Executive Summary

Workshop on Improving ALE Ocean Modeling
NOAA Center for Weather and Climate Prediction
College Park, MD
October 3-4, 2016

A one-and-a-half-day workshop on ocean models that employ the ALE (Arbitrary Lagrangian-Eulerian) method to permit general vertical coordinates was held at NCWCP on October 3-4, 2016. Four different ALE models (GO2, HYCOM, MOM6, and MPAS-Ocean) were discussed. Workshop participants included users and developers of these models, from academia and from seven different national modeling centers (ESRL, GFDL, GISS, LANL, NCAR, NCEP, NRL). A more detailed report on the workshop follows this executive summary, and a workshop agenda with a list of workshop speakers is attached as an appendix.

A number of recommendations for future action were developed during the meeting. The recommendations fall into four broad categories: Code Sharing, Community Building, Code Merger and Performance and Future Development. The recommendations are listed below.

1. Code Sharing

Sharing of common codes for the equation of state, grid generation and remapping, and column physics such as mixing parameterizations, is recommended. The group suggested making vertical remapping/grid generator routines an open source package that can be shared in the manner of CVMix.

a. The development of prognostic equations for the grid generator based upon physical mixing should be considered.
b. We recommend open-source GIT repositories (hosted on code-sharing sites such as Github) for submodules such as CVMix.
c. We recommend the development of common tests for self-consistency, conservation, and known solutions.

2. Community Building

d. The ALE modeling group should consider whether semi-regular meetings similar to this workshop should be undertaken; perhaps via merger/inclusion with the Layered Ocean Model workshop.
e. Following on successful efforts in the atmospheric modeling community, our community should consider developing ocean model development workshops. The long term health of the ocean modeling activity would be significantly enhanced through the sustained sponsorship of integrative activities that develop early career talent and, subsequently, provide cross-model fertilization.

3. **Code Merger**

f. HYCOM and MOM6 have enough similarities to consider a code merger, bringing the strengths of each to a common code base.

4. **Performance and Future Development**

g. The modeling community must confront the physics of boundary layers (at both the top and bottom boundaries) in order to fully exploit the power of ALE models.
h. The group expressed unanimous agreement on the value of funding different approaches to ALE modeling. Impressive results from the DOE MPAS-Ocean and NASA GISS GO2 models demonstrate the need for diversity in ALE modeling. Only by funding a few different ALE streams can we, as a community, be assured that innovations will continue to be developed.
i. To avoid flooding storage silos with large model output datasets, model users should consider, where possible, analyzing runs as they take place rather than afterwards.

Report on "Improving ALE Ocean Modeling Workshop"
NCWCP
October 3-4, 2016

A one-and-a-half-day workshop on ocean models that employ the ALE (Arbitrary Lagrangian-Eulerian) method to permit general vertical coordinates was held at NCWCP on October 3-4, 2016. Four different ALE models (GO2, HYCOM, MOM6, and MPAS-Ocean) were discussed. Workshop participants included users and developers of these models, from academia and from seven different national modeling centers (ESRL, GFDL, GISS, LANL, NCAR, NCEP, NRL). A list of acronyms appears later in this document. A workshop agenda with a list of workshop speakers is attached as an appendix.

The first day included agency overview talks, a discussion of previous attempts to create a unified home for ALE code development, overview talks from all of the modeling groups, discussions of numerical implementation issues, and an overview of applications enabled by ALE models. A few key points that emerged are that (1) ALE is an algorithm for the vertical grid choice, (2) ALE is versatile and permits general vertical coordinates, that include traditional z, isopycnal, and terrain-following coordinates as well as hybrid combinations of the former and other creative treatments yet to be formulated and explored, (3) ALE can eliminate vertical CFL restrictions and minimize velocity errors, and easily enables wetting and drying, (4) a weakness of sorts for ALE models is that grid choice is important, not a simple black box usage, and (5) the MOM6 group has found that model errors have been reduced with the use of a finite-volume pressure gradient force.

The second half-day was set aside for discussions and recommendations. The general topics discussed were numerics and performance, model issues and necessary improvements, code sharing and potential collaborations, and next generation ALE models and future directions. Some key points and recommendations from the discussions held on the second day follow below.

There was some discussion about HPC performance issues that are experienced by all ocean modeling groups including ALE model groups—e.g. limitations on scalability, bit-for-bit multi-CPU reproducibility, refactoring for vector instructions, the low computational intensity of ocean models, and ocean model performance on coming future architectures. In response to a request from David McCarren, workshop participant Alan Wallcraft produced a summary of these issues, attached to this report as an appendix.

The problem of analyzing the massive outputs that large models can generate was discussed. This problem is particularly acute in the case of large high-resolution ensembles. One proposed solution to the large output problem is to analyze model output on the fly as it is running, rather than saving enormous amounts of high-spatial and temporal resolution output for later analysis. For example, in the context of data assimilation (an application for which large ensembles are likely to be used), the capability to output model fields interpolated to specified observation locations would considerably reduce I/O costs. In data assimilation algorithms, the model state must be interpolated to the observation locations in order to compute the analysis adjustments to be made to the model forecast. Without such a feature embedded in the model, 4D-EnKF and 3D-FGAT ("first guess at the appropriate time") methods would require outputting the full model state at or near the observation frequency to achieve similar results. For example, in the context of data assimilation (an application for which large ensembles are likely to be used), the capability to output model fields interpolated to specified observation locations would considerably reduce I/O costs. In data assimilation algorithms, the model state must be interpolated to the observation locations in order to compute the analysis adjustments to be made to the model forecast. Without such a functionality embedded in the model, 4D-EnKF and 3D-FGAT ("first guess at the appropriate time") methods would require outputting the full model state at or near the observation frequency to achieve similar results. For example, an EnKF with a 50-member ensemble assimilating hourly along-track SST data and randomly located in situ profiles over a 24-hour analysis cycle would require 50*24=1200 full model states to be output. With an embedded interpolation scheme, only the model state interpolated to the satellite tracks and sparse in situ profiles would need to be output at those same times, plus the 50 full model states at the analysis time.

This online approach applies most clearly to a reanalysis scenario, in which the observation locations are known in advance. But even in a forecast scenario it may be more efficient to rerun the ensemble of models and supply them with the observation locations to do an online computation of the required interpolations than to output all the model states during the original forecast and then compute the interpolations offline. The observation-minus-forecast and observation-minus-analysis data are the fundamental pieces of information produced in the data assimilation, and these are likely the primary ensemble data to be maintained and stored after the DA processing.

ALE-specific modeling issues were also discussed during the second half-day. For example, the HYCOM group will work to implement the Adcroft et al. 2008 pressure gradient approach, as part of an NRL Nordic Seas project. The HYCOM group will also explore using interface, rather than layer-average, density to guide vertical remapping.

There was much discussion of vertical remapping. Community users of model output generally want output in z-space (or pressure levels for non-Boussinesq configurations). The remapping to z levels has to be done carefully within an ALE model in order to respect conservation laws. An interesting idea that came up is development of prognostic equations for the grid generator based upon physical mixing. It was pointed out that ALE models can be made non-hydrostatic, but this feature has not been implemented yet because ALE models have not been run at horizontal resolutions fine enough to merit this approach. Another important point that came up is that the modeling community must confront the physics of boundary layers (at both the top and bottom boundaries) in order to fully exploit the power of ALE models.

It became clear at the meeting that different groups were using different names for the same type of grid related computation. The term regridding was suggested to be confusing and should be replaced with remapping. Alistair Adcroft suggested normalizing the nomenclature for future clarity. The group also suggested making vertical remapping/grid generator routines a common package that can be shared in the manner of CVMix. The discussions on code sharing and potential collaborations occupied a longer time than the other themed discussions, and led naturally into the final session on next generation ALE models and future directions. The group recommends open-source GIT repositories (hosted on code-sharing sites such as Github) for submodules and common codes to be developed. Sharing of common codes for the equation of state, grid generation and remapping, and column physics such as mixing parameterizations (via CVMix), was recommended. There was widespread support for the development of common tests for self-consistency, conservation, and known solutions.

The idea was floated that this ALE modeling group consider meeting annually or bi-annually, and that one potential format is to absorb such meetings into the Layered Ocean Model (LOM) meetings that have been taking place biannually for many years. The latter idea found favor with some but not with all. An advantage of incorporating ALE workshops into LOM meetings is that it would entail less travel. A potential disadvantage is that the LOM meeting has many presentations on science issues and incorporating the ALE workshops into it may make it too large. The question of whether future ALE workshops should be held, and whether they should be held during LOM meetings, was left unresolved.

One very important point made near the end of the discussions is that the ocean numerical model development community is small and that there is a need to entrain more early-career scientists into the field. The atmospheric modeling community has been very successful in this realm. For example, the Dynamical Core Model Intercomparison Project summer schools (https://goo.gl/7WDpSX) pair international atmophere modeling groups with aspiring graduate students and postdocs. In addition, since the mid 1990s the atmosphere modeling community has nurtured a core capability in computational physics through the bi-annual PDEs on the Sphere workshop. These activities not only develop early-career scientists but also provide an incubator for new ideas. The ocean modeling community lacks similar integrative activities and, instead, has tended toward model-centric workshops (e.g. NCAR OMWG, LOM, ROMS). The long term health of the ocean modeling activity would be significantly enhanced through the sustained sponsorship of integrative activities that develop early-career talent and, subsequently, provide cross-model fertilization.

Another critical point made was that some groups that use similar model codes consider merging their codes for the benefit of both groups. One such potential merger is between the HYCOM and MOM6. A merged model code that incorporated HYCOM's real-time data assimilation packages into a MOM6 dynamical core would partially satisfy the Navy's stated needs to upgrade HYCOM's dynamical core by ~2023, and would satisfy NOAA's desire for a common dynamical core to use in climate modeling, seasonal-interannual forecasting, and near-real-time ocean forecasting. **The group generally agreed that a closer working relationship between NOAA, NRL, and academic partners such as University of Michigan and Florida State University, involving an as-yet-to-be-determined merger between HYCOM and MOM6, was desirable.** It is important to note that such a merger does not prevent any scientist from continuing to use their own versions of HYCOM and MOM6 as they see fit. NRL, GFDL, and Michigan are already working on a MOM6 test run on NRL's 1/12th degree HYCOM grid with atmospheric and tidal forcing. A comparison of results from this MOM6 test run with a twin HYCOM run is expected to be very informative. The group recognized that an inter-agency project to set up and run a model for the benefit of both agencies is a time-consuming endeavor, that will need to be supported with dedicated resources. **The group noted that code sharing is different from code merging. HYCOM and MOM6 have enough similarities to consider a code merger, bringing the strengths of each to a common code base. The grid structure for MPAS-Ocean and GO2 are sufficiently different to preclude a code merger. However, there remain opportunities for code sharing of vertical column physics.**

Finally, the group expressed unanimous agreement on the value of funding different approaches to ALE modeling. Impressive results from the DOE MPAS-Ocean and NASA GISS GO2 models demonstrate the need for diversity in ALE modeling. Only by funding distinct ALE streams can we be assured that innovations will continue to be developed. On a related note, the group recognized that ocean modeling involves people as much as it involves codes; as noted above, the nurturance of young talented model developers is critical for the long-term health of the field.

A number of recommendations for future action were developed during the meeting. The recommendations fall into four broad categories: Code Sharing, Community Building, Code Merger and Performance and Future Development. The recommendations are listed below.

1. Code Sharing

a. The group suggested making re-mapping/grid generator routines an open source package that can be shared in the manner of CVMix. .
b. The group recommends open-source GIT repositories (hosted on code-sharing sites such as Github) for submodules such as CVMix.

c. Sharing of common codes for the equation of state, grid generation and remapping, and column physics such as mixing parameterizations, was recommended.
d. There was widespread support for the development of common tests for self-consistency, conservation, and known solutions.

2. Community Building

a. The ALE modeling group should consider whether semi-regular meetings similar to this workshop should be undertaken; perhaps via merger/inclusion with the Layered Ocean Model workshop.
b. Following on successful efforts in the atmospheric modeling community, our community should consider developing ocean model development workshops. The long term health of the ocean modeling activity would be significantly enhanced through the sustained sponsorship of integrative activities that develop early-career talent and, subsequently, provide cross-model fertilization.

3. Code Merger

a. HYCOM and MOM6 have enough similarities to consider a code merger, bringing the strengths of each to a common code base. The grid structure for MPAS-Ocean and GISS are sufficiently different to preclude a code merger. However, there remain opportunities for code sharing of vertical column physics.

4. Performance and Future Development

a. The development of prognostic equations for the grid generator based upon physical mixing should be considered.
b. The modeling community must confront the physics of boundary layers (at both the top and bottom boundaries) in order to fully exploit the power of ALE models. The group expressed unanimous agreement on the value of funding different approaches to ALE modeling. Impressive results from the DOE MPAS-Ocean and NASA GISS GO2 models demonstrate the need for diversity in ALE modeling. Only by funding a few different ALE streams can we as a community be assured that innovations will continue to be developed.
c. Ocean models and ALE models, in particular, haven't been designed for some of the potential new computer architectures. Ocean models tend to perform little computational work relative to the memory access. Moving ocean models to these new architectures will require substantial development effort.
d. To avoid flooding storage silos with large model output datasets, model users should consider, where possible, analyzing runs as they take place rather than afterwards.

Acronyms:
ALE: Arbitrary Lagrangian-Eulerian
DOE: Department of Energy
ESPC: Earth System Prediction Capability
ESRL: Earth System Research Laboratory
GFDL: Geophysical Fluid Dynamics Laboratory
GISS: Goddard Institute for Space Studies
GO2: Goddard Ocean Model 2
HPC: High Performance Computing
HYCOM: HYbrid Coordinate Ocean Model
LANL: Los Alamos National Laboratory
LOM: Layered Ocean Model meeting
MIT: Massachusetts Institute of Technology
MOM6: Modular Ocean Model version 6
MPAS-Ocean: Model for Prediction Across Scales - Ocean
NASA: National Aeronautics and Space Administration
NCAR: National Center for Atmospheric Research
NCEP: National Centers for Environmental Prediction
NCWCP: NOAA Center for Weather and Climate Prediction
NOAA: National Oceanic and Atmospheric Administration
NRL: Naval Research Laboratory
NWS: National Weather Service
OMWG: Ocean Model Working Group
OSTI: Office of Science and Technology Integration
ROMS: Regional Ocean Modeling System

Workshop organizers:
Brian Arbic, University of Michigan
Robert Hallberg, NOAA GFDL
Avichal Mehra, NOAA NCEP
James Richman, Florida State University

Workshop participants:
Alistair Adcroft, NOAA GFDL
Rainer Bleck, NASA GISS and NOAA ESRL
Frank Bryan, NCAR
Eric Chassignet, Florida State University
James Carton, University of Maryland
Arun Chawla, NOAA NCEP
Gokhan Danabasoglu, NCAR
Darren Engwirda, MIT and NASA GISS
Stephen Griffies, NOAA GFDL
Matthew Harrison, NOAA GFDL
Mary Hart, NOAA NCEP
Robert Helber, NRL

Patrick Hogan, NRL
Max Kelley, NASA GISS
Hyun-Sook Kim, NOAA NCEP
David Lee, LANL
Leonardo Lima, INPE Brazil
Liyan Liu, NOAA NCEP
Carlos Lozano, NOAA NCEP
John Marshall, MIT
Joe Metzger, NRL
Shastri Paturi, NOAA NCEP
Steve Penny, University of Maryland / NOAA NCEP
Mark Petersen, LANL
Todd Ringler, LANL
Ilya Rivin, NOAA NCEP
Jay Shriver, NRL
Luyu Sun, University of Maryland
Shan Sun, NOAA ESRL
Hendrik Tolman, NOAA NWS/OSTI
Alan Wallcraft, NRL
Yan Xue, NOAA NCEP

Participating program managers:

Dan Eleuterio, Office of Naval Research
David McCarren, NOAA
Jesse Carmen, NOAA

Appendix 1. Agenda: Workshop on Improving ALE Ocean Modeling

Venue

NCWCP College Park, MD, Oct, 3-4, 2016

Preamble

In the past few years, a number of research groups in the US have shifted ocean general circulation models from a fixed or single vertical coordinate to an Arbitrary Lagrangian Eulerian (ALE) vertical coordinate system. This shift takes advantage of the superiority of different vertical coordinate systems in different model locations (e.g., shelf vs. deep-ocean, weakly stratified mixed layer versus well-stratified interior). These relative performance advantages have led to a diversity of ocean models. The ALE framework offers the potential for a single model to exploit these relative advantages. However, a convergence in the vertical coordinate system doesn't imply convergence to a single ocean model. The numerical, dynamical and physical implementation and choices of these models differ substantially. Yet, there are aspects of the models that are similar and could be leveraged to improve all ALE models.

Goals of the Workshop

This workshop will gather experts associated with 3 in-use ALE Ocean Models, HYbrid Coordinate Ocean Model (HYCOM), Modular Ocean Model (MOM6) and Model for Prediction Across Scales- Ocean (MPAS-Ocean), as well as ocean modeling and forecasting leaders from academia and national laboratories, to exchange information about the strengths and weaknesses of the ALE models, their numerical, dynamical and physical implementations, and new developments in ALE modeling. The aim of this exchange is fostering improvements in the models by using the strengths of one model to address the weaknesses of another. We will also discuss the applications enabled by ALE models.

Day 1 Monday, October 3

0800 Jim Richman and Brian Arbic: Welcome and goals of Workshop

National Overview

0815 Henrick Tolman: NWS/NGPPS Modeling Plans

0900 Dan Eleuterio: National Earth System Prediction Capability Perspective

0930 Eric Chassignet: Previous attempt to create a unified home for ALE code development (remote presentation)

1000-1030 Coffee Break

ALE Numerics

1030 Rainer Bleck: Historical overview of ALE development

1100 Alistair Adcroft: MOM6 ALE overview + ALE-enabled general coordinates

1130 Darren Engwirda: MIT & GISS High-order accurate reconstructions, re-gridding and pressure gradient evaluation for an ALE ocean model.

1200-1330 Lunch

Model overviews

1330 Max Kelley: GISS The new ALE-enabled GISS ocean model

1400 Mark Petersen and Todd Ringler: Overview of MPAS-Ocean ALE efforts

1430 Alan Wallcraft: Overview of HYCOM ALE efforts

1500-1530 Coffee Break

ALE Applications

1530 Bob Hallberg: Applications enabled by ALE modeling

1600 Future directions/Collaborations

Day 2 Tuesday, October 4 (half day)

0800 Numerics and performance Mehra lead and Shriver rapporteur
Long range simulations require faster performance
~6 min/model day versus current HYCOM ~30 min/model day
Resolution versus high order operators
I/O limitations

0900 Current model issues and necessary improvements Wallcraft/Adcroft lead and Hogan rapporteur
Thermobaric instability in HYCOM
Velocity remapping with ALE operator

1000-1030 Coffee Break

1030 Code sharing and potential collaborations Hallberg lead and Richman rapporteur

1130 Next generation ALE model and future directions Arbic lead and Penny rapporteur

1230 Meeting close

Appendix 2. Computational Aspects of Global Ocean Models

Introduction

Historically structured grid codes have dominated ocean models with some unstructured and semi-structured codes. The models are 3-D but with some of the characteristics of a 2-D problem. Vertical scales of the ocean models are much different from horizontal scales.

HYCOM 1/25th degree fully-global: 9000 x 6595 in horizontal with 41 vertical levels.

Typically, the models use a 2-D domain decomposition.

The vertical dimension is "on-chip" and often treated implicitly.

Ocean models have fast surface gravity waves O(100m/s) which is O(100)x faster than advection and internal gravity wave speeds, motivating a separate 2-D sub-problem with a split-explicit or (less often) semi-implicit time step.

Limits on Ocean Model Scalability

A major problem for ocean models is that little computational work is performed compared to the memory access.

For the 2-D sub-problem, memory accesses are required for the 2-D Halo exchanges and/or 2-D global sums with relatively little computational work between memory accesses. Thus, performance is highly dependent on communication latency.

For the 3-D sub-problem, memory accesses are required for the 3-D Halo exchanges with still relatively little computational work per halo exchange (or per memory access) and still dependent on communication latency.

Typically no overlap between I/O and computations. Thus, I/O eventually limits scalability.

Can get good scalability with a large enough grid or ensembles.

HYCOM 1/25th dgree global tripole (9000 x 6595 x 32): in practice almost exactly 16x faster on 16,000 vs 1,000 cores of Cray XC40 or SGI ICE systems.

Bit-for-Bit Multi-CPU Reproducibility

A requirement for porting to new architectures and different numbers of processors is bit-for-bit reproducibility. Repeating a single processor run produces identical results under this condition. However, repeating a multi-processor run, produces different results, using either OpenMP or MPI. e.g. fastest global sum is non-reproducible unless programmer explicitly avoids non-reproducible operations. We require reproducibility on any number of processors. Then we can test a compiler/system setup once, rather than for every core count.

However, we can't use the highest level of compiler optimization ifort -fp-model precise -no-fma

fp-model precise because vector and scalar operations have different rounding, so the start and end of loop extents can't be scalar if the middle is vector.

fused multiply-add is new with AVX2, it has different rounding and so must be used for all operations in a loop or none.

The Intel compiler is not providing the fastest possible reproducible results.

In some cases this can be worked around with extra coding, but should not be necessary.

Refactoring for Vector Instructions

Earlier generations of ocean models targeted Cray vector shared memory, but recently replaced by models targeting distributed memory (MPI) "scalar" systems.

However, all modern processors either include vector instructions or work well with vector constructs (GPGPUs).

HYCOM is 5% faster on Xeons if its 1st array dimension is a multiple of 8.

The best example of this organization is vertical column physics.

To vectorize the vertical column, we would "push" a horizontal index into the routines. This index is promoted back outside the routine in modern codes.

Taken to an extreme, the single column routine can be very inefficient.

Generally, the best approach is to have a shallow nest of subroutines which a compiler (e.g. on GPGPUs) might be able to in-line into the outer loop to expose the parallelism.

One possible approach to vector refactoring of column physics.

Pack the horizontal dimensions into one index with no land. Have all column arrays aligned for vectorization with exactly the native vector length (pad the length in the last call). Use compile time constants and compiler directives to force maximum vectorization of the column physics routines. The native vector length would be system dependent

Ocean Models on Attached Processors

The low computational intensity of ocean models has been an issue on attached processors. The cost of repeatedly moving arrays from system (host) memory to attached memory is prohibitive.

Only viable approach:

Copy all model arrays to attached memory.

Run MPI across attached processors (without involving the host).

Use the host only for start up and I/O.

I/O includes error reporting, which may
require re-factoring the error handler.

This means that "incremental" approaches to porting won't work. We can't do one subroutine at a time and the attached processor must have enough memory to hold all arrays.

1/25th degree global HYCOM requires 850GB of memory plus tiling overhead. We still must face the low computational intensity bottleneck and may not get good performance without major code re-factoring.

Ocean Models on Future Systems

The memory and programming limitations of attached processors are being reduced over time, which makes host memory more accessible and increase size of "fast" memory. Host-less "attached" processors, with "fast" memory treated as a cache, may provide an option,

Host-less approach involves "more slower cores" vs "fewer faster cores." Currently Intel Knights Landing single socket node with 72 cores per socket vs Intel Xeon dual socket nodes with ~18 cores per socket. Knights Landing has enhanced vector operations but may require more use of hyper-threading for good performance, for example,72 vs 36 cores per node. The question remains which is a) faster per node, b) faster per watt, or c) faster per dollar? In the future ARM server chips with vector extensions will join the "more slower cores" class.

In general, ocean models can scale well (favors more cores) but may need re-factoring to take advantage of vector hardware.
Knights Landing may need hyper-threading for maximum performance.
Increase the number of MPI tasks, or use MPI and OpenMP.

Future Architectures of Interest

Two new architectures are emerging to explore these issues.
IBM Summit/Sierra (2018)
- 140PFlops at 10MW; 14GF/W
  - 3,400 to 4,200 nodes
- Multiple IBM POWER9s and multiple NVIDIA Volta GPGPUs per node
- 512GB RAM per node (high BW memory + DRAM)
- 800GB NVRAM per node (either extension to memory or burst buffer)
- NVLink on-node interconnect (CPUs and GPUs in a common memory space)
- Dual-rail IB-4X EDR (200Gbps) between node interconnect
- GPFS Parallel Filesystem (120 PB; 1TB/s)

Intel/Cray Aurora (2019)
- 180PFlops at 13MW; 14GF/W
  - 50,000+ nodes

1 Intel Knights Hill (3rd generation Phi) per node
128GB+ RAM per node (high BW memory + DRAM)
Intel Omni-Path Gen 2 between node interconnect
Intel SSD burst buffer in each node
Lustre Parallel Filesystem (150 PB; 1TB/s)