to building application codes. In this approach, mature, modular libraries are combined using abstract interfaces and template-based generic programming, resulting in a final code that is verified, scalable, fast and robust, and has access to dozens of algorithmic and advanced analysis capabilities. One recent success story for the components-based code development approach in the area of climate modeling is *Albany/FELIX*² [4, 5], an unstructured grid (Fig. 1, left) finite element land-ice solver written using the *Trilinos* [3] libraries and the *Albany* [2] code base as a part of the SciDAC3 PISCEES³ project for integration into earth system models (ESMs). The component-based code development approach has led to the rapid development (≈ 3 FTEs) of this world-class land-ice model with many sophisticated capabilities. The integration of automatic differentiation into the code has enabled robust nonlinear solves, sensitivity analysis,

¹Detailed in a companion idea paper submitted by A. Salinger entitled "Component-Based Application Code Development, Part 1: The Agile Components Strategy and Albany Code".

²Finite Elements for Land Ice eXperiments.

³Predicting Ice Sheet and Climate Evolution at Extreme Scales.

Contributed ideas paper 3

adjoint-based optimization for ice sheet initialization, and uncertainty quantification (UQ). Ad hoc spin-ups and parameter tuning has been replaced with the use of formal, adjoint-based optimization techniques for generating realistic model initial conditions [6]. In collaboration with experts from the QUEST SciDAC institute, a framework for forward and inverse UQ that incorporates formal UQ methods and computational tools that orchestrate these methods (e.g., *DAKOTA* [7]) has been developed. The *Albany/FELIX* code has demonstrated scalability up to 1 billion unknowns and tens of thousands of cores thanks parallel scalable iterative linear solvers and newly-developed preconditioning methods from *Trilinos* (Fig. 1, right, top). The development of new semi-implicit momentum balance and thickness coupling techniques has led to more stable and efficient time-stepping schemes, expected to reduce substantially run-times for transient land-ice simulations. Finally, using the *Kokkos Trilinos* library and programming model, we were able to create a single code that runs on different architectures (multi-core, many-core, GPU) by merely changing a template parameter (Fig. 1, right, bottom).

Proposed Direction of Work

The proposed direction is to look for ways to equip other climate components (e.g., atmosphere, sea-ice, ocean) and coupled ESMs with the advanced analysis and performance capabilities described above by integrating into these models software libraries and algorithms developed by domain experts. Success of this approach rests strongly on a collaboration model for the development of climate technologies which involves not only glaciologists and climate modelers, but also computational scientists, and is geared towards creating a unified modeling framework that shares software, data standards and tools, and model components. Based on our experience with the *Albany/FELIX* code, the following enhancements in other climate models are conceivable: (1) improved software quality through formal verification studies and regression testing, (2) improved scalability and robustness, (3) improved fidelity (e.g., through the use of unstructured, regionally refined meshes), (4) performance-portability to new and emerging architectures, (5) improved incorporation of data (e.g., through better, optimization-based model initiation techniques), (6) improved validation and UQ methods (e.g., embedded UQ), (7) improved time-evolution algorithms for more stable and faster transient simulations. Some specific ideas worth exploring are embedded UQ for atmosphere, non-linear solvers for sea-ice and implicit/semi-implicit solvers for ocean.

Connections to Math, Comp Sci & and Climate Science

For a successful integration of software libraries with advanced analysis and next-generation capabilities into climate components, close collaborations between climate scientists and mathematicians/computational scientists (e.g. the current applied math and computer science SciDAC3 institutes: FastMath, SUPER, QUEST) are required. The proposed components-based approach would facilitate the development and numerical study of new mathematical approaches and computational algorithms. The use of similar software frameworks and libraries would make state-of-the-art methods readily available to a variety of climate applications.

Potential Impact on the Field

We believe the proposed approach would lead to better climate models, namely models that are more robust and efficient, equipped with advanced analysis capabilities, and capable of running on next generation platforms built towards exascale computing, as described in detail above.

Contributed ideas paper 3

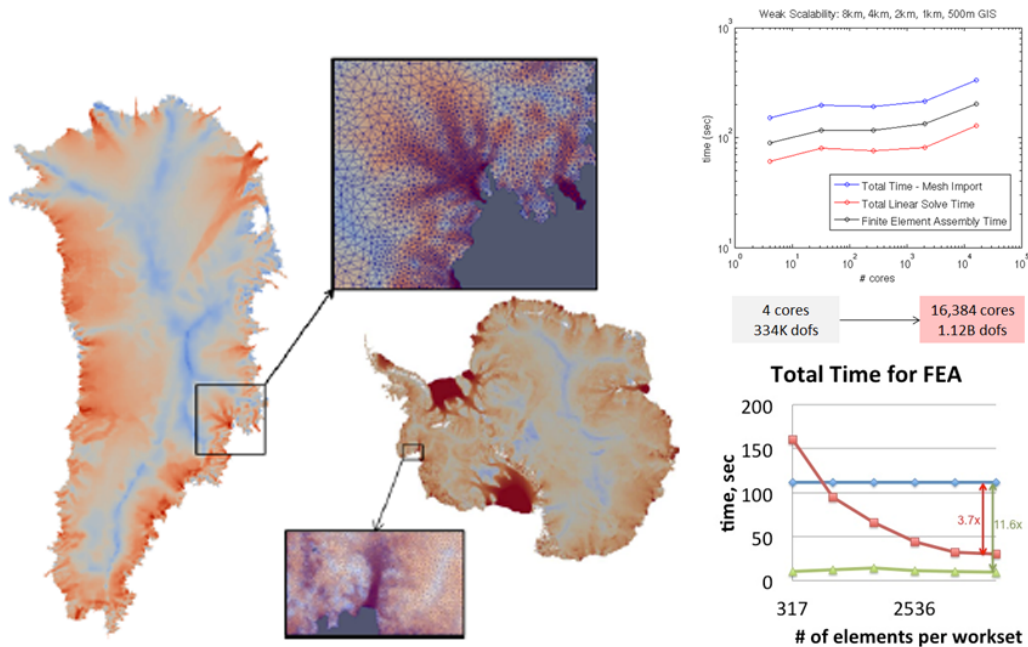


Figure 1: Greenland and Antarctica geometries, discretized by unstructured meshes ((left); weak scaling up to 1.12B unknowns and 16K cores (right, top); performance portability of finite element assembly using *Kokkos*: Serial - blue, OpenMP - green, CUDA - red (right, bottom)

References

- [1] National Research Council of the National Academies, A national strategy for advancing climate modeling. The National Academies Press, Washington, DC, 2012.
- [2] A. Salinger. Albany: A component-based partial differential equation code built on Trilinos. Technical report, Sandia National Labs, 2003.
- [3] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, K. Long, R. Lehoucq, R. Pawlowski, E. Phipps, H. Thornquist, A. Salinger, R. Tuminaro, J. Willenbring, and A. Williams. An overview of trilinos. Technical Report SAND2003-2927, Sandia National Laboratories, 2003.
- [4] I. Tezaur, M. Perego, A. Salinger, R. Tuminaro, and S. Price. Albany/FELIX: A parallel, scalable and robust finite element higher-order stokes ice sheet solver built for advanced analysis. *Geosci. Model. Develop.*, **8**:1–24, 2015.
- [5] I. Tezaur, R. Tuminaro, M. Perego, A. Salinger, S. Price. On the scalability of the Albany/FELIX first-order Stokes approximation ice sheet solver for large-scale simulations of the Greenland and Antarctic ice sheets, *Procedia Computer Science*, **51**:2026–2035, 2015.
- [6] M. Perego, S. Price, G. Stadler. Optimal initial conditions for coupling ice sheet models to Earth system models, *J. Geophys. Res.* **119**(9) 1894–1917.

Contributed ideas paper 3

- [7] B. Adams, L. Bauman, W. Bohnhoff, K. Dalbey, M. Ebeida, J. Eddy, M. Eldred, P. Hough, K. Hu, J. Jakeman, L. Swiler, and D. Vigil. DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis: Version 5.4 User's Manual, Sandia Technical Report SAND2010-2183, 2009.

Contributed ideas paper 4

Embedded Ensembles
Andy Salinger¹
Sandia National Laboratories
Eric Phipps (SNL), Jeffrey Fike (SNL)

Abstract

As the fidelity of climate simulations increase, the need for computing large ensembles of runs will also increase. Computational projections for quantities of interest that are associated either with an extreme event or pertaining to a more localized region will require more fidelity in resolving the probabilistic or uncertainty dimensions. On the computer hardware front, new architectures (particularly GPUs) are rewarding higher compute intensities and fewer communications. By inverting the traditional ensemble generation so that ensembles are moved to the inner loop of a calculation, large efficiency gains may be possible compared to the standard approach of running ensembles as an outer loop around the code. The recent work of Phipps *et al.* [1] has great potential to improve performance of ensembles of calculations on new architectures for climate science applications.

Background/Research to Date

A thorough description of the research underlying this idea has recently been submitted for publication by Phipps *et al.* [1]. Here is an excerpt from this paper, which succinctly communicates the approach and computational benefits.

We explore a new method for implementing sampling methods that simultaneously propagates groups of samples together in an embedded fashion, which we call embedded ensemble propagation. We show how this approach takes advantage of properties of modern computer architectures to improve performance by enabling reuse between samples, reducing memory bandwidth requirements, improving memory access patterns, improving opportunities for fine-grained parallelization, and reducing communication costs. [1]

The method of propagating an ensemble of runs through the code simultaneously requires transforming all scalar data that is dependent on the solution to become a vector. All operations on data become a loop over this vector length. All data that is not a function of the solution, including all mesh and bookkeeping data, does not need to be replicated as a vector. This can significantly reduce both FLOPS and memory access, when a significant portion of the work is amortized over a set of ensembles. Also, when the inner loop (ensemble) size is fixed at compile time, it is nearly assured that this simple (blas-like) loop will be effectively optimized by the compiler. Furthermore, MPI communication calls transfer an ensemble of data at a time, amortizing away much of the latency costs and being more bandwidth limited. For implicit codes that use preconditioned Krylov iterative solvers, there can even be a benefit of calculating a single preconditioner for an ensemble of related matrices, greatly reducing computational costs.

The implementation approach used in Phipps *et al.* [1] is specific to C++ codes, using templated and operator overloading. It requires the code to be written using a template parameter for the data type. When this data type is `double`, then the code performs identically to a version written for propagating a single instance. An `Ensemble` data type has been developed that implements all the basic operations seen in a code (e.g., `*`, `-`, `log`, `exp`, `sqrt`, `sin`) for a vector of information.

¹ The computational science idea behind this proposal is appropriately credited to Phipps and coworkers. Salinger is communicating the idea for its potential impact to climate simulations.

Contributed ideas paper 4

The code correctly handles mixed data types, when a solution-dependent quantity (expressed as a vector) and a mesh-only dependent quantity (expressed as a scalar) are combined in an operation.

There are several codes already written with a template parameter as the data type. This same “interface” to the code is used for automatic differentiation and polynomial propagation, and is detailed in a pair of papers on Template Based Generic Programming [2,3].

The Albany code base [4] (which also has Phipps as an original developer) is written with a template parameter for the scalar type. This code base includes the Albany/FELIX ice sheet velocity solver [5] being used in the ACME earth system model, as well as Albany/Aeras, a research effort on developing a spectral-element atmosphere dycore [6]. Figure 1 shows preliminary results of using embedded ensembles for the shallow water equations in Aeras.

Proposed Direction of work

The proposal is to incorporate ensembles natively into climate models. There are excellent performance gains to be had, particularly on scaling up to exascale platforms.

For C++ codes, the template-based implementation approach of Phipps *et al.*, as in the Albany code base, leads to a very non-invasive implementation. Research is still needed for finding the most effective linear and nonlinear solver approach for codes with implicitness, since there are very large potential gains in this area. Work will continue in matching the implementation to individual hardware and compiler combinations to get optimal performance.

For Fortran codes, we assume the path forward would be to convert data types for solution-dependent quantities to have an extra dimension, and to add an extra loop over the ensemble vector. This would be a natural extension to code transformation tools such as the OpenAD automatic differentiation tool [7].

It is also possible for certain sections of a code to opt-out of this approach, with a loop over the ensemble length around this section of code. This would require some unpacking and repacking of scalar data into a vector layout.

Connections to Math, Comp Sci & and Climate Science

Embedding ensembles as an inner loop in a simulation code opens the door to uncertainty quantification algorithms beyond simple sampling, such as stochastic collocation approaches. The computational science connection is strong, inverting loop orders to provide for higher compute intensities and effective use of exascale-class machines. The effective use of climate projections for decision support, including probabilities of extreme events (e.g. 100-yr floods), will require ensemble size to gain equal importance with spatial resolution.

Potential Impact on the field

The main benefit to the field would be the performance gains of processing an ensemble of information at a time. This is highly dependent on the code, compiler, and hardware, but generally expected to be a factor of 2 speedup for ensembles sizes of order 8 or more. There is some benefit, programmatically, in making effective use of exascale-sized platforms. The scientific gains are to be had when the barrier to computing ensembles is greatly reduced, and extra attention is given to quantifying the uncertainties and probabilities in our high-fidelity climate projections. As the field of climate simulation evolves to include a decision-support role, these dimensions will greatly increase in significance.

Contributed ideas paper 4




Figure 1: Speed-up of the Albany/Aeras code for shallow water equations when using embedded ensembles. In all cases, the timings are for calculating an ensemble of 32 simulations, in batches of 1 through 32 (blue curve), normalized by the time to do 1 forward simulation.

References

- [1] E. Phipps, M. D’Elia, H. C. Edwards, M. Hoemmen, J. Hu, S. Rajamanickam, Embedded Ensemble Propagation for Improving Performance, Portability and Scalability of Uncertainty Quantification on Emerging Computational Architectures, submitted, 2015. <http://arxiv.org/abs/1511.03703>
- [2] R.P. Pawlowski, E.T. Phipps, A. G. Salinger, “Automating embedded analysis capabilities and managing software complexity in multiphysics simulation part I: template-based generic programming,” *Scientific Programming*, 197-219, 2012.
- [3] R.P. Pawlowski, E.T. Phipps, A. G. Salinger, S.J. Owen, C. M. Siefert, M.L. Staten, “Automating embedded analysis capabilities and managing software complexity in multiphysics simulation part II: application to partial differential equations,” *Scientific Programming*, 327-345, 2012.
- [4] A.G. Salinger, R.A. Bartlett, Q. Chen, X. Gao, G.A. Hansen, I. Kalashnikova, A. Mota, R.P. Muller, E. Nielsen, J.T. Ostien, R.P. Pawlowski, E.T. Phipps and W. Sun, “Albany: A Component-Based Partial Differential Equation Code Built on Trilinos”, Sandia National Laboratories Tech Report SAND2013-8430J, 2013.
- [5] I. Tezaur, M. Perego, A. G. Salinger, R.S. Tuminaro, and S.F. Price, “Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis,” *Geosci. Model Dev.*, 8, 1197-1220, 2015.
- [6] Spatz, W. F., Smith, T. M., Demeshko, I. P., & Fike, J. A. Aeras: A next generation global atmosphere model. *Procedia Computer Science*, 51, 2097-2106, 2015.
- [7] J. Utke, U. Naumann, M. Fagan, N. Tallent, M. Strout, P. Heimbach, C. Hill, and C. Wunsch. OpenAD/F: A modular open-source tool for automatic differentiation of Fortran codes *ACM Transactions on Mathematical Software (TOMS)* 34, no. 4, 2008.

Contributed ideas paper 5

New Temporal and Spatial Algorithms for Atmospheric Climate Models (November 23, 2015)

Matthew Norman

Oak Ridge National Laboratory

Abstract

The atmospheric component is currently the most computationally expensive part of a coupled climate model. One of the greatest challenges for atmospheric climate models is that they must run with very little data per node. This is caused by the roughly $2,000$ times real-time throughput constraint coupled with a time step that must be refined as the grid is refined. A 2 times grid refinement leads to at least 2 times less data per node, for instance. Currently, there is so little data per node that data transfer between nodes consumes about 50% of the total runtime, and GPU accelerators are used inefficiently. This will only get worse with grid refinement. We need algorithms that lead to asymptotically larger time steps during spatial refinement.

Also, the semi-discrete time integrators currently used have several deficiencies. They cannot be more than fourth-order accurate and preserve non-oscillatory properties. They rely on repeated calls to the spatial operator and thus are too sensitive to spatial operator cost. Also, they require intermittent stages of parallel data transfer. We need temporal operators that can be any order of accuracy, have less sensitivity to the spatial operator's expense, and require less parallel data transfer.

This idea paper proposes several spatial and temporal operators as alternatives to the current state-of-the-art. The proposed methods provide properties that are more amenable to the coming exascale computer architectures, which will use accelerators, will have a complex memory hierarchy, and will have a large penalty for data transfers.

Background/Research to Date

The current state of the art for atmospheric modeling is the SE operator coupled with RK time integrators[4]. Several studies show that ADER time discretizations have large effective CFL limits than RK integrators, automatically maintain non-oscillatory properties of limited spatial operators, and easily allow any arbitrary temporal order of accuracy without stages or communication with neighboring elements[2]. Most of the overall cost is absorbed in the ADER operator rather than the spatial operator, leaving room for flexibility in using more expensive spatial operators.

The SE operator is very cheap and accurate and has a larger time step than its immediate competitors such as the Discontinuous Galerkin (DG) operator[5]. The nearest competitor besides the DG operator would be the Spectral Finite-Volume method [?], which has a slightly better time step. With the ADER-DT time discretization, SE and DG time steps scale as $\Delta t \propto h^{-1}p^{-2}$, and SFV time steps scale as $\Delta t \propto h^{-1}p^{-1.5}$. There is also the more recently developed BA-MCV operator[3], which scales as $\Delta t \propto h^{-1}p^0$ in 1-D and $\Delta t \propto h^{-1}p^{-0.5}$ in 2-D. The FV operator is more expensive than multi-moment operators because it reconstructs for each Degree of Freedom (DOF), whereas multi-moment operators share the reconstruction for multiple DOFs. However, FV time steps scale as $\Delta t \propto h^{-1}p^0$ for any number of spatial dimensions.

Contributed ideas paper 5

Proposed Direction of work

The BA-MCV and FV operators are great options going forward because there are actually *more* DOFs per node during *p*-refinement, making *p*-refinement a more attractive option than *h*-refinement. There is an additional advantage to the FV method, and that is that Weighted Essentially Non-Oscillatory (WENO) limiting can be performed with very little computational overhead (e.g., 20-40%, depending upon the order of accuracy) because a stencil already has to be retrieved. This means that even though FV costs more per DOF outright, when considering the full scheme (since *some* limiting must be performed), FV might win out.

Connections to Math, Comp Sci & and Climate Science

These BA-MCV, FV, WENO, and ADER-DT methods have the potential to produce accurate climate simulations much more cost-effectively because of larger time steps and greater amounts of data per node (reducing the relative cost of data transfers). This has connections to mathematics because these are new algorithms to solve the Euler equations of fluid motion. These ideas connect to computer science because they are formulated with future exascale hardware and current computational challenges both in mind, and they reduce data transfers, increase time step sizes, and better maintain robustness in limiting than existing algorithms. They connect to climate science because they are targeted toward the challenges currently experienced in real atmospheric climate simulations, which are relatively large MPI overheads, low time steps, and too little data per node to effectively utilize accelerator hardware like GPUs.

Potential Impact on the field

These new operators have the potential to significantly improve our ability to generate science from climate simulations. They allow more effective use of computer resources by achieving similar or better accuracy in a more efficient manner. Also, they potentially enable the climate science field to be awarded larger allocations because of better utilization of accelerators with more data per node due to better time step scaling.

References

- [1] Wang, Zhi Jian. Spectral (finite) volume method for conservation laws on unstructured grids. basic formulation: Basic formulation. *JCP* **178**, 2002, pp. 210-251.
- [2] Matthew R. Norman. A WENO-limited, ADER-DT, finite-volume scheme for emphcient, robust, and communication-avoiding multi-dimensional transport. *JCP*. **274**, 2014, pp. 1-18.
- [3] Matthew R. Norman. Arbitrarily High-Order-Accurate, Hermite-WENO Limited, Boundary-Averaged Multi- Moment Constrained Finite-Volume (BA-MCV) Schemes for 1-D Transport. *Procedia Computer Science*, **51**, 2015, pp. 2688-2697.
- [4] Mark A. Taylor, et al. A mass and energy conserving spectral element atmospheric dynamical core on the cubed-sphere grid. *Journal of Physics: Conference Series*. **78**. IOP Publishing, 2007.
- [5] Ramachandran D. Nair, et al. A discontinuous Galerkin transport scheme on the cubed sphere. *MWR*. **133**, 2005, pp. 814-828.

Contributed ideas paper 6

Fast simplicial finite elements using Bernstein polynomials 24 November 2015

Robert C. Kirby
Baylor University

Abstract

Bernstein polynomials can be defined for arbitrary degree on an arbitrary-dimensional simplex (or tensor product) domain. Unlike many simplicial bases, they admit optimal-complexity algorithms for polynomial evaluation, integration, and differentiation so that they provide a fundamentally new approach to spectral elements on simplices.

Background/Research to Date

The Bernstein polynomials of degree n on the d -simplex are readily defined using barycentric coordinates and multiindex notation. For a nondegenerate simplex $T \subset \mathbb{R}^d$ with vertices $\{x_i\}_{i=0}^d$, let $\{b_i\}_{i=0}^d$ denote the barycentric coordinates. Let $\mathbf{b} \equiv (b_0, b_1, \dots, b_d)$ be a tuple of barycentric coordinates on a simplex. For any multiindex α , we define a *barycentric monomial* by

$$\mathbf{b}^\alpha = \prod_{i=0}^d b_i^{\alpha_i}.$$

We obtain the *Bernstein polynomials* by scaling these by certain generalized binomial coefficients

$$B_\alpha^n = \frac{n!}{\alpha!} \mathbf{b}^\alpha. \quad (1)$$

While inauspicious at first blush, they define a nonnegative partition of unity basis for any polynomial degree in any dimension.

In [5], I proved that the constant-coefficient mass matrix possesses highly detailed structure, admitting a matrix-vector product in only $\mathcal{O}(n^{d+1})$ operations despite its density. Moreover, matrices encoding differentiation for the Bernstein basis are quite sparse, so that stiffness matrices also admit the same order of complexity. I generalized these techniques to include quadrature-based methods for variable coefficients in [7] and $H(\text{div})$ and $H(\text{curl})$ elements in [6]. In Figure 1, we show that at high-order, the matrix free action of operators using Bernstein polynomials vastly outperforms the low-level code generated by FEniCS [8]

Also, in work submitted for publication [4], I have shown that the element mass matrix admits a fast *inversion* algorithm – $\mathcal{O}(n^{d+1})$, of the same order as the forward operations. Although it is not diagonal per orthogonal bases or approximately so per some kind of clever quadrature, it allows optimal-complexity explicit time-stepping in DG methods.

In parallel with these developments is the work of Mark Ainsworth at Brown University [1, 2] who also shows that the Duffy transform [3] tensorializes the Bernstein polynomials so that standard FEM operations using Stroud conical quadrature [10] admit sum-factorization [9].

To sum up, the current state of Bernstein polynomials is that they offer exciting new techniques to the computational science community and are ready for exploration in the context of applications, although along the way certain new fundamental algorithmic research questions will surely appear.

Contributed ideas paper 6

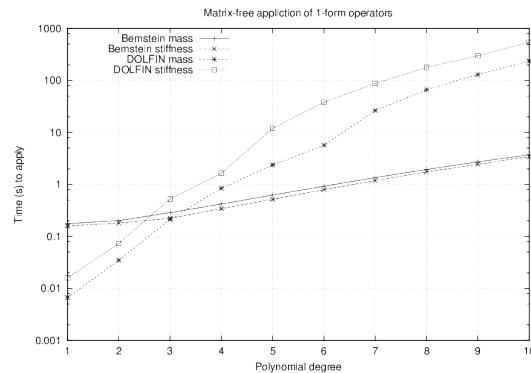


Figure 1: Timing for matrix-free mass and stiffness action for Bernstein polynomials versus DOLFIN on a mesh of 20,000 tetrahedra

Proposed Direction of work

In the future, I hope to perform further investigation of Bernstein polynomials along several fronts:

- Collaborate with domain scientists to foster their application to challenging problems.
- Work with FEniCS and other large-scale software projects to incorporate the highly-specialized algorithms.
- Further fundamental research on finite element algorithms. For example, fast algorithms for the local elimination stage of static condensation or hybridizable DG methods.
- Exploration of fine-grain parallelism in Bernstein algorithms with an eye to GPU implementation.

Connections to Math, Comp Sci & and Climate Science

High-order finite element and spectral methods have long been attractive to the climate community because of their high fidelity, low numerical dispersion, and potential for computational efficiencies. A perennial issue has been the achieving this full efficiency in unstructured geometry owing to the non-diagonal mass matrix. Bernstein polynomials circumvent this problem not by orthogonality, but by highly specialized structure allowing fast inversion of the mass matrix.

Bernstein polynomials also offer an intriguing opportunity for emerging architectures. Unlike many bases, they require very little internal storage since most quantities are calculated on-the-fly during recurrence relations. This means they offer a low-memory, arithmetically-intense yet optimal-complexity way of obtaining high-order approximations.

Potential Impact on the field

Bernstein polynomials offer an opportunity to combine optimal algorithms, low memory requirements, and high-order resolution of problems in unstructured geometry without changing the fundamental mathematical finite/spectral element framework currently in use.

Contributed ideas paper 6

References

- [1] Mark Ainsworth, Gaelle Andriamaro, and Oleg Davydov. Bernstein–Bézier finite elements of arbitrary order and optimal assembly procedures. *SIAM Journal on Scientific Computing*, 33(6):3087–3109, 2011.
- [2] Mark Ainsworth, Gaelle Andriamaro, and Oleg Davydov. A Bernstein–Bézier basis for arbitrary order Raviart–Thomas finite elements. *Constructive Approximation*, 41(1):1–22, 2015.
- [3] Michael G. Duffy. Quadrature over a pyramid or cube of integrands with a singularity at a vertex. *SIAM journal on Numerical Analysis*, 19(6):1260–1262, 1982.
- [4] Robert C. Kirby. Efficient discontinuous galerkin finite element methods via bernstein polynomials. arXiv:1504.03990.
- [5] Robert C. Kirby. Fast simplicial finite element algorithms using Bernstein polynomials. *Numerische Mathematik*, 117(4):631–652, 2011.
- [6] Robert C. Kirby. Low-complexity finite element algorithms for the de Rham complex on simplices. *SIAM J. Scientific Computing*, 36(2):A846–A868, 2014.
- [7] Robert C. Kirby and Thinh Tri Kieu. Fast simplicial quadrature-based finite element operators using Bernstein polynomials. *Numerische Mathematik*, 121(2):261–279, 2012.
- [8] Anders Logg, Kent-Andre Mardal, and Garth N. Wells. *Automated solution of differential equations by the finite element method: The fenics book*, volume 84. Springer, 2012.
- [9] Steven A. Orszag. Spectral methods for problems in complex geometries. *Journal of Computational Physics*, 37(1):70–92, 1980.
- [10] A. H. Stroud. *Approximate calculation of multiple integrals*. Prentice-Hall Inc., Englewood Cliffs, N.J., 1971. Prentice-Hall Series in Automatic Computation.

Contributed ideas paper 7

A physics based iceberg calving model coupled with a global ice-sheet flow model for accurate assessment of sea level rise

November 24, 2015

Haim Waisman (Columbia), Jeremy Bassis (Michigan), Stephen Price (LANL), Ray Tuminaro (SNL), Irina Tezaur (SNL)

Abstract

Mass loss from the Antarctic Ice Sheet occurs primarily through its ice shelves, which are floating platforms of ice that extend into the ocean and surround the Antarctic ice sheet. Because of the cold atmospheric temperatures, most of this mass loss comes from the combination of the detachment of icebergs (calving) and the gradual erosion of ice by submarine melting. One challenge in predicting future ice sheet mass loss, however, is that many of the processes associated with ice shelf demise remain poorly understood. Although progress has been made in our understanding and modeling of submarine melting, a quantitative understanding of iceberg calving from ice shelves remains elusive. This is partly because calving requires that we understand and simulate the three-dimensional initiation and propagation of fractures within the ice over timescales that range from seconds to decades (or longer) [1, 2]. Iceberg calving rates provide boundary conditions to global ice sheet models and thus are extremely important. However, current models are too simplistic and do not demonstrate an ability to mimic observations of ice shelf fracture and calving. Future atmospheric warming may lead to increased melt ponding, “hydrofracture” (as occurred on the Larsen A/B ice shelves), and calving of Antarctic ice shelves (see, e.g., [4]). This process is critical to include in large-scale, predictive models of the Antarctic ice sheet [6, 5]; without accurate models of ice shelf fracture and calving, a critical link between atmospheric forcing and ice shelf decay will be absent in coupled ice sheet and climate models, biasing projections of ice sheet evolution and sea-level rise.

To this end, we propose a physically-motivated ice fracture and calving model (solid mechanics based) coupled with an ice sheet dynamics model (fluid mechanics based), to be incorporated as part of a larger global Earth System Model. The overarching goal is to provide more realistic predictions (including uncertainty ranges) of ice sheet evolution for estimating sea-level rise due to climate warming.

Background/Research to Date

In our previous work, a thermo-viscoelastic model for ice together with a creep damage model was proposed to model ice rheology and ice calving. The model was calibrated using laboratory data of uniaxial creep tests and is currently implemented in the finite element program FEAP [3, 7, 8]. Damage mechanics provides a promising approach to model ice calving and alleviate some of the limitations of traditional fracture mechanics. For example it (1) can model crack initiation and propagation (2) does not require adaptive re-meshing or mesh moving procedures to track fractures and (3) can easily be incorporated with viscoelastic constitutive laws to model the time dependent behavior of ice. In this approach, the usual continuum equations are augmented with an additional internal state variable D representing material deterioration or damage, ranging from $D = 0$ (ice is completely intact) to $D = 1$ (ice is entirely fractured). The evolution of the damage variable D accounts for the progressive accumulation and coalescence of micro-cracks (and also micro-voids) within the ice as it deforms under low loading rates.

In a recent work, the approach described above has also been extended to account for the simultaneous propagation of water filled surface and basal crevasses using a unique poro-

Contributed ideas paper 7




Figure 1: A global ice sheet model coupled to a damage mechanics model at the boundaries

mechanics approach [9] which is important for accurate modeling of the calving physics.

Proposed Direction of work

Coupling a global ice sheet flow model with a calving model at its boundaries is illustrated in Figure 1. This will require significant research to address several open questions related to: (1) the coupling approach: a fully concurrent approach (fluid and solid parts are solved concurrently) versus a staggered approach (fluid and solid parts are solved sequentially) versus an information passing approach (parametric studies are carried out on the solid part to obtain reduced order models or “calving laws” in a functional form, and are then applied to the fluid as boundary conditions); (2) uncertainty quantification across the fluid-solid model and propagation with time; (3) appropriate time stepping methods and solvers for the fluid-solid models, considering a range of time scales and large-scale problems; (4) further improvements to the ice fracture and calving modeling capabilities developed by the PIs (e.g., considering melting and erosion of the ice).

Connections to Math, Comp Sci & and Climate Science

The proposed work leverages past and ongoing DOE-funded efforts (e.g., ISICLES, PISCEES, LDRD, some of which have supported the relevant work discussed herein) in developing next-generation ice sheet models and coupling of those models to DOE-supported Earth System Models. Software and tools developed previously through the DOE BER/ASCR SciDAC initiative as well as the Trilinos solver software will be leveraged in order to solve these large-scale coupled systems on DOE’s high-performance computing platforms. Successful completion of the project will build upon the PIs expertise and collaboration in the disciplines of glaciology, applied mathematics, and computational science.

Contributed ideas paper 7

Potential Impact on the field

Most existing ice sheet models (including DOE's) simulate ice flow dynamics on a large scale, e.g. over all of Greenland and Antarctica, but are currently biased by their inability to simulate fracture or iceberg calving in a physically motivated or verifiable way. Fracture and calving rates, which have a tremendous impact on ice shelf evolution and stability, are also tightly coupled to atmospheric and ocean processes. An advanced capability at including ice sheet models as components of Earth System Models puts DOE in a unique position to take advantage of the proposed improvements to fracture and calving models, in order to provide more realistic predictions and uncertainty ranges on estimates of future sea-level rise due to climate warming.

References

- [1] Bassis JN. *The statistical physics of iceberg calving and the emergence of universal calving laws*. Journal of Glaciology, 57(201):3-16, 2011.
- [2] Bassis JN and Ma Y. *Evolution of basal crevasses links ice shelf stability to ocean forcing*. Earth and Planetary Science Letters, 409:203-211, 2015.
- [3] Duddu R. and Waisman H. *A temperature dependent creep damage model for polycrystalline ice*. Mechanics of Materials, 46:23-41, 2012.
- [4] Trusel L. D., Frey K. E., Das S. B., K. B. Karnauskas, M. P. Kuipers, E. V. Meijgaard, and M. R. Van Den Broeke. *Divergent trajectories of Antarctic surface melt under two twenty-first-century climate scenarios*. Nature Geoscience, 2015.
- [5] Tezaur I., Perego M. , Salinger A., Tuminaro R. and Price S. *Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis*. Geoscientific Model Development, 8:11971220, 2015.
- [6] Evans K. J., Salinger A. G., Worley P. H., Price S. F., Lipscomb W. H., Nichols J. A., White III J. B., Perego M., Vertenstein M., Edwards J., and J. F. Lemieux. *A modern solver interface to manage solution algorithms in the Community Earth System Model*. International Journal of High Performance Computing, 26:54-62, 2012.
- [7] Duddu R. and Waisman H. *A nonlocal continuum damage mechanics approach to simulation of creep fracture in ice-sheets*, Computational Mechanics, 51(6):961-974, 2013.
- [8] Duddu R., Bassis J. and Waisman H. *A numerical investigation of surface crevasse propagation in glaciers using nonlocal continuum damage mechanics*, Geophysical Research Letters, 40(12):3064-3068, 2013.
- [9] Mobasher E. M., Duddu R., Bassis J. and Waisman H. *Modeling hydraulic fracture of glaciers using continuum damage mechanics*. Journal of Glaciology, submitted, 2015.

Contributed ideas paper 8

Unique dynamical structures and multiscale modeling of turbulent geophysical flows

11/24/1025

Qingshan Chen
Clemson University

Abstract

Turbulence is regarded as one of greatest challenges in science and engineering today ([10, 11, 1, 12, 13, 4, 7]). What makes it so hard to solve, both analytically and numerically, is the presence of a wide and continuous spectrum of scales and the nonlinear interactions that they are engaged in. In the coming decade, exascale computing will become reality, and many physical processes and problems that are previously out of reach, such as cloud formation, can now be considered in climate models. But this does not mean that all the scales can be solved. Solving the flows down to the Kolmogorov scale will still be prehistoric, even for exascale computing. Thus, exascale computing makes it possible to consider many more problems and physical processes, but it also presents us an ever wider range of scales to be dealt with. Researchers need a systematic way of treating these unresolved or under-resolved scales. We propose a multiscale modeling program that will exploit the unique dynamical structures of large-scale geophysical flows, and lead to multiscale algorithms that are faithful to the long-term dynamics of the flows, and are more efficient than direct numerical simulations.

Background/Research to Date

Turbulence is an area that both theoreticians and practitioners take an keen interest in ([6]). As a result of these converging interests, there are a diverse body of ideas and approaches towards turbulence modeling. In the CFD community, large-eddy simulation (LES) and the Reynolds Averaged Navier-Stokes (RANS) are the two most common techniques. The LES ([16, 2]) seeks to model the effect of the subgrid eddies on large-scale motions, usually through the Boussinesq hypothesis that eddies are dissipative in the mean. In contrast, RANS ([19]) does not resolve any scales, but keep track of the time-averaged quantities. The literature on these two methods is vast; besides the references already given, [15] provides a survey on these and related techniques. A third approach, the renormalization group (RG) theory, was introduced for turbulence modeling by Orszag and Yaghot ([20]). RG was initially invented by Feynman et al in the study of quantum physics. It is an iterated process that, at each iteration, systematically eliminates the small-scale motions (through averaging, for example), and yields an set of equations for the large-scale motions. With some assumptions and approximations, this method leads to models that has no parameters to tune ([17]). In the geophysical fluid dynamics community, the best known eddy closure scheme is the so-called Gent-McWilliam closure ([8, 9]), which parameterize the eddy transport of passive tracers as a bolus velocity. A notable new development in this area is the work by Ferreira and Marshall ([5]) and by Young ([21]); They show that through a proper averaging of the governing equations, the tracers transport will all be taken care of by the residual mean velocity (bolus velocity does not appear separately). The eddy effect in the momentum equations is captured in the Eliassen-Palm tensor. This approach is compared with the classical GM closure in [14]. In mathematics literature, on turbulence modeling, a very deep and fundamental concept that appeared in the last 20-30 years is invariant manifold (IM,[18, 6]). An IM is a geometric object in the phase space of the dynamical system under study. It has several types, e.g. stable, unstable, center,

Contributed ideas paper 8

and inertial, all of which captures the long-term dynamics of the system. Based on the concept of inertial manifold, and its generalization, the approximate inertial manifold, the nonlinear Galerkin method was proposed, and applied to dynamical, multilevel turbulence modeling ([3]).

Proposed Direction of work

Conceptually, this research program has three essential components, each of which is a task to be undertaken in this project. The three conceptual tasks (CT) are

- A *By means of scale analysis, identify and classify unique dynamical structures for practical turbulent flows.*
- B *Connect unique dynamical structures of practical turbulent flows with the fundamental concept of invariant manifold in dynamical system theories. Analyze the existence of invariant manifolds, and derive the equations for approximate invariant manifolds.*
- C *Develop multiscale algorithms that are efficient and faithful to the long-term dynamics of the system.*

In this project we will deal with three geophysical models, each of which constitutes a Modeling Task (MT). The three models are, in the order of increasing complexity, the shallow water equations (MT1), the multi-layer isopycnal model (MT2), and the Boussinesq hydrostatic primitive equations (MT3). For each of these models, all three of the conceptual tasks will be carried out.

Connections to Math, Comp Sci & and Climate Science

This program is cutting-edge and highly interdisciplinary in that it combines the most advanced and state-of-art concepts and tools from mathematics, geophysics, and computational sciences to tackle one of the greatest challenges in science.

The cornerstone of our multiscale modeling program is its emphasis on the unique dynamical structures exhibited by large-scale geophysical flows, such as geostrophic balance, stratification, planetary boundary layers etc. We classify and quantify these dynamical structures by connecting them with fundamental concepts from dynamical system theories, such as global attractors, invariant manifolds, center manifolds, and inertial manifolds. The multiscale algorithms will be derived based on the invariant manifolds associated with the flows, using an approach similar to how the nonlinear Galerkin methods are derived. Finally, a successful implementation of our multiscale algorithms requires a computational framework that is robust, efficient, and flexible, one that can simultaneously implement various multiscale algorithms in multiple regions, since the dynamical structures that these algorithms are based on are likely localized.

Potential Impact on the field

The multiscale algorithms developed under our program are faithful to the long-term dynamics of the large-scale geophysical flows, and are more efficient than direct numerical simulations. Our program has the potential to change how multiscale modeling is done, and to dramatically improve the outcomes of predictive simulations for geophysical flows. It will help to unleash the power of exascale computing in the coming decade, and sharply increase the fidelity of mid- to long-range climate modeling.

Contributed ideas paper 8

References

- [1] G K Batchelor. Computation of the Energy Spectrum in Homogeneous Two-Dimensional Turbulence. *Physics of Fluids*, 12(12):II-233–II-239, 1969.
- [2] L. C. Berselli, T. Iliescu, and W. J. Layton. *Mathematics of large eddy simulation of turbulent flows*. Springer-Verlag, Berlin, 2006.
- [3] Thierry Dubois, François Jauberteau, and Roger Temam. *Dynamic multilevel methods and the numerical simulation of turbulence*. Cambridge University Press, Cambridge, 1999.
- [4] Gregory Falkovich and Katepalli R Sreenivasan. Lessons from hydrodynamic turbulence. *Physics Today*, 59(4):43–49, April 2006.
- [5] D Ferreira and J Marshall. Formulation and implementation of a “residual-mean” ocean circulation model. *Ocean Modelling*, 13(1):86–107, January 2006.
- [6] C Foias, O Manley, R Rosa, and R Temam. *Navier-Stokes equations and turbulence*, volume 83 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, Cambridge, 2001.
- [7] U. Frisch. *Turbulence: the legacy of AN Kolmogorov*. Cambridge Univ Pr, 1995.
- [8] Peter R. Gent and James C. McWilliams. Isopycnal Mixing in Ocean Circulation Models. *Journal of Physical Oceanography*, 20(1):150–155, 1990.
- [9] Peter R. Gent, Jurgen Willebrand, Trevor J. McDougall, and James C. McWilliams. Parameterizing Eddy-Induced Tracer Transports in Ocean Circulation Models. *Journal of Physical Oceanography*, 25(4):463–474, 1995.
- [10] Andrey Nikolaevich Kolmogorov. Dissipation of energy in locally isotropic turbulence. *Proceedings of the USSR Academy of Sciences*, 32:16–18, 1941.
- [11] Andrey Nikolaevich Kolmogorov. The local structure of turbulence in incompressible viscous fluid for very large Reynolds numbers. *Proceedings of the USSR Academy of Sciences*, 30:299–303, 1941.
- [12] A S Monin and A M Yaglom. *Statistical fluid mechanics: mechanics of turbulence. Vol. I*. MIT Press, Cambridge, MA, 1971.
- [13] A S Monin and A M Yaglom. *Statistical fluid mechanics: mechanics of turbulence. Vol. II*. MIT Press, Cambridge, MA, 1975.
- [14] Juan A Saenz, Qingshan Chen, and Todd Ringler. Prognostic Residual Mean Flow in an Ocean General Circulation Model and its Relation to Prognostic Eulerian Mean Flow. *J. Phys. Oceanogr*, 45(9):2247–2260, September 2015.
- [15] P Sagaut, S Deck, and M Terracol. *Multiscale and multiresolution approaches in turbulence*. London : Imperial College Press ; Hackensack, NJ : Distributed by World Scientific, 2006.
- [16] J Smagorinsky. GENERAL CIRCULATION EXPERIMENTS WITH THE PRIMITIVE EQUATIONS. *MWR*, 91(3):99–164, 1963.
- [17] Leslie M Smith and Stephen L Woodruff. Renormalization-Group Analysis of Turbulence. *Annual Review of Fluid Mechanics*, 30(1):275–310, January 1998.

Contributed ideas paper 8

- [18] Roger Temam. *Infinite-dimensional dynamical systems in mechanics and physics*, volume 68 of *Applied Mathematical Sciences*. Springer-Verlag, New York, New York, NY, second edition, 1997.
- [19] H Tennekes and J L Lumley. *A First Course in Turbulence*. MIT Press, 1972.
- [20] Victor Yakhot and Steven A Orszag. Renormalization group analysis of turbulence. I. Basic theory. *Journal of Scientific Computing*, 1(1):3–51, 1986.
- [21] William R. Young. An exact thickness-weighted average formulation of the Boussinesq equations. *Journal of Physical Oceanography*, 42:692–707, 2012.

Contributed ideas paper 9

July 18, 2016

Smart Proxy; Application of Big Data Analytics in Computational Climate Science

Shahab D. Mohaghegh
Intelligent Solutions, Inc. & West Virginia University

Abstract

Computational climate science is a computationally expensive discipline. High fidelity numerical models in computational climate science take long to run (execute) even on cluster of large number of parallel CPUs. To address this issue, proxy models are developed to decrease CFD execution time. The current paradigm requires sacrificing either physics (solve a simpler problem – reduced physics) or resolution (solve a coarser problem – lower resolution in space and time) in order to gain on speed.

Smart proxy is a complete paradigm shift on how proxy models are developed. As such they do not require either the physics or the resolution of a given CFD model to be compromised to achieve high execution speed. Smart proxies are developed using big data solutions (machine learning and pattern recognition). To build an accurate smart proxy of a CFD model (no matter how complex the model), a small number of simulation runs are made, very large amount of data is extracted from each simulation run, the extracted data is used to train, calibrate and validate a smart proxy that can accurately mimic the functionalities of the CFD model and run in seconds.

The smart proxy that is generate in this fashion has all the functionalities of the original model at the grid block (cell) level. It is far more accurate and runs several orders of magnitude faster than any other proxy model. Smart proxy will significantly increase the utilization of highly complex CFD models that are currently used in computation climate science. Tasks such as quantification of uncertainty associated with model results (that requires tens of thousands of simulation runs) will become practical and routine.

Background/Research to Date

Computational science, addressing numerical solution of complex multi-physic, non-linear, partial differential equations, is on the forefront of engineering problem solving and optimization. These solutions (Reservoir simulation in Petroleum industry and CFD in other industries) are computationally expensive. Proxy models (statistical response surfaces, or reduced physics) attempt to make it practical to use these models by addressing their computational footprint. However these traditional techniques of developing proxy models have limited applications and success.

A relatively new technology that has been successfully applied in the upstream petroleum industry is Surrogate Reservoir Models (SRM) [1], [2], [3]. SRM is a smart proxy of the numerical reservoir simulation model. SRM is a new form of proxy modeling that attempts to observe and learn the detail behavior of a numerical model and then mimic its behavior with accuracy. SRM is a novel implementation of Artificial Intelligence and Data Mining technology in the numerical modeling arena.

Data-Driven Smart Proxies take advantage of the “Big Data” solutions (machine learning and pattern recognition) to develop replicas of the numerical models that are highly accurate, but have very fast response time. Implementation of these proxy models to large and complex numerical

Contributed ideas paper 9

reservoir simulation models have resulted in successful outcome for the oil and gas industry. Figure 1 demonstrates an example of the smart proxy (SRM) in a CO₂ storage numerical model where the SRM has accurately replicated the pressure and saturation distribution throughout the reservoir. The results generate by the SRM in seconds while the actual numerical model take hours.

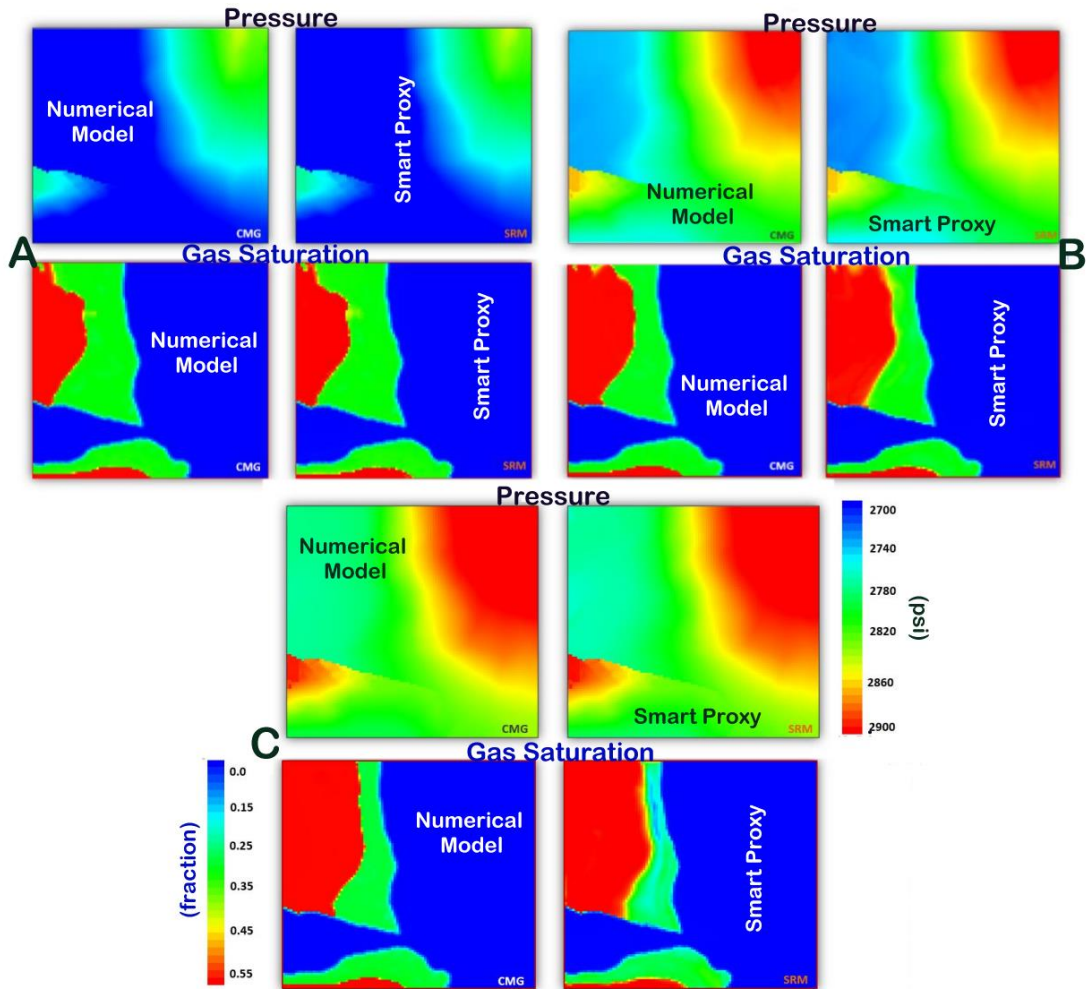


Figure 1. Comparison of the results from Numerical Model and Smart Proxy for Pressure and Gas Saturation Distribution in the first layer of the reservoir, in a Carbon Storage Numerical Model (A) after 1 month of injection, (B) after 8 month of injection, (C) after 16 month of injection.

Figure 2 is an example of our most recent application of the Smart Proxy technology to a multi-phase CFD model. This project was funded by DOE-NETL to explore the possibility of application of the Smart Proxy technology to complex CFD models as a “proof of concept” project.

This figure demonstrates an example of the smart proxy applied to a three-dimensional CFD model simulating a fluidized bed (sand originally at the bottom of the cylinder is

Contributed ideas paper 9

fluidized by the air blowing from the bottom forming transient and turbulent flow). About 4 seconds of the CFD simulation takes about three days of computational time (using DOE's MFIX code) on a cluster of multiple CPUs. The smart proxy that is developed for this particular CFD model replicates the entire process with accuracies shown in Figure 2 in less than 3 minutes.

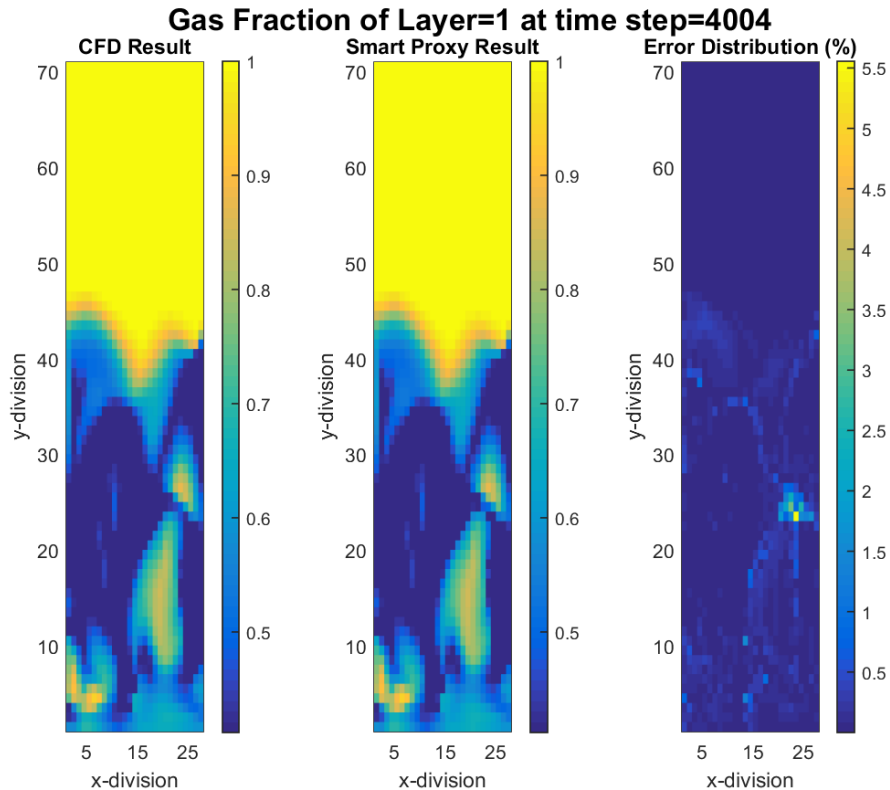


Figure 2. Comparison of the results from a fluidized bed (multi-phase) CFD Model and Smart Proxy for void fraction Distribution. The 3D CFD model that takes several days to execute using a cluster of CPUs is replicated by a smart proxy that completes the CFD execution in less than 3 minutes.

In this presentation, the complete video of accurate replication of the fluidized bed (a single frame – time-step – is shown in Figure 2) will be shown. The audience are encouraged to enter a discussion on the pros and cons of application of this technology to problems that are solved using CFD.

Proposed Direction of work

It is proposed to perform more proof of concept studies on this topic by selecting a reasonably complex climate science model and develop a smart proxy of the model. The objective of the proof of concept study will be to demonstrate that such an idea actually works for models used in computational climate science. In other words, it is proposed to demonstrate that using Big Data

Contributed ideas paper 9

solutions, a smart proxy of a fairly complex (a multi-physics, high resolution model that take tens of hours for a single run) climate model can be trained and validated. Validation of the smart proxy model is performed by having it reproduce (in a very short period of time - seconds) the detail responses of several new simulation runs and compare its results with actual simulation runs that have taken days to accomplish.

Then the smart proxy can be used to perform large number of simulation runs to for example quantify uncertainties with certain types of predictions. Furthermore, methods can be then developed to extract, summarize and quickly analyze large amounts of data that would result from large number of simulation (smart proxy) runs that can now be accomplished a very short period of time.

Connections to Math, Comp Sci & and Climate Science

Smart proxy uses a combination of multiple artificial neural networks and fuzzy systems (including intelligent agents) in order to learn the behavior of a complex numerical models. The math and computer science involved in the development of smart proxy is the same that is used in other artificial intelligence related projects.

Potential Impact on the field

The potential impact of smart proxy in the computational climate science should be assessed by climate scientists. It just seems reasonable to expect that being able to perform very large number of simulations that can examine very large number of possible scenarios and quantify uncertainties associated with our models can be of interest and benefit to the climate science.

References

- [1] Mohaghegh et. al. Converting Detail Reservoir Simulation Models into Effective Reservoir Management Tools using SRMs; Case Study – Three Green Fields in Saudi Arabia. International Journal of Oil, Gas and Coal Technology. Vol. 7, No. 2, 2014, pp.115–131.
- [2] Mohaghegh et. al. Smart Proxy: an Innovative Reservoir Management Tool; Case Study of a Giant Mature Oilfield in the UAE. SPE 177829, ADIPEC - Abu Dhabi International Petroleum Exhibition and Conference. Abu Dhabi, UAE, 9–12 November 2015.
- [3] Mohaghegh et. al. Grid-Based Surrogate Reservoir Modeling (SRM) for Fast Track Analysis of Numerical Reservoir Simulation Models at the Grid block Level. SPE 153844 - Western Regional Meeting. Bakersfield, California, USA, 21-23 March 2012

Contributed ideas paper 10

Modularizing Global Variable In Climate Simulation Software

Tapajit Dey, Yuxing Ma, and Audris Mockus

Department of Electrical Engineering and Computer Science

The University of Tennessee, Knoxville

Abstract

In simulation codes, such as climate models, variables represent a large number of characteristics of earth surface and atmosphere for a single multi-dimensional cell and are distributed over a multitude of cores in the supercomputers where these simulations run. The hundreds of variables allow different parts of simulation representing specific sub-models, for example the photosynthesis, to interact with other sub-models of the simulation.

The scientists of each domain write the simulation code for the sub-model representing their sub-domain. To integrate their code into the entire simulation, they need to deal with hundreds of unfamiliar variables of which only a small subset is relevant to their work. Designing such variables in a modular fashion, so that the scientists could interact only with the variables relevant to their sub-model is likely to increase the productivity of the scientists and to increase accuracy of the simulation codes.

A natural way to group the variables into modules is by using a language feature that group them together, such as, struct construct in C language or a class in C++ language. Each scientist would then need to familiarize themselves with only a small subset of modules that contain variables used in their simulations. For example, Community Earth System Model (CESM) code v1.06 has 51 such modules (structures) that contain 1479 variables. The methods proposed below can be used to assess the modularity of the existing set of structures and to generate alternative modularizations that improve upon it.

In a nutshell, the approach minimizes the number of variables exposed to other domains by the modules used in each domain.

Background/Research to Date

To describe the approach in more detail, we need to define several concepts precisely.

- v_i : is a variable used in one or more domains of the simulation.
- A module m_j is a subset of variables. Modules partition the set of variables.
- A domain d_k is a sub-model representing a specific biological, physical, or other type of process. Each d_k reads from or writes to one or more variables v_i . Unlike for modules, the same variable may belong to multiple domains.
- m_j is *necessary for domain* $d_k \iff \exists v_i : v_i \in d_k \wedge v_i \in m_j$.
- v_i is visible from d_k if it belongs to any of the modules m_j necessary for d_k .

Each partition of v_i s induces a set of modules. We can measure the quality of the partition by counting the variables visible from all domains:

$$\sum_k |\cup_{m_j \text{ necessary for } d_k} v_i| \quad (1)$$

The lower bound for this criteria is simply the number of variables. For example, if the unions of modules necessary for each domain form a partition (do not overlap for any two domains). If, however, any domain uses at least one variable from another domain, this bound may not be achievable. A trivial way to achieve the minimum for the criteria in Equation 1 over all possible partitions given a set of domains is to create a separate module for each variable.

Contributed ideas paper 10

The criteria then becomes: $\sum_k |\cup_{v_i \in d_k}|$. This, however, creates far too many modules defeating the original goal of modularity. It is, therefore, important to note that the criteria should be compared for modularizations that bound the total number of modules or consider Pareto-optimal (non-dominated) solutions defined by minimizing the number of modules and the value of the criteria.

We applied the approach to CESM code (v1.06) and discovered 216 domains (functions) with the optimal (trivial) partition with 1479 modules (each corresponding to one variable) yielding the score (from Equation 1) of 3703. The existing partition of CESM code into the 51 modules (C structures) has the score of 35703. This means that on average, a domain scientist sees ten times more variables than they need to. A partition that contains all variables in a single module, on the other hand, would have the score of 216×1479 or nine times more than the existing partition and 90 times more than the optimal partition.

Proposed Direction of work

The simplest way to resolve the contention between having too many modules and too many variables visible to a domain scientist is to add a penalty term to Equation 1. The penalty would increase with the number of modules used to partition the variables. Initially we will optimize the criteria for each fixed number of modules to find non-dominated solutions and then decide which of them is the most practical one. We plan to use several approaches, for example, a modification of the algorithm used to create independently changeable parts of code (modules) in a very large telecommunications system [1]. The algorithm iteratively redistributes these variables by removing or inserting a variable into module or exchanging a variable between two modules. If the move improves the score it is taken if not it is accepted with a probability p that goes down with iterations (annealing). Using this method, the score decreases from 35703 to 29982 after 10 minutes on a single-core cpu. The probability p and the annealing rate can be optimized to achieve faster convergence. The results, while asymptotically optimal, may yield a different solution with each finite run. The search may be further improved with an initial partition that has no overlaps (where each domain needs only variables from a single partition.) In CESM code we found six non overlapping partitions. Only five exiting modules cross this partition boundary.

Connections to Math, Comp Sci & and Climate Science

The proposed work targets the domain of complex simulation codes, such as ones used in climate science, where multiple phenomena are being modeled with very large numbers of variables used to tie these sub-models together. We plan to apply the approach to several climate simulation codes investigate how data modularity changes over time and could be improved. We plane to interview domain scientist to understand better how data modularity could address their needs and quantify the impact of improved modularity.

Potential Impact on the field

Existing climate models and other-large-scale simulations require contributions of scientist from many different domains, who are familiar with the sub-models relevant to their domain. Existing simulation codes, however, expose them to thousands of variables representing unrelated phenomena. A suitable modularization these data structures may significantly increase the productivity of scientist and lead to more accurate simulations.

References

- [1] Audris Mockus and David M. Weiss. Globalization by Chunking: a Quantitative Approach. IEEE Software, 18(2):30-37, March 2001.

Contributed ideas paper 11

Defining Features in Uncertain Unsteady Flows

December 1, 2015

Hanqi Guo and Tom Peterka, MCS, Argonne National Laboratory
Scott Collis and Jonathan Helmus, EVS, Argonne National Laboratory
Han-Wei Shen and Wenbin He, The Ohio State University

Abstract

The objective of this research is to understand flow transport behavior in uncertain time-varying flow fields. Doing so requires defining stochastic counterparts to traditional deterministic feature metrics. For example, the finite-time Lyapunov exponent (FTLE) and Lagrangian coherent structure (LCS) metrics for deterministic flows can be revamped to U-FTLE and U-LCS to incorporate uncertainty for stochastic flows. The U-FTLE captures the distribution of divergent and convergent flow, and the U-LCS captures the probabilities of finding ridge structures separating the flow into different regions based on divergence. The development of the theory and algorithms for these and other stochastic flow features can help climate scientists understand transport behaviors in ensemble simulations of atmospheric and oceanic processes.

Background/Research to Date

Instead of a single velocity vector at each grid point, an uncertain flow field contains distributions of velocities at each grid point. However, the visualization and analysis of uncertainty still remain grand challenges in the climate science community, particularly uncertainty in time-varying flow fields. For 2D and small scale 3D uncertain datasets, uncertainties can be encoded with glyphs [9, 6] and textures [2]. Vector field topology methods have been extended to uncertain vector fields [7, 8], but they are limited to steady-state data. Recently, we [3] proposed uncertainty metrics for 4D time-varying flow fields by extending a well-established tool—the finite-time Lyapunov exponent (FTLE) [4]—to a probabilistic setting. The FTLE is a scalar field that measures the convergence or divergence rate between neighboring particles in the flow. The ridges of FTLE fields can be used to derive Lagrangian coherent structures (LCS) [5]: the boundaries between attracting or repelling particles in the flow.

Proposed Direction of Work

Traditional deterministic analysis needs to be redefined to accommodate uncertainty, and new algorithms are needed to compute the new quantities. Such stochastic formulations will allow climate scientists to quantify the uncertainty of convergent and divergent transport behaviors that can aid in understanding the uncertainty of derived features such as eddies, flow segmentation, and large-scale teleconnections. For example, we can derive an uncertain finite time Lyapunov exponent (U-FTLE), a statistical representation of FTLE values in uncertain unsteady flow, to capture differences in the distributions of advected neighboring particles. The U-FTLE leads to uncertain Lagrangian coherent structures (U-LCS) that model the probability of belonging to a ridge separating flow regions in the domain. Ridges segmenting the flow field can then be extracted with stochastic ridge finding and density estimation algorithms.

Contributed ideas paper 11




Figure 1: Left: Hurricane Isabel ensemble: (a) ridges of individual ensemble members, (b) U-LCS of all members. Right: WRF ensemble: (a) U-LCS, (b) U-FTLE.

Connections to Math, Computer Science, and Climate Science

The left side of Figure 1 shows the ridge detection results of a dataset modeling hurricane Isabel. In this case, uncertainty results from temporal downsampling of the output data, a common practice. From the U-FTLE, we can see that the FTLE values and their variances are higher near the hurricane eye. We also observe that the U-FTLE values are highly non-Gaussian. The ridges from all FTLE runs are combined into the U-LCS field using curve or surface density estimation. The right side of Figure 1 shows the Weather Research and Forecasting (WRF) model called High Resolution Rapid Refresh (HRRR) [1]. We generated an ensemble from several forecasts of the same point in time, and we modeled the uncertainty of the wind field using U-LCS and U-FTLE. The visualization results show that most of the uncertainty is in areas of steep vertical convection. Even though both of these examples feature weather applications, similar techniques will be needed for future extreme-scale climate simulations. Climate models have similar sources of uncertainty caused by divergence in ensembles and by downsampling errors as do weather models. Hence, we envision that these techniques can be applied to ensembles of high resolution climate models to show uncertainty in features of the general circulation and the errors caused by temporal and spatial sampling.

Potential Impact on the Field

Our preliminary work demonstrates the utility of defining new methods for uncertain unsteady flow and examples of applications in two uncertain unsteady flow datasets from climate and weather models. In the first case, uncertainty resulted because of temporal downsampling. Because time-varying datasets can be extremely large to store, a common practice is to drop time steps for further analysis, but important information can be lost in this process. We were able to detect uncertainty in the separation between the hurricane eye and wall using U-FTLE. In the second case, the uncertainties arose from ensemble simulation runs of weather forecasts, and we visualized the surfaces in storm regions by extracting the U-LCS.

Combined with other visualization techniques, such methods can help scientists analyze the uncertainty of the simulation models. Future research is needed to develop intuitive ways to present the results of metrics such as U-FTLE and U-LCS. Once features are found, for example ridges, a natural next step is to develop stochastic tracking of those features over time. Another area of future work is to parallelize and scale such methods so that they can run in situ with simulations, and to embed these methods into software frameworks such as MPAS.

Contributed ideas paper 11

References

- [1] Curtis Alexander, S. S. Weygandt, D. C. Dowell S. Benjamin, T. G. Smirnova, E. P. James, M. Hu P. Hofmann, J. Olson, and J. M. Brown. The high-resolution rapid refresh: Recent model and data assimilation development towards an operational implementation in 2014. In *Proceedings 26th Conference on Weather Analysis and Forecasting / 22nd Conference on Numerical Weather Prediction*. American Meteorological Society, 2014.
- [2] Ralf P. Botchen and Daniel Weiskopf. Texture-based visualization of uncertainty in flow fields. In *Proc. IEEE Visualization 2005*, pages 647–654, 2005.
- [3] Hanqi Guo, Wenbin He, Tom Peterka, Han-Wei Shen, Scott Collis, and Jonathan Helmus. FTLE and LCS in Uncertain Unsteady Flows. In *To appear in Proceedings of Pacific Vis 2016*. IEEE, 2016.
- [4] George Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D: Nonlinear Phenomena*, 149(4):248–277, 2001.
- [5] George Haller. A variational theory of hyperbolic Lagrangian coherent structures. *Physica D: Nonlinear Phenomena*, 240(7):547–598, 2011.
- [6] Marcel Hlawatsch, Philipp Leube, Wolfgang Nowak, and Daniel Weiskopf. Flow radar glyphs - static visualization of unsteady flow with uncertainty. *IEEE Trans. Vis. Comput. Graph.*, 17(12):1949–1958, 2011.
- [7] Mathias Otto, Tobias Germer, Hans-Christian Hege, and Holger Theisel. Uncertain 2D vector field topology. *Comput. Graph. Forum*, 29(2):347–356, 2010.
- [8] Mathias Otto, Tobias Germer, and Holger Theisel. Uncertain topology of 3D vector fields. In *Proc. IEEE Pacific Visualization 2011*, pages 67–74, 2011.
- [9] Craig M. Wittenbrink, Alex Pang, and Suresh K. Lodha. Glyphs for visualizing uncertainty in vector fields. *IEEE Trans. Vis. Comput. Graph.*, 2(3):266–279, 1996.

Contributed ideas paper 12

Enhancing precision and performance of global and regional climate and environmental models on new parallel architectures

12/01/2015

George Delic, Ph.D. HiPERiSM Consulting, LLC, PO Box 569 Chapel Hill, NC 27514-0569

Abstract

The HiPERiSM experience in working with two major environmental and global models suggests that some opportunities for enhanced parallelism have not been sufficiently explored or developed. Specifically, porting of these models to recent parallel architectures such as add-on accelerators from major vendors (e.g. AMD, Intel, and NVIDIA) is nascent and needs fuller implementation. In addition coding practices and ignorance of compiler features place numerical precision at risk. This proposal seeks DOE support for methods that improve both precision and performance in the aqueous chemistry solvers that are the major compute kernels in these models.

Background/Research to Date

The two models studied by HiPERiSM in recent years are the Community Multiscale Air Quality model (CMAQ) developed at the US EPA, and the Global Climate Model (GCM) developed at Lawrence Livermore National Laboratory, and in use in the Global Modeling Initiative at NASA. Both models have user communities with numbers in the thousands world-wide.

The Chemistry Transport Model (CTM) in CMAQ and GCM has relied on a legacy algorithm (hereafter JSPARSE) originally developed on Cray computers over 20 years ago. These models implement a variety of ODE solvers including Euler Backward Iteration (EBI), Rosenbrock (RB), and Gear (GEAR) with successively increasing runtime cost. The escalation in runtime cost is due to the reliance of RB and GEAR algorithms on a direct LU decomposition and forward/backward solve that typically account for ~60% and ~20%, respectively, of the CTM wall clock time. Because of this, the superior solvers (RB and GEAR) are rarely used in CMAQ and the default choice is typically the EBI algorithm to avoid this added runtime cost. The GCM uses only the LLNL revision of the legacy JSPARSE algorithm. Both models apply MPI for domain decomposition of the entire grid domain. In the case of RB or GEAR algorithms the legacy JSPARSE method organizes grid cells in the domain into blocks of cells, and applies the convergence criterion to the rms error estimate over each cell block. Hence the precision achieved in the legacy algorithm varies as the number of cells in a block is changed.

Work at HiPERiSM [1,2] has replaced the legacy JSPARSE algorithm with a new FSPARSE solver: in the case of CMAQ for RB and GEAR ODE solvers, and in GCM for the GEAR solver. The latter case has resulted in a portable single source code version of the CTM that will execute in parallel on either host CPU's or attached Intel Phi Many Integrated Core (MIC) co-processors, depending on the choice of compiler flags. Fig. 1(a) shows that with MPI processes across two nodes and 4 Intel Phi cards on each, optimal performance occurs when each MPI process has its own dedicated Phi co-processor. Fig. 1(b) shows that observed numerical differences are large and often exceed the value of the species concentrations when comparing the application of the CTM algorithm to 20 cells in a block versus application to a benchmark of each individual cell.

Proposed Direction of work

1. Sparse matrix algorithms

The previous work described above implemented only one (new) sparse matrix algorithm. The direction of proposed work would extend this to the evaluation of other sparse matrix solvers for suitability to the CTM compute kernel of both CMAQ and GCM. These include sparse matrix algorithms available in various libraries such as those

Contributed ideas paper 12

packaged with the compiler (Intel), or the Software Development Kit (SDK with NVIDIA CUDA libraries). Others have become available from academic projects, such as the SuperLU library, or specialty algorithms targeting co-processors (Harwell, UK). The proposed work direction would adapt these for CTM implementation and perform comparative evaluation of execution on either host CPU or attached co-processors.

2. *Stiff ODE solvers*

Here work direction would extend to the evaluation of stiff ODE solvers other than RB or GEAR, and seek methods of improving precision and reducing wall clock time. This requirement is essential as in future the number of chemical species in climate models will increase well beyond the 72 to 193 currently in use.

3. *Limitations and challenges*

The fundamental issue is adapting sparse matrix and stiff ODE algorithms to take the best advantage of the performance potential of new parallel computer architectures. The limitations are due to the pervasive use of scalar code constructs and the challenge is to transform these constructs to parallel code that is portable to current and future commodity computer hardware.

Connections to Math, Comp Sci & and Climate Science

The connection here is with the research into algorithms for sparse matrix operations and stiff ODE solvers, with the emphasis on scalable parallelism. Algorithmic research in these fields has a long history that has not been sufficiently assimilated into Climate Science. Climate modelers are not computer scientist and, in our experience, continue to produce serial code that does not exploit the full potential of modern commodity computer processors.

The connection to computer science is to uncover sparse matrix and ODE solver algorithms that bring a closer approach to peak hardware performance. For example, comparing the Intel CPU and Phi processors, shows that they offer 589Gflop/s and 1208 Gflop/s, respectively. Thus, one MIC card of the first generation has double the peak performance potential of a CPU. However, for performance to approach this value on a MIC card, with a maximum bandwidth of 352 GB/s (44 Gword/s), an algorithm must reach $1208/44=28$ operations per word [3]. These peak values are not reached in the work reported here, but the challenge is to aim higher.

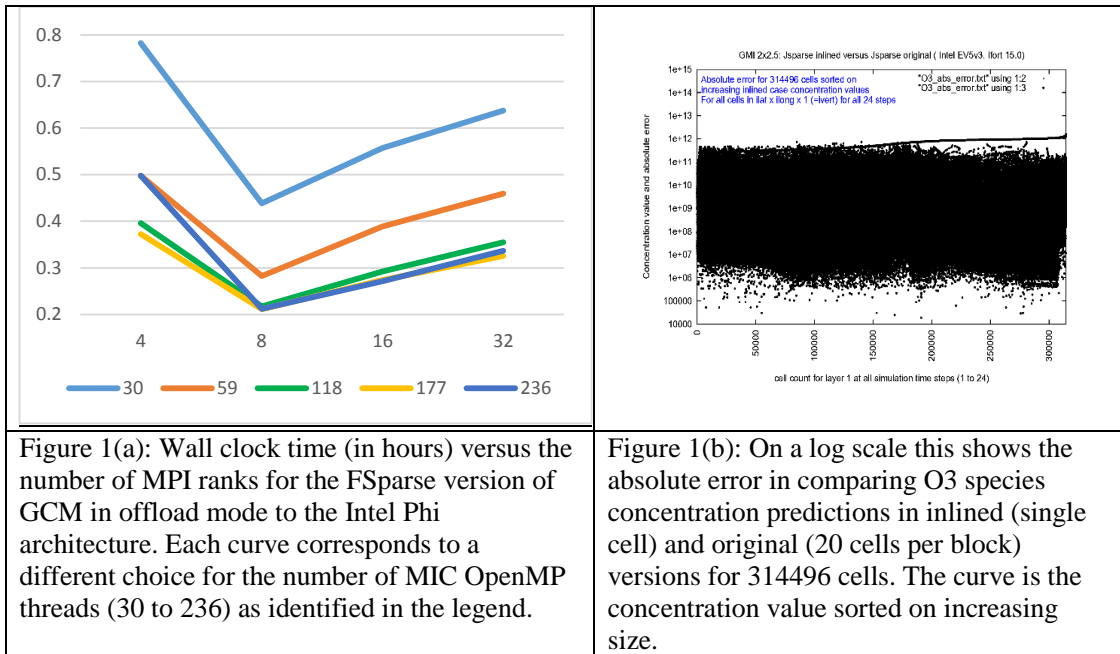
Potential Impact on the field

Prior work demonstrated that performance and precision gains attained so far with these models point the way to an in-depth research study of the relevant algorithms for stiff ODE solvers and sparse matrix libraries. These would bring rewards to the computational climate modeling community with delivery of more efficient and precise code. The outcome of the work proposed here could be to enhance the options available to climate modelers as a user community without expecting them to understand the intricacies of numerical algorithms they would have access to. Such access would be through attached libraries selectable through configuration scripts at the model compile stage.

References

- [1] Delic, G., 2013: contribution to 12th Annual CMAS Conference, Chapel Hill, NC, October 28-30, 2013, <https://www.cmascenter.org/conference/2013/agenda.cfm>
- [2] Delic, G., 2015: contribution to 14th Annual CMAS Conference, Chapel Hill, NC, October 5-7, 2015, <https://www.cmascenter.org/conference/2015/agenda.cfm>
- [3] Albert-Jan N. Yzelman, *et. al.* in High-Performance Parallelism Pearls: Multicore and Many-core Programming Approaches, James Reinders and Jim Jeffers, Morgan (Editors), Kaufmann/Elsevier, Waltham, MA, 2015.

Contributed ideas paper 12



Contributed ideas paper 13

Embedded Urban Framework for ACME Regions of Refined Resolution **Nov 30, 2015**
Melissa Allen, Marcia Branstetter, and Olufemi Omitaomu
Oak Ridge National Laboratory

Abstract

Many minute geospatially located biological processes are incorporated into the ACME model within the C-N-P and hydrological cycles, but high-temporal-resolution geospatially-explicit human activities are not. With micro-simulation modeling, real-time and real-space emissions from traffic; water management decision-making, and impacts to and restoration of cities' critical infrastructure could be tracked, projected, and fed back into the climate system. In anticipation of emerging global urbanization, better understanding and quantification of climate effects on and from human and natural systems in cities are needed, requiring coordinated research in microclimate, global climate, population dynamics and urban development. The proposed Urban Framework for ACME Regions of Refined Resolution (Urban-FAR) can address this need by: quantifying and analyzing the relationships among climatic conditions, urban form, land cover and energy use; and relationships among these activities (at high temporal and spatial resolution) and earth system processes. Research will focus on the effects of city topology, the urban heat island, urban emission and human decision-making on local, regional and global climate. The overarching objective of the project is the computational integration of different approaches across three research areas: earth system modeling; micro-simulation modeling; and impacts, adaptation and vulnerability scenario anticipation.

Background/Research to Date

A vast amount of research, modeling and simulation have been done in the area of 3D microclimate modeling and a variety of tools (using finite difference, finite volume and finite element) are available for its simulation (e.g., [1-3]). There have also been parameterizations for urban microclimate within Global Climate Models (e.g., [4,5]). Additionally, Gill et al. [6] developed a method for linking 3D landscape models to microclimate simulations. Deal and Fournier [7] present a recent culmination of the integration of ecological and engineering approaches to spatially explicit modeling of urban dynamics, and Benson and Torrens [8] integrate a variety of approaches into their "geosimulation." Yet only recently have modelers begun coupling climate modeling with agent-based simulations (e.g., [9]) to achieve high-resolution information about the impacts of human activity on the earth system and vice versa.

Proposed Direction of Work

We propose a high-performance computational dynamical systems framework in which road networks, building topology and land use/land cover are modeled, and over which finite volume (Navier-Stokes based) algorithms for detailed building and city aerodynamics are solved. Inputs will comprise building-level thermal, radiative and topological properties obtained from various LiDAR and multispectral imagery datasets (as in [10]) along with boundary conditions from the Community Land Model Urban portion of the ACME model. Outputs at meter scale capable of being aggregated to the CLM Urban inputs for a two-way coupling will be generated. The framework will also host collectives of autonomous agents within and near cities, who retain demographic attributes (e.g., US Census records by county or city block) and decision-making capabilities for real-time traffic decisions (e.g., data from traffic cams), building energy use (e.g., from EIA summaries disaggregated by building population) and city management of resources such as water and energy (city policy sources). The integrated model will also project spatially explicit (cohort-component) scenario-based changes in this population and probable locations of their commercial and residential development [11] for future decades. Diagnostics will be

Contributed ideas paper 13

developed to capture the transitions between the various scales within the proposed framework as well as measuring the propagation of uncertainties.

- **Key Capability:** This project represents the first coupled system to calculate simultaneously human and natural activities with earth system processes from the global domain to the scale at which humans operate.
- **Attributes:** common architecture for disparate processes at a wide range of scales.
- **Limitations:** computational expense, more variables admit more uncertainty (to be characterized and quantified), limited observational data both for input and validation (although proxies can be generated).

Connections to Math, Comp Sci & and Climate Science

The project will address **three major gaps in the existing computational capability:** i) neighborhood resolution modeling and simulation of urban micrometeorological processes and their effect on and from regional and global climate; ii) effects of city topology, including road networks, building configurations and green spaces, both existing and proposed, on these micrometeorological processes; and iii) impacts of human agent decision-making at various spatial and temporal scales on local to global and seconds to decadal changes in human and earth systems. To enable such a collaborative simulation, we introduce a dynamical systems framework that ingests high-resolution input data from multiple sources, performs a simulation to capture human dynamics and their implications, sends results back to the models as necessary, and continues the simulation with updated inputs from the models (Figure 1).

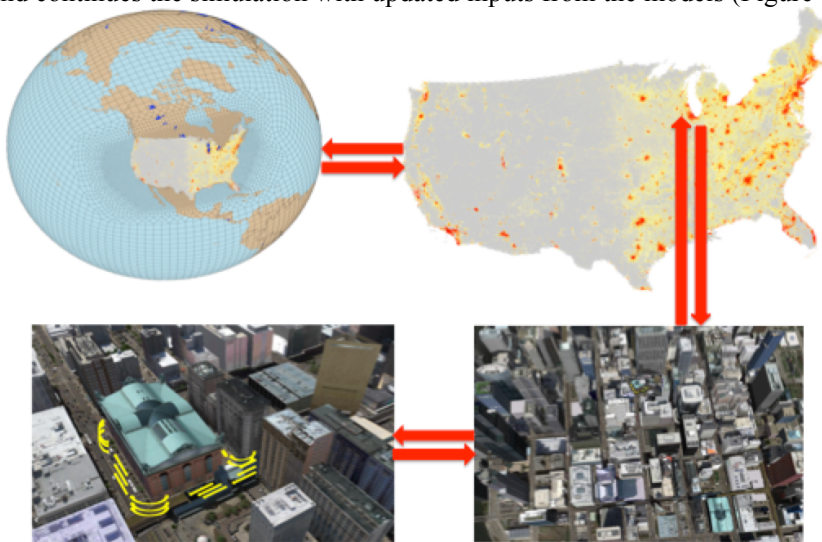


Figure 1: Representation of models and scales involved in a city simulation within an ACME regionally-refined grid cell. Results will be output at second and meter scale, but will also be aggregated up to the regional and global resolution for evaluation of their impact on the earth system as a whole.

Potential Impact on the field

The project has the potential to offer a major shift in earth system modeling because it will represent explicitly the unfolding of human activity within the earth system as the model runs; a true next generation earth system modeling capability. It will allow for deep characterization and investigation of the impact of a much wider variety of human activity on the environment and the climate than has ever before been modeled.

Contributed ideas paper 13

References

- [1] **Allen, M.R.**, A.N Rose, **M.L. Branstetter**, J. Yuan, J.R. New, **O.A. Omitaomu**, T.J. Wilbanks (2015). “Modeling Urban Energy Savings Scenarios using Earth System Microclimate and Urban Morphology.” Abstract GC21E-08 presented at 2015 Fall Meeting, American Geophysical Union (AGU), San Francisco, California, USA, 14-18 December.
- [2] ENVI-met (2015). <http://www.model.envi-met.com/hg2e/doku.php?id=root:start>
- [3] Tunick, A. (2005). Critical assessment of selected urban microclimate model frameworks, Army Research Laboratory, ARL-MR-619, Adelphi, MD 20783-1197.
- [4] Oleson, K.W., G.B. Bonan, J. Feddema, M. Vertenstein, and E. Kluzek (2010). Technical Description of an Urban Parameterization for the Community Land Model (CLMU). NCAR Technical Note NCAR/TN-480+STR, DOI: 10.5065/D6K35RM9.
- [5] Oleson, K.W., G.B. Bonan, J. Feddema, T. Jackson (2011). An examination of urban heat island characteristics in a global climate model, *Int. J. Climatol.*, **31**, 1848-1865.
- [6] Gill, L., E.A. Hathway, E. Lange, E. Morgan and D. Romano (2013). Coupling real-time 3D landscape models with microclimate simulations, *IJEPR* 2.1, 1-19. DOI:10.4018/ijepr.2013010101.
- [7] Deal, B. and D.F. Fournier (2000). Ecological Urban Dynamics and Spatial Modeling, *Energy and Environmental Policy*, 9.59. http://aceee.org/files/proceedings/2000/data/papers/SS00_Panel9_Paper06.pdf
- [8] Benson, I. and P. Torrens (2004). *Geosimulation: Automata-based modeling of urban phenomena*, Wiley, West Sussex, England.
- [9] Park, B.H., **M.R. Allen**, D.A. White, E.M. Weber, J.T. Murphy, M. North, P. Sydelko (2015). MIRAGE: A framework for data-driven collaboration, *Proceedings of the 13th International Conference on GeoComputation Geospatial Information Sciences*, May 20-23, Richardson, Texas, USA.
- [10] Kodysh, J.B., **O.A. Omitaomu**, B.L. Bhaduri and B.S. Neish (2013). Methodology for estimating solar potential on multiple building rooftops for photovoltaic systems, *Sustainable Cities and Society*, **8**, 31-41.
- [11] Glass, J. (2013). In Urban Planning for Climate Change, Pay Attention to Population Dynamics and Smaller Cities, Says Daniel Schensul, NewSecurityBeat November 8, 2013. Accessed April 19, 2015 at: <http://www.newsecuritybeat.org/2013/11/urban-planning-climate-change-pay-attention-population-dynamics-smaller-cities-daniel-schensul>

Contributed ideas paper 14

Compiler-based software analysis toolkit for climate model development 11/30/2015
Dali Wang, Oak Ridge National Lab (ORNL), Oscar Hernandez (ORNL), Graham Lopez (ORNL), Frank Winkler (NCCS/Vampir), Michael Wolfe (NVIDIA/PGI)

Abstract

With decades of rapid developments, complicated software systems have been adopted by climate modeling communities to address scientific questions. Software complexity has become a real barrier that impedes further code development of these scientific codes, including adding new features into the codes, validating the knowledge incorporated in the codes, as well as redesigning the codes for new and emerging computational platforms such as Exascale computers. In the effort, we develop compiler-based innovative methods to (1) better understand existing climate models, (2) modularize complex code, (3) provide analysis and improve instrumentation mechanisms that help users write performance portable codes (4) generate functional testing for key software modules, and (5) add new features for new scientific investigations. This effort will (1) generalize some of the current software engineering practices within the Accelerated Climate Modeling for Energy (ACME) and Interoperable Design of Extreme-scale Application Software (IDEAS) project and (2) co-develop an open-source code analysis toolkit based on the industry-leading compiler: PGI. We believe our methods can be beneficial to broader scientific communities that are facing the challenge of complex codes.

Background/Research to Date

In the past several years, our team has established an innovative methodology to analyze complicated scientific code using a functional testing framework [1-2], tracing, profiling and debugging tools [3-5], interactive visualization tools [6], and source code analyzers [7,8]. Our team also has rich experience on compiler design and compiler-based software analysis and OpenMP implementation [9-11]. We can leverage compiler-based analysis capabilities to extract the key information (such as call-graph and inter-procedural summarization to identify key software metrics that impact performance) on one of the leading climate models (Accelerated Climate Modeling for Energy (ACME)). Further information can also be extracted from the compiler, including data dependence analysis, control flow graphs, alias analysis, and dependence and vectorization analysis. Most recently, DOE labs and the NNSA have teamed with NVIDIA to develop open-source Fortran compiler technologies. It is a golden opportunity to reengineer industry-leading compilers to better support complicated scientific code analysis and to facilitate the code redesign for next generation computers. In the past year, we have developed a virtual observation system for ACME land code, using in-situ communication middleware and interactive explorative data analysis framework [12].

Proposed Direction of work

We will take the advantage of current built-in code analysis capability provided by PGI and GNU compilers, develop necessary functionality within the open source Fortran compiler to extract relevant information on the software structure and data dependency. We will use the ACME model as a case study to demonstrate our automatic instrumentation function to capture and track user defined data streams through the whole simulation. We will also develop the appropriate interface between the PGI-based open Fortran compiler and the open instrumentation and measurement infrastructure (score-P) for large-scale scientific code profiling and tracing. We will also develop module extensions to the PGI-based open Fortran compiler so that data streams can connect to in-situ data communication framework for interactive data exploration and analysis. We have also started a collaboration with Appentra [13], which uses the LLVM compiler to detect semantic parallel patterns, which may help to classify different types of parallelism present in the application and build a classification system for common code patterns.

Contributed ideas paper 14

Connections to Math, Comp Sci & and Climate Science

From the computer science perspective, this effort provides a unique way to reengineer/integrate several key components of modern development tools (including debugger, profiler, and compiler) for scientific code redesign and refactoring. It provides an ambitious and feasible method to deal with large scientific codes where the software structures are too complicated for most existing analysis tools. From the climate science perspective, this effort will provide a unique way to generate science-based software unit testing to improve the model's fidelity and reduce the structure uncertainty via (1) enabling direct model-data comparison using the existing future observational system at different scales and (2) engaging field experimentalists and observational data providers.

Potential Impact on the field

Compiler technologies have been the key component to translate human understandable programs (for example scientific models in high level programming languages) into machine understandable programs. Modern compilers have collected huge amounts of information on programming codes in order to better understand the software structure for code optimization and increased performance. Based on our previous efforts, it is a very reliable and feasible approach to reengineer some part of a compiler system (mostly front end) and repurpose the intermediate information of the software system to provide a deeper understanding of existing scientific codes. It also provides an ambitious but robust approach to deal with general scientific software refactoring and redesign challenges. By integrating further with in-situ data communication middleware and interactive data exploratory capability, it will also provide a unique method for large-scale model validation and verification within the scientific functional testing framework.

References

- [1] Dali Wang, Yang Xu, Peter Thornton, Anthony King, Linghong Gu, Chad Steed, Joseph Schuchart, A Functional Testing Platform for the Community Land Model, *Environmental Modeling and Software*, 2014, Volume 55, Pages 25-31, 10.1016/j.envsoft
- [2] Dali Wang, Wei Wu, Tommy Janjusic, Yang Xu, Colleen Iversen, Peter Thornton, Misha Krassovski, Scientific Functional Testing Platform for Environmental Models: An Application to Community Land Model, *International Workshop on Software Engineering for High Performance Computing in Science*, 37th International Conference on Software Engineering, May 16-24, 2015, Florence, Italy. DOI 10.1109/SE4HPCS.2015.10
- [3] Score-P: <http://www.vi-hps.org/projects/score-p/>
- [4] Vampir: <https://www.vampir.eu>
- [5] DDT: <http://www.allinea.com/products/ddt>
- [6] Chad A. Steed, Daniel M. Ricciuto, Galen Shipman, Brian Smith, Peter E. Thornton, Dali Wang, and Dean N. Williams. Big Data Visual Analytics for Exploratory Earth System Simulation Analysis. *Computers & Geosciences*, 61:71–82, 2013. doi:10.1016/j.cageo.2013.07.025
- [7] Oscar Hernandez, Stephen Poole, Jeffery Kuehn, Siddhartha Jana, Swaroop Pophale, Barbara Chapman. The OpenSHMEM Analyzer. *The 6th Conference on Partitioned Global Address Space Programming Models*, Santa Barbara, CA, 2012.
- [8] Oscar Hernandez, Chunhua Liao, Barbara Chapman. Dragon: A Static and Dynamic Tool for OpenMP. *Shared Memory Parallel Programming with Open MP*. Volume 3349 of the series *Lecture Notes in Computer Science* pp 53-66
- [9] Chunhua Liao, Oscar Hernandez, Barbara Chapman, Wenguang Chen, Weimin Zheng, OpenUH: An optimizing, portable OpenMP compiler, *Concurrency and Computation: Practice and Experience*, Vol 19, Issue 18, pp 2317-2332, 2007

Contributed ideas paper 14

- [10] Oscar Hernandez, Chunhua Liao and Barbara Chapman, Dragon: A static and dynamic tool for OpenMP, Shared Memory Parallel Programming with Open MP, pp 53-66, 2005
- [11] Dali Wang, Wei Wu, Frank Winkler, Wei Ding, Oscar Hernandez, A Feasibility Study on Porting the Community Land Model onto Accelerators using OpenACC, International Journal of Advanced Computer Science and Applications (IJACSA), 5(12), 2014.
<http://dx.doi.org/10.14569/IJACSA.2014.051203>
- [12] A Virtual Observation System for ACME Land Model Simulations, in preparation
- [13] Jacobo Loberias, Manuel Arenaz: Oscar Hernandez. Experiences in extending Parallelware to support OpenACC. Second Workshop on Accelerator Programming using Directives (WACCPD2). Co-located with SC'15 (to appear)

Contributed ideas paper 15

Algorithmic Differentiation for Climate Science

11/30/2015

Sri Hari Krishna Narayanan

Argonne National Laboratory

Daniel Goldberg (University of Edinburgh), Patrick Heimbach (University of Texas at Austin),

David Marshall (Oxford University), Stephen Price (LANL), Carl Wunsch (Harvard University)

Abstract

Computing sensitivities of climate-related quantities of interest (QoI) to very high-dimensional state or parameter spaces in an efficient manner is of considerable interest in climate science. Applications in which sensitivities are propagated in time by means of adjoint models include quantifying controlling factors of the QoI's variations [24, 3], optimal state and parameter estimation ("data assimilation") [25, 26], computing nonnormal transient amplification patterns [27], and formal uncertainty quantification [10, 11]. In an environment of rapid model development algorithmic differentiation (AD) plays a crucial role in supporting adjoint model generation for diverse applications based on up-to-date source codes [9]. With applications ranging from oceanography [19], atmospheric sciences [28], coupled carbon cycle modeling [12] and cryospheric sciences [5, 2], the need for sustained AD tool maintenance and improvements becomes a high priority.

Background/Research to Date

Climate models depend on a large (because spatially varying) parameter set that represents subgrid-scale processes and whose values are inherently uncertain. Similarly, initial and boundary conditions (at the surface or base) are uncertain and will dictate the solution. To understand model sensitivities derivatives are needed for relating the outputs of a model back to its inputs, i.e. control variables. Model calibration, in particular, which seeks to determine optimal estimates of control variables informed by available observations, plays an increasingly important role in attempts to initialize models for prediction. The conventional approach of finite-difference approximation becomes intractable for high-dimensional control spaces and simulations which require days for single-model execution on today's supercomputers. Furthermore, this approach can be inaccurate because of accumulating round-off errors.

An alternative method is the use of adjoint models. For scalar-valued QoI's such as least-squares misfit functions used in model calibration or state estimation, or for sensitivity analysis of climate metrics the adjoint provides an extremely efficient means for computing linear sensitivities. In climate modeling, the challenge, then, is to maintain up-to-date versions of the adjoint model given the evolving nature of the climate model at hand. Writing the adjoint by hand has frequently shown to be cumbersome, and producing code that is increasingly out of date with respect to its parent code. It also is prone to additional coding error.

AD provides a powerful alternative [7]. It exploits the code's composition of elementary arithmetic operations and elementary function evaluations whose analytical derivatives are known. Applying the chain rule of differentiation, derivatives can be propagated from one model variable to another. Various AD tools have been developed based on techniques from source transformation [21, 8, 4, 16] or operator-overloading [22, 17]. Current AD research is directed in part at increasing the efficiency of derivative code for components of large-scale climate models. Reverse-mode AD necessarily uses checkpointing where the application state is determined by source analysis. Many checkpointing schemes are being examined to efficiently

Contributed ideas paper 15

use the available memory and disk while computing adjoints. Preliminary research allows the suspension and restart of the forward and the adjoint model. Adjoinable MPI seeks to differentiate through the MPI calls in the model code. AD advances are exploiting the inherent mathematical properties of the model to allow larger adjoint models to be run. For example, through the special treatment of fixed-point iterations, which differs from a straightforward application of AD software, both computational and memory requirements of the adjoint model are reduced.

Proposed Direction of work

We propose to extend the use and improve the efficiency of adjoint codes of climate model components. At the center of the effort is the open-source AD tool OpenAD [21].

- **New Languages** Many existing models are written in Fortran. As models migrate to C, C++ [14, 20], R, and Python, we propose to apply existing source transformation tools such as ADIC [16] for models written in C. Source transformation AD is not possible for C++ because of features such as templates and overloading. Instead, we propose to apply a mixed method [13] that uses operator overloading AD for most of the application and source transformation AD for the computation-intensive portions that are often written in a C-like fashion.

- **Scaling** An important bottleneck for scaling of adjoint computations is the memory requirement for storing partials and checkpointing of data at the granularity of timesteps. We will apply optimal multilevel checkpointing involving disk storage and memory [1, 18].

- **Robustness** The mean time between failure (MTBF) of leadership-class machines is on the order of hours. Because adjoint computations will often exceed MTBF, we propose to support the restart of the adjoint climate application using existing optimal checkpointing techniques in AD [6]. We will study the implementation of the restart mechanism in the context of multilevel checkpointing.

- **Chaotic Systems** Because of the increasingly nonlinear dynamics with time encountered in climate simulations, the adjoint results can diverge exponentially from the “macroscopic climate sensitivity” [15]. By replacing the initial value problem in these models with the well-conditioned “least squares shadowing (LSS) problem”, it may be possible to compute well-behaved derivatives of climate quantities. [23]. We propose to study the non-straightforward application of AD tools for the LSS approach.

- **Usability and Applicability** Many climate modelers have been reluctant to employ AD. We will therefore strive to make source transformation tools easily applicable to climate codes. Moreover, we will work to demonstrate to the climate community the benefits of AD.

Connections to Math, Comp Sci & and Climate Science

The proposed work covers various fields, including, the solution of (discretized) PDEs, high-performance computing (techniques such as checkpointing and resilience) and different components of climate models (e.g. ice sheets, oceans, atmosphere, carbon cycle).

Potential Impact on the field

Climate modelers who already use AD will be able to run adjoint models at larger scales than they can currently run. Adjoint models that previously could not be executed because of machine failures can be executed with checkpointing support. Climate modelers who are hesitant to migrate to new languages will be able to do so more easily. Success stories may encourage other climate researchers to begin to use AD.

Contributed ideas paper 15

References

- [1] G. Aupy, J. Herrmann, P. Hovland, and Y. Robert. *Optimal Multistage Algorithm for Adjoint Computation*. PhD thesis, LIP-ENS Lyon; Argonne National Laboratory, 2014.
- [2] I. Fenty, D. Menemenlis, and H. Zhang. Global coupled sea ice-ocean state estimation. *Clim. Dyn.*, pages 1–26, Aug. 2015.
- [3] I. Fukumori, O. Wang, W. Llovel, I. Fenty, and G. Forget. A near-uniform fluctuation of ocean bottom pressure and sea level across the deep ocean basins of the Arctic Ocean and the Nordic Seas. *Prog. Oceanogr.*, 134(C):152–172, May 2015.
- [4] R. Giering and T. Kaminski. Applying TAF to generate efficient derivative code of Fortran 77-95 programs. In *Proceedings of GAMM 2002, Augsburg, Germany, 2002*.
- [5] D. Goldberg, P. Heimbach, I. Joughin, and B. Smith. Committed near-future retreat of Smith, Pope, and Kohler Glaciers inferred by transient model calibration. *The Cryosphere Discuss.*, 9(4):4459–4498, 2015.
- [6] A. Griewank and A. Walther. Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Trans. Math. Softw.*, 26(1):19–45, 2000.
- [7] A. Griewank and A. Walther. *Evaluating derivatives. Principles and techniques of algorithmic differentiation. 2nd ed.* SIAM., 2008.
- [8] L. Hascoet and V. Pascual. The Tapenade automatic differentiation tool: Principles, model, and specification. *ACM Trans. Math. Softw.*, 39(3):20:1–20:43, May 2013.
- [9] P. Heimbach, C. Hill, and R. Giering. An efficient exact adjoint of the parallel MIT General Circulation Model, generated via automatic differentiation. *Future Generation Computer Systems*, 21(8):1356–1371, Oct. 2005.
- [10] A. G. Kalmikov and P. Heimbach. A Hessian-Based Method for Uncertainty Quantification in Global Ocean State Estimation. *SIAM J. Sci. Comput.*, 36(5):S267–S295, Sept. 2014.
- [11] T. Kaminski, F. Kauker, H. Eicken, and M. Karcher. Exploring the utility of quantitative network design in evaluating Arctic sea ice thickness sampling strategies. *The Cryosphere*, 9(4):1721–1733, 2015.
- [12] T. Kaminski, W. Knorr, G. Schürmann, M. Scholze, P. J. Rayner, S. Zaehle, S. Blessing, W. Dorigo, V. Gayler, R. Giering, N. Gobron, J. P. Grant, M. Heimann, A. Hooker-Stroud, S. HOUWELING, T. Kato, J. Kattge, D. Kelley, S. Kemp, E. N. Koffi, C. Köstler, P. P. Mathieu, B. Pinty, C. H. Reick, C. Rödenbeck, R. Schnur, K. Scipal, C. Sebald, T. Stacke, A. T. van Scheltinga, M. Voßbeck, H. Widmann, and T. Ziehn. The BETHY/JSBACH Carbon Cycle Data Assimilation System: experiences and challenges. *J. Geophys. Res. Biogeosci.*, 118(4):1414–1426, Oct. 2013.
- [13] K. Kulshreshtha, S. Narayanan, and T. Albring. A mixed approach to adjoint computation with algorithmic differentiation. *IFIP Advances in Information and Communication Technology (IFIP AICT Series): System Modeling and Optimization*, Under Submission.
- [14] E. Larour, H. Seroussi, M. Morlighem, and E. Rignot. Continental scale, high order, high spatial resolution, ice sheet modeling using the Ice Sheet System Model (ISSM). *Journal of Geophysical Research: Earth Surface*, 117(F1):n/a–n/a, 2012. F01022.

Contributed ideas paper 15

- [15] D. J. Lea, M. R. Allen, and T. W. Haine. Sensitivity analysis of the climate of a chaotic system. *Tellus A*, 52(5), 2000.
- [16] S. H. K. Narayanan, B. Norris, and B. Winnicka. ADIC2: Development of a component source transformation system for differentiating C and C++ . *Procedia Computer Science*, 1(1):1845–1853, 2010. ICCS 2010.
- [17] E. T. Phipps, R. A. Bartlett, D. M. Gay, and R. J. Hoekstra. Large-scale transient sensitivity analysis of a radiation-damaged bipolar junction transistor via automatic differentiation. In C. H. Bischof, H. M. Bücker, P. D. Hovland, U. Naumann, and J. Utke, editors, *Advances in Automatic Differentiation*, pages 351–362. Springer, 2008.
- [18] M. Schanen, O. Marin, H. Zhang, and M. Anitescu. Asynchronous two-level checkpointing scheme for large-scale adjoints in the spectral-element solver Nek5000. Technical Report ANL/MCS-P5422-1015, Mathematics and Computer Science Division, Argonne National Laboratory, 2015.
- [19] D. Stammer, M. Balmaseda, P. Heimbach, A. Köhl, and A. Weaver. Ocean Data Assimilation in Support of Climate Applications: Status and Perspectives. *Annu. Rev. Marine Sci.*, 8(1):annurev-marine-122414-034113-28, 2016.
- [20] I. K. Tezaur, M. Perego, A. G. Salinger, R. S. Tuminaro, and S. F. Price. Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis. *Geoscientific Model Development*, 8(4):1197–1220, 2015.
- [21] J. Utke, U. Naumann, M. Fagan, et al. OpenAD/F: A modular open-source tool for automatic differentiation of Fortran codes. *ACM Trans. Math. Softw.*, 34(4):18:1–18:36, 2008.
- [22] A. Walther and A. Griewank. Getting started with ADOL-C. In U. Naumann and O. Schenk, editors, *Combinatorial Scientific Computing*. Chapman-Hall, 2012.
- [23] Q. Wang, R. Hu, and P. Blonigan. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210 – 224, 2014.
- [24] C. Wilson, K. J. Horsburgh, J. Williams, J. Flowerdew, and L. Zanna. Tide-surge adjoint modeling: A new technique to understand forecast uncertainty. *J. Geophys. Res. Oceans*, 118(10):5092–5108, Oct. 2013.
- [25] C. Wunsch and P. Heimbach. Practical global oceanic state estimation. *Physica D*, 230(1-2):197–208, June 2007.
- [26] C. Wunsch and P. Heimbach. Dynamically and Kinematically Consistent Global Ocean Circulation and Ice State Estimates. In *Ocean Circulation and Climate: A 21st Century Perspective*, chapter 21, pages 553–579. Elsevier Ltd., 2013.
- [27] L. Zanna, P. Heimbach, A. M. Moore, and E. Tziperman. Upper-ocean singular vectors of the North Atlantic climate with implications for linear predictability and variability. *Q. J. R. Meteorol. Soc.*, 138(663):500–513, Jan. 2012.
- [28] X. Zhang, X.-Y. Huang, and N. Pan. Development of the Upgraded Tangent Linear and Adjoint of the Weather Research and Forecasting (WRF) Model. *J. Atmos. Oceanic Technol.*, 30(6):1180–1188, June 2013.

Contributed ideas paper 16

Domain Specific Languages for Cloud Resolving Climate Models

Dec. 01, 2015

Pedro Diniz, Robert Lucas, Steve Crago (USC Information Sciences Institute), John Shalf, David Donofrio, Michael Wehner, Leonid Oliker, Samuel Williams, Cy Chan, Brian Van Straalen (Lawrence Berkeley Natl. Lab.), Daniel Quinlan (Lawrence Livermore Natl. Lab.)

Abstract

The computational algorithms involved in climate modeling are very complex and computationally demanding, as they involve a wide range of physical phenomena modeling from temperature, wind, land mass and ocean current to salinity of oceans, ionization of atmosphere. The diversity of the algorithms exacerbates the need for high-performance computing (HPC) resources so that accurate estimates can be found in useful time. Modern general purpose languages use an imperative style of programming where code optimization is co-mingled with the mathematical description of the algorithmics, inhibiting code readability and maintainability. Furthermore, these languages overdefine and over-constrain the definition of an algorithm such that code written for a GPU cannot be easily transformed to perform well on a CPU and vice versa. By contrast, Domain Specific Languages (DSLs) offer an approach to raise the level of abstraction for expressing scientific algorithms by enabling algorithms to be expressed in a concise syntax that is more understandable by the domain scientists and is more amenable to automated code optimization (auto-tuning). Specifically, DSLs enable the generation of architecture specific code using greater semantic content of the DSL. This approach provides performance portability across diverse architectures and enables a viable migration path for existing application codes across a diversity of emerging multiprocessors thus enhancing code evolution and protecting future development investments.

Proposed Direction of Work

The challenge of this decade for scientific computing is to create new programming models and tools that enable concise expression of fine-grained, explicit parallelism that is portable across the diverse computer architectures (GPUs and CPUs – a companion paper entitled “Holistic Approach to Cloud Systems Resolving Climate Models” describes opportunities provided by DSL for the customization of hardware and auto-tuners for performance and energy). We believe domain-specific models and languages to play a key challenges to provide performance and portability with higher productivity than general purpose programming models offer. Using DSLs scientists can describe their computational structures (e.g., stencils, linear solvers) exposing a wide range of climate model parameters and validation criteria for convergence or even resilience in the face of failures. DSL’s can also provide a better interface to auto-tuning frameworks to automatically retarget and retune codes for emerging architectures. Auto-tuning systems search over optimization strategies to find the combination of tunable parameters that maximizes computational efficiency for a given algorithmic kernel. The use of DSLs for describing stencil and grid-based calculations provides an input to auto-tuning technology that can optimize the mapping and compilation onto diverse complex massively parallel hardware. This description is thus translated using a selected set of code templates tuned to the specific nature of the target architecture (possibly and opportunistically exploiting heterogeneity) and coupled with auto-tuners to enhance its performance and/or energy. By doing so, we achieve performance portability using an enhanced toolchain outlined in Figures 1 and 2.

Background/Research to Date

As part of our previous work [Wehner2011] in the analysis of the computational requirements of the Colorado State University icosahedral model at global cloud system resolving scales, we applied a co-design approach, including auto-tuning, to designing synergistic hardware and software technologies. Although the auto-tuning strategy has been successfully applied to libraries, there are still many opportunities for a high impact on the future development of codes for the climate community as we have recently demonstrated in specific case studies, namely:

- We have introduced a generalized stencil auto-tuning framework that takes a straightforward Fortran expression of a stencil kernel and automatically generates tuned implementations of the kernel in C or Fortran to achieve performance portability across diverse computer architectures shown in Figure 3 ([Kamil2009]).

Contributed ideas paper 16

- SEJITS [SEJITS2010] is another approach to packaging DSL's in a conventional language framework by allowing specializers to identify pieces of code for domain-specific parallelization and dynamically rewriting the code to create the parallel implementation.
- Embedded Domain-Specific Languages define a DSL using an existing base-language (perhaps a general-purpose languages; e.g. C, Fortran, C++, Scala) and represent a compromise in the development of DSL. While embedded DSLs limit the design of new syntax (avoid drastic changes from the base-language) they can reuse significant or even the whole base-language compiler infrastructure. We have experimented with this approach in the BISICLES Structured Adaptive Ice Sheet model as depicted in Figure 4.

We have demonstrated the flexibility and generality of a domain-specific auto-tuner framework we developed for several stencil computations with a variety of different computational characteristics, arising from PDE solvers in climate science, image processing, and structural mechanics. Other researches in the climate-modeling community have also develop DSLs to focus on specific kernels of their computations (e.g., [Linardakis2013]). These focused efforts have shown performance of the generated codes to be very close to hand-written codes with the added value of programmer and performance portability.

Connections to Math, Comp Science & and Climate Science

Computer Science: Autotuning technology that has been developed by the CS community requires a path for input. DSLs provide more degrees of freedom for auto-tuning by expressing higher level semantics and preserving domain semantics that would otherwise be destroyed by the compilation.

Applied Mathematics: Defining compact and expressive syntax for describing algorithms requires direct involvement of applied mathematicians and domain scientists. Efficient expression of domain-specific mathematical constructs can improve code readability and long term code maintainability.

Applications/Climate: We have a long track record of working closely with climate scientists (Michael Wehner and with the CSU team as described above) and prototypes have proven very effective. Still, there is a long road ahead to make durable/production-quality approaches for implementing main-line climate codes a reality.

Potential Impact on the Field

This approach combined with automatic code tuners, to automatically map software algorithms to specialized architectures, can lead to orders of magnitude reduction in design time and cost giving us an efficient means to develop a power-efficient HPC system tailored specifically to climate simulation. In particular;

- It enables more efficient expression of new algorithms using a domain-specific syntax.
- Spatial structure and data dependency information made available in a DSL can also be used to optimize data layout across a large machine where data movement becomes the limiting factor. Communication bandwidth and latency can be improved by using approaches such as space-filling curves, graph partitioning, and greedy mapping heuristics.
- Language features that express dependencies between loop iterations could become part of the language, allowing autotuners to explore loop parallelization and vectorization strategies.
- Dependency information can also be used to explore fusion optimizations that compilers cannot currently do automatically.
- Automatic array data layout optimization may require more extensive changes than a compiler and library could achieve by itself. For example, an autotuner can explore tiling sizes, tiling hierarchy, and SoA vs. AoS type decisions.
- A DSL that allows the data dependencies to be extracted at the granularity of function calls and/or loop nests would allow a task DAG to be generated, allowing a runtime to spawn off independent pieces of work to take advantage of another level of parallelism.

The use of domain-specific-languages and target-specific code generations and auto-tuners will enhance the performance and programmer portability of these challenging climate codes on leadership class systems. Overall, the pursued approach will lead, we strongly believe, to a reduction of the sources of inaccuracy and uncertainty in today's climate models to enable policy decisions.

Contributed ideas paper 16

References

- [Linardakis2013] R. Torres, L. Linardakis, J. Kunkel and T. Ludwig, "ICON DSL: A Domain Specific Language for Climate Modeling", in Proc. Of the Third Intl. Workshop on Domain-Specific Languages and High-Level Frameworks for High-Performance Computing, Nov. 18, 2013.
- [Kamil2009] S. Kamil, C. Chan, S. Williams, L. Oliker, J. Shalf, M. Howison, E. Wes Bethel, "A Generalized Framework for Auto-tuning Stencil Computations ", In Proc. of the Cray User Group Conference (research sponsored by the DoE/ASCR Office of Science under contract number DE- AC02-05CH11231).
- [Wehner2011] Wehner et. al., "Hardware/Software Co-design of Global Cloud System Resolving Models", *Journal of Advances in Modeling Earth Systems* Vol3, 2011.
- [SEJITS2010] Bryan Catanzaro, Shoaib Ashraf Kamil, Yunsup Lee, Krste Asanović, James Demmel, Kurt Keutzer, John Shalf, Katherine A. Yelick and Armando Fox, "SEJITS: Getting Productivity and Performance With Selective Embedded JIT Specialization", EECS Department University of California, Berkeley Technical Report No. UCB/EECS-2010-23 March 1, 2010 (<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-23.pdf>)

Figures

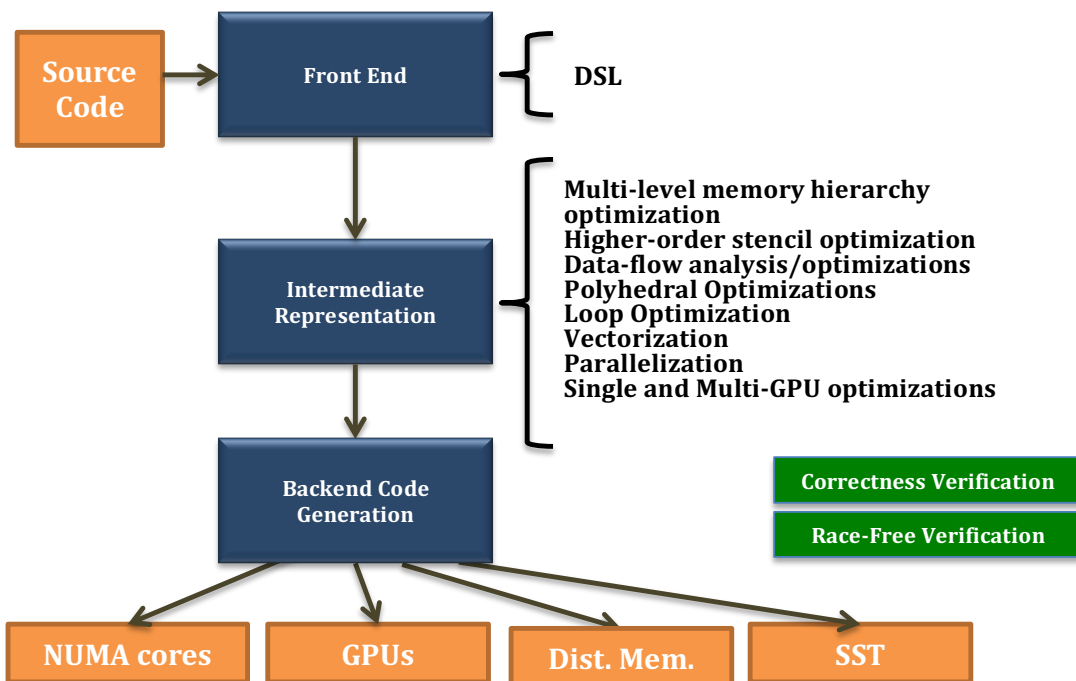


Fig. 1: Use of DSL to provide optimizations and the custom code generation capabilities to different specific types of computer architectures.




Fig. 2: Stencil auto-tuning framework flow. Readable domain-specific code is parsed into an abstract representation, transformations are applied, code is generated using specific target backends, and the optimal auto-tuned implementation is determined via search.

Contributed ideas paper 16

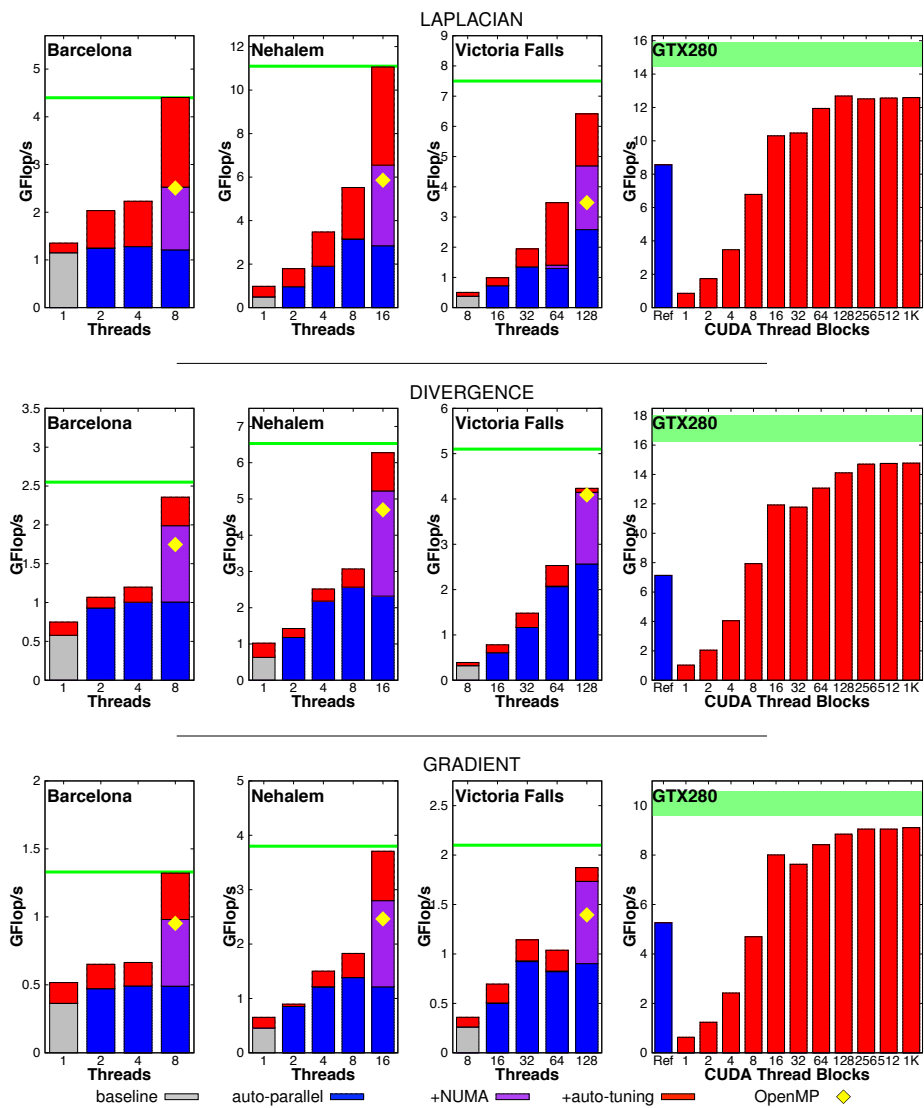


Fig. 3: Laplacian (top row), Divergence (middle row), and Gradient (bottom row) performance as a function of auto-parallelization and auto-tuning — on the four evaluated platforms. Note: the green region marks the range in performance extrapolated from Stream bandwidth. For comparison, the yellow diamond shows performance achieved using the original stencil kernel with OpenMP pragmas and NUMA-aware initialization.

Contributed ideas paper 16

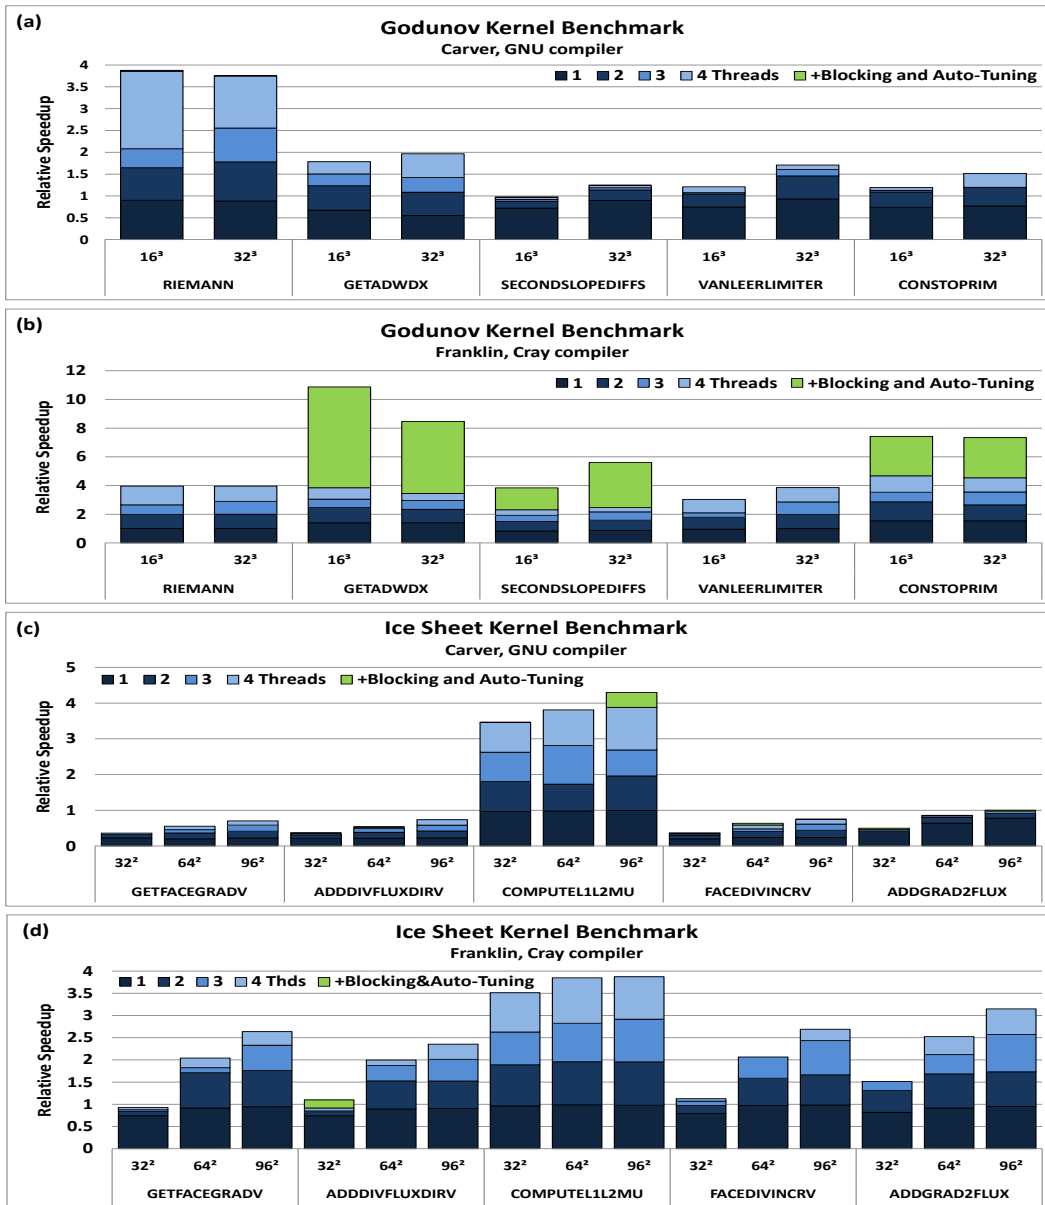


Fig. 4 : Kernel-only benchmarks of the top 5 kernels of the hyperbolic gas dynamics benchmark using one NUMA domain on both Carver and Franklin, for (a-b) Godunov and (c-d) Ice Sheet. Performance is shown relative to the original (unthreaded) version.

Contributed ideas paper 17

Calibration and Comparison of Climate Models: Accounting for Structural and Discretization Error

11/30/2015

Bert Debuschere (SNL, bjdebus@sandia.gov),
Habib Najm (SNL), Khachik Sargsyan (SNL), Kenny Chowdhary (SNL),
Don Lucas (LLNL), Vera Bulaevskaya (LLNL),
Yun Qian (PNNL), Steve Ghan (PNNL),
Daniele Rosa (LBNL), Bill Collins (LBNL)

Abstract

This white paper outlines the need to account for structural and discretization error in climate simulations to improve their predictive fidelity. We propose to apply novel approaches from the uncertainty quantification community for characterizing structural and discretization error, and identify advances in math, computer science, and climate science to make this feasible. The challenges are formidable, but the impact on climate modeling would be enormous by enabling robust tuning of climate models that accounts for all sources of uncertainty and error.

Background/Research to Date

Climate models are complex as they are made up of many discrete formulations of resolved or under-resolved physical processes pertaining to the atmosphere, the land and its ecosystems, the ocean, and other systems. State of the art models generally have hundreds of parameters that need to be tuned for best results [Qian et al., 2015]. Historically, tuning has typically been done by hand, and focused on individual processes [Mauritsen et al., 2012]. Formal methods for model calibration have been developed in the Statistics and Uncertainty Quantification (UQ) communities, often relying on Bayesian concepts, e.g. [Kennedy and O'Hagan, 2001; Yang et al., 2012]. The application of these concepts to coupled climate system models is still in its infancy, however, and is limited by the high computational cost of climate simulations. Moreover, while parameter tuning generally improves the agreement with available observations, the predictive skill of the climate model is often limited by structural error.

To improve the predictive fidelity of climate models, it is essential to take into account both the effect of structural error as well as spatial and temporal discretization errors. Both types of error can be significant, have historically not been accounted for, and can be tightly linked as process parameterizations are strongly influenced by the time and length scales that can be resolved.

Proposed Direction of work

Novel UQ methods are able to account for structural error by embedding statistical model error terms in model parameters, such that the resulting calibrated uncertain model prediction is representative of the impact of structural uncertainty in the model. This approach has been demonstrated, e.g. in the calibration of chemical models [Sargsyan *et al.*, 2015] and tends to be more accurate than additive model error terms, especially for assessing the uncertainty in Quantities of Interest (QoIs) for which no observational data is available. Applying these methods to climate models will require very close collaboration between climate scientists and UQ practitioners to determine the best place(s) to embed model error terms. In this context, it is known that process level (e.g. clouds, convections) calibration is important for model fidelity and it may be beneficial to calibrate the models by constraining the behavior of physical processes at a regional scale. Methods that account for discretization error as a source of uncertainty, and incorporate it accordingly in making uncertain predictions are also emerging in UQ [Conrad *et al.*, 2015]. These methods are still in their infancy, particularly as regards to large-scale computational models, but offer a promising framework. They will require significant development, indeed, to be evaluated and applied to climate models.

Another issue is the high computational cost of climate simulations, which renders statistical

Contributed ideas paper 17

calibration prohibitively expensive. A key approach for mitigating this computational cost is the use of surrogate models that are cheaper to evaluate than the full model [Marzouk *et al.*, 2007, 2009]. Commonly used approaches for surrogate construction are Galerkin projection and regression, often combined with dimensionality reduction. Multifidelity UQ methods, which use ensembles of runs with varying fidelity (*e.g.* low, medium, and high resolution runs), are also useful here as they can deliver significant computational savings by reducing the requisite number of high-fidelity runs, for equivalent accuracy [Ng and Eldred, 2012]. The use of adaptive local surrogates and advanced Markov Chain Monte Carlo methods in calibration can also provide significant savings [Conrad *et al.*, 2015; Cui *et al.*, 2014, 2015].

Cost savings can further be achieved through the use of representative ensembles of shorter simulations to replace the traditional serial-in-time long-term climatology simulations. Wan *et al.* (2015) found this ensemble method can reduce the total computational time by a factor of about 15, and the turnaround time by a factor of hundreds. Another approach uses nudging, such as constraining the simulated wind and temperature fields to reduce the effect of natural variabilities of the large-scale circulation [Zhang *et al.*, 2014]. The risk for this strategy though is that the forcing introduced by nudging can change the basic characteristics of the climate model.

Another key piece in the research to be conducted is the choice of metrics to assess the agreement with observational data. In the presence of structural error, it is not uncommon for a model to do well for one QoI but to perform much worse for metrics that involve other QoIs. One of the research aims will be to investigate how the proper characterization of structural error can reduce such over-fitting with respect to one metric, and allow the consistent use of multiple metrics in the likelihood functions to take advantage of observational data for any of the model QoIs.

Connections to Math, Computer Science and Climate Science

The research directions outlined above are challenging, and will require a multidisciplinary approach for their combined application to a problem as complex and large-scale as climate modeling. In terms of mathematics, advances are needed not just in the ways to represent structural and discretization error; but also in approaches for inferring fields of physics parameters; multifidelity calibration; incorporating data sets with incomplete specification of their uncertainties; optimization, MCMC, Approximate Bayesian Computation, and Density Estimation; dimensionality reduction; optimal experimental design, and others. In terms of computer science, more resilient and efficient approaches for making effective use of extreme scale computing architectures are needed, both for individual runs and large-scale ensembles, including the handling and analyzing of large data volumes. On the climate science side, there is a need for automating the handling of observational data, including proper characterization of their uncertainty; the construction of metrics to evaluate the predictive fidelity of climate models; and the analysis of which model components are most susceptible to structural error.

Potential Impact on the field

The proposed research objectives in this white paper are extremely challenging and the advances needed in math, computer science, and climate science to achieve these objectives are daunting. Some of the proposed advances may easily be 10 years away. However the impact of realizing these objectives, and even the steps on the way to these objectives, on the field of climate modeling would be enormous. Being able to account for structural and discretization error in calibration would allow for automated and robust tuning of climate models in a way that takes advantage of all available observational data in a consistent way, is not disrupted by changes in grid or temporal resolution, can readily handle updates in model components, and is able to readily incorporate new data as it becomes available. Such a framework would be invaluable for determining where models need to be improved and where more data would be the most beneficial for enhancing the overall predictive fidelity of the climate model. Ultimately, this work will enable more rigorously vetted information for analysts and decision makers.

Contributed ideas paper 17

References

- Mauritsen, T., et al. (2012), Tuning the climate of a global model, *J. Adv. Model. Earth Syst.*, 4, M00A01, doi:10.1029/2012MS000154.
- Sargsyan, K., Najm, H. N., and Ghanem, R. (2015). On the Statistical Calibration of Physical Models. *International Journal of Chemical Kinetics*, 47(4), 246–276.
<http://doi.org/10.1002/kin.20906>
- Kennedy, M.C. and O’Hagan, A. (2001). Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63(3): 425–464.
- Conrad, P. R., Girolami, M., Särkkä, S., Stuart, A., & Zygalakis, K. (2015). Probability Measures for Numerical Solutions of Differential Equations. *Annals of Applied Probability*, 23(4), 1437–1468. Retrieved from <http://arxiv.org/abs/1506.04592>
- Zhang K, H Wan, X Liu, SJ Ghan, GJ Kooperman, PL Ma, PJ Rasch, D Neubauer, and U Lohmann. 2014. "Technical Note: On the Use of Nudging for Aerosol-Climate Model Intercomparison Studies." *Atmospheric Chemistry and Physics* 14(16):8631-8645. doi:10.5194/acp-14-8631-2014.
- Wan H, PJ Rasch, K Zhang, Y Qian, H Yan, and C Zhao. 2014. "Short ensembles: An Efficient Method for Discerning Climate-relevant Sensitivities in Atmospheric General Circulation Models." *Geoscientific Model Development* 7(5):1961-1977. doi:10.5194/gmd-7-1961-2014.
- Qian Y, H Yan, Z Hou, G Johannesson, SA Klein, D Lucas, R Neale, PJ Rasch, LP Swiler, J Tannahill, H Wang, M Wang, and C Zhao. 2015. "Parametric Sensitivity Analysis of Precipitation at Global and Local Scales in the Community Atmosphere Model CAM5." *Journal of Advances in Modeling Earth Systems* 7(2):382-411. doi:10.1002/2014MS000354.
- Yang B, Y Qian, G Lin, LYR Leung, PJ Rasch, GJ Zhang, SA McFarlane, C Zhao, Y Zhang, H Wang, M Wang, and X Liu. 2013. "Uncertainty Quantification and Parameter Tuning in the CAM5 Zhang-McFarlane Convection Scheme and Impact of Improved Convection on the Global Circulation and Climate." *Journal of Geophysical Research. D. (Atmospheres)* 118(2):395-415. doi:10.1029/2012JD018213.
- Marzouk, Y.M., Najm, H.N., Rahn, L.A. (2007). Stochastic spectral methods for efficient Bayesian solution of inverse problems. *Journal of Computational Physics* **224**: 560–586.
- Marzouk, Y.M., Najm, H.N. (2009). Dimensionality reduction and polynomial chaos acceleration of Bayesian inference in inverse problems. *Journal of Computational Physics* **228**: 1862–1902.
- Cui, T., Martin, J., Marzouk, Y.M., Solonen, A., Spantini, A. (2014). Likelihood-informed dimension reduction for nonlinear inverse problems. *Inverse Problems* **30**: 114015.
- Cui, T., Law, K.J.H., and Marzouk, Y.M., (2015). Dimension-independent likelihood-informed MCMC, arXiv:1411.3688.

Contributed ideas paper 17

Conrad, P., Marzouk, Y.M., Pillai, N., Smith, A. (2015). Accelerating asymptotically exact MCMC for computationally intensive models via local approximations. *Journal of the American Statistical Association*, in press. arXiv:1402.1694.

Ng, L.W.T. and Eldred, M.S. (2012). Multifidelity uncertainty quantification using nonintrusive polynomial chaos and stochastic collocation. *In Proc. of the 14th AIAA Non-Deterministic Approaches Conference*, AIAA-2012-1852, Honolulu, HI, April 23-26.

Contributed ideas paper 18

Accounting for model error in climate simulations

December 1, 2015

Emil Constantinescu

Mathematics and Computer Science (MCS), Argonne National Laboratory

Mihai Anitescu (MCS), and Rao Kotamarthi (Environmental Sciences)

Abstract

Computational climate simulations are inherently stochastic, with a large source of uncertainty coming from the models themselves. In this short paper we discuss a novel strategy for estimating the structure of the model error, using it within model calibration in a computationally tractable way, and the means of projecting it in decadal simulations. We anticipate that this strategy will significantly improve precipitation projections and cloud processes in climate simulations.

Context

Climate projections are typically carried out by running perturbed deterministic dynamic simulations. An expectation operator over space and time is applied to extract the quantities of interest (QoI) and to measure uncertainty. Because of the climate time scales even relatively small modeling errors can have a significant impact on QoI. These modeling errors arise from unresolved scales, uncertainty in the parameterization of subgrid processes, and numerical solver errors. Stochastic methods have been proposed to alleviate such effects and improve projections and the uncertainty. Ensembles of simulations obtained by using different numerical models, physics, or forcings and stochastic parametrizations are used as proxy for projected probability distributions. However, what distribution the ensembles really represent remains unclear, and often leads to overdispersive results. Moreover, the same data is used to calibrate the physics modules and validate the model simulations, or the physics calibration takes place in a reductionist framework that does not connect well with the rest of the model. Conversely, different ESMs tend to share the same types of physics approximations, leading to overconfidence in the variability across models.

We hypothesize that the modeling error can be derived from model simulations and data via data assimilation sufficiently well to characterize model and system variability. The model error becomes a conditional process with respect to climate simulations and, we believe, will improve significantly the quality of the climate projections.

Model Error Estimation and Projection

Here we briefly sketch the ideas behind developing a stochastic model error framework based on measurements and simulations. These ideas are motivated by our recent work in weather forecasting [1]. This concept is related to model output statistics, where the ensemble prediction is nudged by observations [2]. We propose a stochastic model of the errors conditional on the model output. The key idea here is to avoid inferring the elements that the numerical model represents well, and therefore simplifying the stochastic error process representation and fidelity, as well as delineating the meaning of the distribution.

The model error is developed from the joint distributions of the measurements and model output, and then computing the conditional of the former on the latter. This gives us an estimate of the modeling error and can be used directly in a data assimilation framework aimed at improving the climate model. We use this conditional and climate forecasts to assess the discrepancy and uncertainty when making climate projections. For instance, in [1] we use a space-time Gaussian process to represent modeling errors in wind speed forecasts. In Fig. 1 we show the inferred and the true error (model - reality) as well as confidence intervals.

Computationally Tractable Calibration with Model Error

A major challenge in the case of data assimilation in the presence of model error is the memory usage. Indeed, since the model is imperfect, it can no longer be used to constrain the observations, and every state becomes a free variable in the optimization of the 4DVar functional. Since model error particularly matters in high-resolution cases, where the observations may contradict the model, classical optimization methods are not applicable on emerging architectures that tend to be memory starved. To address these issues, we have developed a new approach that uses the optimality conditions themselves to constrain the state and then estimates the state at each iteration by means of a recursion [3]. While we have demonstrated the memory savings benefits of the approach, it needs to be extended in multiple directions. In particular, recursive procedures are subject to instability. To trade off stability and memory we will investigate a checkpointing-multiple shooting approach as well as preconditioning approaches based on multilevel principles in order to alleviate difficulties converging when used collectively in conjunction with limited memory quasi Newton methods.




Figure 1: Error in the wind field prediction by numerical weather prediction model in January 2012 over a region in the Great Lakes [1].

Potential Field Impact

Even though climate model spatial resolution is increasing rapidly with models having targeted a grid size of 25 km, critical atmospheric processes will remain subgrid and unresolved for the near future. One such process is the representation of clouds in GCMs. Clouds and cloud-related processes are highly parameterized and often tuned to obtain the Earth's radiative balance and global mean temperature. Formation of clouds, their lifecycle, and precipitation are subgrid-scale processes poorly represented in current-generation models and those being developed. As a result, the parameterization of clouds constitutes the single-largest uncertainty in terms of radiative forcings and climate feedbacks in GCMs. Methods for improving the representation of cloud processes and especially interactions with aerosols in high-resolution climate models are definitely a priority. It is not surprising that precipitation patterns are poorly represented in ESMs. The strategy proposed here can be adapted to address directly the precipitation precursors by using high-quality satellite observations.

References

- [1] J. Bessac, E. Constantinescu, and M. Anitescu, Stochastic simulation of predictive space-time scenarios of wind speed using observations and physical models, *Annals of Applied Statistics*, submitted, 2015; <http://arxiv.org/abs/1511.09416>.
- [2] Gneiting, T., Raftery, A. E., Westveld III, A. H., and Goldman, T.: Calibrated probabilistic forecasting using ensemble model output statistics and minimum CRPS estimation, *Monthly Weather Review*, 133, 1098-1118, 2005.
- [3] M. Anitescu, X. Zeng, and E. Constantinescu, A low-memory approach for best-state estimation of hidden Markov models with model error, *SIAM Journal on Numerical Analysis (SIUM)*, 52 (2013), pp. 468–495.

Contributed ideas paper 19

A Holistic Approach to Cloud System Resolving Climate Models December 1, 2015
John Shalf, David Donofrio, Leonid Oliker, Michael Wehner (Berkeley National Laboratory)
Pedro Diniz, Robert Lucas (USC/ISI), Celal Konor, Ross Heikes, David Randall (CSU)

Abstract

Accurately modeling climate change is one of the most important challenges facing computational scientists today, with economic ramifications in the trillions of dollars. While the climate system is well understood, current models still have deficiencies. Specifically, clouds and their feedbacks remain a critical weakness in determining the sensitivity of the climate system to increases in carbon dioxide. Finer model resolutions are needed to quantify local impacts and effectively performing these calculations is predicated on our ability to run high resolution models quickly - the ultimate target being simulations that are capable of resolving climate processes down to the kilometer scale running significantly more than 1000x times faster than realtime. Achieving this target will require a coordinated effort of Climate Scientists, Computer Scientists, and Computer Engineers working together to develop new high-resolution modeling algorithms and computing platforms. However, this target cannot be achieved with conventional roadmaps for leading edge HPC systems. A more holistic approach must be adopted to meet this science target.

Over the past two decades, leading parallel computer systems have been architected for throughput and weak scaling (the traditional path forward), but global cloud system resolving climate models will push requirements very deeply in the strong-scaling regime. To meet these needs will require system design point that focuses more strongly on reducing communication overheads and latencies and the requirements of extreme strong scaling. This can only be accomplished through analysis and projection of climate code requirements to drive specialization of hardware and software environments to underpin the development of higher resolution models.

Background/Research to Date

We have previously analyzed the icosahedral model developed by Professor David Randall's group at Colorado State University to understand the computational and algorithmic requirements for cloud resolving models, and have applied a co-design approach to designing synergistic hardware and software technologies to meet the model's requirements. We have recently shown [Wehner2011] demonstrated how the high computational burden of global cloud system resolving models could be met today using this highly integrated co-design approach. We further demonstrated a prototype implementation of such hardware to meet a performance target of several simulated years per wall-clock day on an existing developmental version of a global cloud system resolving model that would be unachievable using the largest systems available today [Donofrio2009]. The bottom line was that this approach can reach the target performance of 1000x real time simulation at a resolution of 2km by achieving the required ~30PFlops sustained and 2PB total memory (Fig 1). A strawman computer designed to exploit the scalability of the model to 20 million processing cores using a nested domain decomposition was found to be practical today using existing hardware-software co-design and a System-on-Chip (SoC) design approach that uses available commodity "off-the-shelf: (COTS) circuit technologies.

This strategy could be enabled by adopting the practices of the consumer electronics industry (i.e. smartphones, etc.) and reducing wasteful use of transistors through a System-on-Chip design strategy. By leveraging the same enormous commodity IP market for design tools, processors, memory controllers, etc. that has enabled the rapid advances in performance in many consumer technologies it is possible to rapidly, and at reasonable cost, create custom processing

Contributed ideas paper 19

elements tailored to the needs of the scientific community. More recently, experts in semicustom System-on-Chip design representing 5 US agencies, academia, and industry (AMD, Intel, Qualcomm, and ARM) came together (www.socforhpc.org) to discuss the technological and economic challenges of building a semi-custom machine for targeted scientific applications such as climate modeling and laid out a feasible technological path for constructing such a system. This approach of specialization and hardware/software co-design previously enabled semi-custom systems such as the Anton molecular dynamics supercomputers to out-perform HPC systems of all scales by factors of 100x to 1000x for the regime of problems they are designed for. Similarly, if we are to achieve the performance requirements for extreme climate simulation, the only viable approach is to adopt a similar path of specialization.

In contrast to the Anton systems our approach is more affordable and more flexible to the needs of algorithmic changes because we focus on the ability to reuse primarily off-the-shelf circuit designs arranged in a unique and novel ways that do not fall on any current HPC roadmaps to meet our performance targets while *maintaining programmability*. We have demonstrated that it is feasible to address the key requirements of extreme strong scaling through the addition of focused, high-value architectural features. One example of these features is to confront the high communication cost found in many general-purpose cores through efficient direct messaging between processor cores that can reduce the cost of halo exchange from 100s of cycles to < 10 cycles. Similarly, direct support for global address space communication models within each core enables nearly zero overhead exchange of information between cores on-chip and off-chip. Our approach to system specialization also includes the use of domain-specific languages (DSLs) with the goal of making it easier to write complex PDE solver codes, such as the climate code. DSLs provide input to auto-tuners to computer driven automated search to retune for diverse hardware targets. This enables programmers to write loops in styles that are intuitive and once auto-tuned are optimized for consumption by the processor's compiler for construction of the most efficient executable code. A companion whitepaper ("Domain Specific Languages for Cloud Resolving Climate Models") presents the opportunity provided by DSLs and autotuners.

Proposed Direction of work

This approach of specialized hardware combined with code auto-tuners, to automatically map software algorithms to specialized architectures, can lead to orders of magnitude reduction in design time and cost-effective means to develop a power-efficient HPC system tailored specifically to climate simulation. Such an approach could provide climate scientists, within a small number of years, a tool that would radically improve the accuracy of their results and play a central role in policy decisions involving global climate change.

Potential Impact on the field

This approach of high-resolution climate modeling supported on a tailored computing platform is novel and to our knowledge has not been proposed elsewhere within the earth sciences community, and could provide a new avenue to impacting emerging computer architectures for climate. A machine capable of resolving climate models down to the 1-km scale would be a transformational change for the climate modeling community by addressing the largest deficiency in contemporary climate models enabling better policy decisions and has the potential to support the modeling needs of the coming generation of climate scientists. Looking beyond climate, our methodology can be applied in similar ways to multiple areas to solve otherwise intractable problems in the DOE Office of Science that cannot be addressed by conventional roadmaps.

Contributed ideas paper 19

Figures

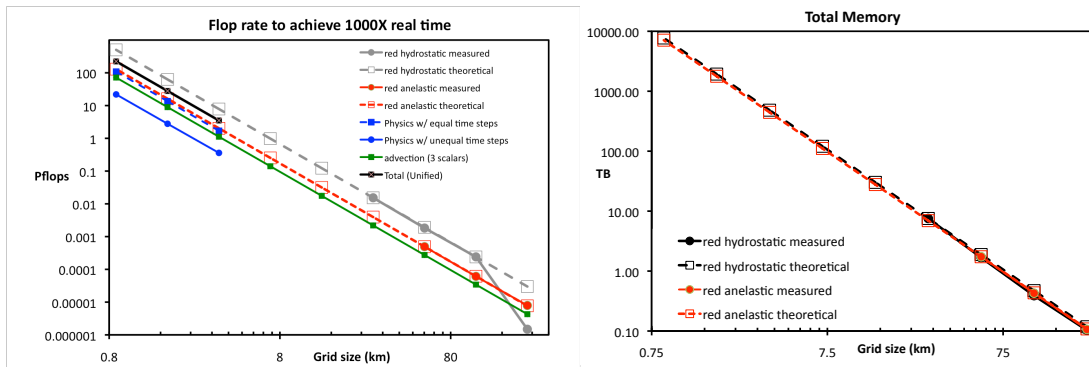


Figure 1: Projected memory scaling and performance scaling requirements for cloud resolving climate models (excerpted from [Wehner2011])

Bibliography

- [Wehner2011] Wehner et. al., "Hardware/Software Co-design of Global Cloud System Resolving Models", *Journal of Advances in Modeling Earth Systems* Vol3, 2011.
- [Donofrio2009] D. Donofrio, L. Olikar, J. Shalf, M. Wehner, C. Rowen, J. Krueger, S. Kamil and M. Mohiyuddin, "Energy-efficient Computer for Extreme-Scale Science", *IEEE Computer*, 42(11), 2009.

Contributed ideas paper 20

Improving predictive capability of land surface models through robust statistical uncertainty quantification techniques

Author: Jaideep Ray, Sandia National Laboratories, Livermore CA

Co-authors: L. Swiler (SNL), G. S. H. Pau, G. Bisht (LBL), F. M. Hoffman (ORNL), M. Huang, Z. Hou and X. Chen (PNNL)

Date: May 24, 2016

Abstract

We advocate for new scalable tools and a mature software base for the robust calibration of climate models. New technologies are required to enable predictive climate simulations with estimates of structural errors in high-resolution, coupled, Earth system models or their individual, computationally demanding, components. Approaches for achieving this goal would extend statistical calibration technologies, develop validated surrogate models for use in many-query applications e.g., statistical inference, global sensitivity analysis, etc., and explore distributed machine learning algorithms (“big data analytics”) frameworks for learning models from voluminous data generated using high fidelity model simulations and field observations. Such tools would yield accurate model parameters and enable calibration and quantification of predictive margins on exascale platforms. These techniques should be demonstrated initially with land surface models for which encouraging preliminary results are available.

Background

Land surface models (LSMs) run at low resolutions serve as components of Earth system models; at high resolutions, they can be used for decision-making at regional and site-scales. This dual-use scenario requires estimation of separate parameter sets from observational data. However, limited data and structural errors in LSMs do not allow parameter estimation with a great deal of certainty. Bayesian calibration, which infers parameters as probability density functions (to quantify uncertainty in the estimates) is an attractive approach but is challenged by the computational cost of invoking the LSM $O(10^4) - O(10^6)$ times. Next-generation *multiscale* LSMs, currently being developed under BER funding, will require joint consideration of multiscale observational data, further increasing calibration difficulties.

The problem of computational cost for Bayesian calibration can be addressed in a number of ways. One can replace the LSM by an emulator (a statistical curve-fit). While common in climate modeling [1-4], emulators have rarely been used for land surface quantities at observational time- and spatial scales [5-7]. In [8], emulators of the Community Land Model (CLM; the codebase for the LSM being developed by the Accelerated Climate Modeling for Energy project) were constructed. They were used to estimate hydrologic parameters and CLM’s structural error, conditional on field observations of latent heat fluxes. However, the method requires the construction of accurate emulators, which may not always be possible [8,9]. Reducing the structural error, e.g. by coupling CLM with a high fidelity, if computationally expensive, model for subsurface processes, can ease estimation of model inputs such as unknown permeability and contaminant fields. Recent work has used PFLOTRAN as the subsurface reactive transport model and addressed the computational cost by using ensemble Kalman filters (EnKF) [10]. Machine learning (ML) methods have also been used in the analysis of observational and simulation data, and for climate projection of mean quantities (e.g., [11]), but their use in the construction of emulators is less common [12,13]. Prediction of numerous LSM variables that are correlated and of varied scales will require new approaches that take the physical relations between variables into consideration.

Ongoing research has four limitations. First, emulators are constructed by processing (reading) outputs of large ensembles of computational simulations. The datasets are too big for processing entirely in memory;

Contributed ideas paper 20

further HPC file systems are not optimized for random reads. Secondly, there are no set methods to select which emulator type should be fitted to a given dataset. The third limitation, in emulator-based Bayesian calibration, arises from uncertainty in construction of an accurate emulator; a robust Bayesian alternative that uses the LSM itself does not exist. Finally, EnKFs use Gaussian assumptions for the states being inferred, posing modeling challenges for non-stationary fields.

Proposed Direction of Work

Advances could take a three-pronged approach to address the limitations listed above. Exploration of multi-chain (i.e., multi-processor) Bayesian calibration techniques, constructed by melding genetic algorithms and multi-chain Markov chain Monte Carlo (MCMC) methods, could address calibration using the ACME Land Model (ALM), rather than an emulator. The method might obtain its scalable characteristics using MPI-2's one-sided, remote memory access to communicate between loosely coupled chains. Computational efficiency on each processor would be achieved by performing a pre-screening of proposed parameter combinations using emulators, which need not be very accurate.

The second prong addresses the construction of surrogate models. A library of types of emulators could be fitted to ALM simulations, and the process of fitting and selecting between them should be automated. To do so, one could draw upon model selection methods from ML that penalize overly complex models. Distributed ML frameworks that have successfully overcome the problem of large reads might be employed to perform the model fitting. Needed is an effort to investigate whether available ML algorithms are sufficient for fitting of emulators or if novel implementations are necessary.

Finally, existing EnKF frameworks, e.g. DART and OpenDA, should be investigated to examine their suitability, and the Gaussian assumptions underlying them, for the estimation of LSM inputs. Challenges lie in the development of accurate prior covariance models for non-stationary fields, and numerical optimizations e.g., covariance localizations, that are necessary for accuracy and stability of EnKFs. Such an effort would inform us whether more sophisticated filters, e.g., particle filters, are required for the Bayesian inference of climate models. We envision these tools being incorporated into the International Land Model Benchmarking (ILAMB) metrics and diagnostics package and utilizing the observational data sets in ILAMB for model calibration.

Connections to Math, Computer Science and Climate Science

These research directions tie to existing BER- and ASCR-funded efforts. An existing prototype multi-chain MCMC implementation, which uses ALM (and not an emulator) for calibration, could be extended to include emulator-based parameter pre-screening. Use of big-data analytic frameworks, e.g. Spark and H2O, to compensate for shortcomings of HPC file systems and making emulators (from simulation data) is novel. Further, statistical modeling of non-stationary fields (for use in EnKFs) would build on ongoing ASCR efforts in random field models and uncertainty quantification. Development of these tools would culminate in statistical inference and data assimilation strategies for LSM parameters with enhanced predictive skill, and deliver the capabilities alongside existing model assessment tools in ILAMB.

Potential Impact on the Field

These tools would aid climate science by allowing estimation of predictive ALM parameters with quantified uncertainty. These data would enable ensemble predictions, which can be used to compute margins and risks for future climate predictions in a data-driven manner. New insights into optimal coupling of models of different computational costs will result. Moreover, the automated emulator-construction and Bayesian calibration packages can be used with other climate component models.

Contributed ideas paper 20

References

- [1] C. Tebaldi and J Arblaster, “Pattern scaling: Its strengths and limitations, and an update on the latest model simulations”, *Climatic Change*. 122:459-71, 2014. doi: 10.1007/s10584-013-1032-9.
- [2] P. B. Holden and N. R. Edwards, “Dimensionally reduced emulation of an AOGCM for application to integrated assessment modelling”, *Geophysical Research Letters*. 37, 2010. doi: 10.1029/2010GL045137.
- [3] P. B. Holden et al, “PLASIM-ENTSem v1.0: a spatio-temporal emulator of future climate change for impacts assessment”, *Geoscientific Model Development*. 7:33-51, 2014. doi: 10.5194/gmd-7-433-2014.
- [4] N. Herger et al, “Improved pattern scaling approaches for the use in climate impact studies”, *Geophysical Research Letters*. 42:3486-94, 2015. doi: 10.1002/2015gl063569.
- [5] G. S. H. Pau et al, “A reduced-order modeling approach to represent subgrid-scale hydrological dynamics for land-surface simulations: application in a polygonal tundra landscape”, *Geoscientific Model Development*, 7:2091–2105, 2014.
- [6] Y. Liu et al, “A hybrid reduced-order model of fine-resolution hydrologic simulations at a polygonal tundra site”, *Vadose Zone Journal*, 2015. doi: 10.2136/vzj2015.05.0068.
- [7] G. S. H. Pau et al, “Accurate and efficient prediction of fine-resolution hydrologic and carbon dynamic simulations from coarse-resolution models”. Submitted, 2015.
- [8] J. Ray et al, “Bayesian calibration of the Community Land Model using surrogates”, *SIAM Journal on Uncertainty Quantification*, 3(1): 199-233, 2015.
- [9] M. Huang et al, “On the Applicability of Surrogate-based MCMC-Bayesian Inversion to the Community Land Model: Case Studies at Flux Tower Sites”, *Journal of Geophysical Research-Atmosphere*, in review, 2015.
- [10] X. Chen et al, “Application of Ensemble-based Data Assimilation Techniques for Aquifer Characterization using Tracer Data at Hanford 300 Area”, *Water Resources Research*, 49, 2013. doi:10.1002/2012WR013285.
- [11] C. Monteleoni et al, “Climate Informatics: Accelerating Discovering in Climate Science with Machine Learning”, *Computing in Science & Engineering*. 15:32-40, 2013. doi: 10.1109/MCSE.2013.50.
- [12] A. Ganguly et al, “Toward enhanced understanding and projections of climate extremes using physics-guided data mining techniques”, *Nonlin Processes Geophys.* 21 (2014) 777-95, doi: 10.5194/npg-21-777-2014. NPG.
- [13] E. B. Goldstein and G Coco, “Machine learning components in deterministic models: hybrid synergy in the age of data”, *Frontiers in Environmental Science*. 3, 2015.

Contributed ideas paper 21

Anomaly detection and data reduction in Climate Simulation Dec. 1 2015
Franck Cappello, (ANL), Robert Jacob (ANL), Sheng Di (ANL)

Abstract

Future exascale climate simulations will generate 100s of PB of data. This unprecedented scale combined with the expected characteristics of exascale systems raise two major challenges related to data characteristics and size that we propose to address: anomaly detection and data reduction.

Anomaly detection

Background/Research to Date: The product of climate simulation may be affected by systematic and non-systematic corruptions induced by bugs in the hardware or a library and silent errors due to radiation. Wrong conclusions could be made potentially leading to wrong decisions. In [1], we list many corruption cases at the hardware and software level. Detection of corruptions relies on spotting anomalies in the data generated by the simulation. This problem is related to Validation and Verification (V&V) as well as to Uncertainty Quantification (UQ). Techniques such as Algorithm Based Fault Tolerance (ABFT) and replication can detect anomalies that would not be spotted by V&V and UQ. However none of these techniques cover both systematic and non-systematic corruptions or are applicable in the exascale context. In [1], we proposed a novel approach called the External Algorithmic Observer (EAO). The main principle is to run along with the simulation a surrogate function of a much lower complexity and to approximately compare the results generated by the simulation and the surrogate function. This approach fits particularly well to time stepping simulations like the ones used in climate simulation.

Preliminary results: We have investigated several types of surrogate functions in a weather model: a) same numerical method with lower order and b) prediction function. Figure 1 shows WRF streamlines computed by two different Runge Kutta numerical methods: RK4 and RK1 (Euler). The color shows the difference between RK4 and RK1. RK4 is the method of choice to compute the solution. RK1 is a low complexity, low order approximation of RK4 that we use here as a surrogate function of the EAO.

Our preliminary results show that under certain conditions, RK1 can be used to detect effectively and efficiently non-systematic and systematic corruptions in RK4. **Proposed Direction of Work:** We believe that the external algorithmic observer is a very powerful principle to detect corruptions in climate simulations. However it needs significant research to become applicable in production. Some of the research needs include: 1) design and development of a low overhead framework that will allow to plug, select low complexity surrogate functions and protect the simulations during execution. 2) study, compare and classify a variety of surrogate functions for the different modules of ACME. 3) study the correlation of the detection performance according to the size of the mesh and time steps. 4) investigate different machine learning algorithms and study conditions on learning sets to provide accurate detection, when detection relies on machine learning techniques. 5) study machine learning techniques to set detection parameters. 6) study the interplay between ensemble computation and the detection from surrogate functions.




Figure 1: WRF streamlines. The blue color show the differences between RK4 and RK1

Data reduction

Background/Research to Date: The size of the data produced by climate simulation is increasing significantly with each generation of systems. However, the bandwidth to the remote storage system has not and will not increase at the same rate. One may even observe

Contributed ideas paper 21

a plateauing of the storage bandwidth. One technique to address this issue is to use burst buffers. These mechanisms will absorb burst of I/O and reduce peaks of stress to the remote storage system. However, for climate simulations there is still a need to store a huge quantity of data. This is why the community has started considering compression as a way to reduce the data size [2]. Lossless compression manages to reduce the data set size by factors typically limited to 1.5-3. This will not be enough to solve the bandwidth mismatch problem in the long run. Higher compression ratios are needed. Recently several researchers have developed lossy compressors for floating point data sets. We believe that only lossy compressors respecting error bounds set by users are relevant for climate simulations. **Preliminary results:** We have developed a novel lossy compressor (SZ) for floating point data sets that outperforms all existing lossy compressors on compression ratio and decompression time [3]. It is also one of the fastest on compression. SZ combines multiple techniques and multiple steps to achieve its compression performance. Figure 2 presents the distribution of the compression ratios (over 1000 time steps) and the distribution of errors (1 time step, others show similar trends) of several ACME sub-models. This early result is very encouraging. We believe that SZ could be improved much further by developing new optimization techniques and compressing across time steps.

Proposed Direction of Work:

We propose to: 1) Study and improve the compression performance of SZ for the data sets produced by all ACME sub-models (improve the compression/decompression time, the compression ratio, and SZ overhead, combining

current compression techniques used by SZ with other techniques such as compression over time steps, entropy based compression, data segments sorting and wavelet modeling for non-smooth portion of variable trajectories, etc.) 2) study the effect of lossy compression with bounded errors on the data analytics performed and conclusions drawn from the compressed data. 3) study climate simulation restart from lossy compressed checkpoints for all sub-models of ACME.




Figure 2: SZ Compression ratios and error (bound of 10^{-6}) for several ACME data sets.

Connections to Math, Comp Sci & and Climate Science

Anomaly detection and data reduction require techniques in Mathematics (surrogate functions, integration methods, etc.) and Computer Science (Machine Learning, Approximate Comparison, etc.). Our research specifically focuses on Climate codes.

Potential Impact on the field

This research will help improve the trustworthiness of climate simulation results and address the critical problem of data size reduction common to many climate science problems.

References

- [1] F. Cappello, et Al., Improving the trust in results of numerical simulations and scientific data analytics, MCS-TM-352, 2015
- [2] A. H. Baker, H. Xu, J. M. Dennis, M. N. Levy, D. Nychka, S. A. Mickelson, J. Edwards, M. Vertenstein, Al Wegener, A Methodology for Evaluating the Impact of Data Compression on Climate Simulation Data, Proceedings of ACM HPDC 2014.
- [3] S. Di, F. Cappello, Fast Error-bounded Lossy HPC Data Compression with SZ, preprint number is ANL/MCS-P5437-1115, 2015

Contributed ideas paper 22

Data Assimilation for ice-sheet models

July 20, 2016

Kody Law

Computer Science and Mathematics Division, Oak Ridge National Laboratory (ORNL), Oak Ridge, TN, USA 37934

Richard Archibald (ORNL), Joseph Kennedy (ORNL).

Abstract

The topic of interest is the development and application of novel data assimilation algorithms to invert ice-sheet observations, such as the surface mass balance (SMB) and surface velocities, to infer model parameters (e.g., the basal sliding coefficient (BSC)) or complete field estimates (e.g., a complete SMB map). The potential impact to the field, in particular to forecasting sea-level rise (SLR), is enormous.

Background/Research to Date

Data assimilation algorithms have had great success in recent years in meteorology [25, 11], oceanography [17, 13], and geosciences [30, 16, 1]. It is only natural to consider their extension to climate models. To our knowledge, thus far the only studies have focused on direct inversion of measured SMB [36, 39, 28, 40] and inversion of surface velocity measurements to infer the initial condition of BSC [3, 35, 34]. The state of the art in data assimilation algorithms being used in practice are still far from fully probabilistic, relying principally on Gaussian approximations being combined with optimization techniques [32, 33, 29]. Meanwhile, theoretical work is laying the foundation for the full probabilistic treatment of such problems with realizable and tractable data assimilation algorithms [27, 38, 2, 21]. In particular, assume the model is given by

$$x_{n+1} = \Psi(x_n) + \xi_n \quad (1)$$

$$y_{n+1} = h(x_{n+1}) + \eta_{n+1}, \quad (2)$$

for $n = 0, 1, \dots$, $\Psi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$, where $\{\xi_n, \eta_{n+1}\}_{n=0}^{\infty}$ are assume i.i.d. and $\xi_n \perp \eta_m$ for all n, m . From the probabilistic perspective, the static (smoothing) problem consists of inferring the distribution $\pi(x_0, x_1, \dots, x_n | y_1, y_2, \dots, y_n)$ for a fixed n , and using this to forecast $\pi(x_{n+k} | y_1, y_2, \dots, y_n)$. The sequential online (filtering) problem consists of inferring the distribution $\pi(x_n | y_1, y_2, \dots, y_n)$, and $\pi(x_{n+k} | y_1, y_2, \dots, y_n)$, as $n \rightarrow \infty$. Inferring the distribution generally means approximating $\mathbb{E}(\varphi(x_j) | y_1, y_2, \dots, y_n)$, for $\varphi : \mathbb{R}^n \rightarrow \mathbb{R}$, which may be done either by an empirical approximation $\frac{1}{N} \sum_{i=1}^N \varphi(x_j^{(i)})$ or a parametric approximation, for example $\pi(x_n | y_1, y_2, \dots, y_n) \approx N(m_n, C_n)$. Operational data assimilation algorithms currently rely on either parametric Gaussian approximations [11] or optimization (i.e. identifying the point $x_j^* = \operatorname{argmax}_{x_j} \pi(x_j | y_1, y_2, \dots, y_n)$) [16, 12].

Proposed Direction of work

There are several tracks worth investigating in the context of data assimilation for ice-sheet models, which are listed below.

- {1} Multilevel Monte Carlo (MLMC) algorithms show great promise for extending statistical models on continuous spaces from a coarse approximation to a fine one without significantly higher computational cost [14, 6]. Indeed it is often possible to deal with uncertainty for the cost of a deterministic formulation. There has been an explosion of work recently in MLMC algorithms

Contributed ideas paper 22

for inference, including both smoothing [2, 26, 20, 24, 7] and filtering [21, 23]. There has been a comparably minute amount of work applying such technology to applications. The only work thusfar in this direction being [10]. In addition to application, further algorithmic development is also necessary. For example, for spatio-temporal dimension higher than 1, it is shown in [18] that a sparse-grid type formulation is optimal. This technology needs to be extended to the inference context. Also, these algorithms are scalable, yet there has yet been no implementation on the emerging architecture of supercomputers, which has potential to leverage the algorithmic technology to heretofore unseen application scales. The recently developed dimension-independent adaptive Metropolis (DIAM) algorithm for inference has exhibited scalability [4], and such strategy will be incorporated into the MLMC algorithms.

{2} Methodology to identify sparsity in a signal based on sparse data [8, 15, 37] shows considerable promise for obtaining the point of maximum probability of the SMB and/or the BSC.

{3} There is opportunity to extend/modify the methodology which combines ensemble and variational techniques following the work [22, 16, 1].

These threads of research concern the development of novel data assimilation techniques. Ensemble alternatives for {2} following from {3} can also be considered. Application to nonlocal models in [24] shows potential to incorporate existing [9] or new nonlocal modeling of fracture and calving.

Connections to Math, Comp Sci & and Climate Science

Contributions to each of the items {1-3} above are fundamental to the progression of computational applied mathematics, and these methods, algorithms, and models are being actively investigated by many groups, following from the citations above. Furthermore it should be noted that there are numerous opportunities to develop novel computer software *adapted to emerging architectures*. In particular, the probabilistic formulation naturally lends itself to a fault tolerant universe and items {1} and {3} can be naturally adapted to asynchronous implementations. Item {2} will in general reduce memory usage through sparse representation.

Contributions to each of the items {1-3} above are also fundamental to the progression of climate science. The significant uncertainties remaining in the ice-sheet contributions to SLR for the 21st century [5] has forced the global climate modeling community to begin incorporating increasingly complex, higher-order models of ice-sheet dynamics run at high resolutions [31]. These models, however, are poorly constrained by a lack of observational data due to the unique difficulties associated with measuring polar ice-sheets. Providing the necessary data fields and parameter values for ice-sheet initialization and SLR prediction relies on robust data assimilation algorithms applied to a variety of sparse (in time and space) datasets.

Potential Impact on the field

Robust estimation of a complete SMB field and ice-sheet parameters, such as the BSC, are necessary to determine the dynamical response of ice-sheets to climatic forcings (e.g., calving fluxes, dynamical thickening/thinning, grounding line retreat) and their contributions to SLR. Understanding of the dynamical components is lacking currently [5], and essential to answering currently open questions such as: whether the rate of SLR is accelerating [39], whether East Antarctica is gaining or losing mass [36, 19, 40], how Antarctica will respond to future climate scenarios [5], and what is the likelihood (and timing) of a possible rapid collapse of West Antarctica [19]. The work proposed here has the potential to provide the models and the methods, for recovering estimates of the parameters currently lacking, and necessary, for a highly credible resolution to these questions and robust quantification of uncertainty.

Contributed ideas paper 22

References

- [1] Sigurd I Aanonsen, Geir Nævdal, Dean S Oliver, Albert C Reynolds, Brice Vallès, et al. The ensemble Kalman filter in reservoir engineering—a review. *Spe Journal*, 14(03):393–412, 2009.
- [2] Alexandros Beskos, Ajay Jasra, Kody Law, Raul Tempone, and Yan Zhou. Multilevel sequential Monte Carlo samplers. *arXiv preprint arXiv:1503.07259*, 2015.
- [3] D. J. Brinkerhoff and J. V. Johnson. Data assimilation and prognostic whole ice sheet modelling with the variationally derived, higher order, open source, and fully parallel ice sheet model Var-GlaS. *The Cryosphere*, 7(4):1161–1184, 2013.
- [4] Yuxin Chen, David Keyes, Kody JH Law, and Hatem Ltaief. Accelerated dimension-independent adaptive metropolis. *To appear in SISC. arXiv preprint arXiv:1506.05741*, 2015.
- [5] J.A. Church, P.U. Clark, A. Cazenave, J.M. Gregory, S. Jevrejeva, A. Levermann, M.A. Merrifield, G.A. Milne, R.S. Nerem, P.D. Nunn, A.J. Payne, W.T. Pfeffer, D. Stammer, and A.S. Unnikrishnan. Sea level change. In T.F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P.M. Midgley, editors, *The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva, Switzerland, 2013.
- [6] KA Cliffe, MB Giles, Robert Scheichl, and Aretha L Teckentrup. Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients. *Computing and Visualization in Science*, 14(1):3–15, 2011.
- [7] Pierre Del Moral, Ajay Jasra, Kody Law, and Yan Zhou. Multilevel sequential Monte Carlo samplers for normalizing constants. *arXiv preprint arXiv:1603.01136*, 2016.
- [8] DL Donoho. Compressed sensing. *IEEE TRANSACTIONS ON INFORMATION THEORY*, 52(4):1289–1306, APR 2006.
- [9] Ravindra Duddu and Haim Waisman. A nonlocal continuum damage mechanics approach to simulation of creep fracture in ice sheets. *Computational Mechanics*, 51(6):961–974, 2013.
- [10] Doaa Elsakout, Mike Christie, Gabriel Lord, et al. Multilevel Markov Chain Monte Carlo (MLM-CMC) for uncertainty quantification. In *SPE North Africa Technical Conference and Exhibition*. Society of Petroleum Engineers, 2015.
- [11] Geir Evensen. *Data assimilation: the ensemble Kalman filter*. Springer Science & Business Media, 2009.
- [12] Pierre Gauthier, Monique Tanguay, Stephane Laroche, Simon Pellerin, and Josee Morneau. Extension of 3DVAR to 4DVAR: Implementation of 4DVAR at the Meteorological Service of Canada. *Monthly weather review*, 135(6):2339–2354, 2007.
- [13] Michael Ghil and Paola Malanotte-Rizzoli. Data assimilation in meteorology and oceanography. *Adv. Geophys.*, 33:141–266, 1991.
- [14] Michael B Giles. Multilevel Monte Carlo methods. *Acta Numerica*, 24:259–328, 2015.
- [15] Tom Goldstein and Stanley Osher. The split bregman method for l1-regularized problems. *SIAM J. Img. Sci.*, 2(2):323–343, April 2009.

Contributed ideas paper 22

- [16] Yaqing Gu, Dean S Oliver, et al. An iterative ensemble Kalman filter for multiphase fluid flow data assimilation. *Spe Journal*, 12(04):438–446, 2007.
- [17] Rodolfo Guzzi. Data assimilation: Mathematical concepts and instructive examples, 2015.
- [18] Abdul-Lateef Haji-Ali, Fabio Nobile, and Raul Tempone. Multi-index Monte Carlo: When sparsity meets sampling. *arXiv preprint arXiv:1405.3757*, 2014.
- [19] Edward Hanna, Francisco J. Navarro, Frank Pattyn, Catia M Domingues, Xavier Fettweis, Erik R Ivins, Robert J Nicholls, Catherine Ritz, Ben Smith, Slawek Tulaczyk, Pippa L Whitehouse, and H Jay Zwally. Ice-sheet mass balance and climate change. *Nature*, 498(7452):51–9, 2013.
- [20] Viet Ha Hoang, Christoph Schwab, and Andrew M Stuart. Complexity analysis of accelerated MCMC methods for Bayesian inversion. *Inverse Problems*, 29(8):085010, 2013.
- [21] Håkon Hoel, Kody JH Law, and Raul Tempone. Multilevel ensemble Kalman filtering. *arXiv preprint arXiv:1502.06069*, 2015.
- [22] Marco A Iglesias, Kody JH Law, and Andrew M Stuart. Ensemble Kalman methods for inverse problems. *Inverse Problems*, 29(4):045001, 2013.
- [23] Ajay Jasra, Kengo Kamatani, Kody JH Law, and Yan Zhou. Multilevel particle filter. *arXiv preprint arXiv:1510.04977*, 2015.
- [24] Ajay Jasra, Kody Law, and Yan Zhou. Forward and inverse uncertainty quantification using multi-level Monte Carlo algorithms for an elliptic nonlocal equation. *arXiv preprint arXiv:1603.06381*, 2016.
- [25] Eugenia Kalnay. *Atmospheric modeling, data assimilation, and predictability*. Cambridge university press, 2003.
- [26] C Ketelsen, R Scheichl, and AL Teckentrup. A hierarchical multilevel Markov chain Monte Carlo algorithm with applications to uncertainty quantification in subsurface flow. *arXiv preprint arXiv:1303.7343*, 2013.
- [27] KJH Law, AM Stuart, and KC Zygalakis. Data assimilation: A mathematical introduction. *Springer Texts in Applied Mathematics*, 2015.
- [28] JTM Lenaerts, MR Den Broeke, WJ Berg, E van Meijgaard, and P Kuipers Munneke. A new, high-resolution surface mass balance map of Antarctica (1979–2010) based on regional atmospheric climate modeling. *Geophysical Research Letters*, 39(4), 2012.
- [29] Andrew C Lorenc, Neill E Bowler, Adam M Clayton, Stephen R Pring, and David Fairbairn. Comparison of hybrid-4DEnVar and hybrid-4DVar data assimilation methods for global NWP. *Monthly Weather Review*, 143(1):212–229, 2015.
- [30] Geir Nævdal, Liv Merete Johnsen, Sigurd I Aanonsen, Erlend H Vefring, et al. Reservoir monitoring and continuous model updating using ensemble Kalman filter. *SPE journal*, 10(01):66–74, 2005.
- [31] F Pattyn, C Schoof, L Perichon, R. C. A. Hindmarsh, E. Bueler, B. de Fleurian, G. Durand, O. Gagliardini, R. Gladstone, D. Goldberg, G. H. Gudmundsson, P. Huybrechts, V. Lee, F. M. Nick, A. J. Payne, D. Pollard, O. Rybak, F. Saito, and A. Vieli. Results of the marine ice sheet model intercomparison project, MISMIP. *The Cryosphere*, 6(3):573–588, 2012.

Contributed ideas paper 22

- [32] Stephen G Penny. The hybrid local ensemble transform Kalman filter. *Monthly Weather Review*, 142(6):2139–2149, 2014.
- [33] Stephen G Penny, David W Behringer, James A Carton, and Eugenia Kalnay. A hybrid global ocean data assimilation system at NCEP. *Monthly Weather Review*, 143(11):4660–4677, 2015.
- [34] Noemi Petra, James Martin, Georg Stadler, and Omar Ghattas. A computational framework for infinite-dimensional Bayesian inverse problems, part ii: Stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.
- [35] Noemi Petra, Hongyu Zhu, Georg Stadler, Thomas JR Hughes, and Omar Ghattas. An inexact Gauss–Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model. *Journal of Glaciology*, 58(211):889–903, 2012.
- [36] Andrew Shepherd, Erik R Ivins, A Geruo, Valentina R Barletta, Mike J Bentley, Srinivas Bettadpur, Kate H Briggs, David H Bromwich, René Forsberg, Natalia Galin, et al. A reconciled estimate of ice-sheet mass balance. *Science*, 338(6111):1183–1189, 2012.
- [37] J. L. Starck, D. L. Donoho, M. J. Fadili, and A. Rassat. Sparsity and the Bayesian perspective. *ASTRONOMY & ASTROPHYSICS*, 552, APR 2013.
- [38] Andrew M Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
- [39] B Wouters, JL Bamber, MR Van den Broeke, JTM Lenaerts, and I Sasgen. Limits in detecting acceleration of ice sheet mass loss due to climate variability. *Nature Geoscience*, 6(8):613–616, 2013.
- [40] H Jay Zwally, Jun Li, John W Robbins, Jack L Saba, Donghui Yi, and Anita C Brenner. Mass gains of the Antarctic ice sheet exceed losses. *Journal of Glaciology*, 2015.

Particle-Resolved Direct Numerical Simulations with Data Mining and Streaming Analysis for Addressing Aerosol-Cloud-Precipitation Challenges

Yangang Liu (Environmental and Climate Sciences Department, Brookhaven National Lab)

Shinjae Yoo (Computer Science Center, Brookhaven National Lab)

Xiaolin Li (Applied Math and Statistics Department, Stony Brook University)

Abstract

Accurately simulating the aerosol-cloud-precipitation system and adequately representing it in climate models continue to pose daunting challenges to the cloud physics and climate communities alike. Particular knowledge gaps exist for vital processes that occur at scales smaller than typical grid sizes of large eddy simulation (LES) model (e.g., 100 m) or cloud-resolving model (CRM, e.g., 1 km), including, but are not limited to, microphysics, turbulent entrainment-mixing between clouds and environmental air, turbulence and their mutual interactions. These processes are either not represented at all or represented rudimentarily in major types of models such as climate models, CRM and LES model, hindering future progress of climate modeling. To address the challenges at the most fundamental level we propose to develop a cross cutting particle-resolved direct numerical simulation (DNS) model with advanced data mining and streaming analysis by forming an multi-disciplinary team from areas of cloud and climate sciences, computer science and mathematics. This new cross-cutting model aims to serve as the most realistic simulation of those gap processes in ambient aerosol-cloud-precipitation by resolving the smallest turbulent eddies, covering a domain of about LES grid size, tracking evolution of individual particles, and containing a package for on-line data mining and stream analysis to timely guide physical understanding and model development and integration with observations.

Background/Research to Date

It is commonly accepted that the accuracy and reliability of climate models in projecting the climate change caused by climate forcing depend heavily on cloud feedbacks and thus on representations (parameterizations) of still poorly understood cloud processes, which have the potential to dampen or enhance changes in essential climatic variables such as temperature and precipitation. The situation worsens when interactions with natural and anthropogenic aerosols are included. Indeed, the latest Inter-governmental Panel on Climate Change continues to assigns “very low confidence” to aerosol- cloud-precipitation interactions, with even the sign of the resulting climate forcing remaining uncertain.

It is not surprising that the problem has been dogging so long, given that the aerosol-cloud-precipitation system represents an extreme example of multi-scale nature and multiple nonlinearities. The system involves physical processes over a wide range of scales from the smallest turbulent eddies (~ 1 mm) to convection (~ km) to global scales in terms of air motion, and from nano-sized aerosol particles to cloud droplets of micrometers to precipitation of millimeter. The system is further riddled with convoluted nonlinearities (turbulence cascade, convection, entrainment, mixing, and microphysical processes)¹ and sensitive thresholds controlling phase changes and cloud-precipitation conversion². Understanding such complex processes and upscaling them to adequate representation in climate models present daunting challenges to the scientific community,

Contributed ideas paper 23

which become more acute for extreme precipitation and weather events. BER has taken up the challenges, esp. through the Atmospheric System Research (ASR) program and the Atmospheric Radiation Measurement (ARM) facility and the new ACME model. Much progress has been made through such and other efforts. However, significant knowledge gaps remain. In particular, many key processes, including microphysics, turbulent entrainment-mixing between clouds and environmental air, turbulence and their mutual interactions, occur at scales smaller than typical grid sizes of large eddy simulation (LES) models (e.g., 100 m) or cloud-resolving models (CRM) (e.g., 1 km) — flagship models for studying cloud processes and informing parameterization development for climate models, and are either not represented at all or represented in very rudimentary ways, seriously hindering progress of future climate models such as ACME³.

Proposed Direction of Work

Fully addressing these vital gaps at the fundamental level calls for a cross-cutting model that not only resolves the smallest turbulent eddies in clouds, but also tracks growth/motion of individual particles. Few attempts^{3, 4} exist along this line, but are limited in scientific scopes, hindering crucial progress. Other important knowledge gaps consist in the understanding and characterization of clustering structures ubiquitously with the aerosol-cloud-precipitation system, and online real-time data process^{5,6,7} that can be used to inform physical understanding and model evaluation.

With the advancements in physics/observation of aerosol-cloud-precipitation, computer sciences, big data mining/streaming^{5,6,7,8}, now is the time to develop a cross cutting particle-resolved DNS model with advanced data mining and streaming analysis to address the aerosol-cloud-precipitation challenges, by forming an inter-disciplinary team and performing extreme scale computations and on-line data mining and stream analysis. This new cross-cutting model is envisioned to resolve the smallest turbulent eddies and cover a domain of about LES grid size, track growth and motion of individual particles, and contain on-line data mining and stream analysis to timely guide physical understanding and model development, and model-observation integration.

Connections to Math, Comp Sci & and Climate Science

Developing such a benchmarking model requires expertise in areas of cloud physics and climate science (e.g., providing correct sets of equations describing the system, physical analysis and understanding of the simulations, and develop parameterizations), computer science (e.g., extreme scale supercomputing and streaming analysis, and advanced mathematics (such as data mining and feature identification), and close interdisciplinary coordination.

Potential Impact on the field

The model expects to change the way we simulate aerosol-cloud-precipitation system at the most fundamental level, and significantly impact at least the following areas: (1) physical understanding of key processes occurring at the scales smaller than CRM or LES scales; (2) advancement of climate models by improving parameterization of these processes in climate models; (3) provide benchmark simulations for approximations at various levels including the state of art spectral bin microphysical model; (4) unprecedented synergy between physical sciences (e.g., cloud physics, fluid dynamics, climate sciences), computer sciences (e.g., extreme scale supercomputing and streaming analysis), and mathematics (such as data mining and feature identification), as well as between models and observations.

Contributed ideas paper 23

References

- [1] Liu, Y., P. H. Daum, and J. Hallett, 2002: A generalized systems theory for the effect of varying fluctuation in cloud droplet size distributions. *J. Atmos. Sci.*, 59, 2279-2289.
- [2] Liu, Y., P. H. Daum, R. McGraw, and M. Miller, 2006: Generalized threshold function accounting for effect of relative dispersion on threshold behavior of autoconversion process. *Geophys. Res. Lett.*, 33. L11804, doi:10.1029/2005GL0255000
- [3] Grabowski, W., and L. Wang, 2013: Growth of Cloud Droplets in a Turbulent Environment. *Annu. Rev. Fluid Mech.*, 45, 293-324.
- [4] Gao, Z., Y. Liu, and X. Lin, 2015: A new particle-resolved direct numerical simulation model for studying processes at sub-meter scales. *Journal of Geophysical Research: Atmospheres* (to be submitted)
- [5] Yoo, S., H. Huang, S. Kasiviswanathan, 2015: "Streaming Spectral Clustering" (to be submitted).
- [6] Huang, H., S. Yoo, S. Kasiviswanathan, 2015: "Unsupervised Feature Selection on Data Streams". ACM International Conference on Information and Knowledge Management (CIKM).
- [7] Huang, H., S. Yoo, D. Yu, H. Qin, 2015: "Diverse Power Iteration Embeddings: Theory and Practice". *IEEE Transactions on Knowledge and Data Engineering (TKDE)* (Accepted).
- [8] Huang, H., S. Yoo, D. Yu, H. Qin, 2014: "Noise-Resistant Unsupervised Feature Selection via Multi-Perspective Correlations", *IEEE International Conference on Data Mining (ICDM)*.

Contributed ideas paper 24

Data-driven stochastic model reduction

December 1, 2015

Fei Lu

Department of Mathematics, UC Berkeley and Lawrence Berkeley National Laboratory
Alexandre J. Chorin (Department of Mathematics, UC Berkeley and Lawrence Berkeley National Laboratory), Kevin Lin (Department of Mathematics, University of Arizona)

Abstract

Prediction for high-dimensional nonlinear dynamic systems is often difficult because such systems are often expensive to solve in full; in addition, one often observes only a few of the variables, so that initial data are not available for the full system. Reduced models for the observed variables are thus needed. The challenges come from the nonlinear interactions between the observed variables and the unobserved variables, and the difficulties in quantifying uncertainties from discrete data.

We address these challenges by developing discrete-time stochastic reduced systems, in which one formulates discrete solvable approximate equations for the observed variables, and uses data and statistical methods to account for the impact of the unobserved variables. A key ingredient in the construction of the stochastic reduced systems is a discrete-time stochastic parametrization based on inference from nonlinear time series. The particular discrete framework in which we chose to identify the discrete models is defined by the NARMAX (nonlinear autoregression and moving average representation with exogenous input); other discrete representations may be useful as well.

Background/Research to Date

We have demonstrated our approach on several model problems, including the two-layer Lorenz 96 system [3] and the Kuramoto-Sivashinsky equation [4]. For the two-layer Lorenz 96 system, we assumed that the goal was to predict the evolution of the large-scale variables when only the large-scale variables could be observed discretely, while the many small-scale variables were unobserved. In the Kuramoto-Sivashinsky equation case, the goal was to predict the evolution of a few observed low wave-number Fourier modes, while the remaining Fourier modes were unobserved. In both cases, we developed discrete-time stochastic reduced systems for the observed variables based on data. The stochastic reduced systems capture the key dynamical features of the variables of interest: their short-term evolution and their long-term statistics. The stochastic reduced systems can be used to advance the predictive understanding of the full system.

In a situation where only some of the variables are observed one may not be able to deduce from data a valid Markovian model such as a differential equation, as demonstrated by the Mori-Zanzig analysis ([2, 7]). An added advantage of our discrete representation is that it can seamlessly represent both Markovian and non-Markovian systems. In addition, as we showed in earlier work, our discrete identification algorithms survive in situations where the data are sparse.

Proposed Direction of work

The use of discrete models simplifies several of the tasks one has to perform to produce a useful reduced predictive model. As we have shown in earlier work, the discrete models make it easier to

Contributed ideas paper 24

estimate and identify model error in the reduced models, and avoids the difficulty of developing high-order schemes for multivariate nonlinear stochastic differential equations. However, one still faces the task of making the resulting reduced models parsimonious enough for practical use. We have begun to address this problem by adapting ideas from differential equations to the analysis of the discrete reduced systems; in particular, we have used ideas related to inertial manifolds to identify the main interactions that need to be taken into account in the modeling, and then identified the best coefficients for the corresponding terms numerically [4]. Our emphasis now is the development of new ideas for producing parsimonious models, using tools such as perturbation schemes, renormalization, and diffusion maps.

We also expect to use our discrete methods to develop efficient algorithms with relatively large time steps or spatial grids, or reduce the dimension of spectral methods, using the data generated by increasingly detailed simulations and experiments. The resulting discrete-time stochastic reduced systems would quantify the uncertainty in the computations, and hence make it possible to estimate the confidence levels in the results.

Connections to Math, Comp Sci & and Climate Science

Stochastic modeling is increasingly being used in weather and climate prediction [1, 5, 6], because of one's limited ability to resolve all the relevant processes and scales in full models. In many practical applications, one is mainly interested in and can observe only the large and potentially predictable scales, and not necessarily the small and fast scales. This leads to the problem of developing stochastic reduced dynamics systems for the large-scale variables from observed data.

The discrete-time representations we use have been developed mostly in the context of computer science and statistics. Our work lies in the intersection of these sciences with mathematics.

Potential Impact on the field

Climate science is necessarily full of uncertainty, with many processes poorly understood and with sparse data. In climate modeling often one must draw conclusions and make predictions on the basis of imperfect models supplemented by data. The models are imperfect because of one's limited understanding of the full system, and because of one's limited ability to resolve all the relevant processes and scales in the full system. We think that the kind of discrete models we advocate will have a major impact on climate science.

References

- [1] M.D. Chekroun, D. Kondrashov, and M. Ghil. Predicting stochastic systems by noise sampling, and application to the El Niño-Southern Oscillation. *Proc. Natl. Acad. Sci. USA*, 108:11766–11771, 2011.
- [2] A.J. Chorin and O.H. Hald. *Stochastic Tools in Mathematics and Science*. Springer, New York, NY, 3rd edition, 2013.
- [3] A.J. Chorin and F. Lu. Discrete approach to stochastic parametrization and dimension reduction in nonlinear dynamics. *Proc. Natl. Acad. Sci. USA*, 112(32):9804–9809, 2015.
- [4] F. Lu, K. Lin, and A.J. Chorin. Data-based stochastic model reduction for the Kuramoto–Sivashinsky equation. *arXiv:1509.09279*, 2015.

Contributed ideas paper 24

- [5] A.J. Majda and B. Khouider. Stochastic and mesoscopic models for tropical convection. *Proc. Natl. Acad. Sci. USA*, 99:1123–1128, 2002.
- [6] TN Palmer, R Buizza, F Doblas-Reyes, T. Jung, M Leutbecher, GJ Shutts, M Steinheimer, and A Weisheimer. Stochastic parametrization and model uncertainty. *ECMWF Technical Memorandum*, 598, 2009.
- [7] R. Zwanzig. Nonlinear generalized Langevin equations. *J. Stat. Phys.*, 9(3):215–220, 1973.

Contributed ideas paper 25

Tuning High-Resolution Coupled Models
Peter Caldwell
Lawrence Livermore National Lab

12/1/15

Abstract

The goal of DOE's Accelerated Climate Model for Energy (ACME) initiative is to develop an ultra-high resolution climate model. In order to adequately *tune* such a high resolution model - which is too expensive to run routinely - a variety of novel approaches will be needed. The goal of this paper is to describe these techniques and to identify areas where mathematical and computational improvements are needed.

Background/Research to Date

So far, high-resolution climate simulations have used tunings from their coarse-resolution counterparts with at most a crude alteration of a single parameter to restore energy balance [1-3]. Because most parameters in climate models are sensitive to resolution, this approach is likely to yield suboptimal results. Over the last few years, DOE has developed several tools that will be useful for more careful tuning at high resolution.

One tool is the ability to run the model in *regionally-refined* mode, where regions of interest are run at high resolution and other regions are run with coarse grid spacing. A parallel effort enables us to run the ACME parameterizations within a limited-area model. Both approaches allow us to test and tune the high-resolution model without incurring the full penalty of a high-resolution model. Questions remain, however, about the validity of these approaches in the face of non-local effects. Tuning each model component (atmosphere, ocean, land, etc) separately can also be used to more efficiently obtain optimal tunings, but these separately tuned components may not sum to an optimal coupled tuning. For example, version 5 of the Community Earth System Model was originally tuned component-wise, but when coupled the model drifted towards an ice age. Better understanding of when and how single-component tuning is appropriate would be beneficial for efficient high-resolution tuning. More efficient approaches to ocean spinup would also be useful for efficient tuning at high resolution.

Another useful tool for high-resolution tuning is the use of short simulations instead of climate-length runs. DOE has been performing these sorts of runs for a long time in the context of the Cloud Associated Parameterizations Testbed (CAPT), where short weather-forecast simulations using a climate model are compared against observations [4]. More recently, model climatologies were shown to be obtainable at a fraction of the cost by running an ensemble of short simulations [5]. Better understanding of when these techniques are informative and where they are misleading would be helpful.

Uncertainty quantification (UQ) techniques are useful for determining which combinations of parameters are best and for better exploring parameter space [6]. Even when tuning is ultimately performed by hand, knowing how much modeled quantities change in response to a unit change in a parameter (which can be readily ascertained from UQ) is handy. Obtaining sufficient samples for UQ is challenging for an expensive model, however. The short simulations described earlier will be useful in this regard, but using coarse-resolution simulations to inform high-resolution tuning will also be extremely useful for efficient high-resolution tuning [7-8]. Using UQ to automatically find optimal tunings would be also very useful. Auto-tuning has not been used operationally in climate models because the goals in tuning a climate model are complex and ill defined. For example, getting a good El Nino pattern may not be a primary objective until your model doesn't have one. It would be interesting and useful to use machine-learning techniques to

Contributed ideas paper 25

develop an objective function by watching the expert hand-tuning process. Developing auto-tuning techniques that can handle complex objective functions is particularly necessary for the promise of UQ auto-tuning to be realized in climate modeling.

Proposed Direction of work

There are many tasks that need to be done before ACME has a robust high-resolution coupled tuning capability in place. Some of these tasks (e.g. regionally-refined mode, CAPT, short ensembles) require only that the ACME team harden existing techniques and develop protocols for using existing tools. Other tasks (e.g. understanding the limitations of short-simulation and regional high-resolution techniques and component-wise versus coupled tuning) are clearly questions that climate scientists need to answer (and are working on). Other tasks (e.g. multi-resolution UQ, auto-tuning with complex objective functions) clearly require mathematical expertise. And finally, several tasks (ocean spin up, coupling questions) require careful statistical understanding, probably as a joint venture between physical and mathematical scientists.

Connections to Math, Comp Sci & and Climate Science

UQ is an obvious area where mathematicians and computer scientists could contribute to improved model tuning. Multi-resolution calibration and auto-tuning with complex objective functions are particularly obvious targets. Statisticians would be helpful in understanding how to optimize ocean spin-up and may be useful for understanding the limitations of short simulations. Experts at process coupling may be able to shed insight on when tuning single components or in a limited area is misleading.

Potential Impact on the field

Thorough tuning is a crucial step in the development of a world-class climate model, yet methods for tuning high-resolution models are still in their infancy. Rapid advances are needed in order for the ACME model to fulfill its promise as the first ultra-high resolution climate model.

References

- [1] Thomas L. Delworth, Anthony Rosati, Whit Anderson, Alistair J. Adcroft, V. Balaji, Rusty Benson, Keith Dixon, Stephen M. Griffies, Hyun-Chul Lee, Ronald C. Pacanowski, Gabriel A. Vecchi, Andrew T. Wittenberg, Fanrong Zeng, and Rong Zhang, 2012: Simulated Climate and Climate Change in the GFDL CM2.5 High-Resolution Coupled Climate Model. *J. Climate*, **25**, 2755–2781.
- [2] T. Jung, M. J. Miller, T. N. Palmer, P. Towers, N. Wedi, D. Achuthavarier, J. M. Adams, E. L. Altshuler, B. A. Cash, J. L. Kinter III, L. Marx, C. Stan, and K. I. Hodges, 2012: High-Resolution Global Climate Simulations with the ECMWF Model in Project Athena: Experimental Design, Model Climate, and Seasonal Forecast Skill. *J. Climate*, **25**, 3155–3172.
- [3] Julio T. Bacmeister, Michael F. Wehner, Richard B. Neale, Andrew Gettelman, Cecile Hannay, Peter H. Lauritzen, Julie M. Caron, and John E. Truesdale, 2014: Exploratory High-Resolution Climate Simulations using the Community Atmosphere Model (CAM). *J. Climate*, **27**, 3073–3099.
- [4] Thomas J. Phillips, Gerald L. Potter, David L. Williamson, Richard T. Cederwall, James S. Boyle, Michael Fiorino, Justin J. Hnilo, Jerry G. Olson, Shaocheng Xie, and J. John Yio, 2004: Evaluating Parameterizations in General Circulation Models: Climate Simulation Meets Weather Prediction. *Bull. Amer. Meteor. Soc.*, **85**, 1903–1915.
- [5] Wan, H., P. J. Rasch, M. A. Taylor, and C. Jablonowski (2015), Short-term time step convergence in a climate model, *J. Adv. Model. Earth Syst.*, **7**, 215–225.
- [6] Guo, Z., M. Wang, Y. Qian, V. E. Larson, S. Ghan, M. Ovchinnikov, P. A. Bogenschutz, A.

Contributed ideas paper 25

Gettelman, and T. Zhou (2015), Parametric behaviors of CLUBB in simulations of low clouds in the Community Atmosphere Model (CAM), *J. Adv. Model. Earth Syst.*, **7**, 1005–1025.

[7] Goh, Joslin, et al. "Prediction and computer model calibration using outputs from multifidelity simulators." *Technometrics* 55.4 (2013): 501-512.

[8] Bulaevskaya, V., D. Lucas. "Climate model calibration across multiple spatial resolutions" Abstract#317429. Joint Statistical Meetings, Seattle, WA, August 8-13, 2015.

Contributed ideas paper 26

UQ and V&V for Climate Modeling

William F. Spatz

Sandia National Laboratories

1 December 2015

Abstract

Uncertainty quantification (UQ), and verification and validation (V&V) are fundamental challenges in climate modeling. This topic entails identifying sources of uncertainty, describing uncertainty associated with input parameters, model comparisons between numerical and/or analytical solutions, evaluating model uncertainty through validation against observations, up-scaling and downscaling, as well as quantifying uncertainty through both forward modeling (sensitivity analyses) and inverse modeling (optimization/calibration) in all components of climate system models at a variety of spatial and temporal scales. These avenues of study properly require massive computing resources not yet available, and approaching these topics will require algorithmic advances in addition to faster hardware.

Background/Research to Date

The InterAcademy Council [1] has reviewed the processes and procedures of the Intergovernmental Panel on Climate Change (IPCC), and found many aspects of the assessment reports lacking with respect to the UQ and V&V of climate models. These deficiencies are understandable, as single climate simulations require massive computing resources and proper UQ and V&V require orders of magnitude more resources: computing cycles that are simply not yet available.

As an example, sensitivity analysis should include many simulations spanning the space of random and uncertain input variables. Each uncertain or random input variable adds a new dimension (the “curse of dimensionality”) to the problem. This multi-dimensional problem is too computationally intensive to be approached in a methodical and comprehensive way. What has been tractable in practice is to use whatever simulations are available as a proxy for the sample space [2]; this space spans different numerical methods, different convection schemes, etc., but is not the complete space of random and uncertain variables.

Proposed Direction of Work

The advent of the Exascale Computing Initiative promises more computing resources will be committed to this problem, but experience with UQ and V&V indicate that exascale machines will not be enough. We propose the following ideas to move climate modeling closer to effective and meaningful UQ and V&V:

Idea: Leverage Advanced Simulation and Computing (ASC) experience to guide the path forward for UQ and V&V in climate modeling. For example, ASC considers four types of errors: producer error (Type I), user error (Type II), solving the wrong problem, and using the solution incorrectly. This is in contrast to current emphasis on “hypothesis testing,” which only addresses producer error [3]. The latter two errors are of importance in policy considerations informed by computational simulation of climate and demand greater formal attention.

Contributed ideas paper 26

Idea: Explore the possibility of multi-fidelity UQ for climate. Utilizing multiple fidelities and resolutions is one way to reduce the order of the model and provide more samples for the ensembles. Additionally, multi-fidelity UQ can utilize lower fidelity results to optimize the sample space for a smaller number of higher fidelity simulations. Thus the efficiency of low-resolution models and the fidelity of high-resolution models can both be exploited.

Idea: Utilize the additional simulations to increase efficiency by running them concurrently. Recent research [4] has demonstrated that a factor of five increase in efficiency is possible (for the specific case studied) when run in this mode. In the context of the multi-fidelity idea above, the lower-resolution simulations will certainly allow for concurrent samples.

Connections to Math, Comp Sci & and Climate Science

This research would directly leverage findings and tools borne not only of the ASC V&V program, but of the QUEST SciDAC Institute as well. Concurrent simulations have been enabled by the Albany project, which in turn leverages leading edge Sandia computational science and has already been applied to climate via the FELIX ice sheet model.

Potential Impact on the Field

The ACME project is justified as a “DOE model on DOE machines for the DOE mission.” Climate modeling for the DOE mission implies that the models should inform energy policy for climate change mitigation and adaptation. This in turn requires that the models enable climate change detection, climate change attribution, risk analysis, and risk management. All of these capabilities are inseparable from uncertainty quantification, verification and validation. UQ and V&V are hard problems that will require both the extreme computing resources promised by the Exascale Computing Initiative, and algorithmic advances as proposed in this idea paper.

References

- [1] InterAcademy Council (2010), “Climate change assessments: Review of the processes and procedures of the IPCC”, available from World Wide Web (https://www.ipcc.ch/pdf/IAC_report/IAC%20Report.pdf).
- [2] Qingyun Duan, “Quantification of Parametric Uncertainty of Large Complex Geophysical Models,” presented at the Workshop on Uncertainty Quantification in Climate Modeling and Projection, Trieste, Italy, July 2015. Available from World Wide Web (<http://indico.ictp.it/event/a14268/session/0/contribution/6/material/slides/>).
- [3] Maarten H.P. Ambaum, “Significance Tests in Climate Science,” *J. Climate*, v. 23, pp. 5927–5932.
- [4] William F. Spitz, Thomas M. Smith, Irina P. Demeshko, Jeffrey A. Fike, “Aeras: A Next-Generation Global Atmosphere Model,” *Procedia Computer Science*, v. 51, International Conference On Computational Science, ICCS 2015 — Computational Science at the Gates of Nature, 2015, pp. 2097–2106.

Contributed ideas paper 27

Dimension Reduction Methods for Improving Weather and Climate Models

November 24, 2016

Erik S. Van Vleck

Department of Mathematics, University of Kansas &
Mathematics and Climate Research Network

Nate Brunzell (Department of Geography and Atmospheric Science, University of Kansas),

Dave Mechem (Department of Geography and Atmospheric Science, University of Kansas)

Abstract

The work proposed here emphasizes computational dynamical systems techniques for dynamic model reduction to improve weather and climate models. In particular, the methods developed expand solutions in terms of a time dependent basis of so-called Lyapunov vectors. These are the analogues of eigenvectors for Lyapunov exponents that measure growth and decay rates of time dependent solutions of differential equations. This work involves developing methods for efficient, accurate computation coupled with assessments of computational error, model error, and quantification of uncertainty for community models of cloud dynamics (WRF-LES), land-atmosphere interaction (CLM), and general circulation (ACME).

Background/Research to Date

This work is an outgrowth of the development and analysis of computational techniques for Lyapunov exponents and related time dependent stability information. Expanding solutions of differential equations in terms of Lyapunov vectors provides for a robust dynamic technique to express systems as time dependent slow-fast systems. The slow part of the system corresponds to directions corresponding to positive, zero, and weakly negative (finite time) Lyapunov exponents while the fast part corresponds to the directions corresponding to strongly negative (finite time) Lyapunov exponents. We have developed effective techniques and long time global error analysis for determining a time dependent orthonormal basis for the Lyapunov vectors [3]. Using this representation for the solutions, we have developed decoupling techniques based upon time dependent inertial manifold ideas and data assimilation techniques based upon assimilation in the unstable subspace. Functionally if we write a standard Galerkin truncation (u are the time and space dependent state variables and α are tuneable system parameters) of a differential or algebraic equation $G(u; \alpha) = 0$ as $PG(Pu; \alpha) = 0$ or in a perturbative framework $PG(u^{(0)} + P\delta; \alpha^{(0)} + \Delta\alpha) = 0$ ($u^{(0)}, \alpha^{(0)}$ best known approximations), then the techniques we have developed and are refining can be written as $0 = PG(Pu + (I - P)u; \alpha) = PG(Pu + \Phi(Pu); \alpha)$ or $0 = PG(u^{(0)} + P\delta + (I - P)\delta; \alpha^{(0)} + \Delta\alpha) = PG(u^{(0)} + P\delta + \Psi(P\delta); \alpha^{(0)} + \Delta\alpha)$. Here we think of a model as corresponding to a differential operator and a variational data assimilation technique corresponding to an algebraic operator. This truncation is performed in space-time and the projection $P = \text{diag}(P_k)$, k denoting time, $P_k = Q_k Q_k^T$ where the columns of Q_k denote the orthonormal basis for the first p Lyapunov vectors at time k . The mappings Φ and Ψ map the slow variables to the fast variables in state space and tangent space, respectively, and are determined using newly developed time dependent decoupling techniques [1, 2] under appropriate conditions. The mappings also provide useful information to assess when or if a standard Galerkin truncation, e.g., $\Phi = 0$, will provide results that are consistent with the model. We have tested these techniques on idealized but increasingly complex model problems and have obtained promising results. Figure 1 shows

Contributed ideas paper 27

results of a sequential shadowing based (PDA) data assimilation [2] applied to Lorenz'96 model with $n = 40$ and $F = 8$ with a reduced dimension of $p = 18$. The codes have been developed so that they can easily be used with any stand-alone codes with a restart capability, such as the community-model codes listed above.

[Figure 1 about here.]

Proposed Direction of work

These new techniques provide an innovative pathway toward more rigorously confronting community models (WRF-LES, CLM, and ACME) with DOE-funded observations from the Atmospheric Radiation Measurement Program (ARM) and Ameriflux towers, which are predominantly point measurements or a soda straw view of the atmosphere. Our initial focus will be on time dependence using single column or two-dimensional time dependent models and state space and parameter space data assimilation. The techniques are applicable in a Bayesian framework to provide PDFs of state variables and system parameters.

- **Attributes:** We employ a full Gauss-Newton framework (linear systems are underdetermined, solved using pseudo inverse for minimum norm Newton updates) that requires approximation of derivatives with respect to state variables and system parameters. This delivers good convergence properties, first order approximations of uncertainty in state variables and system parameters, and applicability to assessment of classical/shadowing global error analysis and use of likelihood functions to assess model error.
- **Limitations:** Full Newton is expensive for high dimensional models and dimension reduction techniques require a gap condition either at the nonlinear level (more difficult to satisfy) or linear level as in perturbative framework (less difficult to satisfy).

Connections to Math, Comp Sci & and Climate Science

The proposed work involves developing and integrating computational dynamical systems techniques based upon time dependent stability theory (Lyapunov exponent theory) for dimension reduction of cloud dynamics, regional climate, and general circulation models. Connections to climate science are achieved through interaction with climate scientists with expertise in cloud dynamics and regional climate modeling and earth system models. The techniques being developed seek to provide explicit information necessary to assess state space and parameter space uncertainties. While the computational expense in doing so is relatively high, codes are being written in a modular form to take advantage of high performance computing capabilities and advances, in particular, numerical linear algebra techniques.

Potential Impact on the field

The proposed work has great potential to provide systematic dimension reduction and data assimilation for state space and parameter estimation combined with uncertainty quantification and assessment of model and computational error. This project requires a combined effort between personnel with expertise in computational dynamical systems techniques and climate scientists. The proposed applications provide realizable stepping stones with further collaboration with high performance computing experts to full scale GCMs and earth system models.

Contributed ideas paper 27

References

- [1] Y.-M. Chung, A. J. Steyer, and E.S. Van Vleck, “Non-Autonomous Inertial Manifold Reduction,” (2015) *submitted*.
- [2] J. Frank, B. de Leeuw, A. Steyer, X. Tu, and E.S. Van Vleck, “Hybrid Shadowing Based Data Assimilation with Assimilation in the Non-Strongly Stable Subspace,” (2015) *in preparation*.
- [3] E.S. Van Vleck, “On the Error in the Product QR Decomposition,” (2010) *SIAM J. Matr. Anal. Appl.* **31** pp. 1775–1791.

Contributed ideas paper 27

List of Figures

- 1 Distance from truth for observations and assimilation in sequential shadowing based (PDA) data assimilation [2] applied to Lorenz'96 model with $n = 40$ and $F = 8$ with a reduced dimension of $p = 18$ 5

Contributed ideas paper 27




Figure 1: Distance from truth for observations and assimilation in sequential shadowing based (PDA) data assimilation [2] applied to Lorenz'96 model with $n = 40$ and $F = 8$ with a reduced dimension of $p = 18$.

Contributed ideas paper 28

Task Parallelism for Complex Multiscale Performance
Phil Jones
Los Alamos National Laboratory

30 November 2015

Abstract

Computer architectures are evolving toward very large numbers of streaming processing elements as a path toward exascale computing. Climate models on current architectures have difficulty exploiting even moderate parallelism using standard domain decomposition techniques. At the same time, climate models are becoming more complex through the inclusion of new processes, the use of variable resolution to resolve a range of spatial scales, and in addressing multiple time scales. Task parallelism may provide an opportunity to exploit the increasing core counts while also managing the complexity of the models being deployed.

Background/Research to Date

Parallel implementations of climate models generally rely on a two-dimensional decomposition of the model's horizontal spatial domain. While this is a natural decomposition that is relatively straightforward to implement, even very-high-resolution climate simulations have difficulty efficiently exploiting systems beyond 10-100k processing elements. Ideas to further exploit parallelism have tended to rely on ensembles rather than revisiting the approach to parallelism. New approaches to parallelism based on task parallelism have (re-)emerged recently with some examples including PaRSEC[1] and Legion [2]. These approaches rely on lightweight task run-time systems and dataflow analyses to efficiently map tasks and their associated data to the underlying hardware. With enough granularity, such systems may provide many advantages in exploiting exascale hardware, including far greater parallelism, load balancing, mapping different tasks to different types of processing elements in hybrid systems, and the management of data. From the climate model perspective, a task-based approach can also be used to manage complexity at various levels. Tasks and associated data can be defined and distributed on the hardware, managing multiple space and time scales based on when results of a task are needed.

Proposed Direction of work

We propose to refactor climate models to take advantage of task parallelism. This refactoring will take place on multiple levels. Frameworks like Legion currently manage tasks at the higher (e.g. component or domain decomposition) levels, distributing work across nodes, managing layouts, data and any required messaging. We can use Legion to explore task layouts and data strategies at these higher levels, in addition to adding new components like analyses or managing separate processes like I/O. At lower levels within components, we can begin to separate different tasks like various tendency terms as kernels and begin to compute these tasks concurrently on different cores or accelerators. This early refactoring will enable us to explore the efficacy of the approach and begin to define the data flow and tasks within our models. As programming models like Legion become more mature, we can rely on more sophisticated data flow analyses to manage a broader range of tasks and exploit further advantages of the approach. Ultimately, we would view the climate model as simply a continuing list of tasks and data dependencies that can be managed and distributed to the hardware with a run-time system.

As described above, a task-based model has several advantages from both the computing and climate science perspectives. Computational advantages include:

- *Exposing additional parallelism.* As future architectures add more processing elements, we will be able to move beyond the limitations of horizontal domain decomposition.
- *Load balancing.* With a fine enough granularity of tasks, idle processing elements or threads will be able to move on to the next work unit.

Contributed ideas paper 28

- *Mapping tasks to most suitable hardware.* For hybrid architectures, tasks that are more suitable to hardware elements (e.g. GPU accelerators, burst buffers, I/O) can be scheduled or mapped to those elements.
- *Optimizing data.* Data management is often the most limiting aspect of current simulations, so explicit data requirements and data flow analyses will help to optimize data flow through the complex memory hierarchy.

For climate science, advantages include:

- *Greater modularity and extensibility.* Adding new processes would involve defining the task and the data dependencies and is less dependent on the overall model framework.
- *Varying complexity.* Within a single model, sub-models of varying complexity could be deployed for different time or space domains. For example, deploying active biogeochemical (BGC) models only in selected regions or changing the complexity of those models based on other criterion. Parameterizing processes in some regions while resolving them in others or deploying an IAV model within a particular region would also be possible. Tasks would simply be launched based on data availability and other criteria and scheduled and load balanced by the run-time system.
- *Time scales.* Because tasks are started and stopped based on analyses of data flow, tasks with different time scales could be launched and run concurrently with shorter time tasks distributed and run while waiting for data from the longer timescale tasks.
- *Simpler programming model.* A task-based programming model would rely more on the task run-time system to distribute work; developers would not need to understand details of a domain decomposition or message passing to contribute to a model.

The above description is obviously an optimistic view and considerations like numerical and physical consistency with other algorithms or parameterizations will impact our ability to define tasks completely independently. Also, a naïve implementation may result in excessive data motion. However, a careful and staged implementation will allow us to explore the effectiveness of this approach.

Connections to Math, Comp Sci & and Climate Science

The methods described above rely on a strong partnership between climate domain scientists and computer scientists. New programming models for task parallelism, advanced memory management and dataflow analysis will be needed from the computer science community, informed by climate science domain expertise in identifying appropriate tasks, dependencies and algorithmic requirements.

Potential Impact on the field

The approach described above should enable more effective use of advanced architectures, enable greater modularity and component design of climate models and allow climate scientist to better integrate models of varying complexity to meet future needs.

References

- [1] George Bosilca, A. Bouteiller, A. Danalis, M. Faverge, T. Herault, J. Dongarra. *PaRSEC: A programming paradigm exploiting heterogeneity for enhancing scalability*, IEEE Computing in Science and Engineering, September 2013.
- [2] M. Bauer, S. Treichler and E. Slaughter and A. Aiken. Legion: Expressing Locality and Independence with Logical Regions, Proceedings of the International Conference on High-performance Computing, Networking, Storage and Analysis, IEEE Computer Society Press, p. 66, 2012.

Contributed ideas paper 29

Realistic Subglacial Hydrology For Improved Ice Sheet-Climate Coupling and Sea Level Prediction

December 1, 2015

Matthew Hoffman

Los Alamos National Laboratory

Luca Bertagna (Florida State University), Max Gunzburger (Florida State University)

Mauro Perego (Sandia National Laboratories), Stephen Price (Los Alamos National Laboratory)

Abstract

Mass loss from ice sheets causes about one-third of current sea level rise and is becoming the dominant contributor [1]. Ice sheet mass loss is the largest uncertainty in sea level change projections; current models are unable to predict rapid changes in ice dynamics in order to put an upper bound on future rates of mass loss [1]–[3]. The sliding of ice sheets over their bed is the primary control on ice flux to the oceans and decades of research have demonstrated that sliding is controlled by the state of a highly dynamic subglacial hydrologic system [4]–[6]. But due to the difficulties involved in modeling this system, the problem of simulating realistic sliding speeds is generally reduced to a physically unrealistic, grid-cell specific, scalar parameter optimization, with the primary drawback that modeled sliding cannot evolve in a realistic way [7]–[10]. Because of these gross over simplifications, ice sheet modelers have emphasized *evolutionary* subglacial hydrology and sliding models as critical missing pieces of current ice sheet models [11]–[13]. We propose to fill this void by adding physically-motivated subglacial hydrology and sliding models to existing models of ice sheet dynamics, for improved sea level projections from Earth System Models (ESMs). The parameters of the hydrology model can be determined via optimization techniques, as with the currently used basal traction parameter, but with the benefit of much reduced parameter dimensionality. In addition to adding more realistic basal physics to ice sheet models, subglacial hydrology models would lead to tighter coupling between ice sheet models and ESMs, through the input of surface melt and the output of subglacial water discharge to the ocean, closing the hydrologic cycle of ice sheets.

Background/Research to Date

Critically, the current generation of large-scale ice sheet models used for predicting sea level change within a coupled, ESM framework completely ignores ice sheet basal processes. Although sliding of the ice over its bed (as opposed to internal deformation) is the primary control on the flux of ice to the oceans, ice sheet models typically tune a “basal traction parameter” at each grid cell to match satellite observations of ice speed and hold those values fixed through future scenarios [7]–[10] (Figure 1). This locks in place a key control on ice sheet evolution and the biases inherited as a result of being the primary tuning knob of the model [7], [9].

Recent years have witnessed a renaissance in standalone process-scale models of subglacial hydrology ([5], [14]–[17]), with the unification of models for inefficient, distributed drainage (for which increasing water flux leads to increased sliding through increased water pressure) and efficient channelized drainage (for which increasing water flux has little or no effect on sliding). Such models have begun to be applied to large-scale ice sheet models capable of running within ESMs [15], [18], but are currently missing physics and, more importantly, the robust numerical and computational implementations required for practical use within large climate simulations.

Proposed Direction of work

Future work would build state of the art subglacial hydrology models into ice sheet models. These models would include the physics of, and automated switching between, inefficient distributed drainage and efficient channelized drainage [16], would be robust across a wide range of forcing [18], and would be flexible at handling disparate spatial and temporal scales between

Contributed ideas paper 29

subglacial drainage and ice flow.

The modeled subglacial drainage system would receive input from surface melt calculated by the ESM's surface energy balance model, and would in turn discharge water to the model ocean or proglacial land components. This would close the modeled ice sheet hydrologic budget and improve the representation of freshwater forcing from the ice sheets to the oceans, which has important impacts on ocean circulation and feedbacks on coupled ice sheet and ocean dynamics (e.g. [19]), both of which are currently only crudely represented in ESMs.

We also propose to perform an estimation of the parameters of the subglacial hydrology model by using a formal optimization strategy to minimize the mismatch with observations (e.g. observed surface ice velocity). The parameters to be estimated (e.g., bed roughness) can be uniform or space dependent, but will not vary at every grid cell as the currently used basal traction parameter does. However, it is reasonable to consider these parameters constant in time for the next few centuries. This is a big improvement with respect to most of the current approaches where parameters like the basal traction coefficient are assumed constant despite observational evidence to the contrary. Quantification of the uncertainty associated to some key parameters is also important for determining the level of confidence in the model predictions.

Connections to Math, Comp Sci & and Climate Science

This work would build off of recent ice sheet modeling advances bringing higher-order, robust, scalable velocity solvers to ice sheet models (through DOE's ISICLES, PISCEES, and ACME projects), as well as physically-motivated basal friction laws that are a function of basal water pressure. In addition to providing the physics of ice motion necessary for coupling to subglacial hydrology models, these past advances have led to mature, robust solver methods (via DOE-developed libraries like *Trilinos* and *Albany*) for the partial differential equations of ice motion. These equations are similar in form to the equations for subglacial water drainage, which would allow us to leverage the successes of past DOE investments. Successful application of subglacial hydrology models will also require parameter estimation and uncertainty quantification techniques. This need will also leverage recent and ongoing work under the PISCEES project and will make use of DOE developed tools like the Trilinos package ROL for optimization and DAKOTA for uncertainty quantification.

Potential Impact on the field

The addition of physically based basal process models for ice sheets would fill one of the most critical caps in current ice sheet models used for sea level change predictions. First, this would add realism to predictions of ice sheet evolution. For the Greenland Ice Sheet, this would allow ice sheet models to explicitly resolve the summer speedup experienced by the ice sheet each summer due to surface meltwater draining to the bed, and which has effects into each subsequent winter. Although the Antarctic Ice Sheet does not currently experience drainage of surface meltwater to the bed, the critical importance of subglacial hydrology in Antarctica has become increasingly evident, with hydrologic-regime switching capable of changing sectors of the ice sheet from negative to positive sources of sea level change [20]. Second, the addition of subglacial hydrology will improve the coupling between ice sheets ESMs. It will allow for a more natural spin-up and initialization of the ice sheet model, and it would close the hydrologic budget within ESMs.

Contributed ideas paper 29

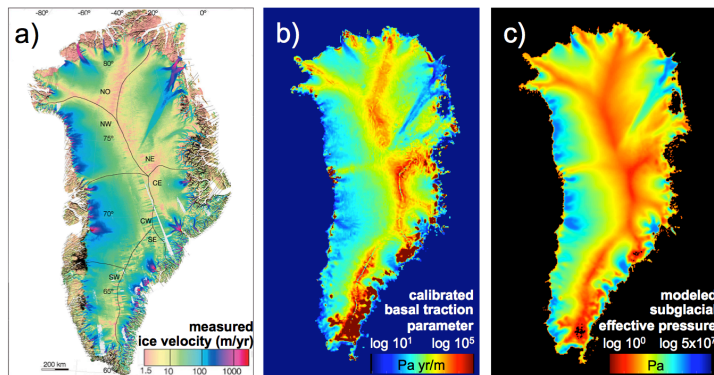


Figure 1: a) Winter ice velocity for the Greenland Ice Sheet measured from satellites[21], [22]. b) The basal traction parameter currently used in most ice sheet models must be **varied at each grid cell** to match observed velocity. c) Preliminary results from subglacial hydrology model with **uniform** parameter values recover spatial patterns in velocity, showing skill of added physics.

References

- [1] J. A. Church, P. U. Clark, A. Cazenave, J. M. Gregory, S. Jevrejeva, A. Levermann, M. A. Merrifield, G. A. Milne, R. S. Nerem, P. D. Nunn, A. J. Payne, W. T. Pfeffer, D. Stammer, and A. S. Unnikrishnan, “Sea Level Change,” in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, T. F. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, and P. M. Midgley, Eds. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2013, pp. 1137–1216.
- [2] J. Walsh, D. Wuebbles, K. Hayhoe, J. Kossin, K. Kunkel, G. Stephens, P. Thorne, R. Vose, M. Wehner, J. Willis, D. Anderson, S. Doney, R. Feely, P. Hennon, V. Kharin, T. Knutson, F. Landerer, T. Lenton, J. Kennedy, and R. Somerville, “Chapter 2: Our Changing Climate,” in *Climate Change Impacts in the United States: The Third National Climate Assessment*, no. October, J. M. Melillo, T. (T. C. . Richmond, and G. W. Yohe, Eds. U.S. Global Change Research Program, 2014, pp. 19–67.
- [3] W. T. Pfeffer, J. T. Harper, and S. O’Neel, “Kinematic constraints on glacier contributions to 21st-century sea-level rise,” *Science (80-.)*, vol. 321, no. 5894, pp. 1340–1343, 2008.
- [4] K. Cuffey and Paterson, *The Physics of Glaciers*, 4th Editio. Amsterdam: Butterworth-Heinneman, 2010.
- [5] G. E. Flowers, “Modelling water flow under glaciers and ice sheets,” *Proc. R. Soc. A Math. Phys. Eng. Sci.*, vol. 471, p. 20140907, 2015.

Contributed ideas paper 29

- [6] G. K. C. Clarke, “Subglacial Processes,” *Annu. Rev. Earth Planet. Sci.*, vol. 33, no. 1, pp. 247–276, May 2005.
- [7] T. L. Edwards, X. Fettweis, O. Gagliardini, F. Gillet-Chaulet, H. Goelzer, J. M. Gregory, M. Hoffman, P. Huybrechts, A. J. Payne, M. Perego, S. Price, A. Quiquet, and C. Ritz, “Effect of uncertainty in surface mass balance-elevation feedback on projections of the future sea level contribution of the Greenland ice sheet,” *Cryosph.*, vol. 8, pp. 195–208, 2014.
- [8] S. R. Shannon, A. J. Payne, I. D. Bartholomew, M. R. van den Broeke, T. L. Edwards, X. Fettweis, O. Gagliardini, F. Gillet-Chaulet, H. Goelzer, M. J. Hoffman, P. Huybrechts, D. W. F. Mair, P. W. Nienow, M. Perego, S. F. Price, C. J. P. P. Smeets, A. J. Sole, R. S. W. van de Wal, and T. Zwinger, “Enhanced basal lubrication and the contribution of the Greenland ice sheet to future sea-level rise,” *Proc. Natl. Acad. Sci. U. S. A.*, vol. 110, no. 35, pp. 14156–14161, 2013.
- [9] M. Perego, S. Price, and G. Stadler, “Optimal initial conditions for coupling ice sheet models to Earth system models,” *J. Geophys. Res. Earth Surf.*, vol. 119, pp. 1–24, 2014.
- [10] R. A. Bindschadler, S. Nowicki, A. Abe-Ouchi, A. Aschwanden, H. Choi, J. Fastook, G. Granzow, R. Greve, G. Gutowski, U. Herzfeld, C. Jackson, J. Johnson, C. Khroulev, A. Levermann, W. H. Lipscomb, M. A. Martin, M. Morlighem, B. R. Parizek, D. Pollard, S. F. Price, D. Ren, F. Saito, T. Sato, H. Seddik, H. Seroussi, K. Takahashi, R. Walker, and W. L. Wang, “Ice-sheet model sensitivities to environmental forcing and their use in projecting future sea level (the SeaRISE project),” *J. Glaciol.*, vol. 59, no. 214, pp. 195–224, 2013.
- [11] C. M. Little, M. Oppenheimer, R. B. Alley, V. Balaji, G. K. C. Clarke, T. L. Delworth, R. Hallberg, D. M. Holland, C. L. Hulbe, S. Jacobs, J. V. Johnson, H. Levy, W. H. Lipscomb, S. J. Marshall, B. R. Parizek, A. J. Payne, G. A. Schmidt, R. J. Stouffer, D. G. Vaughan, and M. Winton, “Toward a new generation of ice sheet models,” *Eos, Trans. Am. Geophys. Union*, vol. 88, no. 52, pp. 578–579, 2007.
- [12] S. Price, G. Flowers, and C. Schoof, “Improving Hydrology in Land Ice Models,” *Eos (Washington. DC)*, vol. 92, no. 19, p. 164, 2011.
- [13] W. Lipscomb, R. Bindschadler, S. Price, E. Bueler, J. Johnson, and D. Holland, “A community ice sheet model for sea level prediction,” *Eos (Washington. DC)*, vol. 90, no. 3, p. 23, 2009.
- [14] C. Schoof, “Ice-sheet acceleration driven by melt supply variability,” *Nature*, vol. 468, no. 7325, pp. 803–6, Dec. 2010.
- [15] M. Hoffman and S. Price, “Feedbacks between coupled subglacial hydrology and glacier

Contributed ideas paper 29

- dynamics,” *J. Geophys. Res. Earth Surf.*, vol. 119, pp. 1–23, 2014.
- [16] M. a. Werder, I. J. Hewitt, C. G. Schoof, and G. E. Flowers, “Modeling channelized and distributed subglacial drainage in two dimensions,” *J. Geophys. Res. Earth Surf.*, vol. 118, no. March, p. n/a–n/a, Oct. 2013.
- [17] I. J. Hewitt, “Seasonal changes in ice sheet motion due to melt water lubrication,” *Earth Planet. Sci. Lett.*, vol. 371–372, pp. 16–25, 2013.
- [18] E. Bueler and W. Van Pelt, “Mass-conserving subglacial hydrology in the Parallel Ice Sheet Model version 0.6,” *Geosci. Model Dev.*, vol. 8, pp. 1613–1635, 2015.
- [19] F. Straneo, P. Heimbach, O. Sergienko, G. Hamilton, G. Catania, S. Griffies, R. Hallberg, A. Jenkins, I. Joughin, R. Motyka, W. T. Pfeffer, S. F. Price, E. Rignot, T. Scambos, M. Truffer, and A. Vieli, “Challenges to Understanding the Dynamic Response of Greenland’s Marine Terminating Glaciers to Oceanic and Atmospheric Forcing,” *Bull. Am. Meteorol. Soc.*, vol. 94, no. 8, pp. 1131–1144, 2013.
- [20] M. Bougamont, P. Christoffersen, S. F. Price, H. A. Fricker, and S. Tulaczyk, “Reactivation of Kamb Ice Stream tributaries triggers century-scale reorganization of Siple Coast ice flow in West Antarctica,” *Geophys. Res. Lett.*, no. 1, 2015.
- [21] E. Rignot and J. Mouginot, “Ice flow in Greenland for the International Polar Year 2008–2009,” *Geophys. Res. Lett.*, vol. 39, no. 11, p. L11501, Jun. 2012.
- [22] I. Joughin, B. E. Smith, I. M. Howat, T. Scambos, and T. Moon, “Greenland flow variability from ice-sheet-wide velocity mapping,” *J. Glaciol.*, vol. 56, no. 197, pp. 415–430, Aug. 2010.

Contributed ideas paper 30

Composing Scalable Nonlinear Solvers for Climate

Nathan Collier

Climate Change Science Institute, Oak Ridge National Laboratory

Matthew Knepley (Rice University)

Date of Submission

1 Dec 2015

Abstract

Scalable nonlinear solvers for partial differential equations (PDEs) are an essential component of climate simulations. Inside these solvers, we find that Newton-Krylov methods are the dominant algorithm because they are generally robust and may be built from existing linear solvers and preconditioners. However, Newton-Krylov methods have their drawbacks. First, the construction and solution of the required linearization causes memory bandwidth and synchronization bottlenecks which limit their scalability. Second, the robustness of Newton's method hinges on an initial guess which is 'close' to the solution. The effect of this limitation is frequently realized as a diminishing time step size in time-dependent, nonlinear PDEs which regularizes the problem yet also increases the frequency of the aforementioned bottlenecks. As earth system models (ESMs) evolve to include more complex and nonlinear physics and chemistry, we suggest that Newton-Krylov methods alone will not be adequate.

However, while Newton-Krylov methods are heavily utilized in simulation, a larger design space for nonlinear solvers exists to complement, improve, or even replace Newton's method. Only a small part of this space has yet been explored either experimentally or theoretically.

Background/Research to Date

A recent paper [1] describes a philosophy for composing nonlinear solvers based on two techniques: nonlinear composite combination and nonlinear preconditioning. The main idea is that while alternative nonlinear methods may have more severe limitations for use as a stand-alone solver, they can serve to precondition the nonlinear system, accelerating overall time to solution. The paper is authored by PETSc developers and also details the implementation of a large collection of nonlinear solvers (nonlinear Richardson, nonlinear GMRES, Newton-Krylov, Quasi-Newton, nonlinear CG) and preconditioning techniques based on decomposition (Gauss-Seidel-Newton, nonlinear Additive Schwarz, Full Approximation Scheme). Their approach enables experimentation with alternative algorithms at runtime, analogous to linear Krylov solvers and preconditioners. The paper demonstrates significant improvements when solving nonlinear elasticity, lid-driven cavity flow, and the p -Laplacian, broadly representing the areas of solid and fluid mechanics. However, the optimal method for a given problem will require experimentation and the development of the techniques themselves.

We find no commensurate capability in other frameworks. There is a subproject [2] as part of the FASTMath Institute which focuses on flexible nonlinear solvers for climate applications. However, the content outlined on this site is a small subset of the solvers that PETSc has at its disposal.

Proposed Direction of Work

We propose an effort to explore techniques of composing nonlinear solvers in climate applications. It is our belief that this will be critical in delivering robust and scalable solvers as we push process representation and resolution. We suggest that the effort should contain:

Contributed ideas paper 30

- An initial exploration phase where the composable family of methods in PETSc may be applied to simpler codes which represent a portion of the physics of a full climate model (for example, subsurface hydrology or atmospheric circulation). The nonlinearities in the example problems of [1] are significantly challenging, but there is no substitute for experimenting on the problem you need to solve. The goal in this phase would be to demonstrate nonlinear approaches which yield faster runtimes than the traditional Newton-Krylov nonlinear approach on a variety of subproblems of interest to the climate community.
- A second phase, contingent on a successful first phase, where the solvers are integrated into ESMs. While some software used in climate modeling in the DOE is based on PETSc, a large portion also utilizes Trilinos or no solver package at all. But this work should not be done if we cannot demonstrate applicability on individual pieces of the overall ESM.

Connections to Math, Comp Sci & and Climate Science

The proposed work is the essence of math and computational science. We are combining lesser-known algorithms in new ways based on a mathematical and computational understanding of the problem to improve run time and enable new science.

Potential Impact on the field

Earth system models are currently expanding in complexity, both in terms of process representation and grid resolution. We believe that this research will enable us to significantly improve runtime and our utilization of leadership computing facilities.

References

- [1] Peter R. Brune, Matthew G. Knepley, Barry F. Smith, and Xuemin Tu. *Composing Scalable Nonlinear Algebraic Solvers*, SIAM Review, 2015, Vol. 57, No. 4 : pp. 535-565
- [2] Andy Salinger, *Flexible Nonlinear Solvers for Climate Applications*, <http://www.fastmath-scidac.org/research/flexible-nonlinear-solvers-climate-applications.html>

Contributed ideas paper 31

Exascale-Ready Programming Models for Climate

1 December 2015

Paul Ullrich

University of California, Davis

Gabriele Jost (Intel Federal RND), Bryce Adelstein Lelbach, Hans Johansen (LBNL)

Abstract

It is widely acknowledged that most modern applications codes are ill-prepared for achieving performance and scalability on pre-exascale computing systems. Modern climate modeling systems are no exception, and in many cases scientific compromises have been made to achieve maximum performance. Unfortunately, part of this is due to the dominance of synchronous parallel programming paradigms, such as Communicating Sequential Processes (CSP) and Fork-Join Threading, typically implemented using the Message Passing Interface (MPI) and OpenMP, respectively. These programming models demonstrate satisfactory scalability for climate applications on current parallel systems, mostly due to static, well-balanced workloads. However, the shift towards multi-core architectures has necessitated the use of both models simultaneously – the “MPI+OpenMP” approach. While this has enabled performance improvements on current systems, the approach suffers from the programming challenges inherent in the compromises of using different programming models together. Additionally, future applications are expected to be more dynamic with a high demand for runtime-orchestrated load balancing, to deal with science-driven variations in work load, as well as system reliability and performance variability, which are expected to worsen in future exascale supercomputers.

To address these issues, we propose investigating parallel runtime systems implementing new, task-based programming models. Out of the available runtimes, we will focus on High Performance ParalleX (HPX) [1], Legion [2,3], Open Community Runtime (OCR) [4] and Kokkos [5]. These models typically are asynchronous, and message- or event-driven. Task parallelism, in combination with oversubscription, facilitates dynamic load balancing, latency hiding, and better fault tolerance. These systems abstract the underlying hardware and provide implicit mechanisms for tolerating performance variability and hardware faults. This paper proposes to explore, assess and compare alternative programming models for climate simulations, by developing a software framework that is adaptable to an asynchronous task-based infrastructure.

Background/Research to Date

Several approaches have been proposed for coping with the challenges of scaling applications to exascale. Expected complexities of large-scale systems include a vastly increased number of cores and heterogeneous memories which are possibly non-coherent. A successful transition to exascale performance must address the following critical issues: runtime-orchestrated asynchronous execution, dynamic load balancing, energy efficiency, and fault tolerance.

To address these issues, task-based programming is undergoing rapid and active research as a potential replacement for the MPI+OpenMP model. The task-based model breaks down the active workload into chunks of work, referred to as tasks. Parallelism is orchestrated by a scheduler which oversees the distribution of tasks to CPU cores, so as to achieve an optimum of load balance and data locality. By abstracting the notion of parallelism from the application developer, the task-based model allows for seamless integration with heterogeneous architectures and implicitly supports fault tolerance and load balancing.

Although the ecosystem of programming models is diverse, there are some key differences between many of the frontrunners that have the potential to dramatically impact performance and ease-of-implementation. In OCR the main programming abstractions are event driven tasks, data

Contributed ideas paper 31

blocks, and events. Events are used for synchronizing task execution. Producer tasks deliver a data block to an event. The data block is consumed by tasks dependent on this event. Models such as HPX support task suspension and communication during task execution. HPX is based on the concurrency model from the C++11 standard, and fully supports C++ data structures. A thorough investigation of the impacts of these design decisions on existing climate model codes is necessary to minimize future investment costs. Some preliminary efforts addressing this need are now underway with a block-structured finite element MPI-enabled C++ framework, Tempest [6], as part of the FastForward 2 initiative (<https://asc.llnl.gov/fastforward/>).

Proposed Direction of Work

The goal of this work would be to assess and compare several modern programming models in the context of global climate modeling. Since all of the above models naturally operate with C/C++, frameworks such as Tempest would be well-suited for supporting the comparison. Once an optimal candidate was selected from the available suite of programming models, this technology could then be adapted for use in other global climate modeling systems. Concerted efforts need to be made to adapt sections of legacy code (such as cloud and radiation parameterizations) to a task-driven model, with less reliance on MPI-segregated shared data.

Connections to Math, Comp Sci & and Climate Science

This project effectively bridges climate science and computer science by investigating how cutting-edge parallel computing technology can improve the performance of climate modeling systems on large-scale parallel computing systems. It further provides valuable feedback to the programming model developers on the capability of application developers to integrate their model into large application codes. It should also be able to demonstrate techniques for testing and validating codes as they evolve from MPI-only to task-driven, shared memory run-times.

Potential Impact on the Field

Supercomputing systems that operate at the exascale are expected to emerge in the next several years, and are an inevitability by the middle of the next decade. The design of climate models that can operate effectively on these architectures is necessary to reach the high model resolutions that are desired for the accurate resolution of meteorological features. In particular, parallel performance, heterogeneous architectures, and fault tolerance are three key issues that will require the adoption of new programming models going forward.

References

- [1] Hartmut Kaiser, Thomas Heller, Bryce Adelstein-Lelbach, Adrian Serio, Dietmar Fey (2014) “HPX – A Task Based Programming Model in a Global Address Space.” PGAS 2014: The 8th International Conference on Partitioned Global Address Space Programming Models.
- [2] Michael Bauer, Sean Treichler, Elliott Slaughter and Alex Aiken (2012). “Legion: Expressing Locality and Independence with Logical Regions.” *Proceedings of the international conference on high performance computing, networking, storage and analysis*. IEEE Computer Society Press, 2012.
- [3] Michael Bauer (2014) *Legion: Programming Distributed Heterogeneous Architectures with Logical Regions*. Diss. Stanford University.
- [4] Tim Mattson, Rob van der Wijngaart, Zoran Budimlic, Vincent Cave, Sanjay Chatterjee, Romain Cledat, Bala Sshasayee and Vivek Sarkar, *OCR the open community runtime interface, version 0.95*, Technical Report, 2015.
- [5] H. Carter Edwards and Daniel Sunderland (2012) “Kokkos Array performance-portable manycore programming model.” *Proceedings of the 2012 International Workshop on Programming Models and Applications for Multicores and Manycores*. ACM.
- [6] Paul A. Ullrich (2014) “A global finite-element shallow-water model supporting continuous and discontinuous elements.” *Geosci. Model Dev.*, Volume 7, pp. 5141-5182, doi: 10.5194/gmd-7-3017-2014.

Contributed ideas paper 32

Skynet for Climate (Computational) Science
Benjamin Mayer
Climate Change Science Institute - Oak Ridge National Laboratory

Dec 1, 2015

Abstract

There are two worlds that we can choose to live in: One of drudgery where we spend our time copying files and performing data transfers repeatedly while trying to squeeze in some original thought while the machines we are beholden to are offline being maintained. The alternative is a world of enlightenment where machines are providing answers in new real time and point out interesting features in the data to be pondered by the great Sages of Climate Wisdom.

Background/Research to Date

There are several key enabling technologies to bring to attention that will be brought together in the Direction of Work Section.

Leading research institutions still conduct model runs and data management manually. Others have been using advanced management techniques for sometime [GFDL], with others moving to automated or semi-automated process and data management techniques [NCAR] [ACME] [MAXPLANCK]. Automation of the process flow, the ability for a system to independently run a simulation while recording the output and provenance is a required prerequisite for the other technologies to work. In itself automated data management is only sufficient to allow more time to be spent on research as opposed to enabling scientists to be much more productive. The more advanced discovery techniques are discussed below.

Disciplines such as wet lab biology have built not only automated machines to carry out experiments devised by humans and are widely available as commercial products [AUTOWET{1,2,3}] but also fully closed loop systems [AUTOBIO]. A closed loop system is one that has domain knowledge, determines an area to further explore, designs an experiment, carries it out, evaluate the results and finally integrates the new knowledge into its system. In this case, all within the computer. This allows a combined logical processing system and physical system to discover new knowledge.

Ontologies are a key technology in the knowledge representation used in the above examples. They are used to represent relationships between concepts in a domain. There are several Climate Ontologies in active use from NASA Climate Model Data Service (CDS), University of Georgia as well as UK Government Department of Energy and Climate Change (DECC).

Another key technology is deciding which theories to explore. In Mathematics the techniques of Automated Theorem Proving systems have been in use for decades, with the combination of the two techniques being used for Computer Aided Mathematics. [MATH {1,2,3}]. We have seen similar technology successfully used in Biology [AUTOBIO].

Proposed Direction of work

Automation of Model Run and Ontology Integration

When a simulation with a new resolution or set of configuration parameters is run the simulation can run into the area of non-sense or not properly tuned. We currently use frequently generated standard diagnostics, such as AMWG, to have a human examine simulation properties, such as RESTOM (energy balance), to make sure the model is behaving properly and possibly make adjustments. In simple cases, such as RESTOM being too high or low, a system trained to examine trends can alert the user of the tendency of the system before it becomes an issue. Further applying an ontology and connection that to the simulation variables a system can be developed which is self tuning for a number of key quality metrics related to simulation spin up.

Contributed ideas paper 32

If we also integrate knowledge about numerical algorithm operations, we could have a system that does real time monitoring of the simulation for unwanted numerical effects. As a simple example of error modes that could be watched are wind speeds approaching the CFL condition. We know that this can crash the model if that bound is crossed. Having a system help monitor would be beneficial by preventing us from wasting compute resources and also reduced time spent computing a bad answer result in faster time to solution.

These are some simple examples of what is possible. In fact there are many aspects of a simulation we should be monitoring, but simply don't have the human resources to be checking constantly. Another possible extension is the yearly production or hero runs we do to evaluate model improvement. This typically consists of four simulations with High and Low resolution and Coupled and Data Ocean being compared against one another. This process consumes a great deal of time to carry out, and likely can be fully automated other than key configuration issues.

Automatic Experiment design

Building on the previous discussion of automation and the existing body of work in Theorem Exploration and computerized closed loop biology systems, let us ponder how this could be applied to Climate. Any experiment design that says "if x is changed there should be an impact y" can be generated and tested. If one provides the ontology with major events that had a climate impact, experiments comparing what the historical impact of the event was can be constructed. If the ability to evaluate the impact is also integrated the system can be closed loop allowing automated learning.

Paper Generation

A randomly generated paper [PAPER{1,2}] has been accepted to a peer reviewed journal. The random part of the paper is that it randomly combined concepts from a database of information. If we use real results and knowledge from a climate ontology, not only can we significant portions of or entire papers could be generated.

Potential Impact on the field

Currently there is a large amount of effort on the technical aspects of run and data management, even for what are seen to be highly automated system in the domain. We can see the ability to make further leaps forward in the involvement of people in the effort to turn the crank on simulations, allowing for more focus on experiment design and analysis of data.

Paper writing is another major sink of time, with being lead author on a paper being a significant investment in ones career. It is possible to reduce the amount of time spent constructing a paper Assistance with presentation and paper construction.

When this system is successful there would be increased demand for computational resources to continue to run simulations and analysis as well as the additional monitoring capabilities.

There are no inherent limitations of this work to the Climate domain, it can be applied to many computational science domains. Additionally we have seen that Biology was able to cross the virtual and reality boarder and was able to carry out physical experiments. While this would be difficult in Climate, there are other domains where this is more possible.

The idea presented here has been one where machines take a leading role in the exploration of knowledge using computational or physical resources. A much more likely outcome as much of the technology is still leading edge is somewhere along the spectrum of humans doing everything or computers doing everything [AID]. Computers and humans should each be used what they are good for, but in that situation if the process is called computer-aided human exploration of the domain or human-aided computer exploration of the domain is yet to be seen.

Contributed ideas paper 32

References

- [ACME] Accelerated Climate Model for Energy (ACME) Proposal, 2013
- [AID] The Computer-Aided Discovery of Scientific Knowledge, Pat Langley
- [AUTOBIO] King , R D , Rowland , J J , Aubrey , W , Liakata , M , Markham , M , Soldatova , L N , Whelan , K E , Clare , A J , Young , M , Sparkes , A C , Oliver , S & Pir , P 2009 , ' The Robot Scientist Adam ' Computer , vol 42 , no. 8 , pp. 46-54 . , 10.1109/MC.2009.270
- [AUTOWET1] <http://www.hudsonrobotics.com/synthetic-biology/>
- [AUTOWET2] <http://www.laboratoryequipment.com/articles/2012/12/streamlining-molecular-biology-through-automation>
- [AUTOWET3] <http://www.sciencemag.org/site/products/robotfinal.xhtml>
- [GFDL] The GFDL Flexible Modeling System Runtime Environment, Amy Langenhorst, NOAA TECH2004
- [MATH1] *AUTOMATHEO*, Workshop on Automated Mathematical Theory Exploration 2009
- [MATH2] Scaling Automated Theory Exploration, Chris Warburton and Ekaterina Komendantskaya
- [MATH3] Learning and Exploration in Automated Theorem Proving Moa Johansson
- [MAXPLANCK] End-to-end optimization potentials in HPC applications for NWP and Climate Research, Luis Kornbluh, 16th ECMWF Workshop on High Performance Computing in Meteorology
- [NCAR] Optimizing workflow for CESM, Sheri Mickelson, CESM Software Engineering Working Group Session 20th Annual CESM Workshop (2015)
- [NLP1]
- [NLP2] Ontology Population via NLP Techniques in Risk Management, Jawad Makki, Anne-Marie Alquier, and Violaine Prince, International Journal of Humanities and Social Sciences 3:3 2009 212
- [PAPER1] SCiGen – An Automated CS Paper Generation <https://pdos.csail.mit.edu/archive/scigen/>
- [PAPER2] *Towards the Simulation of E-commerce* by Herbert Schlangemann, 2008 International Conference on Computer Science and Software Engineering. A random automatically generated paper accepted *after* peer review.

Contributed ideas paper 33

Heavy Physics

David Randall

Department of Atmospheric Science, Colorado State University

December 1, 2015

Abstract

More detailed representations of physical processes can improve the realism of climate simulations. With current architectures, such models are computationally expensive. Concurrent execution of model physics and dynamics can drastically reduce the cost as measured by wall-clock time, but will require new software architectures.

Background

Any global circulation model (GCM) includes a “dynamical core” that solves the equations of fluid motion on the sphere, and a “physics package” that uses physical parameterizations to determine the sources and sinks of the model’s prognostic fields.

In most current models, the dynamics and the physics use comparable amounts of processing power. Exceptions are super-parameterized models [1], in which the physics requires two orders of magnitude more processing power than the dynamics. Over the past 15 years, a large body of work [2] has demonstrated that such “heavy” physics can significantly increase the realism of climate simulations.

The three-dimensional domain of a GCM can be thought of as a collection of vertical columns of grid points, hereafter “grid columns.” Domain decompositions typically keep all grid cells within each grid column on the same node. The dynamics involves interactions among neighboring grid columns. As a result, message passing between nodes is required for the dynamics, and for a given horizontal resolution there is an optimal number of nodes, N_D , that minimizes the wall-clock time needed for execution of the dynamics. Let T_D be the wall-clock time needed to execute the dynamics on N_D nodes.

In contrast, the physics involves interactions only within individual grid columns. For this reason, execution of the physics does not require any message-passing. The physics is embarrassingly parallel, and in principle it can efficiently use as many cores as there are grid columns. Let the optimal number of nodes that minimizes the wall-clock time needed for execution of the physics be N_p , and let the wall-clock time needed to execute the physics on N_p nodes be T_p . Both N_p and T_p are simply proportional to the number of grid columns.

A key point is that $N_D \ll N_p$. The current practice, however, is that the physics and dynamics execute on the same nodes. The number of nodes allocated is approximately N_D , because using more than N_D nodes does not reduce and can actually increase the wall-clock time needed for execution of the full model. Because $N_D \ll N_p$, the wall-clock time needed to run the physics on N_D nodes is much larger than T_p , and as a result the wall-clock time required for execution of the full model is much larger than $T_D + T_p$.

Proposed Direction of work

Much smaller wall-clock times can be achieved if the physics and dynamics run on different sets of nodes, using the optimal choices of N_D nodes for the dynamics and N_p nodes for the physics.

Message passing is then needed for communication between the physics and dynamics, but the message size scales with N_D rather than N_p , and the messages need to be sent only once (in each direction) per physics time step. Ideally the wall-clock time needed for the messages can be

Contributed ideas paper 33

masked by computation.

If the physics and dynamics run sequentially, as in current models, the wall-clock time needed for execution of the full model can, in principle, be made as small as $T_D + T_p$. If the physics and dynamics can be made to run *concurrently*, then the wall-clock time needed for execution of the full model can, in principle can be made as small as T_D , by allocating enough nodes to the physics to reduce its execution time to T_D (or less).

More expensive physics can then be accommodated by allocating more nodes for execution of the physics. Up to a point, this can be done without increasing the wall-clock time needed for execution of the model.

Even greater reductions in wall-clock time may be possible, because the physics is not monolithic. For example, the expensive radiation component of the physics could be run on one set of nodes while the rest of the physics runs concurrently on a different set. Additional message-passing would be required, however, and work is needed to determine whether or not additional efficiencies could actually be realized through this approach.

Connections to Math, Comp Sci & and Climate Science

A heavy-physics GCM would have a software architecture and time-differencing schemes quite different from those of existing models. The construction of such a model would require the participation of climate scientists to develop the physics, computer scientists to design the software architecture, and applied mathematicians to optimize the discretization methods.

Potential Impact on the field

A GCM based on the ideas outlined above could be endowed with highly realistic and detailed parameterizations or super-parameterizations, representing such processes as convection, turbulence, microphysics, and aerosols. It would be capable of efficiently using emerging parallel architectures.

References

- [1] Khairoutdinov, M. F., and D. A. Randall, 2001: A Cloud Resolving Model as a Cloud Parameterization in the NCAR Community Climate System Model: Preliminary Results. *Geophys. Res. Lett.*, **28**, 3617-3620.
- [2] <http://www.cmmap.org/research/pubs-mmf.html>

Contributed ideas paper 34

12/01/2015

Using Design of Experiments to evaluate Numerical Weather Prediction Codes.

Jeffrey A. Smith, Ph.D.

US Army Research Laboratory

Richard Penc, Ph.D. (US ARL through Oak Ridge Affiliated Universities)

Abstract: Numerical Weather Prediction (NWP) generally seeks to forecast weather conditions at some time in the future based on present and past observations using complex mathematical/physical model(s) along with advanced computational techniques. These forecasts can range from hours to months and in some cases years into the future. Spatially, these codes can simulate conditions from small regional microclimates to global climates. Regardless of their application, these codes are a complex assembly of partial differential equations, employing parameterizations that account for sub-grid effects, and which are solved numerically. Given the computational complexity of these codes, it becomes quite difficult to explore the simulation output space in a reasonably efficient manner; however, statistical design of experiments, a technique applied successfully in other areas to large scale simulation models, shows promise in assisting in a structured exploration of these NWP codes.

Background/Research to Date: From the early work of Bjerknes and Richardson [1] and Smagorinsky's [2] landmark paper on the development of a global circulation model, our ability to forecast weather and model the climate have grown increasingly complex. Earth system models with all their attendant assumptions and approximations, along with the computational challenges arising from, for example, ever finer modeling resolutions, interact in a number of ways to introduce even more uncertainty into the model output [3]. This complexity challenges those who rely on these models to support decision making [4] and to assess regional impacts [5, 6] despite ever increasing amounts of data [7]. On possible means of assessing the impact of these interactions and thereby the model results is with statistical experiment design (ED) applied to the simulations codes.

ED reaches back to the work of Fisher [8] who developed the basic techniques for agricultural science. Other researchers extended Fisher's work to industrial process understanding and control [9-11]. Recently, with advances in computational capability, researchers (for example [12]) have begun to apply ED to the study of complex simulation codes. In all, in almost 80 years of research, ED techniques have evolved into a robust and comprehensive collection of methodologies that allow rigorous experimentation in many complex systems far removed from Fisher's initial application; however, there is little evidence to suggest that researchers employ these techniques in the area of numerical weather modeling and climate prediction. Regardless of this absence of direct evidence, some have applied ED techniques to computational fluid dynamics codes [13-15] for engineering design. CFD codes share many of the same complexities exhibited by regional and global weather and climate models, a fact that suggests that ED may prove useful in weather forecasting and the analysis the attendant models.

Proposed Direction of work: The main thrust of this research will be to catalog various existing methods for creating design matrices (the method for creating a structured sampling of the model), applying ED techniques to regional weather models through these matrices, and assessing their strengths and weaknesses.

Connections to Math, Comp Sci & and Climate Science: Certain ED techniques, e.g., those based on nearly orthogonal arrays can be obtained via mixed integer linear programs, which trade-off the independence of orthogonal arrays for more variety in designs. In addition, the use of ED techniques to explore models, especially as they undergo development, can lead to insights that support the software verification process, for example, the use of covering arrays in analysis of highly configurable models [16].

Contributed ideas paper 34

Potential Impact on the field: Though not a direct contribution to climate science, ED techniques are a means of exploring the models that underlay the science. By probing these models in a structured manner rather than by making ‘one at a time changes’ or looking at very narrow cases, we can come to understand how the various assumptions, approximations and interactions within climate models contribute to output uncertainty. Through the understanding gained of these models, we can in turn determine where the science may be underdeveloped, or that methods used to model the effects are not as well developed [3]. By repeatedly applying ED to models as we develop them, we come to a better understanding of the science, and hence can more precisely articulate the strengths and weaknesses in the results. At any stage of the model exploration, understanding these uncertainties will contribute to a better use of the predictions based upon these models for improving policy and planning.

References:

1. Lynch, P., *The Emergence of Numerical Weather Prediction : Richardson's Dream*. 2006, New York: Cambridge University Press.
2. Smagorinsky, J., *General Circulation Experiments with the Primitive Equations*. Monthly Weather Review, 1963. **91**(3): p. 99-164. doi: 10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2
3. Stevens, B. and S. Bony, *What Are Climate Models Missing?* Science, 2013. **340**(6136): p. 1053-1054. doi: 10.1126/science.123755
4. Vecchi, G.A. and G. Villarini, *Next Season's Hurricanes*. Science, 2014. **343**(6171): p. 618-619. doi: 10.1126/science.1247759
5. Hall, A., *Projecting Regional Change*. Science, 2014. **346**(6216): p. 1461-1462. doi: 10.1126/science.aaa0629
6. Schindler, D.E. and R. Hilborn, *Prediction, Precaution, and Policy under Global Change*. Science, 2015. **347**(6225): p. 953-954. doi: 10.1126/science.1261824
7. Overpeck, J.T., et al., *Climate Data Challenges in the 21st Century*. Science, 2011. **331**(6018): p. 700-702. doi: 10.1126/science.1197869
8. Fisher, R.A., *The Design of Experiments*. 1935, Edinburgh, London,: Oliver and Boyde.
9. Box, G.E.P., W.G. Hunter, and J.S. Hunter, *Statistics for Experimenters: An Introduction to Design, Data Analysis, and Model Building*. 1978, New York: Wiley.
10. Montgomery, D.C., *Design and Analysis of Experiments*. 7th ed. 2008, Hoboken, NJ: Wiley.
11. Deming, W.E., *Out of the Crisis*. MIT Press ed. 2000, Cambridge, Mass.: MIT Press.
12. Kleijnen, J.P.C., *Design and Analysis of Simulation Experiments*. International Series in Operations Research & Management Science. 2008, New York: Springer.
13. Berci, M., et al., *Multidisciplinary Multifidelity Optimisation of a Flexible Wing Aerofoil with Reference to a Small Uav*. Structural and Multidisciplinary Optimization, 2014. **50**(4): p. 683-699. doi: 10.1007/s00158-014-1066-2
14. Zhu, B., et al., *Optimization Design of a Reversible Pump-Turbine Runner with High Efficiency and Stability*. Renewable Energy, 2015. **81**: p. 366-376. doi: 10.1016/j.renene.2015.03.050
15. Rahimi, A., T. Tavakoli, and S. Zahiri, *Computational Fluid Dynamics (Cfd) Modeling of Gaseous Pollutants Dispersion in Low Wind Speed Condition: Isfahan Refinery, a Case Study*. Petroleum Science and Technology, 2014. **32**(11): p. 1318-1326. doi: 10.1080/10916466.2011.653701
16. Yilmaz, C., M.B. Cohen, and A.A. Porter, *Covering Arrays for Efficient Fault Characterization in Complex Configuration Spaces*. IEEE Transactions on Software Engineering, 2006. **32**(1): p. 20-34. doi: 10.1109/TSE.2006.8

Contributed ideas paper 35

Optimization and Uncertainty Quantification of Ice Sheet Models

S. F. Price (Los Alamos National Laboratory), M. Perego (SNL), G. Stadler (NYU)

Abstract

Ice sheets are important components of the global climate system but their incorporation as components of coupled Earth System Models (ESMs) presents a number of challenges not yet routinely tackled within the ESM community. A particular set of challenges requiring continued advances in both mathematical and computational methods is the calculation of ice sheet model (ISM) initial conditions that are consistent with observations and can thus be used for prognostic simulations. This requires high-dimensional parameter field optimization and the systematic quantification of uncertainties due to missing information, observational error, and model error.

Background and Research to Date

As Earth’s largest reservoirs of freshwater, ice sheets are important components of the climate system. Their evolution has implications for changes in global sea level, global water mass formation, and ocean circulation. Over the last decade, significant efforts have gone into improving the fidelity, accuracy, and efficiency of ice sheet models (ISMs), including efforts to incorporate them as components of coupled ESMs. Unlike other components of the climate system, ice sheets have equilibrium timescales of 10^3 – 10^5 years and ISMs cannot be easily “spun up” for such a duration (particularly when coupled to ESMs) to remove undesirable transients associated with non-equilibrium initial conditions (e.g., [2,5]). Moreover, in the absence of additional constraints, such a spin up¹ may lead to an ISM state (geometry and velocity) that is a poor representation of the present-day observed state and cannot be used for predictive simulations.

One solution to this problem is the application of formal optimization strategies to minimize an objective function, which targets ice sheet model initial conditions that (i) match today’s observations and (ii) seek an equilibrium (or transient state) that accounts for present-day (or past and present) climate forcing. To date, numerous efforts have targeted one or both of these goals, resulting in advancements in ISM initialization (e.g., [1, 2, 5, 6, 8–13, 15]) and in assigning and propagating uncertainty in model parameters through to uncertainties in model outputs, [7]. Importantly, all approaches successful at targeting (i) and/or (ii) apply methods that take advantage of derivative information—computed efficiently using adjoints—of the forward model. Indeed, the use of first and second derivatives (gradients and Hessians, respectively) is essential for making optimization and uncertainty quantification feasible, particularly in high-dimensional parameter spaces. Despite significant progress made in these fields of research, much work remains to be done in order for ISMs (and other climate model components) to take full advantage of high-dimensional parameter optimization and to make high-dimensional uncertainty quantification (or reasonable approximations of the uncertainty) feasible.

Proposed Direction of work

Specific goals for computing optimal initial conditions include:² (i) Optimization of two and three-dimensional parameter fields of size similar to that of model state variable fields (e.g., the

¹For example, by running the ISM for the duration of a glacial-interglacial cycle.

²Examples are given in the context of ice sheet modeling, but the requirements and methods are applicable to any forward model with uncertain, high-dimensional parameter fields as inputs.

Contributed ideas paper 35

unknown basal friction coefficient field or the geothermal heat flux fields that are part of ISM basal boundary conditions), which will require improved preconditioners for the optimization problem. (ii) Optimization of model state variables (e.g., ice thickness, a prognostic variable), which may be constrained by observations. This requires the development of efficient methods for computing derivatives with respect to state variables. (iii) Optimization of transient models that are constrained by temporal observations over relevant time scales (e.g., ice sheet internal temperatures and layer shapes). This requires the development of efficient techniques for time-dependent data assimilation for high-dimensional observation and model parameter fields.

Optimization is generally the first step of Uncertainty Quantification (UQ). Moreover, it is possible to efficiently compute a Gaussian approximation to the distribution of the parameters using the connection between the covariance matrix of the parameter distribution and the inverse of the Hessian matrix [3, 4, 14]. A particular example of interest (and societal importance) is the quantification of uncertainty associated with predictions of sea level rise from ice sheets in the coming decades and centuries. Although the quantity of interest is low dimensional³, its uncertainty depends on the uncertainty in the model initial conditions, which is typically of high (order 10^6) dimension. In principle, the distribution of initial conditions has to be explored through sampling, which is currently intractable unless appropriate simplifying approximations are made. Approaches we propose to address this high-dimensional UQ problem include reduced-order modeling, a hierarchy of forward models (from cheap and inaccurate to expensive and accurate), a hierarchy of grid resolutions, and Hessian-enabled MCMC sampling.

Connections to Math, Comp Sci and Climate Science

Achieving the goals proposed in this work requires (extremely) efficient, robust, and scalable forward and adjoint model solves. Although many software packages offer robust forward and adjoint solvers, few of them can address problems of the size proposed here (order 10^6 - 10^8 parameters and 10^8 unknowns) while also taking advantage of emerging architectures. The solution of an optimization problem can easily require 100 to 1000 nonlinear forward model solves and these costs are greatly surpassed by those of Bayesian inversion or the forward propagation of an uncertain parameter distribution through to a prediction quantity of interest.

The proposed work would leverage past DOE investments in ice sheet and climate modeling (the ISICLES, PISCEES, and ACME projects) and in computational mathematics and uncertainty quantification (SciDAC's FASTMath, SUPER, and QUEST institutes). The *Albany/Felix* code, developed under PISCEES and included in ACME as part of the Model for Prediction Across Scales (MPAS) Land Ice Model, already features a large-scale optimization capability through the *Trilinos* package ROL and (low to moderate-dimensional) UQ capability through *Dakota*. Given the “embarrassingly parallel” nature of the sampling phase, the solution of the UQ problems noted above would benefit from anticipated new supercomputer architectures.

Potential Impact on the field

Our goals of further developing strategies and tools for optimization of high-dimensional model parameter fields would also benefit other climate model components. The ambitious goal of making high-dimensional uncertainty quantification tractable or of finding tractable approximations represents an extraordinary mathematical and computational challenge, which would have a large impact on not only climate modeling but on many other scientific fields that rely on results from large-scale computer simulations.

³For instance, mass loss and sea level rise from a particular ice sheet at a certain point in the future.

Contributed ideas paper 35

References

- [1] Robert J. Arthern and G. Hilmar Gudmundsson. Initialization of ice-sheet forecasts viewed as an inverse robin problem. *Journal of Glaciology*, 56(197):527–533, 2010.
- [2] D. J. Brinkerhoff and J. V. Johnson. Data assimilation and prognostic whole ice sheet modelling with the variationally derived, higher order, open source, and fully parallel ice sheet model VarGlaS. *The Cryosphere*, 7(4):1161–1184, 2013.
- [3] Tan Bui-Thanh, Omar Ghattas, James Martin, and Georg Stadler. A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
- [4] Pearl H. Flath, Lucas C. Wilcox, Volkan Akçelik, Judy Hill, Bart van Bloemen Waanders, and Omar Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations. *SIAM Journal on Scientific Computing*, 33(1):407–432, 2011.
- [5] F. Gillet-Chaulet, O. Gagliardini, H. Seddik, M. Nodet, G. Durand, C. Ritz, T. Zwinger, R. Greve, and D. G. Vaughan. Greenland ice sheet contribution to sea-level rise from a new-generation ice-sheet model. *The Cryosphere*, 6(6):1561–1576, 2012.
- [6] D. N. Goldberg and P. Heimbach. Parameter and state estimation with a time-dependent adjoint marine ice sheet model. *The Cryosphere*, 7(6):1659–1678, 2013.
- [7] T. Isaac, N. Petra, G. Stadler, and O. Ghattas. Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet. *J. of Comput. Phys.*, 296:348–368, 2015.
- [8] I. Joughin, B. E. Smith, I. M. Howat, D. Floricioiu, R. B. Alley, M. Truffer, and M. Fahnestock. Seasonal to decadal scale variations in the surface velocity of Jakobshavn Isbrae, Greenland: Observation and model-based analysis. *Journal of Geophysical Research*, 117(F2):F02030, May 2012.
- [9] D. MacAyeal. The basal stress distribution of Ice Stream E, Antarctica, inferred by control methods. *Journal of Geophysical Research*, 97(B1):595–603, 1992.
- [10] Laurent Michel, Marco Picasso, Daniel Farinotti, Martin Funk, and Heinz Blatter. Estimating the ice thickness of shallow glaciers from surface topography and mass-balance data with a shape optimization algorithm. *Computers & Geosciences*, 66(0):182 – 199, 2014.
- [11] M. Morlighem, H. Seroussi, E. Larour, and E. Rignot. Inversion of basal friction in Antarctica using exact and incomplete adjoints of a higher-order model. *Journal of Geophysical Research: Earth Surface*, 118(3):1746–1753, 2013.
- [12] M. Perego, S. Price, and G. Stadler. Optimal initial conditions for coupling ice sheet models to earth system models. *J. Geophys. Res.*, 119:1894–1917, 2014.
- [13] N. Petra, H. Zhu, G. Stadler, T. J. R. Hughes, and O. Ghattas. An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model. *Journal of Glaciology*, 58(211):889–903, 2012.

Contributed ideas paper 35

- [14] Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, Philadelphia, PA, 2005.
- [15] A. Vieli, A. J. Payne, Z. Du, and A. Shepherd. Numerical modelling and data assimilation of the Larsen B ice shelf, Antarctic Peninsula. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 364(1844):1815–1839, July 2006.

Contributed ideas paper 36

Climate Model Uncertainty Quantification

Nathan M. Urban

Los Alamos National Laboratory

December 1, 2015

Abstract

Two challenges in climate uncertainty quantification are understanding structural uncertainties in numerical models, and exploring uncertainties using computationally intensive simulations. It may be possible to address both challenges simultaneously by developing emulators of complex models, and synthesizing information across a model hierarchy.

Background/Research to Date

Uncertainty in climate model projections comes from many sources, including uncertainty in initial state (chaos), exogenous forcings or boundary conditions, model parameters, and model structure (choices in numerics, process simplifications, representation of unresolved scales, etc.). To date, much research has focused on initial-value and perturbed-parameter uncertainties, because these are easier to quantify by running ensembles of the same simulation code.

Many uncertainties in climate models are “structural” in nature, such as parameterizations of unresolved sub-grid scale processes, and can be as important as other forms of uncertainty. These structural uncertainties are often *functional* rather than *parametric*, meaning that the uncertainty is in the shape of a function of the state variables (e.g. a closure approximation in the discretized equations), rather than in the value of a scalar parameter. Functional uncertainties are difficult to explore in an automated manner, compared to parametric uncertainties, since functions are high (technically infinite-) dimensional, and the functional form is usually expressed only in source code, a representation not amenable to mathematical manipulation.

Climate models are moving towards exascale computing on DOE leadership-class facilities. Uncertainty quantification typically requires ensembles of simulations to explore the space of uncertainties. While ensemble simulations are often embarrassingly parallel and can exploit unused cores on large systems, it will typically not be possible to run large ensembles of the highest resolution simulations, as the model resolution typically is increased to fill the machine as much as possible in order to gain the greatest system fidelity. Even running dozens of fully-coupled Earth system simulations is a difficult endeavor, and cannot be expected to adequately explore the full space of model uncertainties, particularly when high dimensional structural uncertainties are taken into account.

Proposed Direction of Work

If it is not computationally feasible to sample a large uncertainty space, options include (1) sampling more efficiently, (2) replacing the slow model with a fast emulator, or (3) running a less expensive model. Recent advances in derivative-based and dimension-independent samplers suggest paths to significantly increasing sampling efficiency [1, 12, 14], but may require intrusive methods like adjoint models that are difficult to construct for large-scale codes, and remain fundamentally unlikely to explore very high dimensional spaces with fewer than $O(100)$ simulations. It may be necessary to combine them with other approaches when studying the fully coupled Earth system. Other recent work in probabilistic numerical solvers propagates

Contributed ideas paper 36

time discretization uncertainty throughout a numerical integration [2]; it may be possible to extend this to other forms of uncertainty such as structural uncertainty.

A second option is to efficiently emulate the behavior of the complex model, and then use the emulator to sample uncertainties. This has been done for parametric uncertainties in climate models using response surface techniques. However, parameters in mechanistic emulators have also been tuned to mimic structurally different climate models. This approach can be used to turn a difficult structural uncertainty problem within numerical models into a better-understood parametric uncertainty problem within emulators. The challenge is to identify adequate emulators and the relevant degrees of freedom for capturing structural differences between model dynamics (e.g., feedback factors, response timescales, spatiotemporal covariance structures).

There are many classes of such emulators, including statistical interpolators [3], mechanistic physics models [4], semi-empirical dynamical models [5], state space models [6], dynamical emulators [7], neural networks [8], modal decompositions [9], etc. Intrusive methods for model reduction are common in engineering disciplines, but the high complexity of Earth system models makes intrusive methods difficult to apply. However, it may be possible to combine system identification and intrusive methods [16]. In exascale systems, even non-intrusive methods may need to be applied *in situ* rather than in post-processing, learning a compressed representation of the dynamics online. Nonlinear dimension reduction methods, such as manifold learning [10] or autoencoders [11], may facilitate the automatic generation of reduced system descriptions.

A third option is to run a simpler model, which sacrifices accuracy for a better grasp of uncertainty; this falls under the class of mechanistic emulators above. However, rather than simply tuning the simple model to reproduce a more complex one, it may be possible to exploit the existence of a hierarchy of models of increasing complexity. A large ensemble of simple models may be compared to a smaller ensemble of complex models in order to predict the likely behavior of the complex model in new regions of uncertainty space, including predicting the behavior of new complex model structures from structurally different simple models. In a limiting case, the structure of a partial differential equation solver may be captured by the stencil dynamics evolving the state of an arbitrary grid cell. Parameterizing this local evolution operator using low-rank functional approximation methods may provide a compact mathematical representation of a model structure, which can be related to the behavior of its global solutions.

Connections to Math, Computer Science, and Climate Science

Climate uncertainty quantification for advanced climate models can benefit from scalable high-dimensional statistical inference methods, Monte Carlo samplers, and machine learning techniques. The applied mathematics community also provides relevant methods for dimension and model reduction. Many recent advances have exploited closer relationships between statistics and numerical methods [12, 13, 14, 15, 2, 16].

Potential Impact on the Field

Climate emulators which can be tuned to emulate different numerical models may provide a way to assign uncertainties to outcomes that are not captured by any individual model in the multi-model ensemble. Along with machine learning approaches, emulators may provide efficient ways to tune expensive models, perform sensitivity analysis, and understand what aspects of a model's structure control its large-scale behavior. The methods proposed here could provide a more automated means of exploring model structural uncertainties, breaking the combinatorial barrier of writing new code to explore uncertainty in each subcomponent of a numerical model.

Contributed ideas paper 36

References

- [1] M. Girolami and B. Calderhead, “Riemann manifold Langevin and Hamiltonian Monte Carlo methods”, *J. Roy. Stat. Soc.* **73**, 123 (2011).
- [2] P. Hennig, M.A. Osborne, and M. Girolami, “Probabilistic numerics and uncertainty in computations”, *Proc. Roy. Soc. A* **471**: 20150142 (2015).
- [3] V. Rao, R. Archibald, and K.J. Evans, “Emulation to simulate low-resolution atmospheric data”, *Int. J. Comp. Math.* **91**, 770 (2013).
- [4] M. Meinshausen, S.C.B. Raper, and T.M.L. Wigley, “Emulating coupled atmosphere-ocean and carbon cycle models with a simpler model, MAGICC6 — Part 1: Model description and calibration”, *Atmos. Chem. Phys.* **11**, 1417 (2011).
- [5] C. Penland and P.D. Sardeshmukh, “The optimal growth of tropical sea surface temperature anomalies”, *J. Climate* **8**, 1999 (1995).
- [6] D.W. Gladish and C.K. Wikle, “Physically motivated scale interaction parameterization in reduced rank quadratic nonlinear dynamic spatio-temporal models” *Environmetrics* **25**, 230 (2014).
- [7] S. Conti and A. O’Hagan, “Bayesian emulation of complex multi-output and dynamic computer models”, *J. Stat. Plan. Inference* **140**, 640 (2010).
- [8] A. Wu, W.W. Hsieh, and B. Tang, “Neural network forecasts of the tropical Pacific sea surface temperatures”, *Neural Networks* **19**, 145 (2006).
- [9] J.N. Kutz, X. Fu, and S.L. Brunton, “Multi-resolution dynamic mode decomposition”, e-print arXiv:math/1506.00564 (2015).
- [10] L.J.P. van der Maaten, E.O. Postma, and H.J. van den Herik, “Dimensionality reduction: A comparative review” *J. Machine Learning* (2007); D. Lungu, S. Prasad, M.M. Crawford, and O. Ersoy, “Manifold-learning-based feature extraction for classification of hyperspectral data; A review of advances in manifold learning”, *Signal Processing* **31**, 55 (2014).
- [11] G.E. Hinton and R.R. Salakhutdinov, “Reducing the dimensionality of data with neural networks” **313**, 504 (2006).
- [12] N. Petra, J. Martin, G. Stadler, and O. Ghattas, “A computational framework for infinite-dimensional Bayesian inverse problems, Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems”, *SIAM J. Sci. Comput.* **36**, A1525 (2014).
- [13] T. Isaac, N. Petra, G. Stadler, and O. Ghattas, “Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet” **296**, 348 (2015).
- [14] T. Cui, K.J.H. Law, and Y.M. Marzouk, “Dimension-independent likelihood-informed MCMC”, e-print arXiv:stat/1411.3688 (2015).
- [15] M. Parno, T. Moselhy, and Y. Marzouk, “A multiscale strategy for Bayesian inference using transport maps”, e-print arXiv:stat/1507.07024 (2015).
- [16] H. Chen, “Blackbox Stencil Interpolation Method for model reduction”, S.M. Thesis, MIT (2012).

Contributed ideas paper 37

Automation of parameterization and structure selection of ocean biogeochemical models

Date of Submission: 12/01/2015

Shanlin Wang

Los Alamos National Laboratory (LANL)

Nathan Urban (LANL), Mathew Maltrud (LANL), Alexandra Jonko (LANL)

Abstract

We propose to carry out a sensitivity analysis on parameters in an ocean biogeochemistry (BGC) model, to infer online parameterization from an offline single-column model, and to adaptively learn a minimum sufficient model complexity. The goal of this work is to quantify the oceanic uptake of CO₂ and to predict responses of marine productivity to climate change. We aim to determine a desired representation of BGC and optimal parameter setting for various resolutions of the hydrological model. The results will be applied to online, coupled climate simulations.

Background/Research to Date

Ocean biogeochemical models serve as crucial tools to quantitatively assess and predict the global carbon cycle and marine productivity in the era of unprecedented climate change. Several previous studies investigated the parameter sensitivity of marine ecosystem models, including BEST-NPDZ [1], HadOCC-NPDZ [2], and NEMURO [3], and identified important parameters. Kriest et al. [4] conducted sensitivity analysis of simple ocean biogeochemical models with 2, 3, and 6 state variables, and showed various sensitivities to different variables and model structures (Figure 1a). However, these studies are based on either a small set of observations or on simple model structures. Most Earth system models employ intermediate complexity ocean biogeochemistry models. For example, both CESM and ACME use the BEC (biogeochemical elemental cycling) model as the ocean ecosystem/biogeochemical module, which includes more than 100 parameters (schematic structure shown in Figure 1b). Yet, sensitivity analysis on biological parameters and structure complexity of such models is absent, which lowers the confidence in understanding climate feedbacks involving ocean biogeochemical cycle and predicting marine productivity and the carbon cycle. Furthermore, parameters in biogeochemical models are generally tuned based on expert judgment for different scientific applications and various resolutions of hydrological models that drive the transport of biogeochemical tracers. Information about systematic tuning and key metrics is not provided. This increases the difficulty for broader research community.

Proposed Direction of work

Tuning a biogeochemical model to realistic regional behavior is complicated by the following three factors. First, the real ocean has a far greater biological diversity than is represented in the model. Second, advection of the large number of tracers necessary adds to the model's computational expense. Finally, the unrepresented diversity is both highly heterogeneous within a grid cell and varies between different regions of the global ocean. This leads to spatial and resolution dependence of the model parameters, the tuning of which so far has not been addressed in a global BGC model.

We propose that these challenges can be addressed in the long term by a combination of several approaches:

1. Building a statistical model to predict the behavior of some unrepresented species from represented species (e.g. within the same functional group), in order to reduce the number of tracers needed.
2. Adaptively learning online a minimum sufficient representation, i.e. dynamically varying the complexity (# of represented species) of the model as needed to resolve the behavior of interest.
3. Treating global model tuning as an inference problem, by tuning the parameters of an offline

Contributed ideas paper 37

single-column version of the model, which are statistically related to the parameters of the coupled global model.

4. Representing sub-grid scale heterogeneity by a spatially averaged stochastic process to produce scale-dependent effective parameters within a coarser grid cell.

5. Parameterizing biological diversity informed by a high-complexity model. This involves running an ensemble of offline high-complexity BGC models, with hundreds of species, to learn the dependence of their realistic dynamics on the background physical climate state. A lower-complexity model is then fit to the high-complexity simulations, introducing a distribution of low-complexity parameter fits, conditional on the background climate state, which cause the low-complexity model to reproduce the ensemble of high-complexity behaviors. This distribution can be averaged to give location- (or more accurately, state-) dependent effective parameters, as in (4).

Connections to Math, Comp Sci & Climate Science

This study requires the development of statistical models that can be trained to automatically parameterize unrepresented processes/dynamics and sub-grid scale heterogeneity. The major limitation is the efficiency of ocean BGC models, which require decades to hundreds of model years to reach steady state and increase the hydrological computational cost by 2 - 10 times depending on the model resolution. We will use the fast Newton-Krylov solver [6] for parameter sensitivity tests. Our approach will also include the use of stochastic or random processes, sequential inference, and model emulation and reduction. The research will quantify the uncertainty in BGC models and provide insights into marine ecodynamics.

Potential Impact on the field

The goal of this work is to improve the ease of biogeochemical parameter setting for model users and facilitate the use of ACME by the broader climate research community. The knowledge gained and techniques established can be applied to other ocean and terrestrial biogeochemical models and significantly improve climate simulations. The sensitivity study will reveal potential tipping points in marine ecosystems and provide information commercial fishery managements.

References

- [1] Gibson, G. A., and Y. H. Spitz, 2011, Impacts of biological parameterization, initial conditions, and environmental forcing on parameter sensitivity and uncertainty in a marine ecosystem model for the Bering Sea, *Journal of Marine Systems*, 88(2), 214-231, doi:10.1016/j.jmarsys.2011.04.008.
- [2] Scott, V., Kettle, H., and Merchant, C. J., 2011, Sensitivity analysis of an ocean carbon cycle model in the North Atlantic: an investigation of parameters affecting the air-sea CO₂ flux, primary production and export of detritus, *Ocean Sci.*, 7, 405-419, doi:10.5194/os-7-405-2011.
- [3] Yoshie, N., Y. Yamanaka, K. A. Rose, D. L. Eslinger, D. M. Ware, and M. J. Kishi, 2007, Parameter sensitivity study of the NEMURO lower trophic level marine ecosystem model, *Ecological Modelling*, 202(1-2), 26-37, doi:10.1016/j.ecolmodel.2006.07.043
- [4] Kriest, I., A. Oschlies, and S. Khatiwala, 2012, Sensitivity analysis of simple global marine biogeochemical models, *Global Biogeochem. Cycles*, 26, GB2029, doi:10.1029/2011GB004072.
- [5] Doney, S.C., I. Lima, J.K. Moore, K. Lindsay, M.J. Behrenfeld, T.K. Westberry, N. Mahowald, D.M. Glover, and T. Takahashi, 2009, Skill metrics for confronting global upper ocean ecosystem-biogeochemistry models against field and remote sensing data, *J. Mar. Systems*, 76, 95-112, doi:10.1016/j.jmarsys.2008.05.015.
- [6] Li, X., and F. Primeau 2008, A fast Newton-Krylov solver for seasonally varying global ocean biogeochemistry models suitable for systematic parameter optimization, *Ocean Modell.*, 23, 13–20.

Contributed ideas paper 37

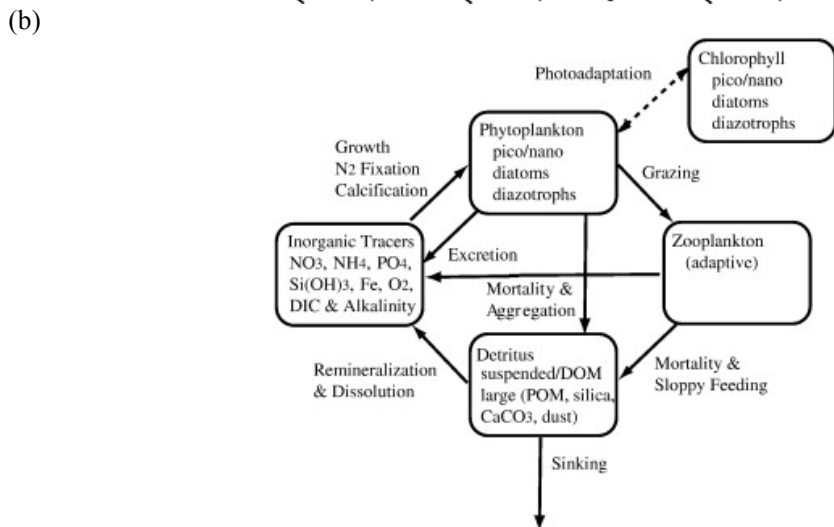
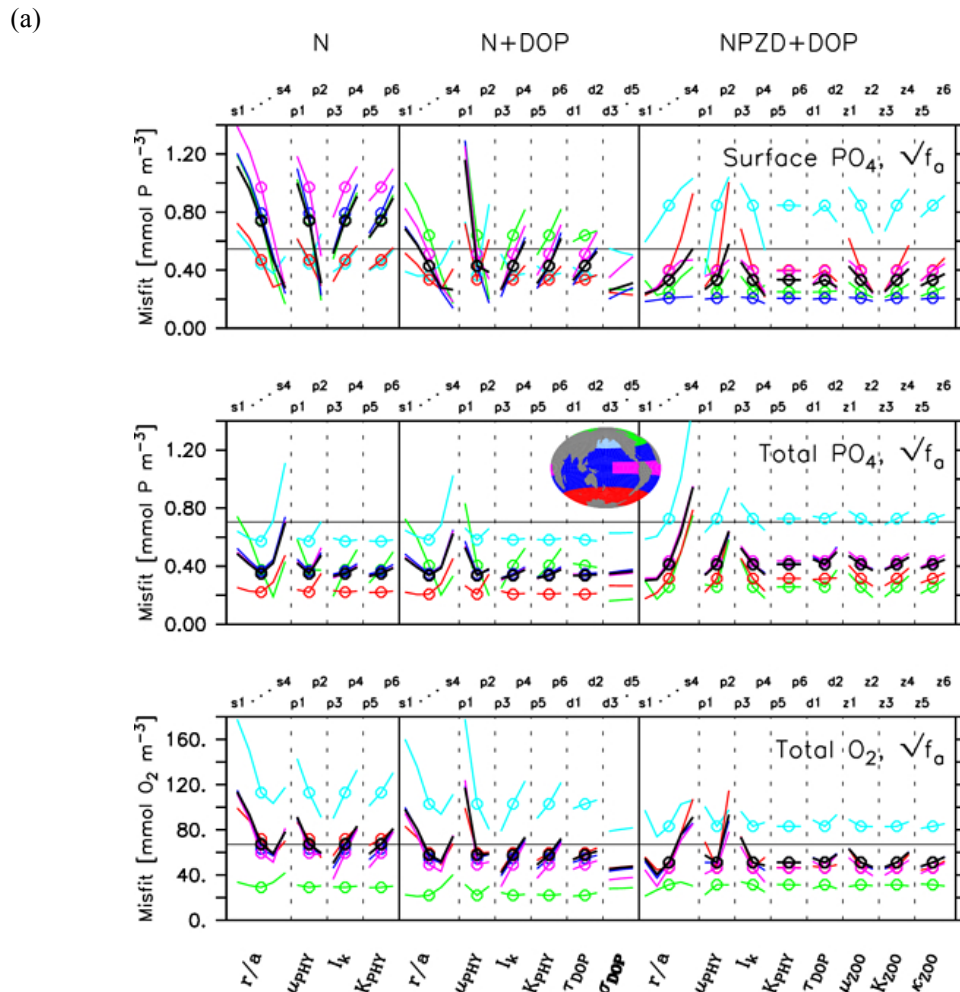


Figure 1: (a) Different global misfit functions for models with different levels of complexity [4]; (b) Processes in the ocean biogeochemical model in CESM and ACME [5].

Contributed ideas paper 38

From Data through Inference to Optimization under Uncertainty: Towards End-to-End Climate Model-Based Decision-Making

Omar Ghattas (University of Texas at Austin) and Georg Stadler (New York University)

Background

In this idea paper we want to make the case for a comprehensive approach that looks beyond constructing (forward, deterministic) climate models to an entire end-to-end mathematical/computational framework aimed at model-based decision-making with quantified uncertainties. By “model-based decision-making” we mean a formal, systematic, and quantitative framework for making inferences about uncertain parameters in complex models of the climate system given observational data, and then optimizing design or control variables under these uncertainties. In particular, the framework encompasses:

1. **Bayesian inverse problem:** First, large observational datasets are assimilated into complex climate models to infer uncertain parameters¹ and reduce their uncertainties via solution of a Bayesian inverse problem.
2. **Stochastic prediction problem:** Second, once the model parameters and their uncertainties have been estimated from the data via statistical inversion, the resulting probability distributions are propagated through the model to yield predictions of quantities of interest with quantified uncertainties.
3. **Stochastic optimization problem:** Finally, given an objective function representing stochastic predictions of quantities of interest and decision variables (design or control) that can be manipulated to influence the objective, an optimization problem governed by the stochastic forward problem is solved to produce optimal decisions. The objective can represent either the performance of the system (in an optimal control problem) or the uncertainty in inferred parameters (in an optimal experimental design problem).

As just one illustration of this end-to-end, “data-to-decisions” framework, consider a subsurface CO₂ injection problem. The model is given by PDEs describing multiphase flow and transport through a porous medium; the observational data are noisy well measurements of pressure and CO₂ concentration; the uncertain parameters describe the permeability field; the optimal experimental design variables are the locations of the observational wells; the optimal CO₂ injection control variables are the locations of the injection wells and the rates of injection; the optimal experimental design objective is the uncertainty in the inferred permeability field; and the optimal control objective is the long-term fate of the injected CO₂.

The vast majority of research in computational science & engineering over the past several decades has addressed the so-called *inner problem*, i.e., the deterministic forward problem. This is entirely appropriate: without accurate, robust (over parameter space), fast, and scalable methods for the inner problem, the *outer problems* of inversion, prediction, and optimization are not viable. *However, for the full power of computational modeling and simulation to be realized for critical societal problems, we must address the ultimate goal—decision-making under uncertainty, and its stepping stones inversion and prediction under uncertainty—within which the deterministic forward problem is merely a “function evaluation.”*

The grand challenge of executing this framework for large-scale complex models, such as those faced in climate science, is that state-of-the-art algorithms and methods for uncertainty-

¹Here we refer to *parameters* in the broadest possible sense, to include not just coefficients of models, but also distributed and boundary source terms, initial conditions, domain shapes, and even parameters representing model structure.

Contributed ideas paper 38

based inversion, forward propagation, optimal experimental design, and optimal control are viable only for low-dimensional and inexpensive models (e.g., ODEs). For complex PDE models (such as those characterizing climate models) with high-dimensional parameter spaces (such as those stemming from discretization of spatial fields), high-dimensional experimental design spaces (such as the spatio-temporal trajectories of observing systems), and high dimensional control variable spaces (such as spatio-temporal distribution of allowable carbon emissions), *current mathematical/computational methods are completely prohibitive.*

Proposed Directions and Potential Impacts

We envision an integrated, model-based, data-driven, end-to-end framework that accounts for uncertainty across all components: from observations, through inverse solution, through optimal observation system design, and ultimately to model-predictive optimal control to support rational model-based decision-making about the integrated climate system. *Indeed, in the long term, the absence of such a framework will prevent the full power of modeling and simulation on extreme scale computers to be brought to bear on the complex decision-making that must be undertaken to understand and mitigate global climate change.*

The key to overcoming the severe mathematical and computational challenges in realizing this framework is to recognize that beneath the apparently high-dimensional inversion, prediction, and optimization problems lurk much lower dimensional manifolds that capture the maps from (inversion/design/control) inputs to outputs of interest. This stems from the fact that: (1) observational data typically inform only low-dimensional manifolds of parameter fields; (2) uncertainties in estimated parameter fields are typically sensitive to only low-dimensional manifolds of optimal sensor location fields; and (3) control objectives are typically sensitive to only low-dimensional manifolds of optimal control spaces.

Thus, the inverse, prediction, and optimal design and control problems are only as large as the information content of the observational data, the sensor spatial modes that most reduce the uncertainty in the inferred parameters, and the controllable modes that most dictate the behavior of the objective of interest—that is, these problems are characterized by their much smaller *intrinsic* dimensions. Black-box methods developed as generic tools are incapable of exploiting the low-dimensional structure of specific inversion, prediction, optimal experimental design, and optimal control operators. To be successful, it is essential that we integrate and cross-fertilize ideas from computer science, numerical analysis, and applied mathematics (PDE-constrained optimization, optimal control theory, Bayesian inference, model reduction, scalable numerical algorithms, randomized algorithms, etc.), while exploiting the structure of the specific inverse, prediction, and optimal design and control operators in the context of target climate models. A critical challenge is to efficiently construct low-dimensional representations of these operators while accessing just the action of these operators in specific directions.

Over the past half-dozen years we have been working along these lines, developing methods for solution, with quantified uncertainties, of large-scale inverse problems [5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19], optimal experimental design problems [1, 3, 4], and most recently stochastic optimal control problems [2] characterized by models of unprecedented size (up to 10^6 uncertain parameters and 10^8 state variables) and complexity (e.g., full nonlinear Stokes ice sheet models). Algorithmic complexity (measured in the number of forward model solves) has been demonstrated to be independent of uncertain parameter dimension, data dimension, and optimization variable dimension. The key has been to exploit derivative information (Jacobians, Hessians) of the maps from stochastic parameters to outputs of interest and their low rank approximations. However significant work remains to overcome the challenges presented by the extremely nonlinear and high-dimensional models characteristic of climate systems.

Contributed ideas paper 38

References

- [1] A. Alexanderian, P. Gloor, and O. Ghattas, On Bayesian A- and D-optimal experimental designs in infinite dimensions, *Bayesian Analysis*, 2015, to appear.
<http://dx.doi.org/10.1214/15-BA969>
- [2] A. Alexanderian, N. Petra, G. Stadler, O. Ghattas, Mean-variance risk-averse optimal control of systems governed by PDEs with random coefficient fields using local approximations, manuscript to be submitted.
- [3] A. Alexanderian, N. Petra, G. Stadler, O. Ghattas, A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0 -sparsification, *SIAM Journal on Scientific Computing*, 36(5):A2122–A2148, 2014.
<http://dx.doi.org/10.1137/130933381>
- [4] A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas, A fast and scalable method for A-optimal design of experiments for infinite-dimensional Bayesian nonlinear inverse problems, *SIAM Journal on Scientific Computing*, 2015, to appear.
<http://arxiv.org/abs/1410.5899>
- [5] T. Bui-Thanh, C. Burstedde, O. Ghattas, J. Martin, G. Stadler, and L.C. Wilcox, Extreme-scale UQ for Bayesian inverse problems governed by PDEs, *Proceedings of IEEE/ACM SC12*, 2012.
<http://dx.doi.org/10.1109/SC.2012.56>
- [6] T. Bui-Thanh and O. Ghattas, Analysis of the Hessian for inverse scattering problems. Part I: Inverse shape scattering of acoustic waves. *Inverse Problems*, 28(5):055001, 2012.
<http://dx.doi.org/10.1088/0266-5611/28/5/055001>
- [7] T. Bui-Thanh and O. Ghattas, Analysis of the Hessian for inverse scattering problems. Part II: Inverse medium scattering of acoustic waves. *Inverse Problems*, 28(5):055002, 2012.
<http://dx.doi.org/10.1088/0266-5611/28/5/055002>
- [8] T. Bui-Thanh and O. Ghattas, Analysis of the Hessian for inverse scattering problems. Part III: Inverse medium scattering of electromagnetic waves in three dimensions, *Inverse Problems and Imaging*, 7(4):1139–1155, 2013.
<http://dx.doi.org/10.3934/ipi.2013.7.1139>
- [9] T. Bui-Thanh and O. Ghattas, An analysis of infinite dimensional Bayesian inverse shape acoustic scattering and its numerical approximation, *SIAM/ASA Journal on Uncertainty Quantification*, 2(1):203–222, 2014.
<http://dx.doi.org/10.1137/120894877>
- [10] T. Bui-Thanh and O. Ghattas, A scalable MAP solver for Bayesian inverse problems with Besov priors, *Inverse Problems and Imaging*, 9(1):27–54, 2015.
<http://dx.doi.org/10.3934/ipi.2015.9.27>
- [11] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems. Part I: The linearized case, with applications to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
<http://dx.doi.org/10.1137/12089586X>

Contributed ideas paper 38

- [12] T. Bui-Thanh, O. Ghattas, and D. Higdon, Adaptive Hessian-based non-stationary Gaussian process response surface method for probability density approximation with application to Bayesian solution of large-scale inverse problems, *SIAM Journal on Scientific Computing*, 34(6):A2837–A2871, 2012.
<http://dx.doi.org/10.1137/110851419>
- [13] H.P. Flath, L.C. Wilcox, V. Akcelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas, Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations, *SIAM Journal on Scientific Computing*, 33(1):407-432, 2011.
<http://dx.doi.org/10.1137/090780717>
- [14] D. Galbally, K. Fidkowski, K. Willcox, and O. Ghattas, Nonlinear model reduction for uncertainty quantification in large-scale inverse problems, *International Journal for Numerical Methods in Engineering*, 81:1581–1608, 2010.
<http://dx.doi.org/10.1002/nme.2746>
- [15] T. Isaac, N. Petra, G. Stadler, O. Ghattas, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet, *Journal of Computational Physics*, 296(1):348–368, 2015.
<http://dx.doi.org/10.1016/j.jcp.2015.04.047>
- [16] C. Lieberman, K. Willcox, O. Ghattas, Parameter and state model reduction for large-scale statistical inverse problems, *SIAM Journal on Scientific Computing*, 32(5):2523–2542, 2010.
<http://dx.doi.org/10.1137/090775622>
- [17] J. Martin, L.C. Wilcox, C. Burstedde, and O. Ghattas, A Stochastic Newton MCMC method for large-scale statistical inverse problems with application to seismic inversion, *SIAM Journal on Scientific Computing*, 34(3):A1460-A1487, 2012.
<http://dx.doi.org/10.1137/110845598>.
- [18] N. Petra, J. Martin, G. Stadler, and O. Ghattas, A computational framework for infinite-dimensional Bayesian inverse problems. Part II: Stochastic Newton MCMC with application to ice sheet flow inverse problems, *SIAM Journal on Scientific Computing*, 36(4):A1525–A1555, 2014.
<http://dx.doi.org/10.1137/130934805>
- [19] H. Zhu, S. Li, S. Fomel, G. Stadler, and O. Ghattas, A Bayesian approach to estimate uncertainty for full waveform inversion with a priori information from depth migration, submitted, 2015.

Contributed ideas paper 39

Dec/1/2015

Transforming Climate Modeling via Scale-Adaptive Computational Techniques

Christiane Jablonowski

Department of Climate and Space Sciences and Engineering, University of Michigan

Jared Ferguson (University of Michigan)

Hans Johansen, Phillip Colella (Lawrence Berkeley National Laboratory)

Abstract

The research advances frontiers in computational fluid dynamics by introducing flexible, solution-adaptive and scale-adaptive techniques into weather and climate models. The goal is to show that innovative computational tools, such as Adaptive Mesh Refinement (AMR) techniques, adaptive numerical approaches and scale-adaptive & stochastic representations of unresolved physical processes, like turbulence and clouds, can transform the field of weather and climate modeling in fundamental ways. The paradigm shift is that computational resources will only be used where needed for accuracy, which is dynamically determined by the geophysical flow field. This guarantees a targeted use of computational resources near “phenomena of interest”. It enables focused scientific investigations with unprecedented small grid spacings in the zoomed areas that could otherwise not be reached within the next decade.

Background/Research to Date

Solution- and scale-adaptive modeling in the climate sciences is an innovative and disruptive technology which is in its early stages today. Fundamental research questions need to be addressed to guarantee numerically and physically consistent, conservative and noise-free down-scaling and upscaling of information. Furthermore, adaptive modeling raises many other open research questions such as the coupling of adaptive atmospheric models to other components like ice, ocean and land models. Very few modeling groups have ever approached adaptive techniques in the atmospheric sciences due to the inherent complexities of 3D weather and climate processes. Besides the open question how to adapt physical parameterizations on variable-resolution grids, complex terrain and rugged coastline data also need to be provided as adaptive boundary conditions. In addition, atmospheric processes are typically modeled in spherical domains with mapped vertical coordinates. This impacts the selection of suitable AMR meshes and raises the question whether 3D adaptations are feasible.

Proposed Direction of work

The science driver for the proposed solution-adaptive research is an improved understanding of multi-scale processes and their non-linear interactions. In particular, the research will advance strategies how to bridge the scale discrepancies between local, regional and global phenomena in weather and climate models without the prohibitive computational costs of global cloud-resolving simulations. This includes adaptations of the computational grids, adjustments of the numerical methods, scale-aware representations of small-scale processes like clouds or precipitation, or the introduction of stochastic modeling techniques that pay tribute to the inherent uncertainty of turbulent processes at all scales. More specifically, the research addresses the following two areas:

Contributed ideas paper 39

Adaptive Mesh Refinement (AMR) The goal of AMR is to refine the grid dynamically in advance of important physical processes that need additional grid resolution, and to coarsen the grid once the system passes. In addition, static adaptations can easily be overlaid. An illustrative example of the author’s prior AMR research is provided in Fig. 1 ([3, 5]). It shows how a block-structured adapted “cubed-sphere” grid tracks features of interest which are the rotational motions (in color) in two 2D shallow water simulations. The examples utilize a fourth-order finite-volume model that is built upon the AMR library Chombo ([1]). The latter has been developed at DoE’s Lawrence Berkeley National Laboratory (LBNL) and supports mapped grids such as spherical cubed-sphere domains. A 3D nonhydrostatic version has also already been built and undergoes idealized testing. The research idea is to advance the AMR capabilities for the climate sciences by (1) gaining in-depth experience with 2D horizontal and fully 3D adaptations, (2) introducing high-order and mixed-order numerical techniques with adaptive diffusive properties, (3) enforcing consistency, conservation and noise-free solutions at refinement boundaries and (4) defining adequate refinement criteria for the targeted capturing of phenomena. The expected outcome is to establish AMR as a mature and flexible computational tool, which represents a major and unprecedented milestone for the climate sciences.

Adaptive representation of subgrid-scale processes and stochastic approaches

Research has become vibrant to advance the scale-awareness of physical parameterizations and, in particular, of cloud and turbulence schemes. Among the new ideas are probability-based and stochastic physical parameterizations that take the impact of the subgrid-scale uncertainty on the resolved scales into account. The stochastic parameterizations introduce structured but random noise, and have not only improved weather prediction simulations but also reduced systematic, long-term climate modeling biases. The research idea is to utilize and advance promising approaches to solution-adaptive subgrid-scale processes, such as CLUBB (Cloud Layers Unified by Binormals, [2]) and the Grell-Freitas stochastic deep convection scheme ([4]), and test them within AMR models. The expected outcome is to identify and mature robust physical parameterizations that behave consistently across all spatial scales without the manual re-tuning of empirical parameters. This will include critical evaluations of the so-called gray zone (with grid spacings $\Delta x \approx 1-10$ km) where convective cloud clusters become partially resolved. The challenge is to avoid the “double-counting” of such processes.

Connections to Math, Comp Sci & and Climate Science

The outlined research is interdisciplinary and combines applied mathematics, computational science and atmospheric science in a synergistic way. Special attention needs to be paid to the computational design of the adaptive approaches, such as the parallel scaling characteristics and computational speed. In particular, it needs to be shown that adaptive modeling is not just feasible, but also more accurate and computationally faster than traditional weather and climate models. This research will lead to unparalleled computational and scientific capabilities that maximize the use of computational resources. It thereby creates the foundation for future-generation weather and climate models.

Potential Impact on the field

The outlined research is a farsighted high-payoff and high-impact research idea that provides the basis for scientific progress in the coming decades. More specifically, adaptive techniques built upon nonhydrostatic dynamical cores are an attractive basis for future-generation weather and climate models that will be enabled to seamlessly capture phenomena across all scales.

Contributed ideas paper 39

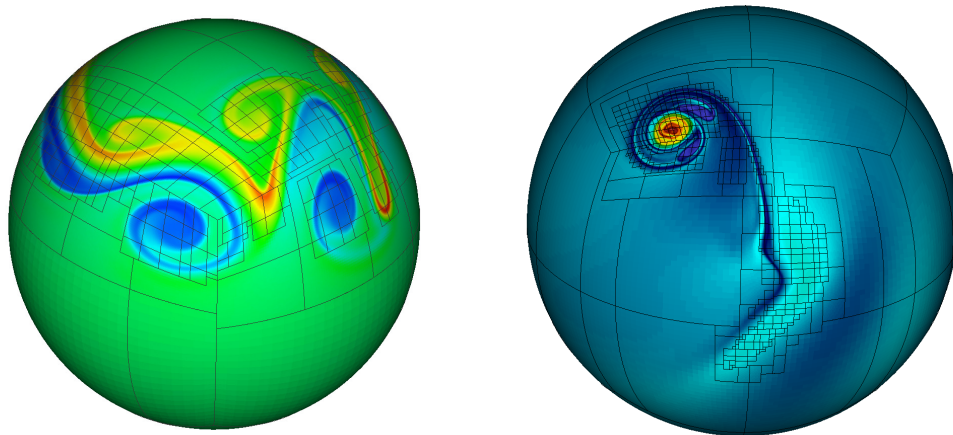


Figure 1: Examples of adaptively refined simulations with the Chombo-AMR model in shallow-water mode ([3]). The refined grids track the rotational motion (depicted in color) of the evolving flow fields.

References

- [1] Adams, M., P. Colella, D. T. Graves, J. N. Johnson, H. S. Johansen, N. D. Keen, T. J. Ligoeki, D. F. Martin, P. W. McCorquodale, D. Modiano, P. O. Schwartz, T. D. Sternberg and B. Van Straalen, 2014. *Chombo Software Package for AMR Applications - Design Document*. Lawrence Berkeley National Laboratory Technical Report LBNL-6616E, Available from <https://commons.lbl.gov/display/chombo/Chombo+-+Software+for+Adaptive+Solutions+of+Partial+Differential+Equations>
- [2] Bogenschutz, P. A., A. Gettelman, H. Morrison, V. E. Larson, C. Craig and D. P. Schanen, 2013. *Higher-order turbulence closure and its impact on climate simulations in the Community Atmosphere Model*. *J. Climate*, 26, 9655-9676
- [3] Ferguson, J., C. Jablonowski, H. Johansen, P. McCorquodale, P. Colella and P. A. Ullrich, 2016. *Assessing the Chombo-AMR Adaptive Mesh Refinement approach for geophysical flows*. In preparation
- [4] Grell, G. A. and S. R. Freitas, 2014. *A scale and aerosol aware stochastic convective parameterization for weather and air quality modeling*. *Atmos. Chem. Phys.*, 14, 5233-5250
- [5] Jablonowski, C., R. C. Oehmke and Q. F. Stout, 2009. *Block-structured adaptive meshes and reduced grids for atmospheric general circulation models*. *Phil. Trans. R. Soc. A*, 367, 4497-4522

Contributed ideas paper 40

Next generation implicit solvers and analysis algorithms for ice sheet modeling 11/20/2015

Mauro Perego (Sandia National Laboratories),
Stephen Price (Los Alamos National Laboratory),
Andrew Salinger (Sandia National Laboratories)

Abstract

Dynamics of ice sheets are typically modeled in an explicit or semi-implicit fashion. In this work we advocate for fully implicit modeling of ice sheets based on interpreting the ice sheet evolution equation as an obstacle problem governed by variational inequalities (as recently proposed by E. Bueler and G. Jovet). Advantages of this approach include the ability to use significantly larger time steps to evolve the ice geometry and to compute steady-state solutions in a single solve, rather than performing millennial-scale transient simulations (currently not possible due to model complexity and computational cost). The latter capability will also allow for the prediction of inherent instabilities (bifurcations) of the coupled ice dynamic and evolution system through parameter continuation.

Background/Research to Date

The computational modeling of ice sheets has been significantly advanced during the last decade, resulting in accurate, fast, robust, and scalable solvers for the numerical solution of ice sheet dynamics and evolution. These models include BISICLES¹[2], Elmer/Ice²[3], FELIX [13, 9, 10], ISSM³ [8], PISM⁴[14], and Ymir [5]. While the diagnostic solution of the momentum balance equations are always done implicitly, for historical and practical reasons the evolution of the ice geometry is typically modeled explicitly (e.g., using a forward Euler scheme) with very simplistic numerical treatments of ice advance and retreat. This leads to (1) a requirement of very small time time-steps (sub-annual to daily depending on spatial resolution) to ensure stability and, (2) low accuracy with respect to the simulated position of the ice sheet margin and/or grounding line.

Land ice flow can be interpreted as a free boundary problem which can be naturally modeled using variational inequalities. In [7], E. Bueler and G. Jovet interpreted ice thickness evolution as an obstacle problem and derived a model based on variational inequalities. The main advantage of this formulation is that the evolution of grounded ice sheet margins are implicitly handled and do not need an *ad hoc* discretization. Their formulation is limited, however, in that it does not include ice shelves (the floating extensions of ice sheets) or calving margins and is limited to a low-order approximation flow model (i.e., the “shallow ice approximation”).

To our knowledge, no ice sheet models under active development have the capability to allow for the prediction of inherent instabilities in ice sheet dynamics or to perform bifurcation analysis.

Proposed Direction of work

We propose to model the complete ice flow problem (dynamics and evolution) as a fully implicit obstacle problem governed by variational inequalities. The challenge will be to extend the work of E. Bueler and G. Jovet [7] to higher-order flow models (e.g. nonlinear Stokes or “Blatter-Pattyn” type models) and to include in this same implicit formulation the flow of ice shelves and mass loss at calving margins. The resulting inequality should be discretized with variational methods such

¹<http://bisicles.lbl.gov>

²<http://elmerice.elmerfem.org/>

³<https://issm.jpl.nasa.gov/>

⁴<http://pism-docs.org/wiki/doku.php>

Contributed ideas paper 40

as the finite element method. The resulting discrete variational inequalities can be solved using semi-smooth Newton methods for the equivalent complementarity problem (see e.g. [6] and [1]).

We also propose to use parameter continuation to predict and analyze possible instabilities inherent to a particular choice of ice dynamics model. An example scenario of interest to climate scientists is as follows: For a constant-in-time climate forcing (e.g. surface mass balance), what is the resulting steady state for the ice sheet as a function of one or more ice sheet or climate model parameters (e.g. average ocean surface temperature, average near-surface atmospheric temperature, or average surface mass balance). Of particular interest would be understanding whether or not a specific perturbation to the average ocean temperature could trigger an instability like the *West Antarctica ice sheet instability* (see e.g. [12]), resulting in significant mass loss from the ice sheet (and subsequent sea-level rise). This sort of analysis could be used to narrow down and identify unstable regimes as a function of particular climate forcings, which could then be followed up by a more systematic and detailed set of transient runs, in order to better understand and constrain the impact of those instabilities in the future.

A natural software environment for developing the proposed work is the Albany/FELIX ice sheet code [13], which is included in the Accelerated Climate Model for Energy (ACME) as part of Model for Prediction Across Scales (MPAS) Land Ice Model. It is strongly integrated with several Trilinos packages, including LOCA [11], a library of continuation algorithms that can be used for performing bifurcation analysis on large-scale simulations.

Connections to Math, Comp Sci & and Climate Science

The proposed direction of work presents several mathematical and computational challenges. The problem features several nonlinearities and is characterized by large and complex geometries that would require efficient and robust solvers. This would stimulate and leverage research in numerical and computational methods for solving problems based on variational inequalities. Despite a number of available numerical methods, we currently know of no solvers that could robustly handle the large-scale problems (order 100 million unknowns) this would present while simultaneously taking advantage of DOE's emerging computer architectures.

While the discussion above specifically focuses on ice sheet modeling, the concept of applying continuation methods for finding and analyzing instabilities is generic to many other other aspects of climate science. For example, a further application with the potential for greatly impacting our understanding of earth's climate would be a stability study of the *Atlantic meridional overturning circulation* [4].

The development of the proposed work can clearly leverage DOE investments in ice sheet climate modeling (ISICLES and PISCEES SciDAC institutes and the ACME program) as well as in computational mathematics (FASTMath and SUPER SciDAC institutes).

Potential Impact on the field

The proposed fully implicit approaches for modeling ice sheet dynamics and evolution will significantly improve the accuracy, robustness, and efficiency of ice sheet simulations. Additionally, they will provide a prognostic tool for efficiently computing steady-state solutions, something that does not currently exist due to model complexity and the very long timescales required for ice sheets to come into equilibrium with climate forcing (order tens to hundreds of thousands of years). Moreover, the ability to predict instabilities in steady-state ice sheet configurations will significantly change how we investigate ice sheet evolution and increase our understanding for what the most sensitive parameters of the system are in terms of specific quantities of interest, such as future sea-level rise.

Acknowledgments Mauro Perego thanks Ed Bueler and Jed Brown for fruitful discussions.

Contributed ideas paper 40

References

- [1] S. J. Benson and T. S. Munson. Flexible complementarity solvers for large-scale applications. *ArXiv Mathematics e-prints*, July 2003.
- [2] S. L. Cornford, D. F. Martin, D. T. Graves, D. F. Ranken, A. M. Le Brocq, R. M. Gladstone, A. J. Payne, E. G. Ng, and W. H. Lipscomb. Adaptive mesh, finite volume modeling of marine ice sheets. *J. Comput. Phys.*, 232(1):529–549, January 2013.
- [3] O. Gagliardini, T. Zwinger, F. Gillet-Chaulet, G. Durand, L. Favier, B. de Fleurian, R. Greve, M. Malinen, C. Martín, P. Råback, J. Ruokolainen, M. Sacchetti, M. Schäfer, H. Seddik, and J. Thies. Capabilities and performance of elmer/ice, a new-generation ice sheet model. *Geoscientific Model Development*, 6(4):1299–1318, 2013.
- [4] M. Hofmann and S. Rahmstorf. On the stability of the atlantic meridional overturning circulation. *Proceedings of the National Academy of Sciences*, 106(49):20584–20589, 2009.
- [5] T. Isaac, G. Stadler, and O. Ghattas. Solution of nonlinear stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics. *SIAM Journal on Scientific Computing*, 37(6):B804–B833, 2015.
- [6] K. Ito and K. Kunisch. Semismooth newton methods for variational inequalities of the first kind. *ESAIM: Mathematical Modelling and Numerical Analysis*, 37:41–62, 1 2003.
- [7] G. Jovet and E. Bueler. Steady, shallow ice sheets as obstacle problems: Well-posedness and finite element approximation. *SIAM Journal on Applied Mathematics*, 72(4):1292–1314, 2012.
- [8] E. Larour, H. Seroussi, M. Morlighem, and E. Rignot. Continental scale, high order, high spatial resolution, ice sheet modeling using the ice sheet system model (issm). *Journal of Geophysical Research: Earth Surface*, 117(F1):n/a–n/a, 2012. F01022.
- [9] W. Leng, L. Ju, Y. Xie, T. Cui, and M. Gunzburger. Finite element three-dimensional stokes ice sheet dynamics model with enhanced local mass conservation. *Journal of Computational Physics*, 274:299 – 311, 2014.
- [10] M. Perego, M. Gunzburger, and J. Burkhart. Parallel finite-element implementation for higher-order ice-sheet models. *Journal Of Glaciology*, 58(207):76–88, 2012.
- [11] A. G. Salinger, E. A. Burroughs, R. P. Pawlowski, E. T. Phipps, and L. A. Romero. Bifurcation tracking algorithms and software for large scale applications. *International Journal of Bifurcation and Chaos*, 15(03):1015–1032, 2005.
- [12] C. Schoof. Ice sheet grounding line dynamics: Steady states, stability, and hysteresis. *Journal of Geophysical Research: Earth Surface*, 112(F3):n/a–n/a, 2007. F03S28.
- [13] I. Tezaur, M. Perego, A. Salinger, R. Tuminaro, and S. Price. Albany/FELIX: A parallel, scalable and robust finite element higher-order Stokes ice sheet solver built for advance analysis. *Geoscientific Model Development*, 8(4):1197–1220, 2015.
- [14] R. Winkelmann, M. A. Martin, M. Haseloff, T. Albrecht, E. Bueler, C. Khroulev, and A. Levermann. The Potsdam Parallel Ice Sheet Model (PISM-PIK) Part 1: Model description. *The Cryosphere*, 5:715–726, 2011.

Contributed ideas paper 41

Iterative Parallel In Time Methods for Climate Simulations **Dec. 1, 2015**

Michael Minion

Lawrence Berkeley National Lab

Abstract

As global climate and earth system models increase in both physical fidelity and spatial resolution, the computational challenge of integrating these equations over long time scales appears daunting. Not only does the coupling of many physical processes on multiple time scales require custom time integration methods, but increasing spatial resolution means increasing the number of time steps and hence runtime using traditional sequential time stepping methods. The main “idea” in this idea paper is that recently developed temporal integration methods that are iterative in nature could be adopted in the climate modeling community to ease the construction of accurate integrators for multi-scale and multi-physics systems and to provide an opportunity to exploit parallelism in the time direction.

Background/Research to Date

The problem of numerically modeling the earth’s climate requires the temporal integration of a highly complex coupled system of partial differential equations (PDEs) for physical time periods of decades or longer. The underlying physical problem consists of multiple characteristic domains (ocean, atmosphere, land, ice), multiple physical processes, and multiple spatial and temporal time scales. Hence specialized temporal integration methods for single or coupled systems are often desirable.

The current spatial resolution of global climate models is 50-100km with time step sizes of 10-20 minutes. Increasing the spatial resolution is clearly desirable to both reduce numerical error and in some cases remove the need for parameterizations of subgrid-scale physics. In most instances however, increasing the spatial resolution requires a corresponding increase in temporal resolution due to stability and/or accuracy constraints. Research into parallel methods for time dependent PDEs has traditionally focused on parallelization in the spatial domain by distributing spatial degrees of freedom across multiple processors. Temporal integration, on other hand, is still typically done by sequential time stepping using a variety of methods depending on the system under consideration.

Since the early 1990s, the fastest supercomputers have employed an ever larger number of standard processors, which enabled the number of spatial degrees of freedom on the largest climate simulations to increase as well. The wall clock cost of increasing time steps was masked for many years by the increased clock speed of individual processors. Due to the limitations of Denard scaling, the exponential increase in clock speeds has ended and is not expected to resume without a fundamental technology change. The inconvenient outlook is that even with perfect spatial parallel scaling, increasing spatial resolution of climate simulations on future supercomputers will mean increasing run-time unless some parallelization in the time direction can be realized.

Attempts to construct efficient parallel temporal integration methods date back to at least the 1960s [1], but are still not generally employed in large scale PDE simulations in any domain science. The reason for this is partly due to the previous success of weak spatial scaling mentioned above and partly due to the lack of algorithms with decent parallel scaling for more than a few processors. However, in the last 15 years iterative time parallel algorithms such as

Contributed ideas paper 41

parareal [2], PITA [3], PFASST [4], and RIDC [5], have sparked renewed interest in parallel temporal integration and promise better parallel scaling in the time direction.

The PFASST algorithm is based on the iterative temporal method called spectral deferred corrections (SDC) [6]. SDC methods have a convenient feature in that higher-order accurate methods can be readily constructed from first-order methods, which enables the construction of higher-order methods that can treat different terms in the equation independently as in traditional splitting schemes, implicitly or explicitly depending on stiffness, and even with different time steps or integrators [7, 8, 9, 10, 11]. One apparent disadvantage of SDC methods is that the number of function evaluations per time step is considerably higher than standard lower-order methods. Some of this apparent extra cost can be amortized by using larger time steps for the same accuracy, however methods to reduce the cost of SDC by using a multi-level approach (MLSDC) have shown recent promise [12]. The basic idea is to replace some SDC iterations on the full resolution spatial grid with iterations on a coarser version of the problem. Coarsening can take the form of coarsening of spatial and/or temporal resolution, lower-order spatial discretizations, reduced accuracy implicit solves, or even reduced fidelity governing equations. The PFASST algorithm can be thought of as a way to perform MLSDC iterations on many time intervals in parallel, with improvements to the initial conditions on each time step passed forward during the MLSDC sweeps.

Proposed Direction of work

To date, the PFASST algorithm has not been evaluated on any of the standard test cases in the climate community. First, a popular set of tests from Williamson, et. al. [13] will be implemented and tested over a range of parameter choices (see also related work [14, 15]). These tests all involve the evolution of the shallow water equations on the sphere and hence do not present the opportunity to exploit some of the advantages of MLSDC methods listed above. Nevertheless, success on these test cases is necessary to convince researchers in the climate community that pursuing parallel in time methods is worthwhile.

If the standard tests are successful, we would like to next demonstrate the use of multi-physics coarsening in MLSDC and PFASST. Ideally this could be done using a coupled atmospheric flow and chemistry package, where reduced spatial resolution is coupled with a reduced chemistry representation on coarse levels. Ultimately, the ideal course of action would lead to the evaluation of parallel in time strategies for a coupled ocean-atmospheric code. Finally, we would like to provide some guidance through analysis and numerical experiments of the type of problems in climate science that do not seem readily parallelizable in the time direction using current technology.

Potential Impact on the field

Depending on the point of view, the research described above could potentially:

1. Allow climate simulations with increased spatial resolution without increasing run time.
2. Facilitate a more robust and accurate temporal coupling of different climate components.

Contributed ideas paper 41

References

- [1] Martin J. Gander. 50 years of Time Parallel Time Integration. In *Multiple Shooting and Time Domain Decomposition*. Springer, 2015.
- [2] J L Lions, Y Maday, and G Turinici. A "parareal" in time discretization of PDE's. *Comptes Rendus De L Academie Des Sciences Serie IMathematique*, 332(7):661–668, 2001.
- [3] Charbel Farhat and Marion Chandesris. Time-decomposed parallel time-integrators: theory and feasibility studies for fluid, structure, and fluid-structure applications. *International Journal for Numerical Methods in Engineering*, 58(9):1397–1434, 2003.
- [4] M Emmett and M Minion. Toward an efficient parallel in time method for partial differential equations. *Comm. App. Math. and Comp. Sci*, 7:105–132, 2012.
- [5] Andrew J Christlieb, Colin B Macdonald, and Benjamin W Ong. Parallel High-Order Integrators. *SIAM Journal on Scientific Computing*, 32(2):818, 2010.
- [6] A. Dutt and L. Greengard. Spectral deferred correction methods for ordinary differential equations. *BIT Numerical Mathematics*, 40(2):241–266, 2000.
- [7] Michael L. Minion. Semi-implicit spectral deferred correction methods for ordinary differential equations. *Communications in Mathematical Sciences*, 1(3):471–500, September 2003.
- [8] Anne Bourlioux, Anita T. Layton, and Michael L. Minion. High-order multi-implicit spectral deferred correction methods for problems of reactive flow. *Journal of Computational Physics*, 189(2):651–675, August 2003.
- [9] Anita T. Layton and Michael L. Minion. Conservative multi-implicit spectral deferred correction methods for reacting gas dynamics. *Journal of Computational Physics*, 194(2):697–715, March 2004.
- [10] A Nonaka, J B Bell, M S Day, C Gilet, A S Almgren, and M L Minion. A deferred correction coupling strategy for low Mach number flow with complex chemistry. *Combustion Theory and Modelling*, 16(6):1053–1088, December 2012.
- [11] Elizabeth L. Bouzarth and Michael L. Minion. A multirate time integrator for regularized Stokeslets. *Journal of Computational Physics*, 229(11):4208–4224, June 2010.
- [12] Robert Speck, Daniel Ruprecht, Matthew Emmett, Michael Minion, Matthias Bolten, and Rolf Krause. A multi-level spectral deferred correction method. *BIT Numerical Mathematics*, 2014.
- [13] D.L. Williamson, J.J. Hack, R. Jakob, P.N. Swarztrauber, and J.B. Drake. A standard test set for numerical approximations to the shallow water equations in spherical geometry. Technical report, Oak Ridge National Laboratory (ORNL), Oak Ridge, TN, August 1991.
- [14] T. Haut and B. Wingate. An asymptotic parallel-in-time method for highly oscillatory PDEs. *SIAM Journal on Scientific Computing*, 36(2):A693–A713, 2014.
- [15] Jun Jia, Judith C. Hill, Katherine J. Evans, George I. Fann, and Mark A. Taylor. A spectral deferred correction method applied to the shallow water equations on a sphere. *Monthly Weather Review*, page 130208070720007, February 2013.

Contributed ideas paper 42

Adaptive multiresolution strategies for coupled climate simulations

Nov. 24, 2015

Vijay Mahadevan

Argonne National Laboratory

Robert Jacob (ANL)

Abstract

Climate modeling is at the strong scaling limit, and scientists will not be able to continually increase resolution on future architectures while retaining a 5-years/day simulation speed. Climate models for multidecadal simulations will instead need new approaches such as nonlinear regression-based adaptive multiresolution strategies that maximize computational efficiency on emerging architectures without sacrificing coupled solution accuracy.

Motivations

Understanding the Earth's climate through robust and accurate modeling of the intrinsically complex, coupled ocean-atmosphere-land-ice-biosphere models requires extreme-scale computational power [1]. As clock rates and core counts increased steadily in the 90's and early 2000's, climate modelers were able to improve their solution quality somewhat by increasing the resolution of the dynamics and adding more detailed representation of nonfluid processes in the separate models, while subjecting them to rigorous validation [2]. But improved predictability in next-generation climate models can only come from careful treatment of the multiphysics and multiscale nature of the climate system [3].

Previously, investigations seeking the optimal numerical scheme for tackling multiphysics behavior have relied on operator splitting between component models, with regime-specific approximations [4, 5] that may be violated over long-term integration. In order to preserve the numerical accuracy and stability of the coupled multimodel simulations [6], adaptation in algorithmic choices depending on the dominant dynamical scales of interest will be essential. Additionally, in several atmospheric flow and ocean model solvers, consistent and conservative remapping strategies [7] determine the propagation of discretization accuracy between models. The recent use of spectral advection schemes has allowed use of larger step sizes ($CFL > 1$) while still maintaining positivity for tracer advection. Even such schemes, however can yield diffusive solution propagation and entail high computational cost in imposing mass conservation or incompressibility preserving constraints.

Proposed Direction of Work

The next generation of multidecadal predictive climate solvers will need efficient algorithms to determine sensitivities and correlation between **coupled** solution fields in order to quantify uncertainty in the models. Additionally, further research is necessary to better understand subgrid-scale parameterizations that result from multiscale interactions occurring in subsurface flow, tropical convection, and dynamic ocean-ice coupling problems.

The coupling methods for multimodel climate problems can be improved by utilizing adaptivity in the coupling schemes that can account for the wide characteristic variations in computational complexity and approximation properties among the component models. Achieving such a goal will require understanding sources of errors due to spatiotemporal discretization,

Contributed ideas paper 42

linearization choices of coupled terms, remapping between moving unstructured grids, and evaluating the scalability metrics of the algorithms as a function of memory, cost and data movement. The parametric sensitivity of the solution fields of interest (say global atmospheric and ocean temperature) on nonlinear coupling strength between models can be used to create nonlinear regression models. Such models can effectively drive the adaptive selection of computationally efficient algorithms on current and next-generation architectures, with minimal impact on numerical accuracy.

Adaptive block-structured (cubed-sphere) or unstructured (MPAS-based) mesh refinement of models [8] can resolve local geophysical fluid dynamics efficiently. Further investigations of conservative, local flux-preserving remapping schemes [9] can effectively reduce multimesh errors in advection of tracers such as water vapor in the atmosphere along with conservation of momentum and energy. Such an effort, however, also requires the availability of different linearizations in addition to remapping schemes under one flexible framework, such as the coupler/driver of the Accelerated Climate Model for Energy, in order to understand the pros and cons of the adaptive multiresolution choices [10].

Connections to Math, Computer Science & and Climate Science

Individual low- and intermediate-resolution of models are already at their scalability limits on existing machines, and leveraging the exponential increase in computing power by the traditional method of increasing resolution will not be possible without high-throughput numerical algorithms that show no signs of arriving. Maximizing computational efficiency while accurately resolving complex interactions in multimodel climate problems through algorithmic adaptivity requires detailed investigation beyond what is available in the current literature. Spatial adaptivity for multiresolution of atmospheric flows [11] and ocean modeling [12] has been pursued separately in the past, but general extensions to the full climate system requires extensive research.

Climate models are currently static in their numerical schemes, grids, and coupling strategies and may require new frameworks to make them more dynamically adaptive. The availability of such a unified climate modeling framework can also alleviate the generation of ensemble calculations that comprehensively captures scale variations in the coupled models in various problem regimes. Such Earth-system models can then utilize deterministic, adaptive selection of scalable algorithms through machine learning and nonlinear regression analysis, derived from importance-based sampling of the parametric space for ensemble calculations.

Potential Impact on the Field

For two decades, climate modelers have been able simply to “turn up the crank” on the resolution and process models, to improve the fidelity of climate simulations. Because future architectures offer more, but weaker, cores, this approach will no longer provide easy benefits. Instead, adaptive strategies must be developed that result in a simulation platform where numerical errors from discretization, linearization, and interpolation can be quantified consistently. This will be a crucial step forward toward realizing scalable, predictive computing on emerging architectures for climate applications.

Contributed ideas paper 42

References

- [1] W Washington, DB Bader, B Collins, et al. Challenges in climate change science and the role of computing at the extreme scale. In *Proc. of the Workshop on Climate Science*, 2008.
- [2] David A Randall, Richard A Wood, Sandrine Bony, et al. Climate models and their evaluation. In *Climate Change 2007: The physical science basis. Contribution of Working Group I to the Fourth Assessment Report of the IPCC (FAR)*, pages 589–662. Cambridge University Press, 2007.
- [3] Mu Mu, Duan Wansuo, and Wang Jiacheng. The predictability problems in numerical weather and climate prediction. *Advances in Atmospheric Sciences*, 19(2):191–204, 2002.
- [4] Jon Reisner, Andrzej Wyszogrodzki, Vincent Mousseau, and Dana Knoll. An efficient physics-based preconditioner for the fully implicit solution of small-scale thermally driven atmospheric flows. *Journal of Computational Physics*, 189(1):30–44, July 2003.
- [5] Brian J Soden and Isaac M Held. An assessment of climate feedbacks in coupled ocean-atmosphere models. *Journal of Climate*, 19(14):3354–3360, 2006.
- [6] Florian Lemari, Eric Blayo, and Laurent Debreu. Analysis of ocean-atmosphere coupling algorithms: Consistency and stability. *Procedia Computer Science*, 51:2066–2075, 2015.
- [7] Paul Aaron Ullrich, Peter Hjort Lauritzen, and Christiane Jablonowski. Some considerations for high-order ‘incremental remap’-based transport schemes: Edges, reconstructions, and area integration. *International Journal for Numerical Methods in Fluids*, 71(9):1131–1151, July 2012.
- [8] Julia Slingo, Kevin Bates, Nikos Nikiforakis, et al. Developing the next-generation climate system models: challenges and achievements. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 367(1890):815–831, 2009.
- [9] Jeffrey Grandy. Conservative remapping and region overlays by intersecting arbitrary polyhedra. *Journal of Computational Physics*, 148(2):433–466, January 1999.
- [10] Thomas Reichler and Junsu Kim. How well do coupled models simulate today’s climate? *Bulletin of the American Meteorological Society*, 89(3):303–311, 2008.
- [11] Jörn Behrens. Adaptive atmospheric modeling: Key techniques in grid generation, data structures, and numerical operations with applications. *Lecture Notes in Computational Science and Engineering*, 2006.
- [12] MD Piggott, GJ Gorman, CC Pain, et al. A new computational framework for multi-scale ocean modelling based on adapting unstructured meshes. *International Journal for Numerical Methods in Fluids*, 56(8):1003, 2008.

Contributed ideas paper 43

Title: Exascale-Computing Enabled Reparameterization of Biosphere Processes to Deliver High-Resolution, Multi-Scale, Fully-Coupled Climate Models

Date of Submission: Tuesday, December 1, 2015

Main Author Name: Powell, Amy J.

Main Author Affiliation: Sandia National Laboratories, Albuquerque, NM

Additional Authors: William F. Spotz, Mark A. Taylor (Sandia National Laboratories);

Abstract

As global economic activity increases, so do anthropogenic emissions, and effective mitigation steps altering this trajectory have yet to be developed, due in part to low resolution in climate models. To achieve global participation in mitigation regimes, models must give a clearer, more reliable picture of the interaction of the changing climate and the biosphere. Currently, biotic processes, represented as biogeochemical (BGC) cycling, are crudely parameterized for terrestrial and ocean model components. To develop multiscale, high-resolution models accurately detailing and predicting the interaction of biochemical networks and climate, we propose combining state-of-the-science measurements and bioinformatics resources with advanced linear algebra methods to estimate metabolic fluxes at multiple scales, towards the goal of reparameterizing biosphere processes robustly. Innovation resulting from this work will include multiscale, high-resolution methods for estimating ecosystem-level metabolic fluxes coupled to climate models, identification of biochemical pathways that respond to climate change pressures and exascale computing paradigms for investigating biosphere-climate interaction scenarios.

Background

Multi-scale, high-resolution components (land, sea-ice, ocean and atmosphere) of biosphere processes for Earth System Models, such as the Department of Energy's (DoE) Accelerated Climate Model for Energy (ACME), are among the most pressing scientific needs of the early twenty-first century, yet current versions do not include any of the last three decades' advances in the biological sciences, such as inexpensive, high-throughput genome-based analyses and the concomitant development of sophisticated open-source bioinformatics and computational biology resources that have yielded disruptive insights to numerous fields, including medicine, agriculture and computing [1]. Cross-cutting insights include the staggering diversity, abundance and ubiquity of microbial (heterotrophic) life on the planet, and their fundamental control of BGC cycling (*e.g.*, nitrogen transformation, carbon turnover, integration of terrestrial primary-producer communities, trophic structuring of marine ecosystems).

Current methods for estimating BGC responses and feedbacks focus on carbon (C), nitrogen (N) and phosphorous (P) flows, using high-level, arbitrary "functional types" with fixed C:N:P signatures to represent biologically-mediated nutrient cycling and physical-biological hydrologic processes for land and oceans components [2]. This hyper-reductive, dated approach for modeling biotic processes that potentially interact with climate is problematic, and possibly accounts for significant biases in land-to-air CO₂ flux and ocean uptake, and known weak responses to the El Niño Southern Oscillation [2].

Proposed work

To develop state-of-the-science, fully-coupled biological parameterizations for ACME, the climate-modeling community should use genome-based measurements and publicly-available bioinformatics resources to inventory biochemical pathways in an environment, and estimate fluxes through ecosystem-scale biochemical networks to identify and predict network sections potentially interacting with climate. The proposed work is feasible because biochemical networks, the full complement of which represents an organism's metabolic potential, are readily expressed in the language of linear algebra; stoichiometric matrices, containing biochemical reactions by metabolite data, can be used to compute flux vectors at individual pathway to meta-metabolism-scales [3]. Metabolic flux, expressed mathematically as serial sparse matrix vector products, can be made performance-portable over threads, GPUs and diverse next-generation architectures using DoE tools and expertise. Estimating metabolic flux for a single organism is routine with "omics" datasets, but the scope and complexity of scaling such analyses for ecosystems, represent a multi-phased challenge. Other phases include robustly integrating new ecosystem-scale metabolic flux data with climate models, and exercising the re-parameterized models under perturbation scenarios in exascale computing environments.

Improving Numerics for Parameterized Physics

December 1, 2015

Hui Wan

Pacific Northwest National Laboratory

Phil Rasch (Pacific Northwest National Laboratory), Peter Caldwell, Carol Woodward, and David Gardner (Lawrence Livermore National Laboratory)

Abstract

We advocate systematic evaluation and improvements to numerical methods for the parameterization of sub-grid scale processes in climate models through close collaboration between climate scientists and applied mathematicians. The effort would address a substantial source of structural uncertainty in current models, and provide more a solid basis for the development of better parameterization schemes.

Background/Research to Date

State-of-the-art atmosphere/ocean general circulation models are numerical solvers for very complex sets of integro-differential equations. The representation of sub-grid scale processes, often referred to as the “physics parameterization”, is known to strongly affect characteristics of the discrete solutions, but very limited attention has been paid to the numerical accuracy in this part of the models. The traditional mindset was that approximations introduced in the analytic development of the physics parameterization overwhelm the inaccuracies introduced by the discretization itself [1]. While this may be the case in some models and circumstances, our analysis [1] has suggested that discretization errors can be sufficiently large that the solutions can easily be dominated by this error (Figure 1). Studies have also shown that time integration methods and/or time step length can substantially alter the simulated cloud properties and climate sensitivity [e.g. 3, 4]. Awareness of the issue of numerical errors in parameterized physics has been increasing [4], but more concrete and coordinated efforts are needed to address this very complex problem.

Proposed Direction of work

We believe a comprehensive approach is needed to evaluate and improve time integration methods for sub-grid scale processes and their interactions in state-of-the-art climate models. The first component of the effort would be to better understand the properties of the mathematical equations. Are the equations deterministic or stochastic? Are there terms in the equations that make the system stiff? What kind of convergence properties should be expected of the discrete solutions? The second component would be to establish methods to isolate, quantify, and attribute time stepping errors in current models, in contexts that can facilitate understanding, but




Figure 1. Time integration error in the CAM5 model and its dependence on time step length. The error is defined as the root-mean-square difference in temperature (unit: K) after a 1-hr simulation, calculated against a reference solution obtained using a 1-s time step. Different colors correspond to different model configurations. “Full Model”: standard CAM5 with the spectral element dynamical core and real-world land-sea mask and topography. All other simulations: Dynamical core (Dyn) plus one parameterization in an aqua-planet setup. “St Cld”: stratiform cloud macrophysics and microphysics; “Sh Cu” and “Dp Cu”: shallow and deep cumulus convection; “Rad”: radiation. From [2].

Contributed ideas paper 44

meanwhile are relevant to the modeling of the real world (e.g. using models of a hierarchy of complexity). It would be beneficial to address both components concurrently, since progress in one component may provide helpful insights for the other. The ultimate goal is to develop computationally efficient solution methods that produce small time integration errors at a cost that allows meaningful climate simulations to be feasible. (Current “operational” time steps in global models range from a few minutes to an hour, which allows multi-year and multi-century simulations to be feasible.)

Connections to Math, Comp Sci & and Climate Science

The proposed work requires direct collaboration between climate modelers and applied mathematicians. On the math side, we would be addressing the topic of time integration methods for complex, nonlinear, and possibly stiff systems. On the climate science side, we would detangle complex dependencies between physical assumptions and numerical implementation, and reveal/remove compensating errors. Computer scientists could help in the design of solvers with desirable load balance and scalability.

Potential Impact on the field

Accurate numerical methods will ensure that the resulting climate simulations are correct solutions of the intended mathematical equations. This not only will help reduce numerical uncertainties in current models, but also will provide a better platform for climate scientists’ efforts in developing advanced, scale-aware parameterizations. An exemplary collaboration between climate modelers and applied mathematicians is likely to motivate stronger future interactions between the two fields. In the long run this will help achieve and maintain numerical robustness for climate models and simulations.

References

- [1] Beljaars, A. (1991): Numerical schemes for parametrizations. Proceedings of the ECMWF seminar on numerical methods in atmospheric models, 1–42. European Centre for Medium-range Weather Forecasts, Reading, United Kingdom.
- [2] Wan, H., P. J. Rasch, M. A. Taylor, and C. Jablonowski (2015), Short-term time step convergence in a climate model, *J. Adv. Model. Earth Syst.*, 7, 215–225, doi:10.1002/2014MS000368.
- [3] Wan, H., P. J. Rasch, K. Zhang, Y. Qian, H. Yan, and C. Zhao (2014), Short ensembles: An efficient method for discerning climate-relevant sensitivities in atmospheric general circulation models, *Geosci. Model Dev.*, 7, 1961–1977, doi:10.5194/gmd-7-1961-2014.
- [4] Gettelman, A., H. Morrison, S. Santos, P. Bogenschutz, and P. H. Caldwell (2014a), Advanced two-moment microphysics for global models. Part II: Global model solutions and aerosol-cloud interactions, *J. Clim.*, doi:10.1175/JCLI-D-14-00103.1.
- [5] Gross, M., S. Malardel, C. Jablonowski and N. Wood (2015), Bridging the (Knowledge) Gap between Physics and Dynamics, *Bull. Amer. Meteorol. Soc.*, doi:10.1175/BAMS-D-15-00103.1, in press.

Contributed ideas paper 45

Space-Time Adaptivity For Exascale Climate

December 1, 2015

Daniel F. Martin

Lawrence Berkeley National Laboratory (LBNL)

Hans Johansen (LBNL), Phillip Colella (LBNL)

Abstract

Many problems in climate exhibit large variations in spatial and temporal scale which can be addressed by local refinement in space and time. Adaptive Mesh Refinement has already shown promise in climate-related applications such as ice sheet modeling and atmospheric dynamics. As progress toward higher resolution and exascale computing continues, algorithmic innovations such as higher-order discretizations coupled with framework-driven software improvements promise to make AMR more mainstream for computational climate science.

Background/Research to Date

Block-structured adaptive mesh refinement (AMR) has proven useful in many scientific applications exhibiting large variations in scale. Originally developed for compressible flows, it was soon extended to a variety of applications, including mesoscale atmospheric modeling [12], incompressible flows [1, 8], combustion [5], and magnetohydrodynamics [7, 11]. In the climate community, perhaps the biggest AMR success to date is in ice sheet modeling. Ice sheets are particularly well-suited for AMR because they demand extremely high (sub-kilometer) resolution near rapidly-evolving grounding lines, while remaining relatively quiescent over much of the remaining continental-scale domains. The DOE-funded BISICLES ice sheet model, based on the Chombo AMR framework, is able to fully resolve the dynamics of the Antarctic ice sheet at a computationally-tractable cost (Figure 1) [2, 3]. We have also developed high-order finite volume methods on mapped grids [4, 10], with ongoing development of applications to gyrokinetic modeling in tokamaks [6] and atmospheric fluid dynamics [9].

Proposed Direction of work

We plan to extend our work in efficient high-order adaptive finite-volume methods in several climate-related directions. These include the use of higher-order discretizations in space and time, which include a more accurate treatment of orography (using cut cells), better-resolved planetary boundary layer interactions, and high-resolution tracking of extreme weather events. We plan to implement these capabilities in our existing climate-related efforts in ice sheet modeling (where cut cells are critical for representing grounding lines) and cubed-sphere finite-volume atmospheric models. Future opportunities for this technology could be surface / subsurface hydrology and oceanography, among others. At the same time, much work will be needed on the computer-science side ensuring that these methods remain computationally efficient at scale on emerging architectures. Issues that will need to be addressed include: efficient implementations of mixtures of regular and irregular calculations as they arise in multiblock, AMR, and cut-cell discretizations; the effective use of hierarchical parallelism; and performance portability across many-core and GPU-based architectures. To address these issues, we will need to collaborate with the computer science community to effectively use and customize more general software solutions.

Contributed ideas paper 45

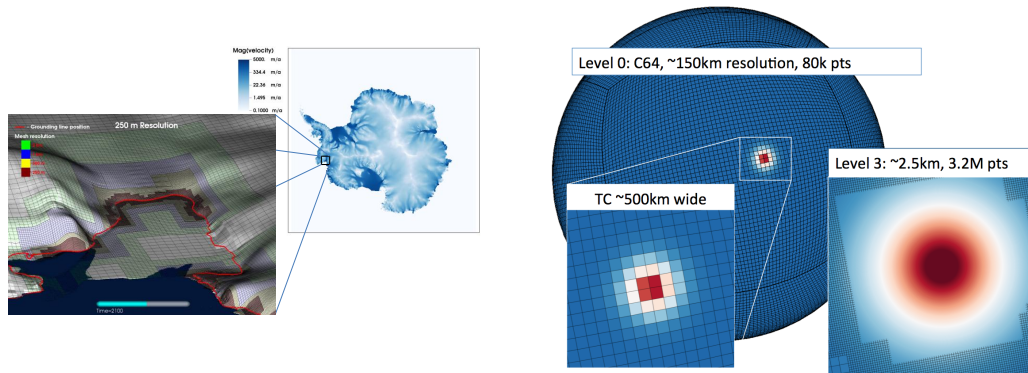


Figure 1: Examples of AMR in climate applications. *(left)* BISICLES-computed full-continent Antarctic vertically-averaged ice velocity field. Inset shows placement of mesh refinement for the Pine Island Glacier. Finest spatial resolution is 250m near grounding lines, while the coarsest mesh resolution is 8 km. *(right)* Coarse resolution for a tropical cyclone benchmark for atmospheric simulations are severely under-resolved. 3 levels of AMR refinement fully resolves the TC, at a savings of 1,000x the number of points compared to a uniform mesh (3.2B points at 2.5km). Time stepping and other multi-scale physics improvements could lead to a factor of 10,000x or more improvement in efficiency.

Connections to Math, Comp Sci & and Climate Science

ASCR-sponsored development has been at the heart of most AMR success stories, including BISICLES and AMR climate (both SciDAC). This has taken two forms: mathematical algorithmic development, including higher-order methods with better floating-point intensity, as well as CS-based software framework development efforts (Chombo, BoxLib). Partnerships with BER like the ISICLES and PISCEES projects (for ice sheet modeling) and the Multiscale climate SciDAC FASTMath partnership have proven invaluable for making the connection between these development efforts and climate science.

Potential Impact on the field

The movement toward the next-generation computing promises a revolution in our ability to fully resolve the processes which define our climate. Paradoxically, this will increase the importance and usefulness of AMR methods. In many applications using uniform refinement, the entire solution domain is under-resolved, due the limits of computational resources. Given adequate resources, higher-order methods combined with AMR can play an essential role in ensuring that computational resources are used as effectively as possible to fully-resolve the relevant physics, while dynamically adapting to changes in the system.

References

- [1] Almgren, et al (incl. P. Colella), A Conservative Adaptive Projection Method for the Variable Density Incompressible Navier-Stokes Equations,” *Journal of Computational Physics.*, 142, pp. 1-46, 1998

Contributed ideas paper 45

- [2] Cornford, et al (incl. D.F. Martin), "Adaptive mesh, finite volume modeling of marine ice sheets", *Journal of Computational Physics*, 232(1), pp. 529-549, 2013.
- [3] Cornford, et al (incl. D. F. Martin), "Century-scale simulations of the response of the West Antarctic Ice Sheet to a warming climate", *The Cryosphere*, August 18, 2015.
- [4] P. Colella, M.R. Dorr, J.A.F. Hittinger, and D.F. Martin, "High-Order Finite-Volume Methods in Mapped Coordinates", *Journal of Computational Physics*, 230(8), pp. 2952-2976 2011.
- [5] Day M.S. and J. B. Bell Numerical Simulation of Laminar Reacting Flows with Complex Chemistry," *Combust. Theory Modelling* 4(4) pp.535-556, 2000.
- [6] Dorf, et al (incl. P. Colella, D.F. Martin), Numerical modelling of geodesic acoustic mode relaxation in a tokamak edge. *Nuclear Fusion*, 53(6), 063015, 2013.
- [7] Li, et al (incl. D.F. Martin). A stable, accurate methodology for high mach number, strong magnetic field MHD turbulence with adaptive mesh refinement: resolution and refinement studies. *The Astrophysical Journal*, 745(2), 139, 2012.
- [8] Martin D., Colella, P., "A Cell-Centered Adaptive Projection Method for the Incompressible Euler Equations", *Journal of Computational Physics*,.163 (2000), pp. 271-312, 2000.
- [9] McCorquodale, et al (incl. H. Johansen, P. Colella). An adaptive multiblock high-order finite-volume method for solving the shallow-water equations on the sphere, *Communications in Applied Mathematics and Computational Science*, 10(2), 121-162, 2015.
- [10] McCorquodale, et al (incl. P. Colella), "High-order finite-volume methods for hyperbolic conservation laws on mapped multiblock grids", *Journal of Computational Physics*, 288, pp. 181-195, 2015.
- [11] Miniati, F., and Martin, D. F. .Constrained-transport magnetohydrodynamics with adaptive mesh refinement in CHARM. *The Astrophysical Journal Supplement Series*, 195(1), 5, 2011.
- [12] Skamarock, W. C., and Klemp, J. B. Adaptive grid refinement for two-dimensional and three-dimensional nonhydrostatic atmospheric flow. *Monthly Weather Rev*, 121(3), pp. 788-804, 1993.

Contributed ideas paper 46

Expanding the Utility of High-Resolution Global Climate Models via Short Ensembles

Dec. 1, 2015

Salil Mahajan

Oak Ridge National Laboratory

Katherine J. Evans and Mathew R. Norman (Oak Ridge National Laboratory)

Abstract

Conducting long simulations with high resolution climate models is a very time-consuming and expensive endeavor. The cost of obtaining stationary noise statistics of the system makes high-resolution climate model configurations unusable for traditional climate studies - where long control and perturbed simulations are the norm. Using short simulation ensembles to quantify natural variability on shorter time-scales of interest could enhance the utility of high-resolution climate models for predictability or process studies, for e.g. to study the impact of anthropogenic emissions on the near-term climate. We discuss approaches using short ensembles to answer critical climate science questions with high-resolution climate models.

Background

There are two main purposes of running long simulations with a fully coupled global climate model. The first is to propagate the shock from the ill-posed initial conditions of different components or cold starts, which take hundreds of simulated years because of the low frequency variability associated with oceans, ice and biogeochemistry cycles. The second purpose is to quantify the natural variability of the system as a stationary process (stationary noise statistics). Traditionally, climate models are run for hundreds of simulated years after they appear to have reached equilibrium to capture the natural variability on a large spectrum of timescales relevant to the climate system. The benefits of running a long simulation at equilibrium are (1) large sample sizes for relatively high-frequency variability (2) capture some cycles of low-frequency variability (e.g. time period of the ocean meridional overturning circulation is about 1000 years).

The current high-resolution climate models are computationally expensive, for example the ACME v0 quarter degree configuration requires about 500K core-hours per simulated year. Further, the wall-time to integrate these high-resolution model configurations is also long (10 simulated years/month). Thus, conducting long simulations with the model to evaluate its climate statistics is a very expensive endeavor. In fact, it may not be feasible to allow the fully coupled model to reach an equilibrium state. Practically, with the current computational resources, one should only expect short simulations with high-resolution configurations of climate models. But, the hurdles in obtaining stationary noise statistics of the system would make such configurations unusable for climate studies using traditional methods - where long control and perturbed simulations are the norm.

Operational weather and climate forecast systems like the initialized weather forecasting systems, seasonal prediction systems, decadal prediction systems; and even regional atmosphere-ocean coupled simulations, routinely use an ensemble of short simulations. Model biases, deficiencies and inaccurate initial conditions tend to drift the weather/climate model state that would move towards a different attractor if the model were allowed to run for long periods. Thus, models in such situations are only integrated for a short length whence the model-drift is not a significant contributor to the model state evolution.

Contributed ideas paper 46

Proposed Direction of work

Implementing such a strategy of short ensembles in global climate model studies that focus on predictability or process studies, for e.g. to study the impact of anthropogenic emissions on the near-term climate, would expand the usability of high-resolution models. Recent work has shown that running short ensembles instead of a long simulation allows for an efficient use of computing resources for the same number of simulated years by packing more computational work on each compute node (Norman et al. 2015, in prep.). Additionally, short ensembles also allow for a faster throughput as the individual runs can be run simultaneously (in parallel).

If an equilibrium model state is indeed achieved, an ensemble of short runs may be used to quantify natural variability at shorter timescales of interest. Baker et al. (2015) recently used 1-yr short ensembles for model verification, albeit for a low-resolution model (100 km). It may also be worthwhile to investigate strategies to utilize a drifting model state, much like prediction systems. We elaborate on two possible applications of short ensembles below:

- **Short Perturbed Ensembles:** Traditionally, to isolate the impact of a perturbation/forcing on the climate, a long single-forcing simulation is compared with a long control simulation. The long simulations ensure that the effects of low-frequency variability are similar between the two runs. For questions addressing shorter time-scales, say the impact of dust aerosols on Atlantic hurricanes, alternatively one could integrate short forced runs starting from various points on the control run trajectory, say integrating from May-Dec. with perturbed dust aerosols initialized from several May states from the control simulation. Initializing from the trajectory of the control integration ensures that the climate state would evolve similarly without the perturbation. These short simulations do not allow the low frequency signals to damp climate response signals thus enhancing the signal-to-noise ratio when compared to the corresponding segments of the control run. The advantage of such short runs is that they require small sample sizes (fewer simulated years) and result in a quick throughput as they can be run in parallel. Recently, Verma et al. 2016 (in prep.) used such an approach to isolate the impact of sulfate aerosols on the tropical Pacific from that of ENSO variability with a low-resolution model.
- **Using a Model with a Drifting Model State:** In the absence of an equilibrated model state, the possible trajectories of the model from a drifting model state (not in equilibrium) can instead be quantified for a short simulation length of interest with an ensemble. Each short simulation will start from the same point on the drifting model trajectory, but with slightly different start conditions (machine precision perturbations). The short ensembles would include both the natural variability of the system at the shorter time-scales and the trend associated with the drift. An ensemble of forced simulations, say with modified aerosols, beginning at the same model state quantifies the impact of the forcing, the impact of the drift plus the non-linearities and feedbacks associated with the interaction of the forcing and the drift. Ideally, in the absence of non-linearities, detrending the forced ensemble with the trends in the control ensemble would allow for quantification of the impact of the forcing on shorter time-scales. Nonetheless, the impact of non-linearities could be crudely estimated by conducting the control and forced experiments with a model configuration that is at equilibrium, say for a low-resolution model for scale-independent interactions.

The use of short ensembles can be applied to understand tele-connections, climate extremes, impact of short-lived and long-lived forcing on short time-scales, etc. in high resolution global climate models. For variability on longer (say decadal) timescales, the length of the short-runs could be extended to several years based on the time-scale of interest, which would still be beneficial over single long simulations.

Contributed ideas paper 46

References

- [1] Baker, A. H., Hammerling, D. M., Levy, M. N., Xu, H., Dennis, J. M., Eaton, B. E., Edwards, J., Hannay, C., Mickelson, S. A., Neale, R. B., Nychka, D., Shollenberger, J., Tribbia, J., Vertenstein, M., and Williamson, D. (2015): A new ensemble-based consistency test for the Community Earth System Model, *Geosci. Model Dev. Discuss.*, 8, 3823-3859, doi:10.5194/gmdd-8-3823-2015
- [2] Norman, M. R., A. L. Gaddis, V. G. Anantharaj, K. J. Evans, S. Mahajan, P. H. Worley (2015): Computational Benefits of an Ensemble-Based Approach to Climate Modeling and Testing at Scale (*in preparation*)
- [3] Verma, T., S. Mahajan, R. Saravanan, P. Chang (2016): Oceanic Feedback in Regional Climate Response to Sulfate Aerosol Forcing (*in preparation*)

Contributed ideas paper 47

Sea Ice Modeling Across Scales at Exascale and Beyond

December 1, 2015

Wieslaw Maslowski¹, Andrew Roberts¹, Elizabeth Hunke², Michal Kopera¹, Frank Giraldo¹

¹Naval Postgraduate School, ²Los Alamos National Laboratory

Abstract

As the climate and Earth system modeling community approaches the capability of exascale computing and beyond, the overall trend is to resolve physical processes and resulting feedbacks within and across individual model components. This trend increases the urgency of addressing the current limitations of sea ice models in representing both the sub-floe scale (i.e. 1 km or less or sub-kilometer) and multi-floe or continuum scale (i.e. 10 km or more). We outline a possible new approach to modeling sea ice dynamics across these scales and discuss potential advantages and limitations of combining highly scalable mathematical methods, such as Discontinuous and Continuous Galerkin (DG and CG) methods, for use in the next generation of sea ice models.

Background/Research to Date

The Arctic is undergoing rapid and some of the most coordinated climate changes currently occurring anywhere on Earth, with Arctic sea ice cover changes exceptional in at least the last 1400 years [1] and related surface temperature extremes unusual in at least the past 600 years [2]. Historical model reconstructions of the high north from global climate and Earth System models (GC/ESMs) are in broad agreement with these changes; however, the rate of change in the GC/ESM forecasts remains outpaced by observations [3]. Reasons why models may not be able to simulate rapid environmental change in the Arctic stem from a combination of coarse model resolution, inadequate parameterizations, unrepresented processes and spatio-temporal scales, and a limited knowledge of physical and other real world interactions.

Model limitations are hindering our ability to predict future state of Arctic sea ice and its impacts on regional and lower latitude climate. These problems are important, because a reduction of perennial sea ice cover exposes open water to direct interactions with the atmosphere, which in turn influences regional atmospheric circulation patterns and temperature profiles, especially along the seasonal marginal ice zone. Sea ice thickness variability in space and time modifies the Arctic-wide atmospheric circulation and appears to impact the troposphere-stratosphere coupling [4]. A more realistic representation of time-dependent conditions of the Arctic sea ice cover and their effect on air-sea interactions is necessary in models and it requires coupling of the respective model components.

Contrary to atmospheric and oceanic models, sea ice models are currently unable to universally simulate the frozen ocean from large scale (10-1000 km) down to very small, sub-kilometer or sub-floe scale. Continuum sea ice model dynamics are designed to operate at or above the so-called multi-floe scale, i.e. at resolutions of 10 km or more [5]. Meantime discrete element models, resolving individual floes and their interactions, continue needing refinements to simulate mechanical constraints between floes [6]. Neither class of models has capability to simulate sea ice successfully both at the sub-floe scale and the multi-floe or continuum scale across entire basins and on climatic timescales. However, development of ocean and atmospheric models using quasi-uniform grids with local refinement demand the use of multi-scale sea ice physics in models, to soon span the sea ice floe-scale divide.

Proposed Direction of work

While there is currently a strong focus in sea ice modeling research on resolving sea ice at the sub-floe scale [7,8], little attention has been given to the possibility that beneath horizontal scales of about 1km, sea ice may often be non-hydrostatic. Recent work (in preparation by Roberts and Hunke) indicates that there is an analytic finite non-hydrostatic horizontal length scale limit for sea ice, and this may be the most important consideration in modeling sea ice dynamics across scales, rather than the horizontal length-scale of floes. This morphological approach links small and large-scale sea ice dynamics, and may facilitate multi-scale sea ice dynamics models.

Contributed ideas paper 47

In addition, finite element or discrete granular Lagrangian models [6,7,8] are often very expensive computationally, as they have to search and follow collisions between particles, which leads to highly irregular communication patterns and significantly reduces computational scalability. Instead, we argue that the above morphological approach when combined with highly scalable discontinuous Galerkin (DG) methods [9] will be more suitable for use of the exascale computing capability in the next generation of sea ice models. In particular, DG methods could be useful for a very accurate interpolation to an Eulerian mesh or adaptive mesh tracking as well as for the solution of sea ice velocity field and transport.

Connections to Math, Computational and Climate Science

DG methods were successfully applied to atmospheric [9] and shallow water [10,11] models, including extended the capabilities of those models to non-conforming dynamically adaptive meshes [12]. Scalability study of the atmospheric model NUMA [13] showed strong scaling of the CG and DG method for up to 32k processors, up to 3.14 million MPI threads and sustained performance of 1.2 Petaflops, which allowed a 3 km uniform global resolution simulation of a baroclinic instability test case within the time frame required for operational weather prediction. In addition to modeling sea ice at multi-scale in the Arctic and Southern Ocean, there is another important requirement for this capability, related to the sea ice role at the ice-sheet/ocean interface. A recently funded DOE project, involving some of the authors of this paper, is aiming to leverage the NUMA technology for constructing a geometrically flexible variable resolution model for the Greenland Fjords with the goal to improve the representation of the ice-sheet/ocean interactions in coupled climate models.

Potential Impact on the field

Advanced modeling of multi-scale sea ice physics will (i) advance understanding of physical processes and feedbacks involved in polar amplification of global climate change, (ii) facilitate studies of sea-level rise from Greenland and Antarctic ice sheet mass loss and (iii) improve understanding and potentially reduce uncertainty in prediction of polar climate change at seasonal to decadal scales.

References

- [1] Kinnard, C., Zdanowicz, C. M., Fisher, D. A., Isaksson, E., de Vernal, A. and Thompson, L. G. 2011. Reconstructed changes in Arctic sea ice over the past 1,450 years, *Nature*, 479(7374).
- [2] Tingley, M. P., and Huybers, P. 2013. Recent temperature extremes at high northern latitudes unprecedented in the past 600 years, *Nature*, 496(7444), doi:10.1038/nature11969.
- [3] Stroeve, J. C., Kattsov, V., Barrett, A., Serreze, M., Pavlova, T., Holland, M., and Meier, W. N. 2012. Trends in Arctic sea ice extent from CMIP5, CMIP3 and observations, *Geophys. Res. Lett.*, 39(16), doi:10.1029/2012GL052676.
- [4] Krinner, G., Rinke, A., Dethloff, K., Gorodetskaya, I.V. 2010. Impact of prescribed Arctic sea ice thickness in simulations of the present and future climate. *Climate Dynamics* 35:4, 619-633.
- [5] Hunke, Elizabeth C., Lipscomb, William H., Turner, Adrian K. 2010. Sea-ice models for climate study: retrospective and new directions. *J. Glaciol.*, V. 56, N. 200, pp.1162-1172(11)
- [6] Hopkins, Mark A. 2004. A discrete element Lagrangian sea ice model", *Engineering Computations*, Vol. 21 Iss 2/3/4 pp. 409 - 421 (<http://dx.doi.org/10.1108/02644400410519857>)
- [7] Bouillon, S. and Rampal, P. 2015. Presentation of the dynamical core of neXtSIM, a new sea ice model, *Ocean Model.*, 91, 23–37, 5890, 5891, 5892, 5914.
- [8] Baohui, Li, Li Hai, Liu Yu, Wang Anliang, Ji Shunying. 2014. A modified discrete element model for sea ice dynamics. *Acta Ocean. Sinica*, 33(1): 56–63, doi: 10.1007/s13131-014-0428-3.
- [9] Restelli, M., and Giraldo, F.X. 2009. A Conservative Semi-Implicit Discontinuous Galerkin Method for the Navier-Stokes Equations in Nonhydrostatic Mesoscale Atmospheric Modeling, *SIAM J. Sci. Comp.*, Vol. 31, 2231-2257.
- [10] Giraldo, F.X., and Restelli, M. 2009. High-Order Semi-Implicit Time-Integration of a Triangular Discontinuous Galerkin Oceanic Shallow Water Model, *Int. J. Num. Meths. Fluids*, Vol. 63, 1077-1102.
- [11] Nair, R. D., Thomas, S. J., and Loft, R. D. 2005. A discontinuous Galerkin global shallow water model. *Monthly Weather Review*, 133(4), 876-888.
- [12] Kopera, M.A., and Giraldo, F.X. 2014. Analysis of Adaptive Mesh Refinement for IMEX Discontinuous Galerkin Solutions of the Compressible Euler Equations with Application in to Atmospheric Simulations, *J. Comput. Physics* Vol. 275, 92-117.
- [13] Kelly, J.F., and Giraldo, F.X. 2012. Continuous and Discontinuous Galerkin Methods for a Scalable 3D Nonhydrostatic Atmospheric Model: limited-area mode, *J. Comp. Phys.*, Vol. 231, 7988-8008.

Contributed ideas paper 48

New Algorithms for Tracer Transport
Philip Cameron-Smith
LLNL

2015-12-1

Abstract

Tracers are critical to simulating climate, evaluating climate against observations of various tracers, answering questions about how parts of the Earth system affect other parts, and evaluating the effectiveness of climate mitigation strategies. The problem is that tracer transport typically requires a lot of communication between grid-cells at small timesteps, which inhibits the scalability of climate simulations to large processor counts. And the computational challenge increases as global models move to higher resolution, so that tracer transport will be the largest computational requirement of climate models in the future. Dramatic algorithmic improvements are therefore required to fully realize the value of tracers to climate science on future compute platform.

Background/Research to Date

There have been many tracer transport schemes over the years, which generally fall under the categories of Lagrangian (following particles), Eulerian (fluid flow through a fixed grid) or some hybrid of the two. There are additional numerical constraints that are challenging to satisfy at the same time, conservation of mass, positive definiteness, concentration mixing, and relationships between different tracers.

Much of the computational work is focused on making marginal improvements to existing algorithms on new computer architectures. The main new algorithmic scheme I am aware of is the CSLAM semi-Lagrangian scheme. However, these do not seem to offer the algorithmic breakthrough needed to avoid the tracer transport bottleneck as resolution continues to increase.

Proposed Direction of work

CSLAM offers a good solution to scaling to large numbers of tracers, however it doesn't resolve scaling associated with shrinking timesteps and inter-processor communication. However, some of the limitations of CSLAM can be avoided by allowing Courant numbers larger than 1.

There may also be other advection schemes that have not yet been invented. The history of the field is that new and improved ideas are continually being invented, but it takes many years to bring an idea to fruition.

Connections to Math, Comp Sci & and Climate Science

New algorithmic schemes, that mathematically satisfy the conservation properties of tracer advection and mixing, as well as scale well to large tracer numbers, higher resolution, and larger computers, requires that Math, Comp Sci, and Climate Science work together..

Potential Impact on the field

See above.

Contributed ideas paper 49

Adaptive Cartesian Mesh Methods in Climate Modeling December 1, 2015

Donna Calhoun
Boise State University

Abstract

Dynamically adaptive Cartesian grid methods have been used for over 30 years as a means of making Cartesian grid calculations practical and scalable. These methods have been widely successful in several areas related to natural physical systems, including astrophysics, where most modern large scale codes use Cartesian adaptivity (FLASH, Castro, Ramses, [12, 1, 26]), tsunami modeling, flooding, landslides and storm-surges (GeoClaw [4]), and mantle convection (p4est [9, 18]). However, for several reasons, adaptive Cartesian grid methods have not been widely adopted by computational climate scientists, despite several efforts over the years to promote their use. However, given technical advancements in adaptive grid software, as well as increased computing power, time may be right to revisit these methods and promote their potential use in computational climate modeling.

Adaptive Mesh Refinement (AMR) : Background

Cartesian grid Adaptive Mesh Refinement (AMR) can be traced back to the original work of Berger and Oliger [5], and Berger and Colella [3]. This early work was largely used for solving hyperbolic conservation laws, particularly those involving localized spatial features such as shocks. Following closely on Berger's work, Skamarock, Oliger, Klemp and others applied AMR techniques to numerical weather prediction (NWP) [21, 20, 22]. These early efforts effectively demonstrated the potential for AMR to greatly improve computational efficiency in NWP, but also pointed out potential shortcomings. In particular, it was noted that abrupt transitions in resolution between coarse and fine grids could trigger physical instabilities in the model equations, resulting in noticeable differences between solutions computed on uniformly refined grids and those computed on AMR grids . A more fundamental concern was that grid-scale dependent parameterizations were now a potential source of non-deterministic behavior brought about by non-smooth dynamically varying grid resolutions and related truncation errors. Where AMR did a superb job of capturing relatively stable shocks in aerodynamic applications, it was now seen as a liability in problems with large scale marginally stable or unstable features.

These early efforts may have discouraged widespread adoption of dynamic grid adaptivity, but the computational benefits that variable resolution meshes can provide have been largely recognized by computational weather and climate scientists. Weather codes, such as WRF or the Finite-Volume dynamical core (FV3) from GFDL make use one-way or two-way nesting, where a fine grid over a region of interest is embedded in a larger global coarse grid [24, 14]. The two climate codes, MPAS and FV3, chosen for Phase 2 assessment of the NOAA sponsored program to identify their Next Generation Global Prediction System (NGGPS) use smoothly varying static meshes, even if fully dynamical adaptive capabilities may not be available or practical in either code [23, 17].

Towards more widespread use of adaptivity in climate modeling

While doubt may persist as to how suitable AMR techniques are for accurately resolving features of large-scale climate models, there can be little doubt that full dynamic adaptivity is

Contributed ideas paper 49

extremely computational efficient. Underscoring this point, the 2015 ACM Gordon-Bell Prize committee recently awarded J. Rudi et al. for an extreme scale implicit solver for PDEs modeling highly heterogeneous flow in the Earth's mantle. The recipients of this award, given each year in recognition of outstanding achievement in high performance computing, achieved optimal algorithmic performance in large part through aggressive octree adaptivity available in the P4EST library [18, 9, 8].

The DOE, through their SciDAC program has been enormously supportive of scientific software development, and in particular general purpose software for Cartesian block-structured AMR (CHOMBO, BOXLIB). What these software projects have demonstrated is that, with active involvement from core developers, adaptive methods can be used to solve large-scale problems central to the DOE mission. What is now needed is a concerted effort to get more adaptive codes into the hands of the climate scientists with the computational expertise to begin to address the issues first observed in [22].

One approach to making adaptive codes more appealing and accessible to these climate modelers is to make the separation between numerics and grid management more transparent, without hiding the details of the adaptive time stepping algorithms from potential users. This separation can be achieved through use of grid management libraries such as P4EST [9]. One effort at developing a new adaptive code base is the FORESTCLAW project.

In the FORESTCLAW project, Calhoun, in collaboration with C. Burstedde, one of the chief developers of P4EST, has developed a very light-weight, but scalable general purpose multi-block, Cartesian AMR code suitable for the kinds of investigations needed to address concerns about using adaptivity in climate models [7, 10, 11]. With suitable input from climate scientists, a program of addressing specific mathematical and computational issues regarding the use of AMR in climate modeling could be laid out. More importantly, these issues could be investigated in a concrete manner by using FORESTCLAW to write mini-application programs targeted at isolated concerns.

Connections to Math, Comp Sci & and Climate Science

J. Boyd, in a recent SIAM Newsletter article, summed up rather colorfully the challenges that “computational missionaries” face when they parachute into climate modeling. “Clever adaptive algorithms that work for smooth, straight shocks disintegrate into computational anarchy when flayed by gravity waves, assaulted by moist convective instability, battered by highly temperature-sensitive photochemistry, and coupled to the vastly different time and space scales of the ocean.” [6]. But developing and using AMR codes is a daunting challenge, and one that computational mathematicians and climate scientists alike are justifiably reluctant to tackle. However, the availability of lightweight, easy-to-use, highly scalable code base such as FORESTCLAW may provide the inspiration necessary encourage an “all-hands-on-deck” approach to tackling some of the concerns raised about used AMR methods in climate science.

Potential impact on the field

The US Air Force recently made the decision use the UK Met Office's Unified Model in favor of WRF [19]. One of the reasons given for this decision is that the UK model provides forecasts across a wide range of scales. Ironically, world-class expertise in the adaptive methods that could be used to achieve comparable capabilities in US codes are working in DOE laboratories. In light of the need to maintain US competitiveness in climate modeling, it is inconceivable that we would leave on the table a technology as successful as dynamic AMR.

Contributed ideas paper 49




Figure 1: Tracer transport on the sphere. This benchmark was designed by P. Lauritzen et al. [16] run using FORESTCLAW [7, 10]

References

- [1] A. S. Almgren, V. E. Beckner, J. B. Bell, M. S. Day, L. H. Howell, C. C. Joggerst, M. J. Lijewski, A. Nonaka, M. Singer, and M. Zingale. CASTRO: A New Compressible Astrophysical Solver. I. Hydrodynamics and Self-gravity. *Astrophys. J.*, 715(2):1221, 2010.
- [2] Marsha J. Berger, Donna Calhoun, Christiane Helzel, and Randall J. LeVeque. Logically rectangular finite volume methods with adaptive refinement on the sphere. *Phil. Trans. R. Soc. A*, 367(1):4483–4496, 2009.
- [3] Marsha J. Berger and Phillip Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.*, 82(1):64–84, May 1989.
- [4] Marsha J. Berger, David L. George, Randall J. LeVeque, and Kyle T. Mandli. The GeoClaw software for depth-averaged flows with adaptive refinement. *Advances in Water Resources*, 34(9):1195–1206, September 2011.
- [5] Marsha J. Berger and Joe Oliger. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.*, 53:484–512, 1984.
- [6] John Boyd. Multiscale numerical algorithms for weather forecasting and climate modeling: Challenges and controversies. *SIAM News*, 41(9), November 2008.
- [7] Carsten Burstedde, Donna Calhoun, Kyle T. Mandli, and Andy R. Terrel. FORESTCLAW: Hybrid forest-of-octrees AMR for hyperbolic conservation laws. In Michael Bader, Arndt Bode, Hans-Joachim Bungartz, Michael Gerndt, Gerhard R. Joubert, and Frans Peter, editors, *Parallel Computing : Accelerating Computational Science and Engineering (CSE)*, volume 25 of *Advances in Parallel Computing*, pages 253–262. IOS Press, 2014.
- [8] Carsten Burstedde, Omar Ghattas, Georg Stadler, Tiankai Tu, and Lucas C. Wilcox. Towards Adaptive Mesh PDE Simulations on Petascale Computers. In *Proceedings of Teragrid '08*, 2008.

Contributed ideas paper 49

- [9] Carsten Burstedde, Lucas C. Wilcox, and Omar Ghattas. `p4est`: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM J. Sci. Comput.*, 33(3):1103–1133, 2011.
- [10] Donna Calhoun and Carsten Burstedde. ForestClaw : A parallel solver for finite volume schemes on adaptive, multiblock Cartesian meshes. *in preparation*, 2015.
- [11] ForestClaw. <http://www.forestclaw.org>.
- [12] B. Fryxell, Kevin M. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, and Peter MacNeice. FLASH: An adaptive mesh hydrodynamics code for modeling astrophysical thermonuclear flashes. *Astrophys. J. Suppl*, 131(1):273–334, 2000.
- [13] Francis X. Giraldo, J. F. Kelly, and E. M. Constantinescu. Implicit-explicit formulations of a three-dimensional nonhydrostatic unified model of the atmosphere (numa). *SIAM J. Sci. Comput.*, 35(5):B1162–B1194, 2013.
- [14] Lucas M. Harris and Shian-Jiann Lin. A Two-Way Nested Global-Regional Dynamical Core on the Cubed-Sphere Grid. *Mon. Wea. Rev.*, 141(1):283–306, 2015/11/30 2012.
- [15] Christiane Jablonowski, Robert C. Oehmke, and Quentin F. Stout. Block-structured adaptive meshes and reduced grids for atmospheric general circulation models. *Philos. T. Roy. Soc. A*, 367(1907):4497–4522, 2009.
- [16] Peter H. Lauritzen, Paul A. Ullrich, Christiane Jablonowski, Peter A. Bosler, Donna Calhoun, A. J. Conley, T. Enomoto, L. Dong, S. Dubey, O. Guba, A. B. Hansen, E. Kaas, J. Kent, J.-F. Lamarque, M. J. Prather, D. Reinert, Vladimir V. Shashkin, William C. Skamarock, B. Sorensen, Mark A. Taylor, and M. A. Tolstykh. A standard test case suite for two-dimensional linear transport on the sphere: results from a collection of state-of-the-art schemes. *Geosci. Model Dev.*, 7:105–145, 2014.
- [17] MPAS Tutorial practice guide. <http://www2.mmm.ucar.edu/projects/mpas/tutorial/UK2015/>.
- [18] Johann Rudi, A. Cristiano I. Malossi, Tobin Isaac, Georg Stadler, Michael Gurnis, Peter W. J. Staar, Yves Ineichen, Costas Bekas, Alessandro Curioni, and Omar Ghattas. An Extreme-Scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth’s Mantle . In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, November 15-20 2015.
- [19] Jason Samenow. Air Force plans to drop US forecast system for U.K. model draws criticism. *Washington Post*, April 20th 2015.
- [20] William Skamarock, Joseph Oliger, and Robert L Street. Adaptive grid refinement for numerical weather prediction. *J. Comput. Phys.*, 80(1):27 – 60, 1989.
- [21] William C. Skamarock. *Adaptive Grid Refinement for Numerical Weather Prediction*. PhD thesis, Stanford University, October 1987.
- [22] William C. Skamarock and Joseph B. Klemp. Adaptive Grid Refinement for Two-Dimensional and Three-Dimensional Nonhydrostatic Atmospheric Flow. *Mon. Wea. Rev.*, 121(3):788–804, 1993.

Contributed ideas paper 49

- [23] William C. Skamarock, Joseph B. Klemp, Michael G. Duda, Laura D. Fowler, Sang-Hun Park, and Todd D. Ringler. A Multiscale Nonhydrostatic Atmospheric Model Using Centroidal Voronoi Tessellations and C-Grid Staggering. *Mon. Wea. Rev.*, 140(9):3090–3105, 2012.
- [24] William C. Skamarock, Joseph B. Klemp, Jimy Dudhia, David O. Gill, Dale M. Barker, Michael G. Duda, Xiang-Yu Huang, Wei Wang, and Jordan G. Powers. A description of the advanced research wrf version 3. NCAR Technical Note NCAR/TN-475+STR, NCAR, June 2008.
- [25] Julia Slingo, Kevin Bates, Nikos Nikiforakis, Matthew Piggott, Malcolm Roberts, Len Shafrey, Ian Stevens, Pier Luigi Vidale, and Hilary Weller. Developing the next-generation climate system models: challenges and achievements. *Philos. T. Roy. Soc. A*, 367(1890):815–831, 2009.
- [26] Roman Teyssier. Cosmology Hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *Astronomy and Astrophysics*, 385:337–364, 2002.
- [27] Fred Toepfer. Next Generation Global Prediction System (NGGPS), August 2015. Presentation at NOAA UMAC Meeting.

Contributed ideas paper 50

Solvers for Implicit Solution of Climate Dynamics

Dec. 1, 2015

Carol S. Woodward, Lawrence Livermore National Laboratory

David J. Gardner (LLNL) and Katherine J. Evans (ORNL)

Abstract

Efficient solutions of global climate models require effectively handling disparate length and time scales. Implicit solution approaches allow time integration of the physical system with a step size governed by accuracy of the processes of interest rather than by stability of the fastest time scales present. Implicit approaches, however, require the solution of nonlinear systems within each time step, and it is unclear as to the best way to solve these systems. Due to the expense of implicit methods, semi-implicit or IMEX schemes are also considered. These methods still require a nonlinear solve, but the system may be smaller or easier to solve. This paper presents two options and the related research necessary for efficient nonlinear solvers within climate dynamics.

Background/Research to Date

Currently, global climate models use explicit or semi-implicit time-stepping integration methods within their dynamical cores (Evans2013, ECMWF2012). However, increased interest in atmospheric models configured with overall finer and/or regionally refined spatial discretization is causing these methods to become intractably expensive due to the CFL restriction of the time step size. Implicit time integration methods become a feasible alternative when there is a significant separation between the scales of interest to resolve and the smallest scales of the system, with the exact crossover point being problem and optimization dependent. Using finer spatial grids solving the same flow balance equations increases this separation.

The solution of the nonlinear and subsidiary linear systems required at each time step of an implicit approach forms the critical bottleneck affecting the overall scalability of the implicit dynamical core (Evans2009). Most preconditioner development for atmospheric climate models has been focused on developing Helmholtz solvers that arise in semi-implicit formulations of the shallow-water equations (Thomas2003, Thomas2002). Recently, Yang et. al (Yang2010) developed a domain decomposition based preconditioner for the fully implicit shallow-water equations based on a finite volume discretization.

However, the default dynamical core in the ACME model is the spectral element dynamical core developed originally in the HOMME code. An implicit version of this core was first analyzed in (Evans2010). The implicit solution approach employed in this work used an unpreconditioned Newton-Krylov method with a finite difference Jacobian-vector multiply. While this approach was effective in evaluating the efficacy of an implicit approach, it was not efficient. Even efforts to improve the data structures and memory use cannot overcome the extra computational burden of the method and the solver library infrastructure, so scalability and expense of the preconditioner is a key aspect for the success of the method. Follow-on work (Lott2015) developed an algorithmically scalable preconditioner for the Newton-Krylov method. This work confirmed that an effective preconditioner can be developed.

Proposed Direction of work

However, the Lott2015 work showed that an investment in an implementation of explicit Jacobian entries would be required before a truly efficient preconditioner could be attained with a Newton-Krylov method. One significant benefit of spectral elements is that the effect of differential operators can be implemented only through matrix vector multiplies and explicit entries into the matrix need not be formed. However, with a preconditioned iterative method, these operators must be applied multiple times, and the re-evaluation of them is quite expensive.

Contributed ideas paper 50

As a result, standard implementations of spectral element operators cannot be used. An existing effort is currently developing explicit Jacobian entries for the 3D primitive equations for climate dynamics in CAM-SE. These matrix entries can be used both for efficient fully implicit and semi-implicit solution approaches. This existing effort will apply these entries to develop a preconditioned Newton-Krylov method. However, we advocate for additional work into efficient alternate preconditioning approaches such as multilevel methods, semi-implicit methods as preconditioners, etc. with the goal of developing an effective toolbox for solving a wide range of locally refined climate problems.

In addition, we advocate for research into other nonlinear solvers for implicit systems within either fully implicit or partitioned (IMEX) approaches to climate dynamics. In particular, the use of Anderson accelerated fixed point methods for solving nonlinear systems within each time step should be explored. Fixed point methods have been found to be overly slow for solving stiff time evolution systems due to their linear rate of convergence. However, recent work has shown that Anderson acceleration can speed these methods up (Anderson1965, Walker2011, Gardner2015). The potential benefit is that a fixed point iteration can be formulated such that each iteration is approximately equivalent to the cost of an explicit update. By accelerating the fixed point solver, we expect to see a larger time steps allowed due to the implicit formulation. If we can ensure that the step is sufficiently larger (factor of N) than an explicit step while the iteration requires less than N iterations, then we expect a benefit.

Connections to Math, Comp Sci & and Climate Science

Research in the area of nonlinear solvers for climate dynamics problems will leverage strong mathematical expertise as well as application expertise in order to most effectively approximate the dominant physics within the preconditioners or implicit parts of split (IMEX) systems. It will be necessary to develop better interfaces with solver libraries so that they are easier and more flexible for application scientists to implement and alter as needed for both the Newton-Krylov and Anderson acceleration methods.

Potential Impact on the field

With future climate science questions focusing on topics that require additional vertical levels, more complete and multiscale parameterizations, and regional refinement in the major components of Earth System models being developed in ACME and other ESM's, a solver option that removes the CFL restriction, and also enables the use of advanced architectures could lead to significant speedups of the overall simulations.

References

- (Anderson1965) Anderson D G 1965 Iterative procedures for nonlinear integral equations *J. Assoc. Comput. Mach.* **12** 547–60
- (ECMWF2012) ECMWF: IFS Documentation - Cy38r1, Part III: Dynamics and Numerical procedures. Tech. Rep. cy38ra, European Centre for Medium-Range Weather Forecasting (2012). URL <http://www.ecmwf.int/research/ifsdocs/CY38r1/IFSPart3.pdf>.
- (Evans2009) Evans, K.J., Rouson, D.W.I., Salinger, A.G., Taylor, M.A., Weijer, W., III, J.B.W.: A Scalable and Adaptable Solution Framework within Components of the Community Climate System Model, Lecture Notes in Computer Science, vol. 5545, pp. 332–341. Springer (2009).
- (Evans2010) Evans, K.J., Taylor, M.A., Drake, J.B.: Accuracy analysis of a spectral element atmospheric model using a fully implicit solution framework. *Monthly Weather Review* 138,

Contributed ideas paper 50

3333–3341 (2010).

(Evans2013) Evans, K., Lauritzen, P., Neale, R., Mishra, S., Taylor, M.A., Tribbia, J.: AMIP simulation with the CAM4 spectral element dynamical core. *Journal of Climate* 26, 689–709 (2013). DOI 10.1175/JCLI-D-11-00448.1.

(Gardner2015) Gardner, D. J., et al. "Implicit integration methods for dislocation dynamics." *Modelling and Simulation in Materials Science and Engineering* 23.2 (2015): 025006.

(Lott2015) Lott, P. Aaron, Carol S. Woodward, and Katherine J. Evans. "Algorithmically scalable block preconditioner for fully implicit shallow-water equations in CAM-SE." *Computational Geosciences* 19.1 (2015): 49-61.

(Thomas2003) Thomas, S.J., Dennis, J.M., Tufo, H.M., Fischer, P.F.: A Schwarz preconditioner for the cubed sphere. *SIAM Journal on Scientific Computing* 25(2), 442–453 (2003). DOI <http://dx.doi.org/10.1137/S1064827502409420>.

(Thomas2002) Thomas, S.J., Loft, R.D.: Semi-implicit spectral element atmospheric model. *Journal of Scientific Computing* 17, 339–350 (2002).

(Walker2011) Walker H F and Ni P 2011 Anderson acceleration for fixed-point iterations *SIAM J. Numer. Anal.* **49** 1715–35.

(Yang2010) Yang, C., Cao, J., Cai, X.C.: A fully implicit domain decomposition algorithm for shallow water equations on the cubed sphere. *SIAM Journal on Scientific Computing* 32(1), 418–438 (2010).

Contributed ideas paper 51

Full-system evaluation of Earth system models

December 1, 2015

Joseph H. Kennedy

Computer Science and Mathematics Division, Oak Ridge National Laboratory (ORNL), Oak Ridge, TN, USA 37934

Bert Debusschere (Sandia National Laboratories), Khachik Sargsyan (Sandia National Laboratories), Forrest M. Hoffman (ORNL), Katherine J. Evans (ORNL)

Abstract

Policy makers are increasingly reliant on Predictions from Earth system models (ESMs), especially in the case of sea-level rise (SLR) and extreme events [6]. Necessarily, policy makers require extremely high levels of model credibility and scientific confidence in model predictions. Accordingly, confidence building activities, such as verification and validation (V&V) and uncertainty quantification (UQ), are maturing on a component-by-component basis within today's ESMs. However, within the next decade, the scientific community will need to adopt full-system model evaluation (FSME) approaches to authenticate fully-coupled model simulations and predictions. Developing FSME will require investigation into novel evaluation techniques and building the required infrastructure to tie individual component diagnostics packages together.

Proposed Direction of work

Currently, many disparate component-evaluation packages are being actively developed within the ESM community (e.g., the AMWG diagnostics package for the CESM atmospheric component [1], LIVVkit for the CESM/ACME land-ice component [7], and ILAMB for land models [8]). These evaluation packages are typically run by their respective component-model's working group (WG), and have their own sets of associated data, statistical algorithms, and quality measures. Due to this fractionation, and the associated large technological and logistical burden of multiple-component evaluations, it is only possible to evaluate coupling between contiguous model components. How model changes propagate throughout the full system to non-contiguous model components cannot be captured pragmatically.

A FSME framework, however, could be designed such that all the individual component evaluation algorithms, quality measures, and associated data are available through a central portal, allowing any related scientist or developer to evaluate how changes propagate throughout the entire model system. Within a FSME, validation tests can be run where quality measures are evaluated across multiple model components, which enables direct UQ of quantities that depend on the full Earth system, such as SLR.

Math, Comp. Sci. & Climate Science Challenges

Developing a FSME framework will require the ESM community to overcome many mathematical, computational, and scientific challenges. The initial challenge will be developing the necessary infrastructure to connect current component-model evaluation packages, and to develop a data repository to contain all the associated benchmark, observation, and reanalysis data needed for the model evaluations. The framework will have to efficiently transform data between individual component projections and conservatively interpolate the data onto different component grids for comparisons.

Contributed ideas paper 51

Additionally, new metrics will need to be developed to take advantage of the FSME framework. Because the calculation of many of today’s component-model validation metrics are computationally expensive and being run at Leadership Computing Facilities (LCFs), it is expected that full-system model metrics will require exascale computing facilities and robust data analysis techniques. For example, because of the probabilistic nature of many model components, many metrics will require the evaluation and comparison of several model ensembles.

The development of a formal and automated way to compare ensemble runs is an especially challenging problem given the natural and acceptable variability within the ensembles. To answer this question, two components are needed: a) a metric for measuring the “distance” between the two ensembles, and b) a means to decide on equivalence given this metric.

For a rigorous ensemble comparison, an appropriate measure is the Kullback-Leibler (KL) divergence, which takes into account the full probability distribution function (PDF) of the ensemble [3]. As to whether a small number of realizations are likely to be part of a benchmark ensemble, we will rely on the Continuous Rank Probability Score (CRPS) [5, 9]. This test will be useful for quick, ‘on the fly,’ testing of single runs or ensembles in progress.

We can consider ensembles to be statistically similar for values of the distance metric below a specified threshold. This decision can be formulated as a binary classification problem, and we can use tools from the machine learning community to quantify the confidence in this decision based on the nature of the ensembles (in terms of size, simulation length, output variables, and perturbation types). The classifier performance will be characterized with Receiver Operating Characteristic (ROC) curves, which plot the classifier true positive rate versus the false positive rate, for different values of the distance metric threshold [4]. The area under the ROC curve is a measure of the performance and robustness of the classifier [2]. Larger ensembles with longer simulation times will generally have less noise in the ensemble PDF of the output quantities, and will therefore result in a more robust classification. Constructing ROC curves for the associated classifiers will quantify this robustness and guide the optimal selection of the ensembles.

Potential Impact on the field

Development of a FSME framework will allow for an unprecedented level of scientific confidence in ESM predictions and ESM credibility. Changes to ESM components, parameters, or input data can then undergo system-wide V&V tests, and direct UQ of quantities that depend on the full system can be preformed. Through the development of the new metrics and comparisons that can leverage the FSME framework, we will have a better understanding of how the entire Earth system behaves, and is likely to behave in the future.

References

- [1] Atmosphere Model Working Group. The AMWG diagnostic package. <https://www2.cesm.ucar.edu/working-groups/amwg/amwg-diagnostics-package>, 2014. Version 140804.
- [2] A. P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997.
- [3] Sung-Hyuk Cha. Comprehensive survey on distance/similarity measures between probability density functions. *International Journal of Mathematical Models and Methods in Applied Sciences*, 1(3):300–307, 2007.
- [4] Tom Fawcett. An introduction to roc analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.

Contributed ideas paper 51

- [5] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [6] Maria Carmen Lemos and Richard B. Rood. Climate projections and their impact on policy and practice. *Wiley Interdisciplinary Reviews: Climate Change*, 1(5):670–682, 2010.
- [7] LIVVkit. The Land Ice Verification and Validation toolkit. <https://github.com/LIVVkit/LIVVkit>, 2015. v1.0.0.
- [8] Y. Q. Luo, J. T. Randerson, G. Abramowitz, C. Bacour, E. Blyth, N. Carvalhais, P. Ciais, D. Dalmonech, J. B. Fisher, R. Fisher, P. Friedlingstein, K. Hibbard, F. Hoffman, D. Huntzinger, C. D. Jones, C. Koven, D. Lawrence, D. J. Li, M. Mahecha, S. L. Niu, R. Norby, S. L. Piao, X. Qi, P. Peylin, I. C. Prentice, W. Riley, M. Reichstein, C. Schwalm, Y. P. Wang, J. Y. Xia, S. Zaehle, and X. H. Zhou. A framework for benchmarking land models. *Biogeosci.*, 9(10):3857–3874, October 2012.
- [9] C. Safta, D. Ricciuto, K. Sargsyan, B. Debusschere, H. N. Najm, M. Williams, , and P. Thornton. Global sensitivity analysis, probabilistic calibration, and predictive assessment for the data assimilation linked ecosystem carbon model. *Geoscientific Model Development*, 8:1899–2015, 2015.

Contributed ideas paper 52

Machine-learning guided, multi-resolution approaches to high-fidelity representation of global hydrology in ESMs

December 1, 2015

Richard Tran Mills, Intel Corporation and Forrest M. Hoffman, ORNL

Spatial resolutions of land surface models (LSMs) have historically been dictated by the resolutions at which a coupled weather or climate model is computationally feasible to run (about $O(10)$ km laterally for what is currently considered a very high-resolution global model). In recent years, however, compelling arguments in favor of running LSMs (with more realistic hydrologic process models) at dramatically higher resolutions [1] have been made. Such “hyperresolution” models would enable representation of several important processes related to carbon and nitrogen cycling. For example, they would allow representation of CO_2 outgassing from river channel systems that are too fine to represent in current resolution models; these fluxes can be very significant, e.g., in the Amazon basin an estimated 0.5 Gt per year of carbon is outgassed [2], much of it from small rivers [3].

Recent advances have seen the coupling of sophisticated surface–subsurface hydrology and reactive transport models—needed for accurate representation of surface/subsurface water dynamics and biogeochemical processes at high resolutions—to LSMs [4, 5, 6]. The ability of some of these models to fully utilize leadership-class supercomputers [7, 8, 9, 10, 11, 12] will enable dramatically higher-resolution simulations of global hydrology and associated biogeochemistry in LSMs as supercomputers continue to increase in power, but *even with all foreseeable advances in computing power and solver algorithms, the achievable resolutions for global LSM simulations will be unable to resolve some important processes that are difficult to represent via sub-grid parameterizations.*

We consider an example—permafrost-affected Arctic regions—that particularly illustrates some of the challenges involved in scaling up the effects of fine-scale land surface processes to the regional or global scale. Assumptions of linear scaling employed in LSMs are poorly suited to the particular organization of fine- and intermediate-scale features in these landscapes. Figure 1(a) illustrates: At the resolution of a typical coupled climate model cell (100 km), there are many small-scale features and extreme heterogeneity in the land cover and microtopography of the cell, and we continue to see heterogeneity as we zoom in: the distinct, geometrical organization of such *patterned ground* remains apparent to sub-meter resolution. Rates of, for instance, organic matter decomposition (and thus greenhouse gas emission) are strongly controlled by (among other factors) the local distribution of soil moisture, and simply averaging up to the global grid cell does not yield meaningful predictions. The NGEE–Arctic project [13] is addressing this by constructing mechanistic, process-resolving models that are accurate at the small scale, and then using these to conduct a series of simulations of permafrost-affected landscapes using a sequence of nested computational domains ranging from fine to global climate-grid resolutions. This scaling approach is described in much more detail in [14]. We note that the iterative upscaling-downscaling sequence is an off-line one: An ensemble of simulations will be conducted at a finer scale and then analyzed to produce parameter values for a coarser scale; simulations at the coarser scale will then be run and analyzed to produce parameters generated for yet coarser simulations, or to produce improved boundary conditions for repeating finer-scale runs.

The NGEE–Arctic scaling approach cannot rely solely on simulations; it must be constrained by observations. It is therefore desirable to develop systematic sampling strategies that can maximize coverage while objectively representing environmental variability at the scales of interest. We have recently applied a quantitative methodology for stratifying sampling domains, informing site selection, and quantifying the representativeness of candidate measurement sites in the State of Alaska [15]. We applied a geospatiotemporal clustering procedure to downscaled general circulation model results and observational data for the State of Alaska at a nominal resolution of 4 km^2 to define a set of eco-climatic regions possessing relatively homogeneous attributes, at multiple levels of division across two decadal time periods, and examined their present distributions across Alaska and how they are expected to shift as a result

Contributed ideas paper 52

of projected climate change. Using a metric of “representativeness” based on Euclidean distances in the n -dimensional state space employed in the cluster analysis, optimal sampling locations (the “realized centroids”) that are most representative of the conditions in each eco-climatic region were obtained, and candidate locations for future NGEE–Arctic measurement sites near towns in Alaska were identified (see Figure 1(b)). In addition to guiding selection of candidate measurement sites, the cluster analysis and n -dimensional data space regressions offer a basis for up-scaling and extrapolating measurements to land areas within and beyond the sampling domains, and provide a down-scaling approach to the integration of models, observations, and process studies.

To circumvent the tremendous computational expense of fully-resolved ultra-high resolution models, we propose an approach inspired partly by the “super-parameterization” [16, 17] approaches developed by some members of the atmospheric modeling community. These approaches avoid the tremendous computational expense of cloud system-resolving models while still realistically representing small-scale and mesoscale processes, by embedding a small-scale 2D or quasi-3D model inside each large scale (we will refer to this as the “global scale” for convenience, even though these may be regional-scale) cell. The embedded small-scale models generally employ periodic lateral boundary conditions, and do not interact with each other except indirectly through the fluxes on the global-scale grid. This approach essentially consists of two models with two sets of model variables: the large-scale models and variables of the coarse, global grid, and the fine-scale models and variables of the embedded, small-scale models. The two sets of model variables are coupled by enforcing the property that the horizontal average of a small-scale variable is exactly equal to the value of the corresponding large-scale value.

Although 2D models are attractive from the standpoint of computational expense, in many landscapes, the hydrologic processes work in an intrinsically 3D space. Embedding finely resolved 3D small-scale grids in every global-scale grid is not computationally feasible, however, even if the individual small-scale grids do not communicate with each other. However, *we believe that the same geospatiotemporal clustering-based techniques that we have used for choosing measurement sites in ecological sampling networks can be used to choose a sparse set of “measurement” sites where fine-scale, 3D models are to be embedded.* The basic idea is to periodically run an on-line clustering using the parameters and state variables from all cells in the global grid (likely with ancillary fine-scale data such as high-resolution topographic or land-cover data); group cells with similar states, properties, and forcings together into clusters; choose a representative subset of cells from each cluster, and run an embedded fine-scale model—coupled to the global scale in a manner very similar to that used in atmospheric super-parameterization models. Quantities of interest, e.g., biogeochemical reaction rates, from the sparse collection of fine-scale models are then mapped back to the other global cells that are members of the same (or similar) clusters that do not have an embedded fine-scale model. The simplest approach is to “paint by numbers”: if only one member of a cluster has an embedded model, assign the same upscaled value from that member to all others of the cluster; if there are multiple members of the cluster that run embedded fine-scale models, assign other members the value from the closest (in data space) member or some average. Such an approach might work surprisingly well: geospatiotemporal clustering has been used to dramatically reduce the dimensionality of global CLM simulation output [18] while maintaining high accuracy (Figure 1(c)). Alternatively, one could construct model response surfaces from the ensemble of fine-scale simulations within a cluster. Such an approach was used by Luxmoore et al. in [19] and [20] to scale the results of small-scale plant physiological responses to various environmental factors up through a hierarchy of models to the regional ecosystem scale.

We believe that the machine-learning guided approach outlined presents a viable way to adaptively and parsimoniously select locations in which embedded ultra high-resolution model patches should be run, but there are many open questions: What level of division (i.e., the number of clusters) should be employed? How many different levels of scale must be used? What is the best approach to prescribing boundary conditions for the embedded fine-scale models? We hope to explore these and other questions with the climate science and computational mathematics communities.

Contributed ideas paper 52

References

- [1] E. F. Wood, J. K. Roundy, T. J. Troy, L. P. H. van Beek, M. F. P. Bierkens, E. Blyth, A. de Roo, P. Dill, M. Ek, J. Famiglietti, D. Gochis, N. van de Giesen, P. Houser, P. R. Jaff, S. Kollet, B. Lehner, D. P. Lettenmaier, C. Peters-Lidard, M. Sivapalan, J. Sheffield, A. Wade, and P. Whitehead, “Hyperresolution global land surface modeling: Meeting a grand challenge for monitoring earth’s terrestrial water,” *Water Resources Research*, vol. 47, no. 5, pp. n/a–n/a, 2011. [Online]. Available: <http://dx.doi.org/10.1029/2010WR010090>
- [2] J. E. Richey, J. M. Melack, A. K. Aufdenkampe, V. M. Ballester, and L. L. Hess, “Outgassing from amazonian rivers and wetlands as a large tropical source of atmospheric CO₂,” *Nature*, vol. 416, pp. 617–620, 2002.
- [3] M. de Ftima F. L. Rasera, M. V. R. Ballester, A. V. Krusche, C. Salimon, L. A. Montebelo, S. R. Alin, R. L. Victoria, and J. E. Richey, “Estimating the surface area of small rivers in the southwestern Amazon and their role in CO₂ outgassing,” *Earth Interactions*, vol. 12, pp. 1–16, 2008.
- [4] R. M. Maxwell and N. L. Miller, “Development of a coupled land surface and groundwater model,” *Journal of Hydrometeorology*, vol. 6, no. 3, pp. 233–247, 2005.
- [5] S. J. Kollet and R. M. Maxwell, “Capturing the influence of groundwater dynamics on land surface processes using an integrated, distributed watershed model,” *Water Resources Research*, vol. 44, no. 2, pp. n/a–n/a, 2008, w02402. [Online]. Available: <http://dx.doi.org/10.1029/2007WR006004>
- [6] G. Tang, F. Yuan, G. Bisht, G. E. Hammond, P. C. Lichtner, J. Kumar, R. T. Mills, X. Xu, B. Andre, F. M. Hoffmand, S. L. Painter, and P. E. Thornton, “Using reactive transport codes to provide mechanistic biogeochemistry representations in global land surface models: CLM-PFLOTRAN 1.0,” *To appear in Geoscientific Model Development Discussions*, 2015.
- [7] R. T. Mills, C. Lu, P. C. Lichtner, and G. E. Hammond, “Simulating subsurface flow and transport on ultrascale computers using pflotran,” in *Journal of Physics: Conference Series*, vol. 78, no. 1. IOP Publishing, 2007, p. 012051.
- [8] R. T. Mills, G. E. Hammond, P. C. Lichtner, V. Sripathi, G. K. Mahinthakumar, and B. F. Smith, “Modeling subsurface reactive flows using leadership-class computing,” in *SciDAC 2009 Scientific Discovery through Advanced Computing*, ser. J. Phys.: Conf. Ser., vol. 180, no. 1, 2009.
- [9] R. T. Mills, V. Sripathi, G. K. Mahinthakumar, G. E. Hammond, P. C. Lichtner, and B. F. Smith, “Engineering PFLOTRAN for scalable performance on Cray XT and IBM BlueGene architectures,” in *SciDAC 2010 Scientific Discovery through Advanced Computing*, ser. J. Phys.: Conf. Ser., 2010.
- [10] G. E. Hammond, P. C. Lichtner, R. T. Mills, and C. Lu, *Ground Water Reactive Transport Models*. Bentham Science Publishers, 2010, ch. PFLOTRAN: Reactive Flow & Transport Code for Use on Laptops to Leadership-Class Supercomputers.
- [11] G. E. Hammond, P. C. Lichtner, and R. T. Mills, “Evaluating the performance of parallel subsurface simulators: An illustrative example with pflotran,” *Water Resources Research*, vol. 50, no. 1, pp. 208–228, 2014. [Online]. Available: <http://dx.doi.org/10.1002/2012WR013483>
- [12] S. J. Kollet, R. M. Maxwell, C. S. Woodward, S. Smith, J. Vanderborght, H. Vereecken, and C. Simmer, “Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources,” *Water Resources Research*, vol. 46, no. 4, pp. n/a–n/a, 2010, w04201. [Online]. Available: <http://dx.doi.org/10.1029/2009WR008730>

Contributed ideas paper 52

- [13] N.-A. team, “NGEE–Arctic webpage,” 2015. [Online]. Available: <http://ngee-arctic.ornl.gov/>
- [14] P. Thornton, B. Riley, C. Koven, F. Hoffman, R. T. Mills, and the NGEE Arctic Science Team, “Developing a hierarchical scaling framework for modeling a dynamic arctic landscape in a changing climate,” Tech. Rep., 2011. [Online]. Available: <http://ngee-arctic.ornl.gov/system/files/NGEE%20Arctic%20Scaling%20White%20Paper%20%28Part%201%29.pdf>
- [15] F. M. Hoffman, J. Kumar, R. T. Mills, and W. W. Hargrove, “Representativeness-based sampling network design for the State of Alaska,” *Landscape Ecol.*, vol. 28, no. 8, pp. 1567–1586, Oct. 2013.
- [16] W. W. Grabowski, “An improved framework for superparameterization,” *Journal of the atmospheric sciences*, vol. 61, no. 15, pp. 1940–1952, 2004.
- [17] D. Randall, M. Khairoutdinov, A. Arakawa, and W. Grabowski, “Breaking the cloud parameterization deadlock,” *Bulletin of the American Meteorological Society*, vol. 84, no. 11, pp. 1547–1564, 2003.
- [18] J. Kumar and F. M. Hoffman, “Personal communication,” 2013.
- [19] R. J. Luxmoore, W. W. Hargrove, M. L. Tharp, W. M. Post, M. W. Berry, K. S. Minser, W. P. Cropper Jr, D. W. Johnson, B. Zeide, R. L. Amateis *et al.*, “Signal-transfer modeling for regional assessment of forest responses to environmental changes in the southeastern united states,” *Environmental Modeling & Assessment*, vol. 5, no. 2, pp. 125–137, 2000.
- [20] R. J. Luxmoore, W. W. Hargrove, M. Lynn Tharp, W. Mac Post, M. W. Berry, K. S. Minser, W. P. Cropper, D. W. Johnson, B. Zeide, R. L. Amateis *et al.*, “Addressing multi-use issues in sustainable forest management with signal-transfer modeling,” *Forest Ecology and Management*, vol. 165, no. 1, pp. 295–304, 2002.

Contributed ideas paper 52

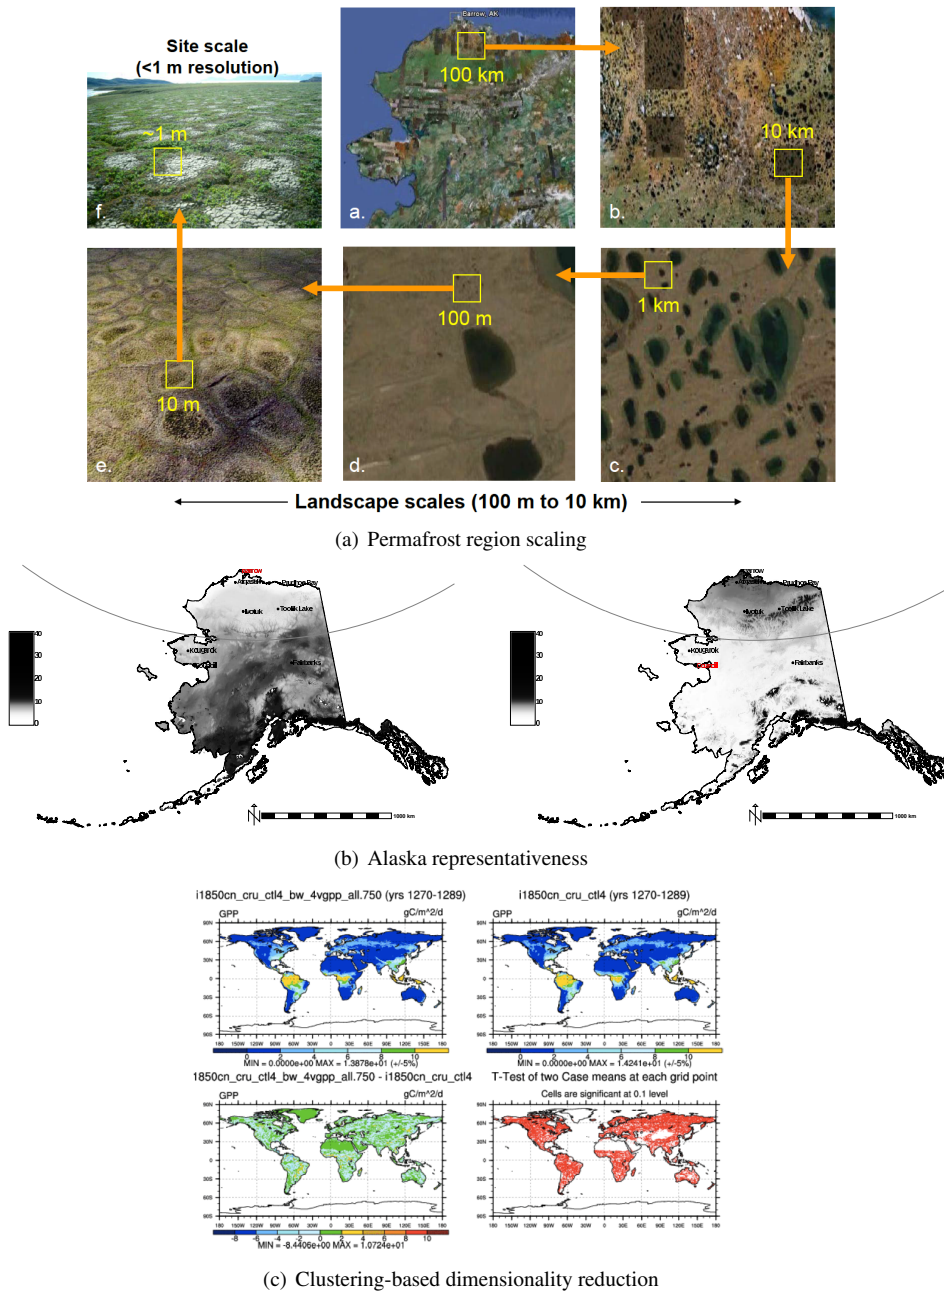


Figure 1: Different aspects of spatial scaling and representativeness in landscapes. a) A zoom-in over patterned ground on the North Slope of Alaska, near Barrow, AK, illustrates distinct (and geometric) small-scale features. Figure courtesy of Peter Thornton, ORNL. b) A comparison of representativeness (closeness in terms of Euclidean distance in the data-space of eco-climatic variables used in the cluster analysis) for present-day conditions for two of the eight potential NGEE–Arctic sites considered in [15] (Barrow on the left, Council on the right). White to light gray land areas are well-represented by the site, while dark gray to black land areas are poorly represented by the site. Our analysis suggests that, because Barrow has been selected as the first sampling site, Council is a strong candidate for a second site in the northern half of the state because of its dissimilarity to Barrow. c) A comparison of Gross Primary Productivity (GPP) from an approximately 60,000 cell CLM simulation (top right) and from a reduced version (top left) in which geospatiotemporal clustering has been used to identify 750 clusters; the GPP values on the top left map have been filled in assigning the GPP value associated with the centroid of a cluster to all cells that are members of the cluster. Figure courtesy of J. Kumar and F. M. Hoffman of ORNL, from as yet unpublished work.

Contributed ideas paper 53

Higher Standards on the Control of Numerical Accuracy

2015-12-01

Jed Brown

University of Colorado Boulder

Abstract

The relatively low current standard for control of numerical accuracy is hurting climate models in many ways. It is crucial for reliability, trust, and the efficient use of scientist time that we become more systematic about verification—quantifying the extent to which we solve the stated continuum equations, especially as climate models transform into engineering tools that advise policy and investment decisions.

Background/Research to Date

The 1986 Editorial Policy Statement for the Journal of Fluids Engineering [5] states

A professional problem exists in the computational fluid dynamics community and also in the broader area of computational physics. Namely, there is a need for higher standards on the control of numerical accuracy. [...] It [is] impossible to evaluate and compare the accuracy of different turbulence models, since one [cannot] distinguish physical modeling errors from numerical errors related to the algorithm and grid. This is especially the case for first-order accurate methods and hybrid methods.

The statement goes on to mandate systematic evaluation of numerical errors. This policy has been influential in defining the standards of the computational engineering community, for whom validation is readily available because products are manufactured and tested. Models that do not ascribe to this standard are widely viewed as less reliable and lower quality.

Climate models are being sold to policy makers as engineering tools, but with much lower verification standards and with no comparable validation process (due to the time scales involved and practical infeasibility of avoiding overfitting due to ad-hoc calibration process). Indeed, many prominent models have grave deficiencies.

- MPAS [7] uses an inconsistent discretization of Coriolis on non-orthogonal grids [6], including at the pentagons that are necessary when tiling the sphere with hexagons.
- CICE relies on the EVP formulation which introduces artificial parameters that materially affect solutions [3, 4] and efficiency.
- CAM column physics uses a splitting scheme that is not temporally convergent and significantly affects computed solutions [2, 1].

Models are calibrated to compensate for the systematic numerical errors produced by such schemes, thus making it difficult to distinguish the properties of the continuum model from that of the discretization. More concerning is that experimentation with improved discretizations requires expensive, ad-hoc recalibration involving significant domain expertise.

Contributed ideas paper 53

Proposed Direction of work

New and existing models must be reevaluated to distinguish modeling error from numerical error. Numerical errors should be controlled by improving discretizations, addressing deficiencies in decreasing order of prominence. In cases where it is computationally infeasible to control numerical error to a similar scale as modeling error, the parameter calibration procedure must be formalized such that it can be rapidly tuned for a new discretization. A more formal calibration process should be encouraged throughout the model, but it is especially important in such cases, for scientific understanding, to establish confidence, and to encourage improvements. Note that it is unacceptable to evaluate numerical methods in idealized scenarios with artificial viscosity (e.g., hyperviscosity) turned off, then use it with such terms present.

Specific solutions to fix non-convergent methods while maintaining efficiency will likely involve increased use of implicit methods and less reliance on naive splitting methods. While possibly more technical to develop, these techniques should improve algorithmic scalability and better expose fine-grained parallelism.

Connections to Math, Comp Sci & and Climate Science

The first goal for a mathematician working on numerical methods is usually verification, demonstrating that a particular algorithm solves the given continuum equations to some accuracy. This concept is relatively easy to explain and is accessible even to those who lack domain expertise. If the concept of convergence and metrics for evaluation is made prominent, it creates an environment in which domain scientists can effectively communicate with applied mathematicians and computer scientists and spur research likely to have a direct impact on the field. In the current system, it is far too common for researchers to spend time on methods that will never be viable for poorly specified reasons, resulting in dead-end abstract work and/or moving on to research in other disciplines. A more systematic approach to discretization errors will likely encourage refactoring models to more precisely define the equations, thus improving software modularity and maintainability and making it easier to utilize general-purpose numerical libraries.

Potential Impact on the field

A systematic approach to numerical accuracy should lead to adoption of convergent numerical methods yielding higher quality predictions and reducing the effort required to recalibrate a model when the discretization is changed. It will aid communication and make research on improved numerical methods more appealing.

References

- [1] Jed Brown. Time integration for atmospheric physics. Presented at SIAM CS&E, 2015.
- [2] Peter Caldwell, Bereket Lebassi Habtezion, Hui Wan, and Phil Rasch. Process integration/timestep sensitivity. Slides, 2013.
- [3] Jean-François Lemieux, Dana Knoll, Bruno Tremblay, David M Holland, and Martin Losch. A comparison of the Jacobian-free Newton–Krylov method and the EVP model for solving the sea ice momentum equation with a viscous-plastic formulation: A serial algorithm study. *Journal of Computational Physics*, 2012.

Contributed ideas paper 53

- [4] Martin Losch, Dimitris Menemenlis, Jean-Michel Campin, Patick Heimbach, and Chris Hill. On the formulation of sea-ice models. Part 1: Effects of different solver implementations and parameterizations. *Ocean Modelling*, 33(1):129–144, 2010.
- [5] P.J. Roache, K.N. Ghia, and F.M. White. Editorial policy statement on the control of numerical accuracy. *Journal of Fluids Engineering*, 108:2, 1986.
- [6] J Thuburn, CJ Cotter, and T Dubos. A mimetic, semi-implicit, forward-in-time, finite volume shallow water model: comparison of hexagonal–icosahedral and cubed-sphere grids. *Geoscientific Model Development*, 7(3):909–929, 2014.
- [7] John Thuburn, Todd D Ringler, William C Skamarock, and Joseph B Klemp. Numerical representation of geostrophic modes on arbitrarily structured C-grids. *Journal of Computational Physics*, 228(22):8321–8335, 2009.

Contributed ideas paper 54

Advanced Time Integration Methods for Atmospheric Physics

Dec. 1, 2015

David J. Gardner, Lawrence Livermore National Laboratory

Carol S. Woodward (LLNL), Peter M. Caldwell (LLNL), Katherine J. Evans (ORNL)

Abstract

The parameterized atmospheric physics component of global climate models comprises a significant portion of the computational cost of large-scale climate simulations. This aspect of the climate model is composed of numerous physical processes acting on a wide range of spatial and temporal scales with each process sequentially or parallel split from the others. With the large timestep sizes utilized for efficiency in production scale runs, the splitting error from decoupling the different physics processes can become large. Modern time integration methods offer a better understanding of splitting errors and provide a means for improving process coupling and reducing time stepping errors. Incorporating these new techniques into production climate models will require close collaboration among climate scientists, numerical/applied mathematicians, and solver library developers.

Background/Research to Date

Coupled global climate models are complex multiscale multiphysics systems encompassing a broad set of physical processes. A great deal of effort in the atmospheric component of climate models has focused on the numerical solution of the fluid flow and much less attention has been paid to the numerical treatments of parameterized sub-grid scale processes [Beljaars1991, Collins2015]. As a result, errors due to modeling approximations can be surpassed by numerical errors [Beljaars1991] that lead to decreased convergence rates in the Community Atmosphere Model (CAM) [Wan2015].

CAM utilizes sequential splitting and first order explicit time updates so that each process advances the system state serially. At large timestep sizes, this approach can result in instabilities and large errors, for example, when coupling the cloud macrophysics and microphysics. Moreover, the microphysics itself is comprised of several parallel split processes, any of which may attempt to remove more moisture than is available in the system. This triggers a conservation limiter to ensure non-negativity of the solution and introduces additional error into the system. Another challenge demonstrated by the microphysics is that processes are not necessarily smooth. In particular they rely on ‘if’ statements to handle different physical events, such as freezing/melting, and are thus not always differentiable.

Integrating such a system in a stable, consistent, and accurately coupled manner while ensuring efficient simulations and maintaining code modularity presents a significant challenge. Analysis of idealized multirate systems with split processes in [Dubal2004] and [Dubal2005] demonstrates the need for careful coupling with low order methods and some level of implicitness to accurately integrate fast processes with large timesteps. Advances in partitioned time integration methods offer new approaches for constructing accurate splitting methods of arbitrary orders of accuracy.

Proposed Direction of work

Of particular interest are the potential benefits of implicit-explicit (IMEX) methods for integrating atmospheric processes. In this splitting approach the right-hand side is divided into slow terms that are integrated explicitly and fast terms that are handled implicitly. With an IMEX method the splitting error can be made to match the order of the individual methods in the scheme. This creates the possibility for highly accurate coupling between split processes and should lead to greater consistency and stability than is currently possible with low order splitting techniques. Furthermore, there exist a variety of IMEX approaches including linear multistep

Contributed ideas paper 54

methods [Frank1997], Runge-Kutta methods [Ascher1995, Kennedy2003], and extrapolation methods [Constantinescu2010]. It is unclear how to select the best IMEX method to achieve a good balance between accuracy and efficiency for any given system. Thus a fair amount of exploration is necessary to understand what IMEX properties and nonlinear solver choices are most advantageous for atmospheric physics.

Solver libraries such as SUNDIALS [Woodward2015] and PETSc [Balay2015] provide a number of advanced integration methods including Runge-Kutta based IMEX schemes. Utilizing time integration packages provides a good starting point for exploring the impact of IMEX schemes in atmospheric physics. Interfacing with these libraries also enables access to robust nonlinear solvers that can be leveraged for implementing IMEX schemes not supported by the time integration packages. Additionally, time integration, nonlinear solver, or computational enhancements made in the library become readily available to the climate model code base. For example, time integration in climate models relies on fix time step sizes while many of the time integration methods in solver libraries, allow for temporal error estimation and time step adaptivity. Such approaches could be useful to avoid applying conservation limiters. However, such approaches are likely not a good fit for climate simulations at this time but may become important in future models.

Connections to Math, Comp Sci & Climate Science

Incorporating advanced time integration methods into the atmospheric physics component of a global climate mode will require detailed knowledge of the process parameterizations and their derivations to identify the time scales of interest and to modify process formulations as necessary to maximize accuracy and performance. This domain knowledge will need to be paired with expertise in time integration and nonlinear solvers to understand the costs and benefits of various IMEX schemes and how different process formulations could impact the time integrators. In order to effectively utilize solver packages it will be important to work with the library developers to improve software documentation, usability, and flexibility to simplify code portability and maintenance.

Potential Impact on the field

Current approaches to process spitting and time integration of multirate sub-grid scale parameterizations rely on first order explicit updates with sequentially split processes. At large time step sizes, this approach can introduce significant errors in to the model that will need to be addressed. Modern IMEX methods for splitting processes provide a means to efficiently deal with the multiscale nature of climate models and produce more accurate simulations.

References

[Ascher1995] Ascher, U. M., et al., (1995) Implicit-explicit methods for time-dependent partial differential equations, *SIAM J. Numer. Anal.*, 22 (3), 797-823.

[Balay2015] Balay, S., et al., PETSc, (2015), <http://www.mcs.anl.gov/petsc>

[Beljaars1991] Beljaars, A. (1991), Numerical schemes for parameterizations, *Proceedings of the ECMWF seminar on numerical methods in atmospheric models*, 1–42. European Centre for Medium-range Weather Forecasts, Reading, United Kingdom.

[Collins2015] Collins, W. D., et al. (2015), Progress in fast, accurate multi-scale climate simulations, *Procedia Comp. Sci.*, 51, 2006-2015.

Contributed ideas paper 54

[Constantinescu2010] Constantinescu, E.M. and A. Sandu, (2010), Extrapolated implicit-explicit time stepping, *SIAM J. Sci. Comp.*, 31(6), 4452-4477.

[Dubal2005] Dubal, M., N. Wood, and A. Stantforth (2005), Mixed parallel-sequential-split schemes for time-stepping multiple physical parameterizations, *Mon. Wea. Rev.*, 133, 989–1002.

[Dubal2004] Dubal, M., N. Wood, and A. Stantforth (2004), Analysis of parallel versus sequential splittings for time-stepping physical parameterizations, *Mon. Wea. Rev.*, 132, 121–132.

[Frank1997] Frank, J., W. Hundsdorfer, and J. G. Verwer (1997), On the stability of implicit-explicit linear multistep methods, *App. Num. Math.*, 25, 193-205.

[Kennedy2003] Kennedy, C. A. and M. H. Carpenter (2003), Additive Runge-Kutta schemes for convection-diffusion-reaction equations, *App. Num. Math.*, 44, 139-181.

[Wan2015] Wan, H., et al. (2015), Short-term time step convergence in a climate model, *J. Adv. Model. Earth Syst.*, 7, 215–225.

[Woodward2015] Woodward, C. S., et al., (2015), SUNDIALS, <http://computation.llnl.gov/casc/sundials>

Contributed ideas paper 55

Functional Representation of Climate Simulation Output December 1, 2015

George Ostrouchov

Oak Ridge National Laboratory

Mark Ainsworth (Brown University), Scott Klasky (Oak Ridge National Laboratory)

Abstract

The increasing gap between available processor cycles and speed of data movement cannot be solved with business as usual and a disruptive change is needed. One such disruptive solution that already has a strong mathematical, statistical, and computational tool base is to transition away from storing grid data representations to storing functional data representations. From a theoretical viewpoint, this is a transition from representing solutions as points on a grid in Euclidean space (size depends on number of grid points) to functional representations in Hilbert space (size depends on function smoothness and regularity). Functional data analysis tools already exist for common climate science analytics such as variability attribution, principal components, canonical correlation, clustering, and many other multivariate methods.

Background/Research to Date

Computational Climate Science data serve very large and diverse communities of researchers. The diversity across many disciplines and many uses within disciplines results in access patterns of the data across many different scales and in many different configurations. This diversity of access patterns produces a combined pressure on computational climate science to continue providing more fidelity across nearly all of its components.

Computational climate models respond to this pressure by providing more detail and fidelity as they take advantage of larger computational resources. As the spatial and temporal resolution of the models increases, the model output that has more resolution and fidelity also increases. More physics and chemistry components in the models also add new output components.

At the same time, developments in computer hardware and architecture continue to increase the gap between processor speed and data movement speed, putting more and more pressure on computational climate science to slow output data increases relative to simulation model fidelity improvements.

These conflicting pressures will not cease and require a solution with disruptive concepts that make conceptual changes to current practice. An opportunity lies in the bright side of this trend that brings low relative cost of compute cycles, which favors more sophisticated treatment of the output data before moving it to disk.

A common thread throughout much of simulation science that is fundamentally true in climate simulation: we simulate many cells of a discretized domain using first-principles/small-scale/local relationships described by equations for the purpose of discovering their combined/large-scale/non-local evolution. While we have equations for the propagation and equilibrium of physics and chemistry within cells and between pairs of cells, we do not have equations for the global simulation results. The results necessarily require an empirical treatment that is flexible enough to represent the data to a specified fidelity while also not introducing unnecessary artifacts.

Nonparametric empirical models such as B-splines, wavelets, and other basis functions, provide a rich set of models (functions) to represent any data on a continuous spatial domain. Further, techniques such as crossvalidation and regularization provide tools to prevent the in-

Contributed ideas paper 55

roduction of artifacts. These techniques are the tools of nonparametric statistics. Functional data analysis (FDA) evolved from combining nonparametric statistics with multivariate statistical analysis of data over continuous domains. The maturity of functional data analysis is evident in several books that have been published over the last decade or two [1, 2, 3, 4].

Applications of FDA to climate and weather data are already numerous as one can see by entering a web search on the terms "*functional data analysis*" AND *climate*.

Proposed Direction of work

We propose the conversion of grid based climate output data to functional representations before moving these representations to disk. This conversion does demand a considerable amount of initial computation but the result is a greatly reduced data representation that can be sampled for visualization at any resolution, even beyond what was originally provided by the simulation.

Tools already exist for computing functional representations in R [5] and elsewhere although none of them are currently used specifically for data reduction such as that described here. Their current use is for data analysis over continuous domains. For this use, many logistical and representation issues as well as scalability of FDA itself would need to be addressed.

Ideally, FDA representations would be provided to the climate simulations as a service accessible via various in-situ mechanisms such as those available in ADIOS [7] that would require minimal changes to the climate simulation codes.

Scalability of the extensive set of FDA techniques in R can be addressed with the use of infrastructure and in-situ capability of the pbdR project [6]. Compelling solutions may be further speeded up with judicious use of compiled language components. While initial experiments at scale may be realized as early as within two years, the complete paradigm shift in data representation will trigger other changes in the computational climate science data analysis toolchain that may take 10 years or more to resolve.

Connections to Math, Comp Sci & and Climate Science

FDA is a recent development in statistics that has mathematical and computational components.

Potential Impact on the field

Scalability of climate data output will no longer depend directly on simulation discretization and will grow much more slowly with fine scale variability.

References

- [1] James Ramsay and B. W. Silverman 2005. *Functional Data Analysis, Second Edition*. Springer.
- [2] Frédéric Ferraty and Philippe Vieu 2006. *Nonparametric Functional Data Analysis: Theory and Practice*. Springer.
- [3] James O. Ramsay, Giles Hooker, and Spencer Graves 2009. *Functional Data Analysis with R and MATLAB*. Springer.
- [4] Tailen Hsing and Randall Eubank 2015. *Theoretical Foundations of Functional Data Analysis, with an Introduction to Linear Operators* Wiley.

Contributed ideas paper 55

- [5] R Development Core Team (2008). *R: A language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0, <http://www.R-project.org>.
- [6] George Ostrouchov, Wei-Chen Chen, Drew Schmidt, and Pragneshkumar Patel 2012. *Programming with Big Data in R*. <http://pbd.r.org>.
- [7] Norbert Podhorszki, Qing Liu, Jeremy Logan, Jingqing Mu, Hasan Abbasi, Jong-Youl Choi, Scott A. Klasky 2015. *ADIOS 1.9 User's Manual*. ORNL/TM-2009/100.

Contributed ideas paper 56

Co-design of New Architectures for Large Scale Climate Data Analytics

December 1, 2015

Forrest M. Hoffman and Jitendra Kumar (Oak Ridge National Laboratory), William W. Hargrove (USDA Forest Service), and Richard T. Mills (Intel Corp.)

Abstract

The volumes of climate and other Earth science data are rapidly growing as model resolutions increase and observing networks and satellites collect data at higher spatial and temporal resolutions. New data analytics approaches are required on high performance computing platforms to synthesize and analyze these data. Described here is an effort to develop and extract key analytics methods useful in climate research, to optimize these methods for existing Leadership Computing platforms using large climate data sets, and to develop benchmark problems for co-design of future data analytics platforms.

Background/Research to Date

Observational and modeled data acquired or generated by the various disciplines within the realm of the Earth sciences encompass temporal scales of seconds to millions of years (10^0 – 10^{13} s) and spatial scales of microns to tens of thousands of kilometers (10^{-6} – 10^7 m). Because of rapid technological advances in sensor development, computational capacity, and data storage density, the volume, complexity, and resolution of Earth science data are increasing at a commensurate rate. Moreover, combining, integrating, and synthesizing data across Earth science disciplines offers new opportunities for scientific discovery that are only beginning to be realized. The rise of data-intensive scientific pursuits, in Earth sciences and other disciplines, has led some visionaries to proclaim it the fourth paradigm of discovery alongside the traditional experimental, theoretical, and computational archetypes (*Hey et al.*, 2009).

The promise of scientific advances in sustainability and environmental change research has stimulated an enormous increase in the volume of model and observational data. Earth system model (ESM) simulations are being conducted at progressively higher resolutions, generating growing volumes of data and requiring improved hardware and software infrastructure and more efficient post-processing and data mining analysis techniques. Organized global climate modeling activities, like the Coupled Model Intercomparison Project (CMIP) that coordinates simulations in support of the United Nations' Intergovernmental Panel on Climate Change (IPCC) assessment reports, can generate tens of terabytes to several petabytes of simulation results in raw form (*Overpeck et al.*, 2011). Observational data pose their own challenges. Satellite remote sensing data tend to be very large and their size has grown as spatial and temporal resolutions have increased.

While great strides are being made in the quantitative assessment of model fidelity through comparison with benchmark observational data sets (*Randerson et al.*, 2009; *Luo et al.*, 2012) and uncertainty quantification (UQ) methods (*Ricciuto et al.*, 2012), today's large and complex Earth science data often cannot be synthesized and analyzed using traditional methods or on individual workstations. Data mining, machine learning, and high performance visualization approaches are increasingly filling this void and can often be deployed only on parallel clusters or supercomputers. However, supercomputer architectures designed for compute-intensive simulations, usually containing large numbers of cores with high speed interconnects between nodes, are not typically optimal for large scale analytics. Instead, such applications demand large and fast on-node memory, high bandwidth input/output (I/O), and fast access to large local disk volumes. Most Earth scientists are ill-equipped to develop analytics codes for these architectures, while system vendors have largely focused on compute-intensive applications, and must acquire representative analytics benchmarks and scientific expertise to design systems for geospatial big data analytics.

Contributed ideas paper 56

Approaches for Advancing Climate Research

To realize the promise of new scientific discovery from very large, long time series Earth science data, a distinct balance of increasing computational, storage, and bandwidth capacity from high performance computing resources is required. Traditional analysis methods and algorithms are insufficient for analyzing and synthesizing such large data sets, and those algorithms rarely scale out onto distributed-memory parallel platforms. Therefore, new analysis techniques and scalable algorithms and software tools must be developed to enable analysis, exploration, and visualization of today's Earth science data. Highly scalable parallel k -means cluster analysis (*Hoffman and Hargrove, 1999*) and global domain traversal using a parallel MapReduce approach (*Kendall et al., 2011*) are two techniques used for spatiotemporal analysis that exploit massively parallel, high performance computing resources. These algorithms are generalized, flexible, and increasingly employed in Earth science, making them good candidates for analytics cores in a co-design process with system architects and integrators directed at large scale climate data reduction and analysis.

Hoffman and Hargrove (1999) developed a parallel k -means clustering algorithm, which they implemented on an early Beowulf-style parallel cluster computer they constructed from surplus PCs (*Hargrove et al., 2001*), useful for segmentation, feature extraction, network analysis, change detection, model inter-comparison, and model-data comparison in a number of Earth science applications (*Hoffman et al., 2008*). Recent improvements to that code, including adoption of a triangle-inequality-based acceleration technique and “warping” of unassigned/empty cluster centroids, have significantly reduced the time to solution (*Hoffman et al., 2008*), and a new technique for initial centroid determination has improved the statistical performance of the clustering result. These enhancements have enabled the analysis of large satellite data sets for identification of forest disturbances, a key component in a national-scale early warning system for detection of threats to forest health (*Hargrove et al., 2009; Hoffman et al., 2010*). This well-instrumented code could be encapsulated and combined with similar machine learning methods in a library, then optimized on DOE's existing Leadership Computing platforms for real, large scale climate and satellite remote sensing analytics. Moreover, benchmark tests and data sets could be designed for exercising and interrogating the performance of different potential future system designs, possessing various balances of CPUs, GPUs, memory hierarchies, I/O capacities, and volatile and non-volatile storage options.

Connections to Math, Computer Science, and Climate Science

Development and application of data analytics methods connects mathematicians and statisticians, performing algorithm design, with computer scientists, implementing and optimizing algorithms for target hardware platforms, with climate scientists, applying the tools to answer new questions and test hypotheses. Standalone and flexible implementations of these algorithms could be used by computer scientists and computer system vendors to inform the design of future Leadership Computing platforms suitable for the large scale data analytics needs of climate researchers.

Significance

The data analytics methods described above have been applied in climate-related studies and hold promise for future development in conjunction with a growing body of Earth science data. Standard analysis metrics, drawn from model-data benchmarks and UQ methods, can be developed using data analytics approaches and incorporated into the post-processing sections of simulation workflow systems. Moreover, these scalable analytics algorithms can be used in conjunction with appropriate data sets and evaluation metrics for the co-design of future data analytics hardware platforms for DOE.

Contributed ideas paper 56

References

- Hargrove, W. W., F. M. Hoffman, and T. Sterling (2001), The do-it-yourself supercomputer, *Sci. Am.*, 265(2), 72–79.
- Hargrove, W. W., J. P. Spruce, G. E. Gasser, and F. M. Hoffman (2009), Toward a national early warning system for forest disturbances using remotely sensed phenology, *Photogramm. Eng. Rem. Sens.*, 75(10), 1150–1156.
- Hey, T., S. Tansley, and K. Tolle (Eds.) (2009), *The Fourth Paradigm: Data-Intensive Scientific Discovery*, 284 pp., Microsoft Corporation, Redmond, Washington, USA.
- Hoffman, F. M., and W. W. Hargrove (1999), Multivariate geographic clustering using a Beowulf-style parallel computer, in *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA '99)*, vol. III, edited by H. R. Arabnia, pp. 1292–1298, CSREA Press.
- Hoffman, F. M., W. W. Hargrove, R. T. Mills, S. Mahajan, D. J. Erickson, and R. J. Oglesby (2008), Multivariate Spatio-Temporal Clustering (MSTC) as a data mining tool for environmental applications, in *Proceedings of the iEMSs Fourth Biennial Meeting: International Congress on Environmental Modelling and Software Society (iEMSs 2008)*, edited by M. Sánchez-Marrè, J. Béjar, J. Comas, A. E. Rizzoli, and G. Guariso, pp. 1774–1781.
- Hoffman, F. M., R. T. Mills, J. Kumar, S. S. Vulli, and W. W. Hargrove (2010), Geospatiotemporal data mining in an early warning system for forest threats in the United States, in *Proceedings of the 2010 IEEE International Geoscience and Remote Sensing Symposium (IGARSS 2010)*, pp. 170–173, doi: 10.1109/IGARSS.2010.5653935, invited.
- Kendall, W., J. Wang, M. Allen, T. Peterka, J. Huang, and D. Erickson (2011), Simplified parallel domain traversal, in *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11*, pp. 10:1–10:11, ACM, New York, NY, USA, doi: 10.1145/2063384.2063397.
- Luo, Y. Q., J. T. Randerson, G. Abramowitz, C. Bacour, E. Blyth, N. Carvalhais, P. Ciais, D. Dalmonech, J. B. Fisher, R. Fisher, P. Friedlingstein, K. Hibbard, F. Hoffman, D. Huntzinger, C. D. Jones, C. Koven, D. Lawrence, D. J. Li, M. Mahecha, S. L. Niu, R. Norby, S. L. Piao, X. Qi, P. Peylin, I. C. Prentice, W. Riley, M. Reichstein, C. Schwalm, Y. P. Wang, J. Y. Xia, S. Zaehle, and X. H. Zhou (2012), A framework for benchmarking land models, *Biogeosci.*, 9(10), 3857–3874, doi:10.5194/bg-9-3857-2012.
- Overpeck, J. T., G. A. Meehl, S. Bony, and D. R. Easterling (2011), Climate data challenges in the 21st century, *Science*, 331(6018), 700–702, doi:10.1126/science.1197869.
- Randerson, J. T., F. M. Hoffman, P. E. Thornton, N. M. Mahowald, K. Lindsay, Y.-H. Lee, C. D. Nevison, S. C. Doney, G. Bonan, R. Stöckli, C. Covey, S. W. Running, and I. Y. Fung (2009), Systematic assessment of terrestrial biogeochemistry in coupled climate-carbon models, *Global Change Biol.*, 15(10), 2462–2484, doi:10.1111/j.1365-2486.2009.01912.x.
- Ricciuto, D. M., M. L. Branstetter, D. Wang, J. Mao, J. Kumar, F. Hoffman, H. Najm, K. Sargsyan, C. Safta, and P. E. Thornton (2012), CSSEF progress towards uncertainty quantification in CLM-CN, DOE BER Poster Session, Oak Ridge National Laboratory, Oak Ridge, Tennessee.

Contributed ideas paper 57

Increasing orders of complexity in Earth System Models

Jean-François Lamarque

National Center for Atmospheric Research

Abstract

I discuss here some significant areas of research that would significantly benefit from increases in complexity (which include resolution) over the next decade, focusing on new areas of science that are unreachable at present.

Background/Research to Date

1) Climate change and evolution

In a recent paper [Veilleux et al., *Nature Climate Change*, 5, 1074–1078, 2015] the strong link between climate change and genetics was presented through experiments that identified changes in the genome that led to fish adaptation within 2 generations. Such experiments are clear identifiers that there is a level (or lack) of adaptability of the biological world that is totally unexplored and un-modeled in the current or even net generation of Earth System Models (ESMs). Along the same lines, the traditional approach in ESMs lacks direct microbial control over soil carbon dynamics [Wider et al., *Nature Climate Change*, 3, 909–912, 2013]. In addition to adaptation by genetic modification, competition is a key element to evolution [Vallina et al., *Nature Communications*, 5, 2014].

2) Climate change and human decision process

The first attempt of providing an interactive coupling between human decision-process and climate change has been recently published [Collins et al., *Geosci. Model Dev.*, 8, 2203–2219, 2015]. This highlights the possibilities and limitations in the overall approach. In particular, little to none of actual human psychology and decision process is included in the present coupling.

3) Representation of process in atmospheric boundary layer

The chemical and physical processes occurring in the few tens of meters above the Earth surface actually represent a grey zone in modeling. Certain processes are taken into account in the land model, others in the atmosphere. In reality there is a continuum of processes that are critical to the representation of the chemical state of the atmosphere (e.g. [Kaser, L., et al. (2015), *Chemistry – turbulence interactions and mesoscale variability influence the cleansing efficiency of the atmosphere*, *Geophys. Res. Lett.*, 42, doi:[10.1002/2015GL066641](https://doi.org/10.1002/2015GL066641)]). This would obviously extend to air quality issues and representation of other canopy processes.

Proposed Direction of work

- 1) requires the consideration of many interacting and competing species (i.e. tracers), and introduction of additional information (genome) that need to involve over time.
- 2) requires of an iterative process (projection, identification of climate impact) that will need to rely on large ensembles for significance
- 3) requires rethinking the way models are interacting at their boundary, similar in approach to what is being used for simulating the interaction between ocean and land-ice.

Connections to Math, Comp Sci & and Climate Science

Overall ability to simulate multiple temporal scales (no mention was made of the obvious need of

Contributed ideas paper 57

representing atmosphere and ocean dynamics and physics at the 1-10 km scale) with sufficient throughput and very extensive I/O for comparison with process-based observations.

Potential Impact on the field

Huge impact on the ability of representing adaptation in ESMs, in addition to enabling the representation of a continuum of scales of importance to the human and biological environment.

Building Computational Bridges Across the Water, Ecosystem, and Soil Biogeochemistry Disciplines

Li Li, Yuning Shi, Chris Duffy
Penn State University

Abstract. We propose to develop a code that integrates mechanistic understanding of the below ground processes (soil carbon decomposition and stabilization through C/N – microbe – roots – water – mineral interactions) with terrestrial ecosystem and hydrological models at the watershed scale. Models that integrate mechanistic process understanding across relevant disciplines do not yet exist. Such development will offer tools for predictive understanding of soil C–climate change feedbacks and watershed biogeochemical response to changing climate.

Background/Research to Date. Existing models typically focus on processes of interests in individual research communities / disciplines. The connection / coupling among processes across disciplines, however, lacks a mechanism-based approach. For example, the water-rock-microbe-roots interactions can significantly influence C decomposition and stabilizations, as well as the cycling and fluxes of other elements (N and metals). Mechanistic understanding of C decomposition and stabilization requires understanding of water-rock-microbe-roots interactions across disciplines including hydrology (soil moisture determined by water partitioning into runoff, subsurface water, and evapotranspiration), microbiology (microbial community and dynamics), soil biogeochemistry (microbe-mediated C decomposition, mineral-associated C stabilization), ecosystem biogeochemistry (roots, litter fall, leaf growth), among others. Understanding the complexity of process coupling and predicting emergent behavior in such systems requires an integrated approach representing the current state of art within different disciplines. Such connections across disciplines, however, are not yet established in existing models. For example, most Earth System Models (ESMs) have relatively simple representation of soil microbial processes (Wieder et al., 2015) with bacteria response to warming (temperature), water content however not to water composition and roots. C stabilization through soil C sorption has only recently been explicitly modeled (Tang and Riley, 2015; Wieder et al., 2014). Inclusion of these processes in ESM have led to numerical observation of weaker soil carbon-climate change feedback (Tang and Riley, 2015), as well as prediction of the global soil C response to climate change that is different from that of traditional ESM without microbial mechanisms and soil C sorption processes (Wieder et al., 2013).

Climate change can lead to different precipitation patterns and abundance, therefore leading to shifts in soil moisture conditions, the extent of water interacting with dissolving soils (through mineral dissolution and precipitation, and ion exchange), which can have significant impacts on microbe C decomposition efficiency and the abiotic C stabilization (sorption to mineral surface /clays), as well as how hydrodynamics and soil biogeochemical processes response to climate change. A cross-disciplinary effort is needed in model development to integrate hydrologic, ecosystem, and soil biogeochemistry process coupling to address questions that cannot be addressed within individual disciplines.

Proposed Direction of work. We propose a cross-disciplinary model development effort that integrates hydrologic, ecosystem, and soil biogeochemistry (Reactive transport) processes, as shown in Figure 1. The proposed work will develop a code bioRT-Flux-PIHM-BGC based on existing ecosystem biogeochemistry code Biome-BGC and newly developed hydrologic-reactive transport code RT-Flux-PIHM with new extensions to include belowground biological processes (microbial and root processes).

The developed model can be used to answer questions that cannot be previously answered. For instance, how do vegetation and roots affect soil moisture and microbe spatial distribution, therefore influencing soil C decomposition? How does soil mineralogy influence mineral-associated C stabilization, control C residence time, C fluxes and C cycle feedbacks to climate change? How do vegetation and roots alter soil biogeochemical rates and fluxes (e.g., Ca,

Contributed ideas paper 58

Mg, Na, K, Si, Fe, heavy metals) that are important controls of water quality? How do watershed biogeochemistry processes respond to changing climate?




Figure 1. Representative processes from different disciplines (different colors) in different modules that can be used to integrated into bioRT-Flux-PIHM-BGC. The integration among processes from mechanistic based understanding in the model will allow systematic understanding at the watershed scale. The watershed is based on the Susquehanna Shale Hills (SSHCO), an NSF supported critical zone observatory.

Attributes. The integrated model will allow process coupling across disciplinary boundary, identification of emergent behavior, answers to questions that cannot be answered within disciplines. We recently developed the code RT-Flux-PIHM, the first code that couples subsurface reactive transport with surface hydrology and land-surface interactions. Within the proposed framework, the code will be augmented to include 1) the belowground water-rock-microbe-roots interactions; 2) C and N cycling processes in Biome-BGC. This will ultimately offer unprecedented capabilities to understand and predict the coupling among energy, water, C, N, and biogeochemistry under changing climate. **Limitations.** An integrated model with many different processes is computationally expensive and numerical challenging. Parameterization will also require intensive data sets.

Connections to Math, Comp Sci & and Climate Science. The proposed work stands at the intersection of multiple science disciplines (hydrology, ecosystem sciences, biogeochemistry, climate science), as well as math and computational sciences. Detailed, mechanistic representation of processes means solve many partial different equations over time and space (plant type, C fluxes, soil C composition, water quantity and composition, concentrations of aqueous species). This presents significant mathematical and computational challenges.

Potential Impact on the Field. This proposed model development will fill an important gap in understanding soil carbon-climate change feedback. Reduce uncertainties in future climate prediction. It will also help understand how earth systems, or watershed biogeochemistry (microbe, soil composition, water quality) response to climate change, and the connections among geosphere, biosphere, and atmosphere.

Contributed ideas paper 58

References

- Tang, J. and Riley, W.J. (2015) Weaker soil carbon-climate feedbacks resulting from microbial and abiotic interactions. *Nature Clim. Change* 5, 56-60.
- Wieder, W.R., Allison, S.D., Davidson, E.A., Georgiou, K., Hararuk, O., He, Y., Hopkins, F., Luo, Y., Smith, M.J., Sulman, B., Todd-Brown, K., Wang, Y.-P., Xia, J. and Xu, X. (2015) Explicitly representing soil microbial processes in Earth system models. *Global Biogeochemical Cycles* 29, 1782-1800.
- Wieder, W.R., Bonan, G.B. and Allison, S.D. (2013) Global soil carbon projections are improved by modelling microbial processes. *Nature Clim. Change* 3, 909-912.
- Wieder, W.R., Grandy, A.S., Kallenbach, C.M. and Bonan, G.B. (2014) Integrating microbial physiology and physio-chemical principles in soils with the Microbial-MIneral Carbon Stabilization (MIMICS) model. *Biogeosciences* 11, 3899-3917.

Contributed ideas paper 59

Improved Topography in ACME through Immersed Boundaries July 6, 2016

David M. Hall (Department of Computer Science, University of Colorado Boulder)

Mark Taylor (Sandia National Lab)

Abstract

The current representation of topography is one of the greatest sources of error in the ACME atmospheric dynamical core, particularly at high resolutions. As refinement increases, resolved terrain becomes rougher and steeper. The current terrain-following approach handles steep slopes poorly, producing tilted, distorted coordinate surfaces throughout the lower atmosphere. This induces imbalances between pressure-gradient and geopotential-gradient forces, resulting in spurious velocity fields, nonphysical gravity waves, and a poor representation of low level clouds. As an alternative, we propose to investigate direct forcing immersed boundary techniques commonly employed in aerospace for the representation of complex geometries. This approach is able to represent topography of arbitrary complexity without terrain-following coordinates and without modifying the underlying mesh. We are confident in this approach, as it has been applied successfully in the WRF weather forecast model to represent both topography and urban environments. Furthermore, the immersed boundary method uses only local information, preserving the excellent scaling properties of the atmospheric dynamical core.

Background/Research to Date

While commonly used in aerospace [1], immersed boundary methods are largely unfamiliar to the atmospheric modeling community, with the notable exception of Lundquist et al.[2, 3] who used immersed boundaries to simulate orography and urban environments in the WRF model. Although there are many variants, the direct forcing approach of Mohd Yussaf [4] is particularly simple and effective. In this approach, values are assigned to discrete node points beneath the terrain surface, which induced the desired boundary conditions at the interface, as illustrated in Fig. 1. For example, for a normal wind velocity of 1 m/s just above the surface, assigning -1 m/s at the same distance just below the surface induces a value of zero at the surface, thereby creating a no-flux boundary condition. In practice, the velocity values are obtained by mirroring the immersed points along the surface normal and interpolating field values at the image points. With this technique, the terrain surface need not coincide with the grid points, and neither the mesh nor the coordinate system need be modified to fit it.

Proposed Direction of work

We propose to investigate low and high-order variants of the direct forcing immersed boundary method in the atmospheric dynamical core. Impacts on climatology, gravity wave production, and precipitation would be evaluated using both the primitive equation model and the nonhydrostatic dynamical core currently under development. Implementation would require replacing the hybrid vertical coordinate with constant pressure surfaces and construction of a pre-processing step wherein surface normals are computed and image points are located within each element. In some cases, the image points will lie in a neighboring element, and data at these points will be collected for transmission during the edge exchange step. Performance analysis should demonstrate that this method maintains current levels of throughput and preserves the dynamical core's strong scaling advantages. Optionally, a hybrid approach may be employed

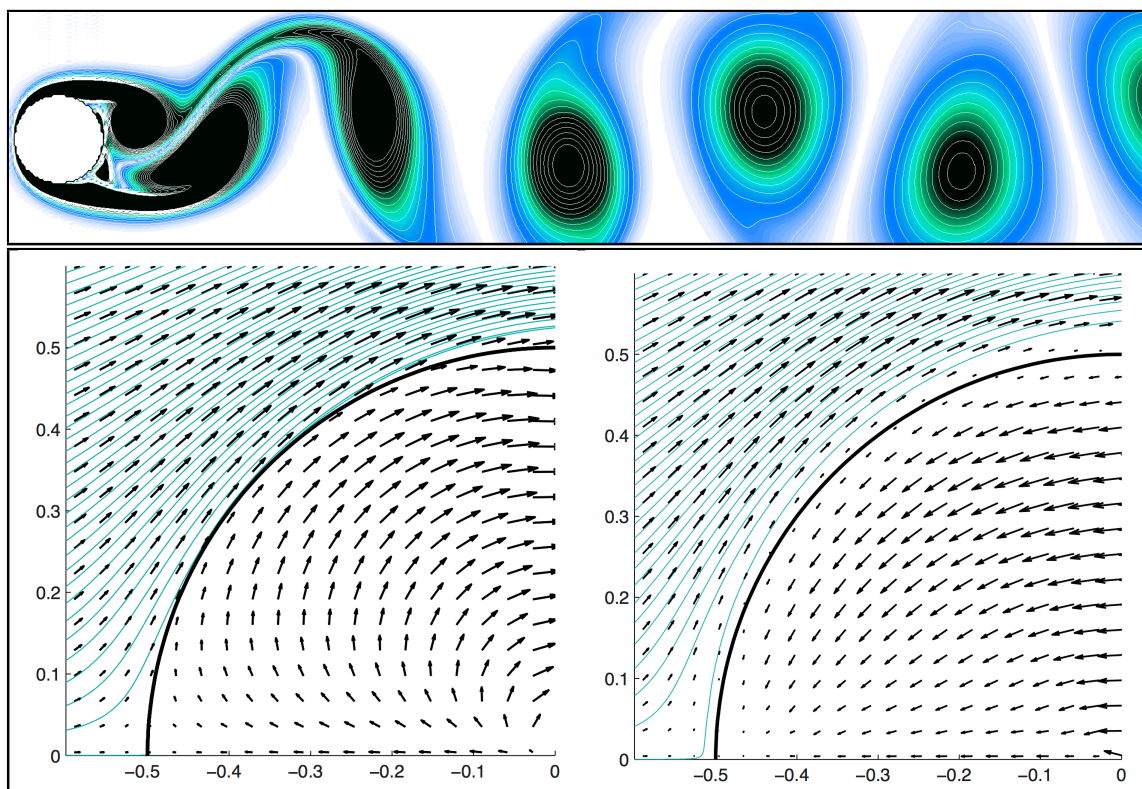


Figure 1: The direct forcing immersed boundary technique assigns values to subsurface nodes in order to induce the desired surface properties. **(top)** Immersed boundary flow past a fixed cylinder at Reynolds number 1000. **(left)** Reflection of the surface-normal velocity component creates a no-flux boundary condition. **(right)** Reflection of both normal and tangential velocity components creates a no-slip boundary condition.

wherein the existing terrain-following coordinates are used to capture smooth continent scale features while the immersed boundary technique is used to resolve local mountainous terrain.

Connections to Math, Comp Sci & and Climate Science

Mathematically, the proposed immersed boundary technique represents an alternative formulation of the boundary condition on the partial differential equations of motion. Computationally, the method implements complex computational geometry with a high degree of locality and parallel-efficiency in distributed computing environments. It is relevant to climate science, as the quality of the topography directly impacts predicted wind fields, wave production, tracer transport, cloud production, and patterns of precipitation.

Potential Impact on the field

This technique has the potential to remove a dominant source of error from the atmospheric dynamical core, thereby improving the quality of all coupled simulations. Furthermore, its importance grows over time as increased levels of refinement are desired and topography becomes better resolved. This technology should have a particularly large impact on predictions of water availability in mountainous regions.

Contributed ideas paper 59

References

- [1] Rajat Mittal and Gianluca Iaccarino. Immersed boundary methods. *Annu. Rev. Fluid Mech.*, 37:239–261, 2005.
- [2] Katherine A Lundquist, Fotini Katopodes Chow, and Julie K Lundquist. An immersed boundary method for the weather research and forecasting model. *Monthly Weather Review*, 138(3):796–817, 2010.
- [3] Katherine A Lundquist, Fotini Katopodes Chow, and Julie K Lundquist. An immersed boundary method enabling large-eddy simulations of flow over complex terrain in the wrf model. *Monthly Weather Review*, 140(12):3936–3955, 2012.
- [4] J Mohd-Yusof. methods for complex geometries. *Annual Research Briefs-1998*, page 325, 1998.



U.S. DEPARTMENT OF
ENERGY

Office of Science



 **OAK RIDGE**
National Laboratory

Managed by UT-Battelle
for the US Department of Energy