DART Perspectives on Infrastructure Challenges of High Resolution and Coupled Data Assimilation

Kevin Raeder (NCAR:CAM+DART development)

Contributors:

Alicia Karspeck (NCAR:CESM+DART,POP+DART)

NCAR:DART software:

Jonathan Hendricks, Helen Kershaw

Tim Hoar, Nancy Collins,

Jeffrey Anderson (NCAR: DART PI)

Data Assimilation Research Testbed

- Ensemble Kalman Filter

- Research
    - Dozens of model and observation interfaces
    - Good parallel performance on many machines (not tuned like operational)
    - Education via tutorials

- Software Engineering Focus
    - Standardized, small interface to any sized model
    - Flexibility; choice of model, any relevant observations, filter
    - Tools which make EnKF work well with large models, ‘small’ ensembles.
    - Make DA accessible to non-DA experts

CESM + DART:
NCAR's current framework for global assimilation applications

CESM+DART

- Community Earth System Model (multi-component earth system model designed for climate studies)
- Ensembles of CESM run in any supported configuration (i.e. coupled-model, high-resolution, etc) can interface with the DART system.

Relevant to this workshop → there is active work at NCAR on resolving challenges associated with

- coupled ocean-atmosphere assimilation
- high resolution 1/10° global ocean assimilation (eddy resolving)

Challenges associated with CESM global assimilation (and large-state space models in general)

1) **Data motion**
   - within DART (internal)
   - Between CESM and DART
   - Between DART and the file system (DA diagnostics)
2) **Memory and Parallelism**
   - How best to distribute the state and ensemble information?
3) **Model Design**:
   - The CESM architecture was not designed with ensemble DA in mind
4) **System Robustness:**
   - more processors leads to more machine failures.
5) **Storage and access:**
   - huge output data sets.

Resources and funding are needed to address these challenges

Status of coupled assimilation with CESM+DART

- Multi-component assimilation (global ocean/atmosphere at nominal 1° resolution) has been operable for 2+ years.
  - Quality of reanalysis is very promising, however
  - Infrastructure work is needed to make the system viable (faster and more parallelizable)

- Cross-component assimilation is on the horizon

  Open questions:
  - how to localize across components?
  - Is there a detectible benefit in analysis quality over multi-component coupled assimilation

- To date, this work has largely been accomplished with non-base funds:
  - NSF EASM
  - NSF+DOE supercomputing
  - Funding beyond 2016 unclear

- Recent allocation of NCAR base funds towards making CESM interact more efficiently with DART

Our Terminology:
"Coupled": *model* is ocean +atmosphere+…
"Multi-component DA": DART can assimilate 'native' observations into *each* component.
"Cross-component" = any observation can influence any model variable

Significantly more software engineering is needed to adapt large, complex state vectors in DART to existing computing environments.

Assimilating with the 1/10° ocean model: challenges associated with large state vector*

- State vector (T,U,…) size
  - > largest 32 bit integer: convert integers to 64 bit
  - = 17 Gb (2/3 of a Yellowstone node)

- New "RMA" capabilities allow each state vector to be distributed across multiple nodes, but requires:
  - internode communication for many observation forward operators
  - book-keeping of which node has which parts of the state vector

- On the horizon: distribute static memory across nodes

*Funded through NSF EASM3

Assimilating with the 1/10° ocean model: challenges associated with large state vector

About the new "RMA" version of DART:

- Eliminates the need to hold the entire state on any one processor, making it possible to assimilate to models with otherwise prohibitively large state-spaces
- Forward observation operators can use state variables that are distributed in memory essentially randomly.
  - The state is distributed across all processors.
  - Using MPI one-sided communication a task can get the state elements it needs from any other process. This happens asynchronously.
  - *Asynchronously is the key point – you don't have to calculate which processors need what from where and when they need it*

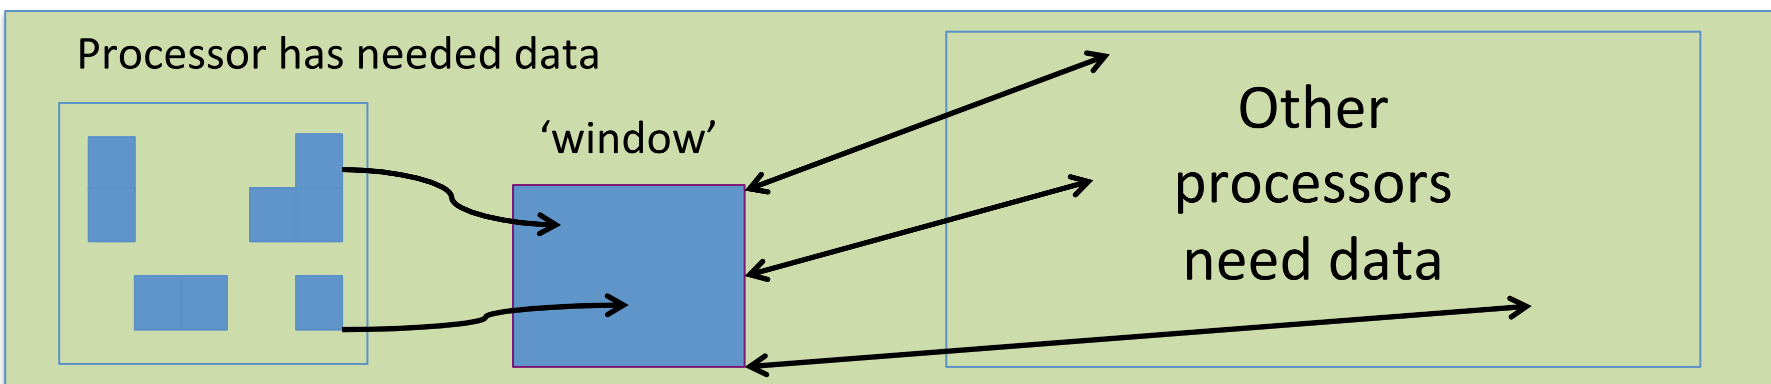

Assimilating with the 1/10° ocean model: challenges associated with intensive I/O

- 30 members -> 516 Gb of state vectors read and written each assimilation.
- 4 Gb/s for read + movement to where it's needed.
- I/O is a significant fraction of the assimilation time, and it doesn't scale well.

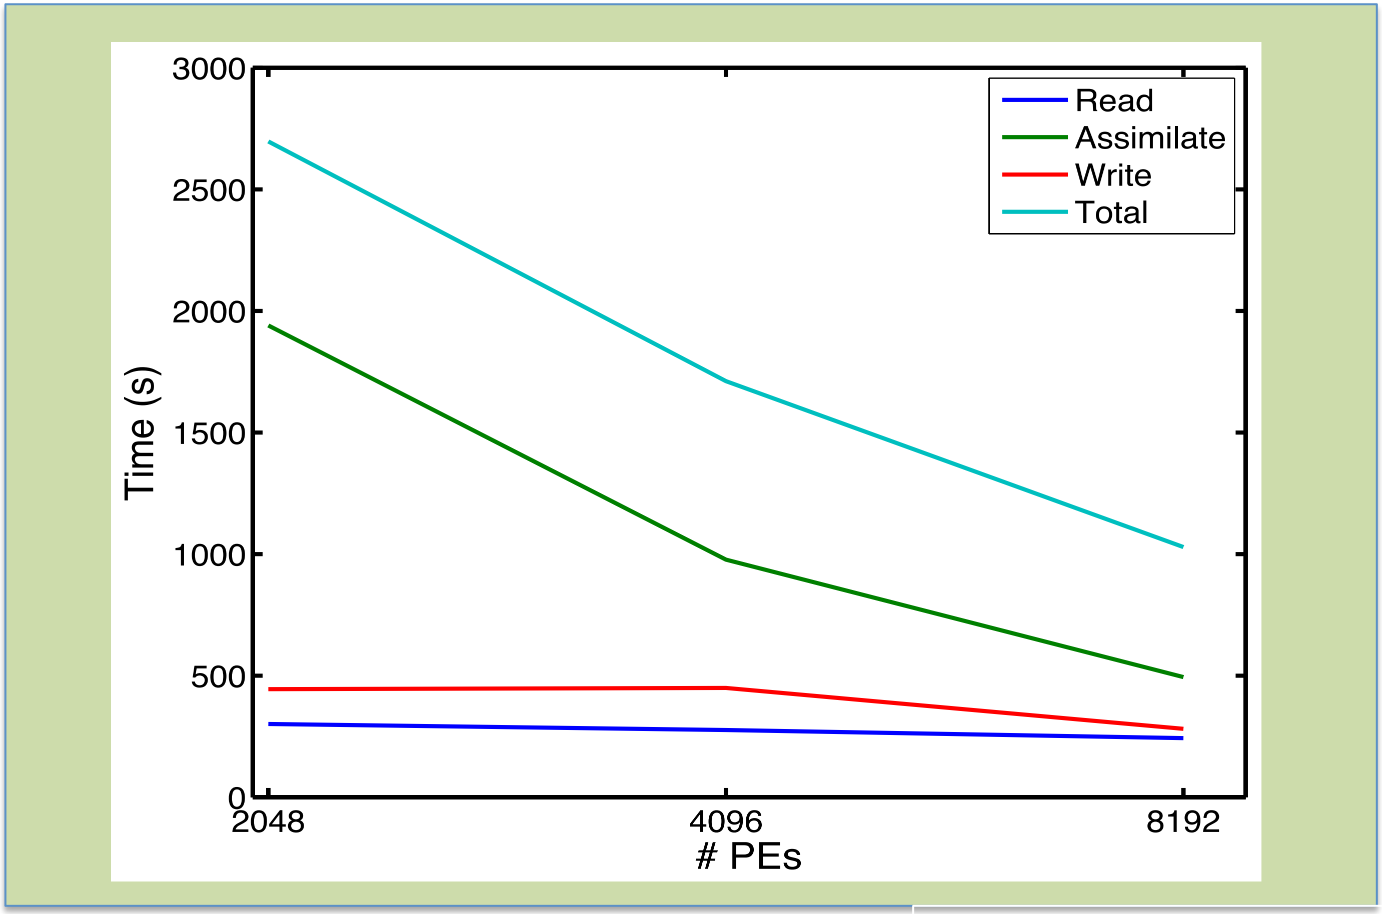

Assimilating with the 1/10° ocean model: challenges associated with intensive I/O

Potential accelerations given current I/0 limitations

- Read in only subset of state values which will be effected by assimilation.
- POP2: Run ensembles with a 1° model to constrain the planetary scales and use a stationary ensemble OI to constrain the mesoscale variability.

Rethinking how CESM generates ensembles and how CESM interfaces with DART

New ideas and research horizons

- Make CESM more parallel
- Make CESM more resilient to machine failures
- Create a true ensemble coupled model (?)
  - Separate model members are replaced by a new dimension ("member") in each state variable. A given computation would be performed for all members at the same time. That would move larger chunks of data fewer times.
- Compute forward operators within the model, rather than afterwards in the DA.
- Pass state between model and DA via memory, rather than files.

Summary and Conclusions

- For 'high' resolution, significant software infrastructure development in CESM and DART is needed to use even currently available computing resources effectively.

- The trends of increasing processor count, decreasing shared memory/node, and increasing data set sizes will require ongoing software infrastructure development.

- Large model development in CESM+DART benefits greatly from true software engineers, rather than (talented) DA scientists.