

MPAS-O: Plan for 2013

Todd Ringler

Theoretical Division
Los Alamos National Laboratory

Staffing, Effort and Focus

Core (i.e. reallocatable, ~1.0 FTE each):

Todd Ringler: manager, salesman and part time modeler

Mark Petersen: algorithm development, implementation, testing, model evaluation

Doug Jacobsen: performance evaluation, optimization and testing

Leveraged (i.e. deliverables with non-empty intersection with MPAS-O needs, ~0.5 FTE each)

Rob Lowrie: high-order transport

Qingshan Chen: mesoscale eddy parameterization

Nathan Urban: data assimilation

Consultants (i.e. answer the door if we knock loud enough, charge per question.)

Mat Maltrud: the man with 10 hands one for each pot.

Phil Jones: COSIM lead and full time cat herder.

Changes in CY12:

Matt Pratola (MPAS-O/DART) and Jonathan Graham (LES) left.

Nathan Urban arrived.

Quick Look at 2012: Goals and Outcomes

Quick Look at 2012: Goals and Outcomes

CY2012 Goal #1 -- Paper to introduce MPAS-O to Community:

- + Submitted to Ocean Modeling in September
- + In review
- + Revised version available here: <https://www.dropbox.com/s/7619np99h6njvkp/multiResolutionOceanRI.pdf>

Quick Look at 2012: Goals and Outcomes

CY2012 Goal #1 -- Paper to introduce MPAS-O to Community:

- + Submitted to Ocean Modeling in September
- + In review
- + Revised version available here: <https://www.dropbox.com/s/7619np99h6njvkp/multiResolutionOceanRI.pdf>


CY2012 Goal #2 -- A quality MPAS-O simulation produced through the CESM:

- + Design MPAS-O/CESM forcing module
- + Run MPAS-O with data atmosphere forcing
- + Present results at OMWG
- Do 100 year simulation
- Implement KPP
- Implement GM


Manuscript: A Multi-Resolution Approach to Global Ocean Modeling

Ringler, T., Petersen, M., Jacobsen, D.L, Higdon, R. L., Jones, P. W., & Maltrud, M. (2012).
A Multi-Resolution Approach to Global Ocean Modeling. *Ocean Modeling*.

SSH RMS
uniform 15 km global



SSH RMS
variable resolution 15 km in NA




MPAS-O within the CESM system

MPAS-O and CPL are talking.

Configuration, interpolation and coupling issues are mostly resolved.

Should take the opportunity to rethink the ocean model / coupler design.



fresh-water flux

Priorities for 2013

1. Beta release on Sourceforge before June 3, 2013
2. Four peer-reviewed papers
 - a. evaluation of vertical coordinate
 - b. comparison to POP et al.
 - c. high-order transport
 - d. variable resolution GM
3. Improve performance by 50%

Priorities for 2013

1. Beta release on Sourceforge before June 3, 2013
2. Four peer-reviewed papers
 - a. evaluation of vertical coordinate
 - b. comparison to POP et al.
 - c. high-order transport
 - d. variable resolution GM
3. Improve performance by 50%

Purpose of 2013 plan is to be ready for CORE and coupled CESM simulations in 2014.

Goal #1: Beta release on Sourceforge before June 3, 2013
(note: release coordinated with MPAS-A)

Goal #1: Beta release on Sourceforge before June 3, 2013 (note: release coordinated with MPAS-A)

Purpose of initial release is the following:

1. make us clean up code, test cases, documentation
2. engage developers who are interested in the framework
3. provide baseline code for doing science.

Goal #1: Beta release on Sourceforge before June 3, 2013 (note: release coordinated with MPAS-A)

Purpose of initial release is the following:

1. make us clean up code, test cases, documentation
2. engage developers who are interested in the framework
3. provide baseline code for doing science.

Included in the initial release will be:

1. tagged source code
2. scripted test-case suite
3. suite of real-world meshes and restarts

Goal #1: Beta release on Sourceforge before June 3, 2013 (note: release coordinated with MPAS-A)

Purpose of initial release is the following:

1. make us clean up code, test cases, documentation
2. engage developers who are interested in the framework
3. provide baseline code for doing science.

Included in the initial release will be:

1. tagged source code
2. scripted test-case suite
3. suite of real-world meshes and restarts

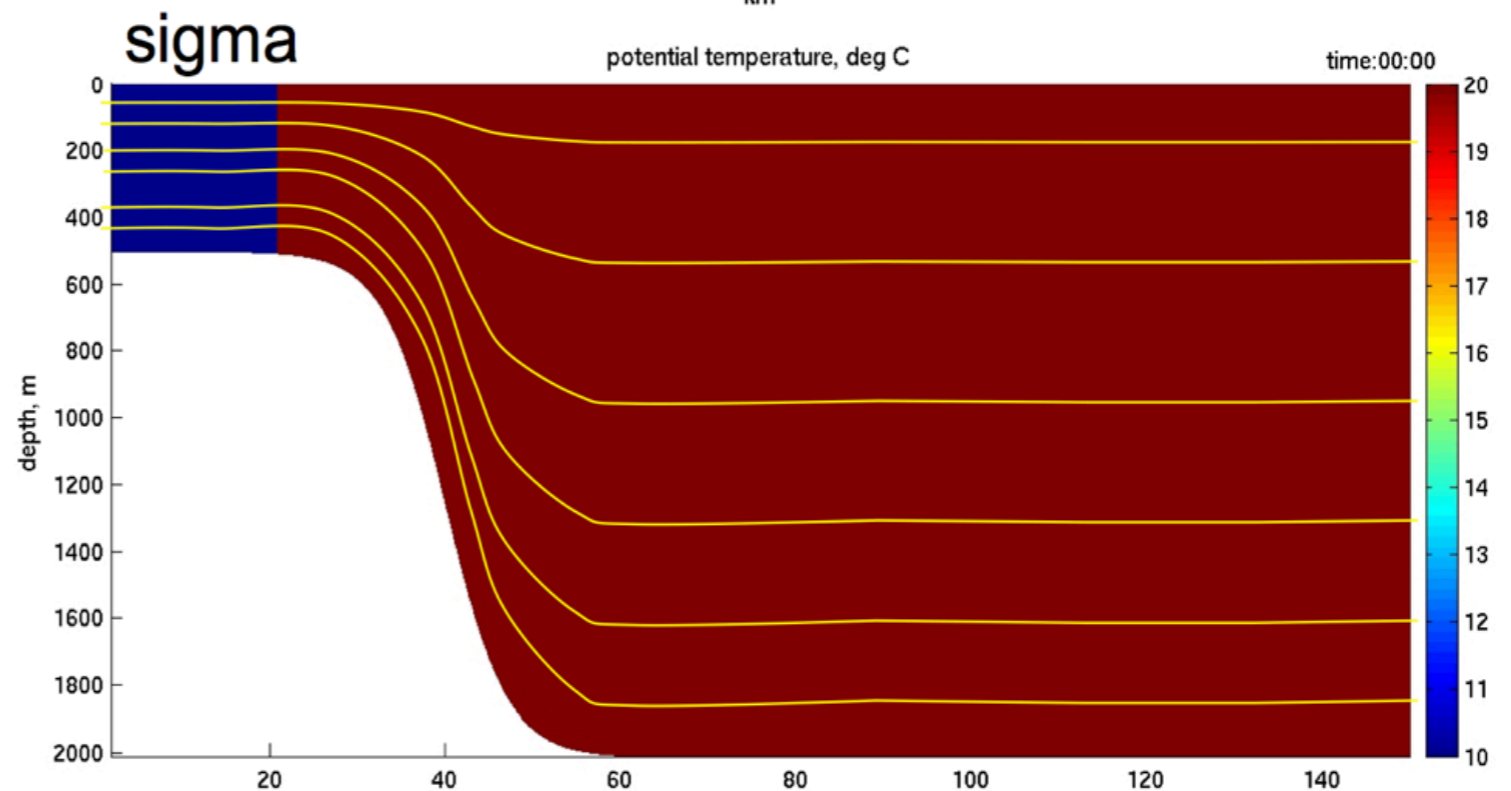
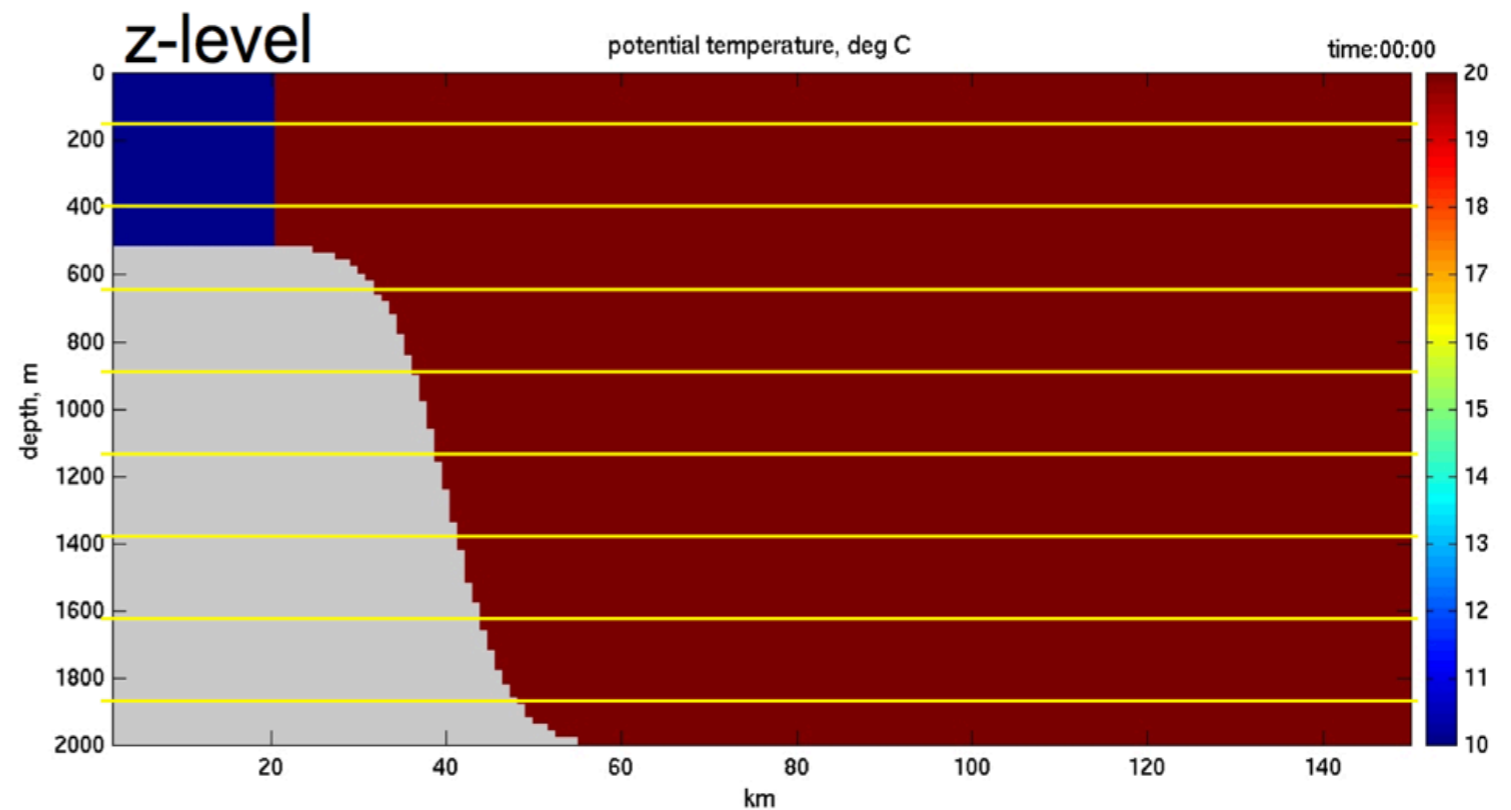
Once code is on Sourceforge, we can point the CESM build script to the open source svn repository.

Goal #2, Paper #1: Evaluation of vertical discretization

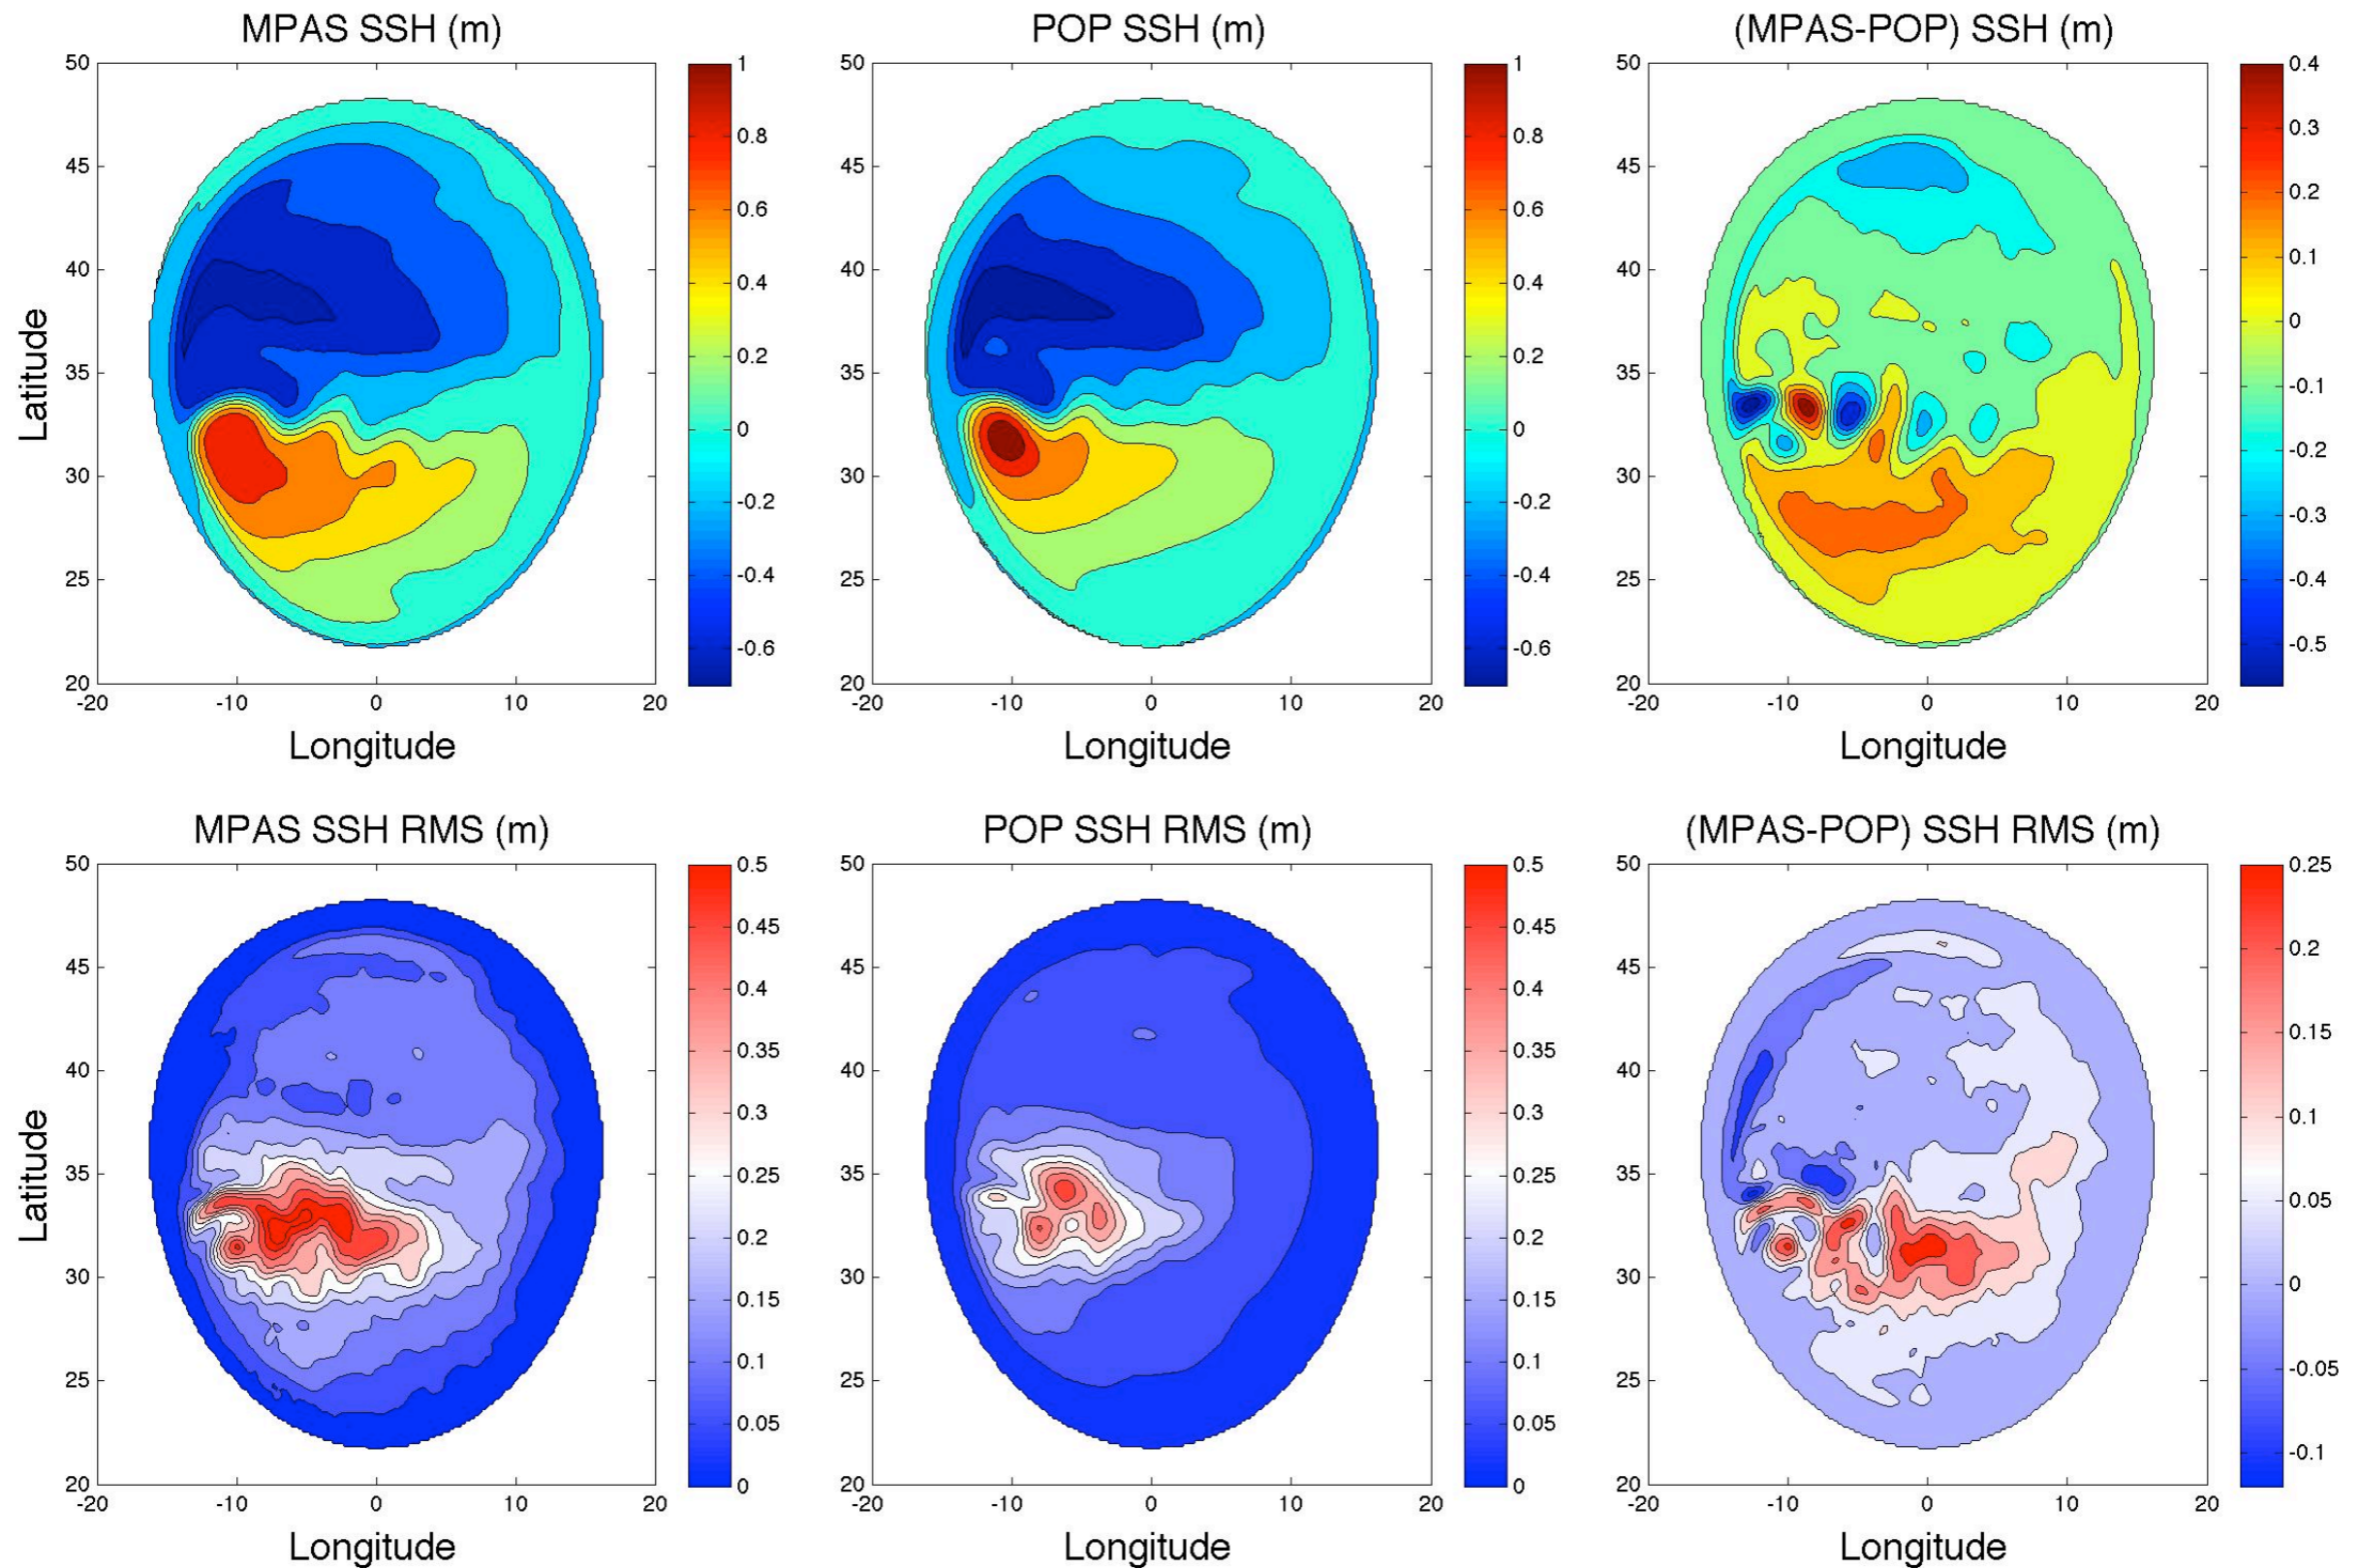
The model has the capability to track arbitrarily positioned (non-zero thickness) layer interfaces.

Paper will compare z-level, z-star, z-tilde, sigma and (hopefully) hybrid grids in terms of spurious dipycnal mixing.

Outcome will be a default vertical grid specification for the 2014 CORE and coupled model simulations.



Goal #2, Paper #2: Comparison to POP (and other OGCMs)



Draft specification of test case: <https://www.dropbox.com/s/ww8e1fj2cf26fdq/SOMA.pdf>

Goal #3: Improve performance by 50%

Current scheme is a two time-scale system

```
do time step
  do barotropic subcycling
    solve 2d SWM system
  end do

  get target layer interfaces

  solve 3D baroclinic system

  tracer update with FCT

  solve for derived quantities
end do
```


Goal #3: Improve performance by 50%

Current scheme is a two time-scale system

```
do time step
  do barotropic subcyling
    solve 2d SWM system
  end do
  get target layer interfaces
  solve 3D baroclinic system
  tracer update with FCT
  solve for derived quantities
end do
```

In addition, the proposed scheme limits diffusion through its Lagrangian treatment of the vertical and limiting the number of times per simulation that tracer transport is conducted.

Proposed scheme is a three time-scale system

do tracer time step

```
do baroclinic time step
  do barotropic subcyling
    solve 2d SWM system
  end do
  no vertical transport
  solve 3D baroclinic system
  accumulate horizontal mass flux
  solve for derived quantities
end do (baroclinic)
get target layer interfaces
tracer update with FCT
  solve for derived quantities
end do
```

Other Important Tasks that are on the List

Link, test, evaluate CVMix within MPAS-O system

Code, test, evaluate GM with MPAS-O system.

Finish off CESM / MPAS-O coupling design and implementation.

Implement rate-of-strain tensor computation

ACC / MOC deep ocean stratification

Coupling MPAS-O to ice shelves

Thoughts on building an durable analysis toolbox

What is the appropriate “model” for ocean analysis given the following:

1. OGCMs will be typically be using 1000s to 10000s processors.
2. Disk is moving away from coprocessor so fast “it is red-shifted.”¹
3. The burden of storing, retrieving and moving data continues to grow.

¹ attributed to Rich Loft

Thoughts on building an durable analysis toolbox

What is the appropriate “model” for ocean analysis given the following:

1. OGCMs will be typically be using 1000s to 10000s processors.
2. Disk is moving away from coprocessor so fast “it is red-shifted.”¹
3. The burden of storing, retrieving and moving data continues to grow.

We all recognize the growing need to “process on the fly.”

¹ attributed to Rich Loft

Thoughts on building an durable analysis toolbox

What is the appropriate “model” for ocean analysis given the following:

1. OGCMs will be typically be using 1000s to 10000s processors.
2. Disk is moving away from coprocessor so fast “it is red-shifted.”¹
3. The burden of storing, retrieving and moving data continues to grow.

We all recognize the growing need to “process on the fly.”

Proposal: Leverage all of the effort put into the OGCM.

1. Upside: MPI, netCDF, PIO, model operators are free.
2. Downside: Simulation and Analysis are typically separate enterprises.

¹ attributed to Rich Loft

Thoughts on building an durable analysis toolbox

What is the appropriate “model” for ocean analysis given the following:

1. OGCMs will be typically be using 1000s to 10000s processors.
2. Disk is moving away from coprocessor so fast “it is red-shifted.”¹
3. The burden of storing, retrieving and moving data continues to grow.

We all recognize the growing need to “process on the fly.”

Proposal: Leverage all of the effort put into the OGCM.

1. Upside: MPI, netCDF, PIO, model operators are free.
2. Downside: Simulation and Analysis are typically separate enterprises.

```
make ifort CORE=swmodel  
make ifort CORE=nhyd_atmos  
make ifort CORE=ocean
```

¹ attributed to Rich Loft

Thoughts on building an durable analysis toolbox

What is the appropriate “model” for ocean analysis given the following:

1. OGCMs will be typically be using 1000s to 10000s processors.
2. Disk is moving away from coprocessor so fast “it is red-shifted.”¹
3. The burden of storing, retrieving and moving data continues to grow.

We all recognize the growing need to “process on the fly.”

Proposal: Leverage all of the effort put into the OGCM.

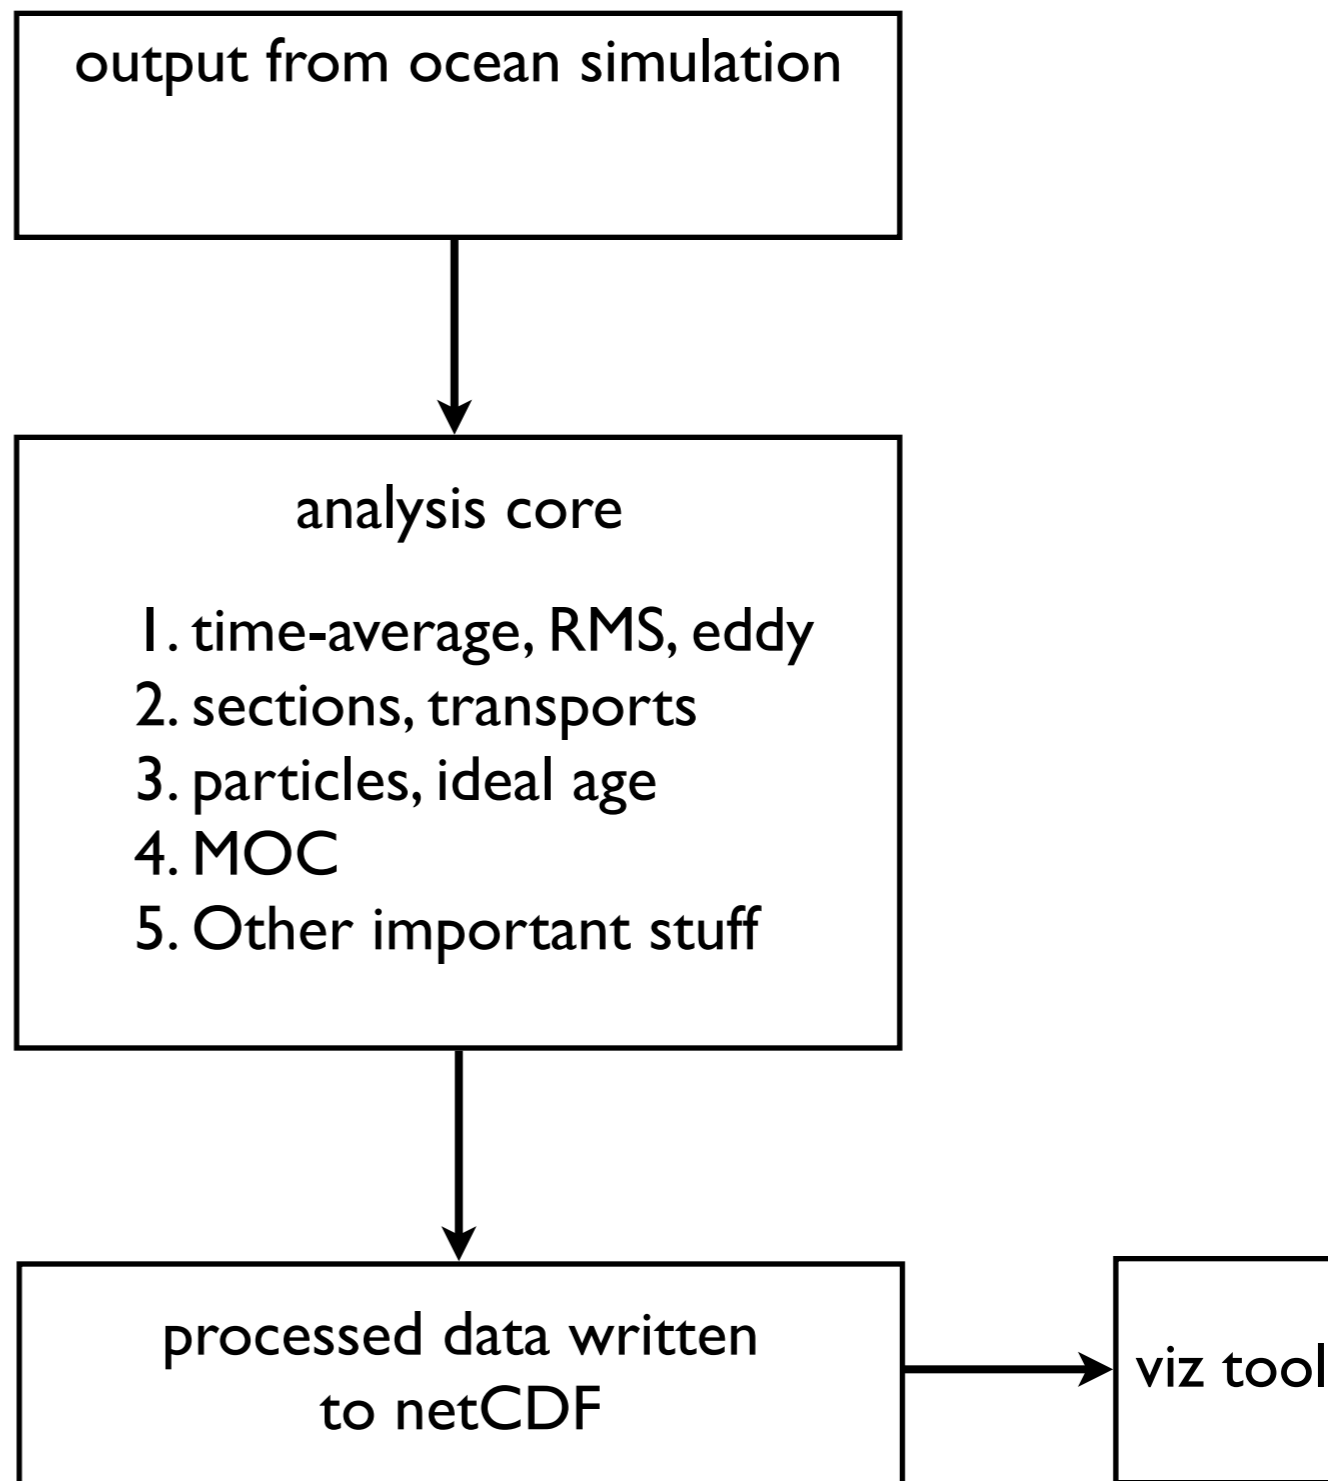
1. Upside: MPI, netCDF, PIO, model operators are free.
2. Downside: Simulation and Analysis are typically separate enterprises.

```
make ifort CORE=swmodel  
make ifort CORE=nhyd_atmos  
make ifort CORE=ocean
```

why not make ifort CORE=ocean_analysis ?

¹ attributed to Rich Loft

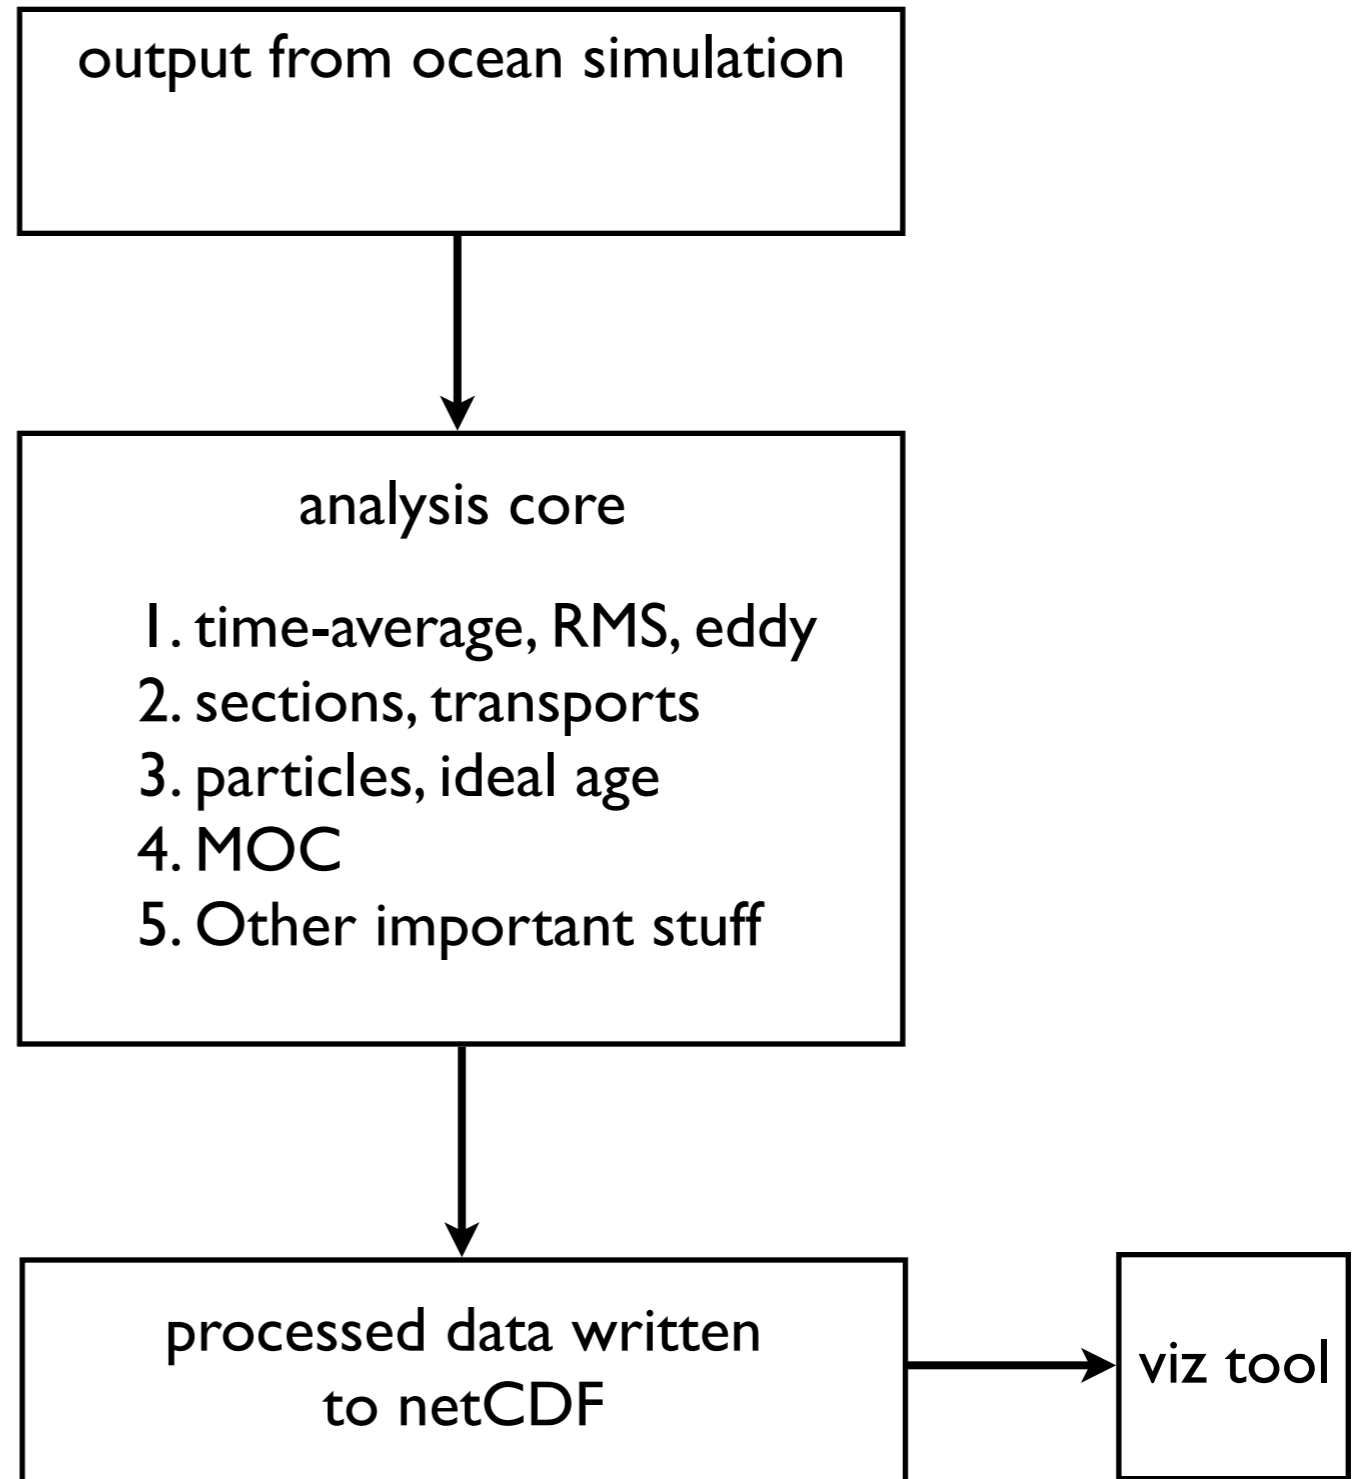
make ifort CORE=ocean_analysis



make ifort CORE=ocean_analysis

Scenario #1

Scientist prototypes analysis tools in desktop environment, reading in simulation data, writing FORTRAN code.



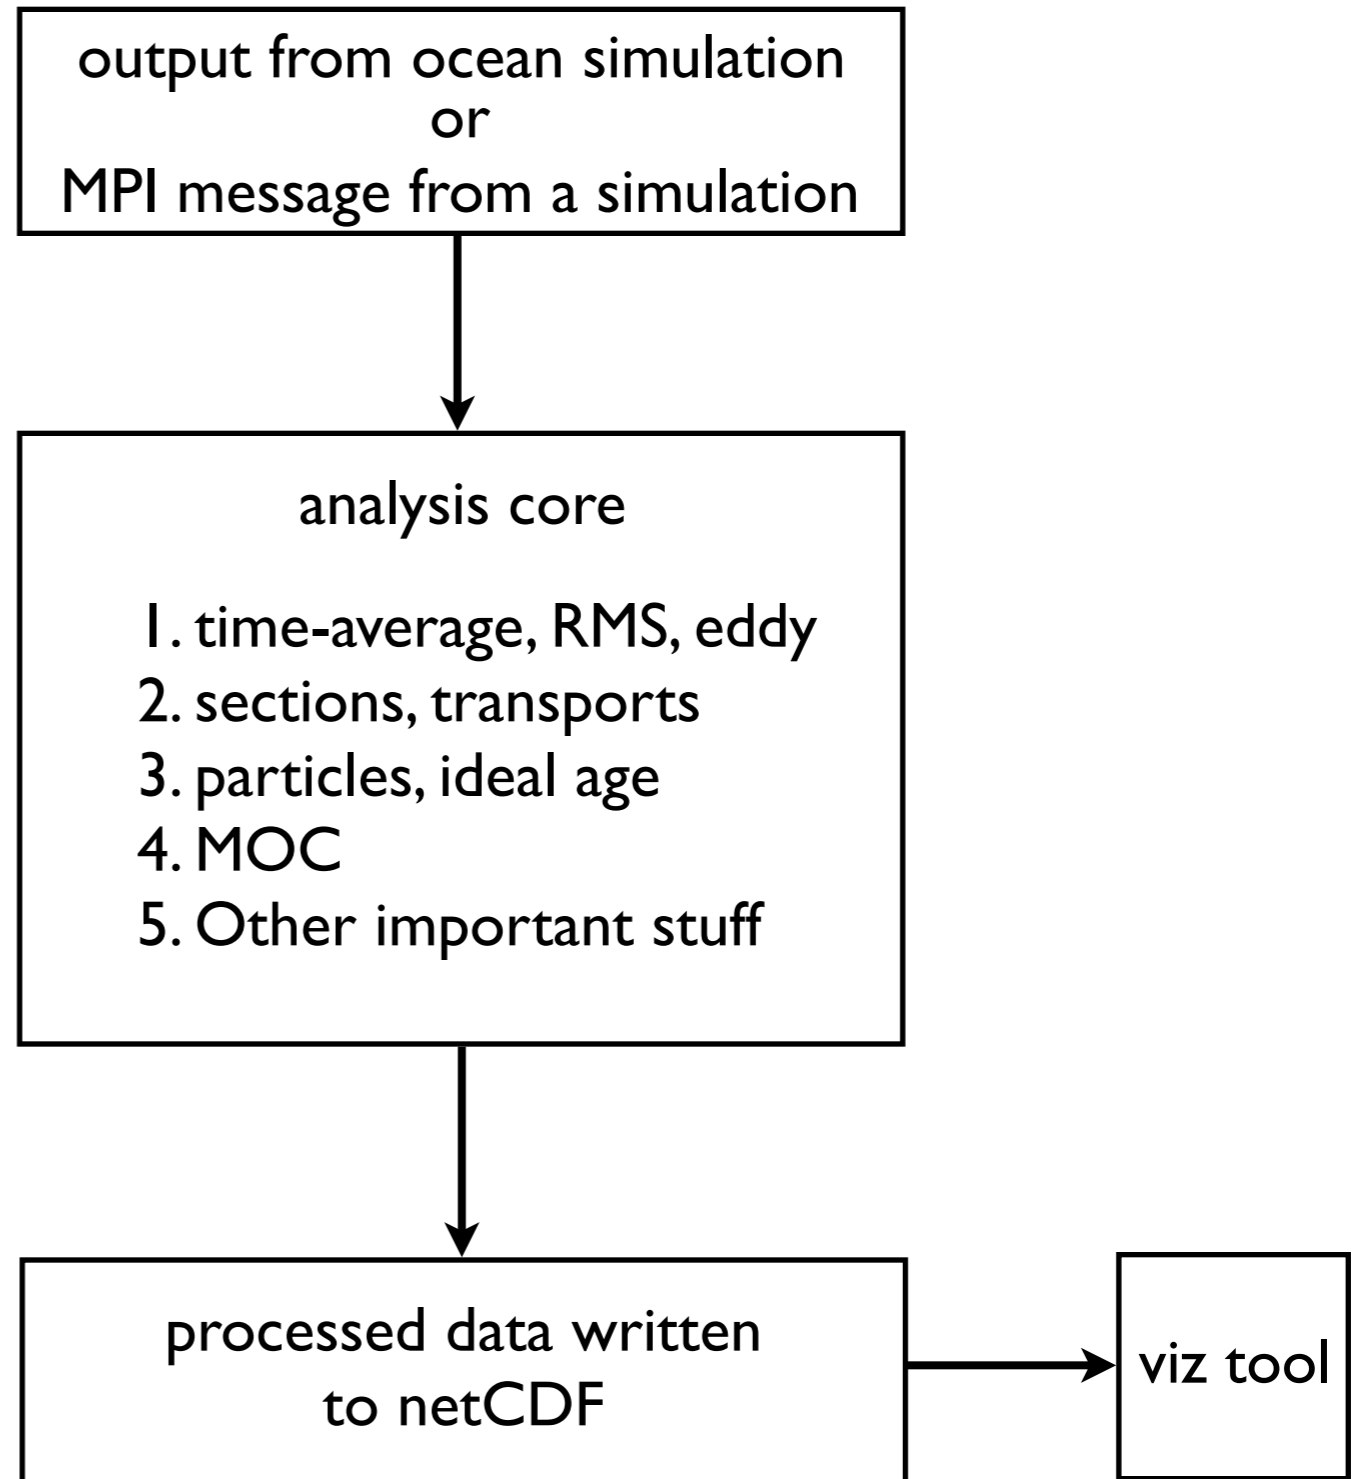
make ifort CORE=ocean_analysis

Scenario #1

Scientist prototypes analysis tools in desktop environment, reading in simulation data, writing FORTRAN code.

Scenario #2

Exact same source code developed in Scenario #1 is called during a simulation, thus reducing the amount of post processing.



make ifort CORE=ocean_analysis

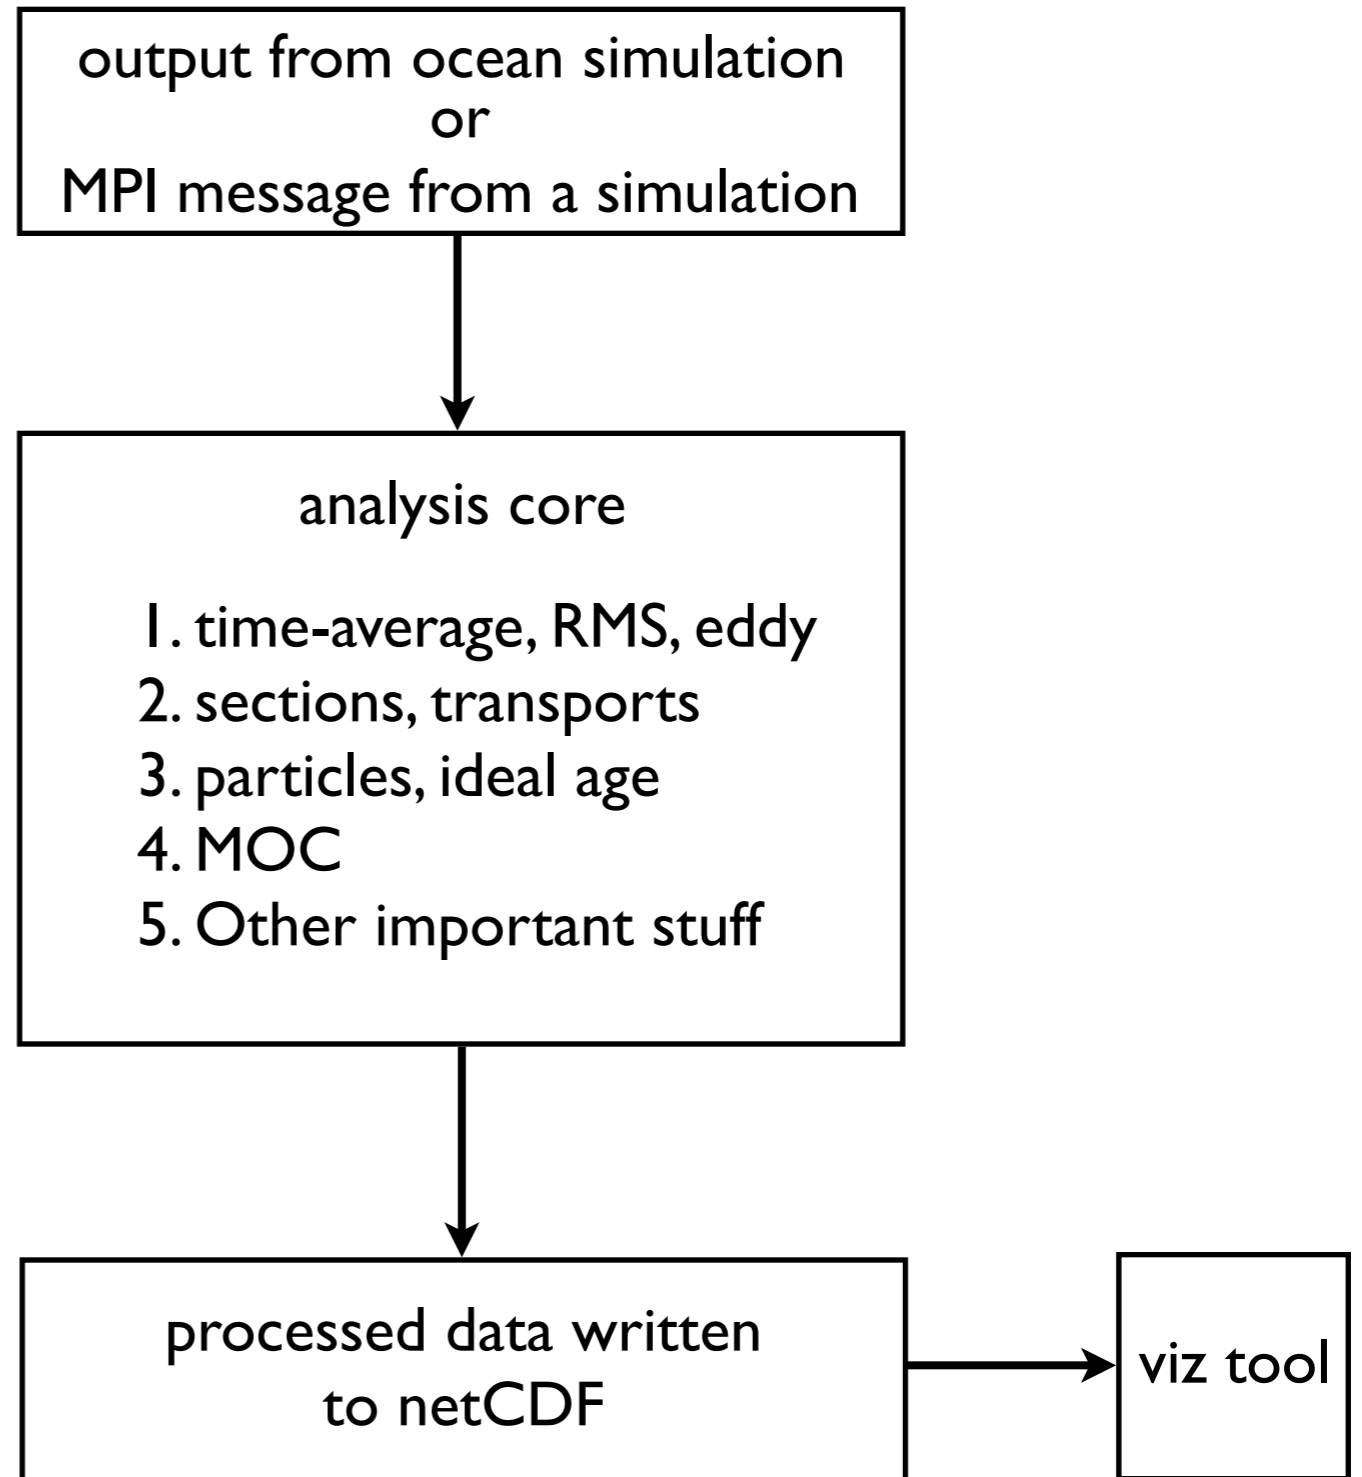
Scenario #1

Scientist prototypes analysis tools in desktop environment, reading in simulation data, writing FORTRAN code.

Scenario #2

Exact same source code developed in Scenario #1 is called during a simulation, thus reducing the amount of post processing.

By putting the analysis and the forward model into the same framework, we increase our chances of 1) getting modelers to analysis and 2) get analysts to do modeling.



make ifort CORE=ocean_analysis

Scenario #1


Scientist prototypes analysis tools in desktop environment, reading in simulation data, writing FORTRAN code.

Scenario #2

Exact same source code developed in Scenario #1 is called during a simulation, thus reducing the amount of post processing.

By putting the analysis and the forward model into the same framework, we increase our chances of 1) getting modelers to analysis and 2) get analysts to do modeling.

We need help here!



Thanks!

Discussion Items:

1. How do we entrain OMWG / Ocean Section into this process?
2. Target horizontal/vertical resolution for CORE and CESM simulations?
3. Thoughts on analysis framework?
4. BGC (can we abstract the RHS sources to a separate model?)
5. PIO, it is required to build MPAS components.
6. wrt CESM coupling, how backward compatible do we want to be?
7. Relative maturity of MPAS-CICE and MPAS-Ocean
8. Other