ACME
Accelerated Climate Modeling for Energy

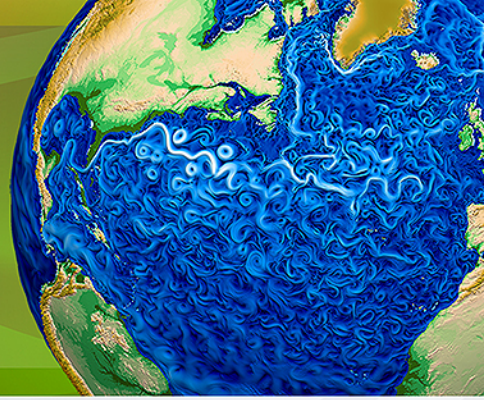

ACME: Status and Outlook on New Architectures

Mark Taylor

ACME All Hands Meeting

June 5, 2017

Bolger Center

U.S. DEPARTMENT OF ENERGY

ACME v1-v3
Target Resolutions

ACME v1

- “Low-Resolution”: 110 km atm/lnd, 60-30 km ocn/ice
  - Workhorse configuration - model development, CMIP/DECK type science campaigns, O(1000) years of simulation
- High-Resolution: 27 km atm/lnd, 18-6 km ocn/ice
  - ACME v1: baseline simulations O(100) years
  - ACME v2: Climate science campaigns on pre-exascale or exascale systems 2019-2023

ACME v2/v3+

- RRM (Regionally Refined Model)
  - ACME v1: All components capable of running on RRM meshes
  - ACME v2: Affordable high resolution configurations for CONUS, Arctic
- Ultra-high resolution:
  - 1 km: Cloud resolving, Coastal modeling/inundation/ice shelves, Arctic embayments
  - 100m resolution (LES regime, boundary layer mixing, sub-watershed resolution)
- ACME-MMF: Global cloud resolving at 5 SYPD
  - Via superparameterization and GPU acceleration

ACME Computational Resources

ANVIL: ACME Dedicated Resource

- 120 nodes (soon 240 nodes!)
  - Intel Xeon, 36 cores/node
  - 80M core hours/year
  - Fastest per-node performance and best core-hours per simulated year
- ACME v1 low-res:
  - 7 SYPD, 20K core-hours per simulated year
- Compare to NERSC Edison
  - Similar (but older) Intel Xeon architecture
  - 2x more expensive due to charge factor
  - 2017 NERSC usage: 30M out of 120M allocation

ANVIL: ACME Dedicated Resource

- Exascale version of Anvil?
  - Too much power, uncertain architecture roadmap
- ACME Mission:
  - Develop an Earth system modeling capability for next-generation computing architectures

ACME v1 on Today's DOE LCF's O(10) petaflop

OLCF Titan

- 19K nodes
  - 16 core AMD + NVIDIA GPU
- Good machine for ACME:
  - ACME v0 high-res: 2 SYPD, 1.5M/year
  - ACME v1 high-res: 0.5 SYPD, 2.5M/year
  - We've used 50M core-hours used under ALCC, 70M used under INCITE. 70M remaining.
- ACME v1 high resolution
  - 15% of our code (and growing) can make use of the GPU. Insufficient GPU utilization to be competitive in INCITE.
  - Exception: ACME-MMF
  - Looking for more performance staff

ALCF Mira

- Mira: 49K nodes (16 core BG/P)
- ACME v0 high-res: 1 SYPD, 0.7M/year
- ACME v1 high-res: .4 SYPD, 8M/year
- Great machine for ACME v0 high-res
  - 127 year pre-industrial control
  - 6x40 present day ensembles
  - 240M core-hours available in CY17
- ACME v1 - too expensive
  - Performance work needed to fix this – but end-of-life machine so focusing on KNL architecture higher priority

KNL: Cori and Theta

- Intel KNL (68 cores per node)
- NERSC Cori-KNL 9145 nodes
- ALCF Theta: 3624 nodes
- Most promising architecture for ACME v1 high-res:
  - 0.9 SYPD, 2.1M/year
  - Some technical issues remain. Should be able to achieve 1.3 SYPD
  - Should be able to perform O(100) years worth of simulation in CY2017.

KNL: Cori and Theta

- KNL Architecture:
  - Coupled model currently runs slower on KNL than on Xeon
  - Components which vectorize well can run faster (e.g. atm dycore).
  - Need staff to work on vectorizing
- Charge factor at NERSC impacts cost. ACME v1 low-res:
  - 20K/year Anvil
  - 40K/year Edison
  - 100K/year Cori-KNL (for now…)

DOE Pre-Exascale Machines

ACME v1 on Aurora

- ALCF Aurora 2019
  - 50K Nodes, 3rd gen Intel Phi
- Larger system & faster nodes as compared to Cori-KNL and Theta
- Promising machine for ACME v1/v2 high-resolution model

ACME v1 on Summit

- OLCF Summit 2019
  - 3400 Nodes
  - Multiple IBM power9 and NVIDIA GPUs
- Harder to estimate performance
- Potential for larger performance gains
- More labor required to port code
- Most GPU work currently focused on ACME-MMF

Thanks!