Asynchronous Many-Task Programming Models for the Earth System: Couplers and Ocean/Ice

Phil Jones, Irina Demeshko, Jonathan Pietarila-Graham, David Ringo (LANL)

CANGA

Motivation

Programming Models for Complex Earth System Models

Earth System models like the Energy Exascale Earth System Model (E3SM) have typically utilized standard domain decomposition approaches with an MPI+OpenMP programming model. Unfortunately, the E3SM and other models have difficulty exposing enough parallelism In this approach and load-balancing the system is difficult. The Coupling Approaches for Next-Generation Architectures (CANGA) project is working to exploit a number of advantages of task-parallel programming models.

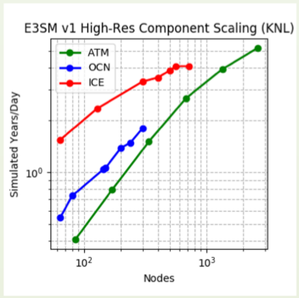

Computational Advantages

- Exploit additional parallelism
- Improve load balancing
- Fault tolerance
- Map tasks to hardware elements

Scientific Advantages

- Better manage complexity
- Improved extensibility (add tasks)
- Ability to couple at process level, not component
- Manage multiple time, space scales

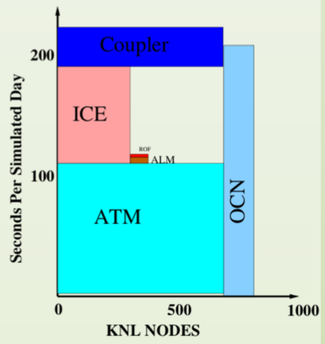

Approach

Create a prototype coupled E3SM model in a task-parallel paradigm

FleCSI and Legion:

Flexible Computation Science Infrastructure (LANL)

- C++ framework
- Supports tasking runtimes (Legion, HPX, CHARM++)
- Provides control, execution and data models

Legion (Stanford, LANL, NVidia)

- Based on logical regions (intersection index/field spaces)
- Task-based run-time with efficient DAG generation
- Apparently sequential semantics

FleCSI

https://github.com/laristra/flecsi/

http://legion.stanford.edu/

Legion architecture

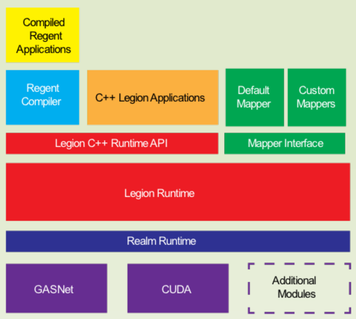

Coupler Level (Top-down):

- Create prototype coupler
- Replicate existing component coupling
- Push downward to expose more processes

Component Level (Bottom Up):

- Create task-parallel versions of ocn, ice, land
- Break up further into finer scale tasks
- Generate task hierarchy for coupler layer

Progress

Prototype Coupler:

- Created Regent/Legion prototype w/ stub models
- Demonstrated dependency graph and parallel execution
- FleCSI prototype reading data models in progress

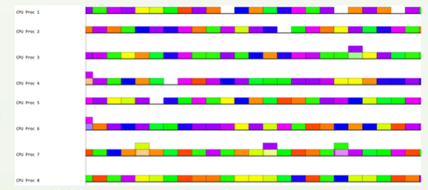

Parallel (8-core) execution of prototype coupler using Regent/Legion. Colors represent tasks.

Prototype Ocean Model:

- Completed Regent prototype with all existing ocean function interfaces
- Generated initial dependency graph based on interfaces
- FleCSI specialization layer for MPAS-Ocean framework and data types in progress

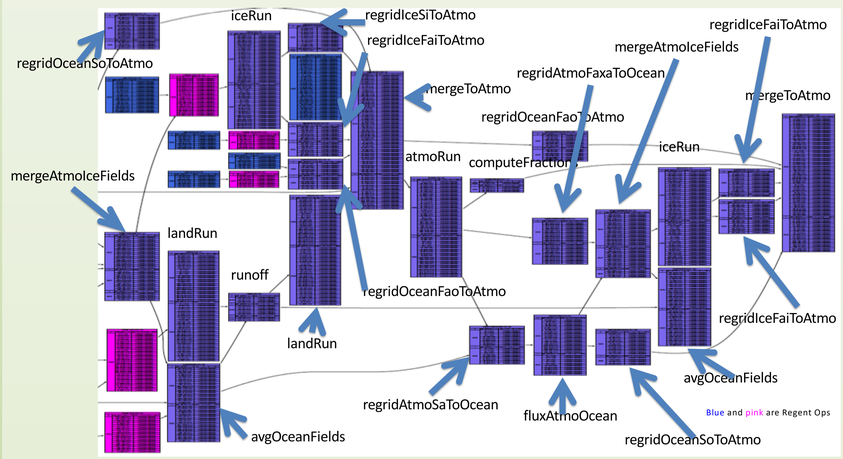

Data flow and dependencies in current Regent prototype coupled system meant to replicate current coupling at component level. Once task-parallel components are developed (see below), further parallelism can be exposed.

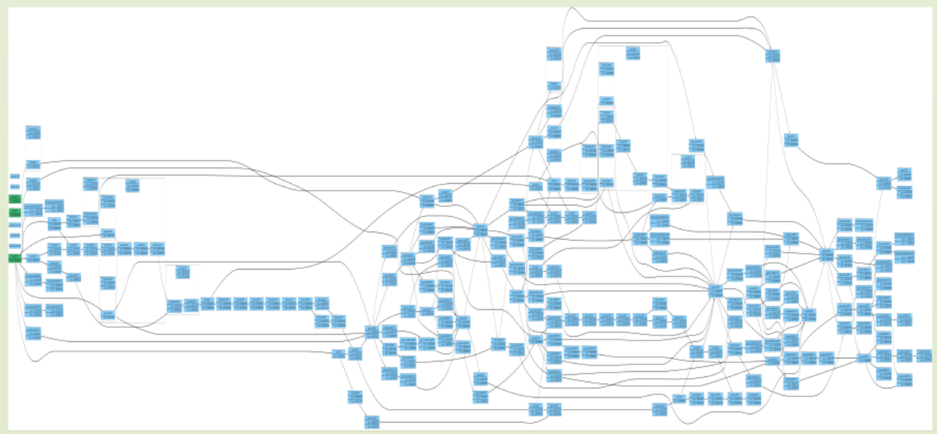

Ocean distributed acyclic graph (DAG) generated by Legion based on current ocean interfaces, showing both existing task parallelism and serial dependencies that must be further broken down to achieve independent execution.

Contact Info:

Phil Jones, Los Alamos National Laboratory, pwjones@lanl.gov
CANGA: http://canga-scidac.org
E3SM: http://e3sm.org

SciDAC Scientific Discovery through Advanced Computing
Los Alamos National Laboratory EST. 1943
Sandia National Laboratories
Argonne National Laboratory
Oak Ridge National Laboratory
UC Davis University of California
Stony Brook University
The Ohio State University
U.S. Department of Energy | Office of Science