EarthInsights: Parallel Clustering of Large Earth Science Datasets on the Summit Supercomputer

Sarat Sreepathi[1], Jitendra Kumar[1], Forrest M. Hoffman[1], Richard T. Mills[2], Vamsi Sripathi[3], William W. Hargrove[4]

[1]Oak Ridge National Laboratory
[2]Argonne National Laboratory
[3]Intel Corporation
[4]USDA Forest Service

1st Workshop on Leveraging Artificial Intelligence (AI) in the Exploitation of Satellite Earth Observations & Numerical Weather Prediction(NWP)
NOAA, College Station, MD
April 23-25, 2019

Motivation

- Rapid proliferation of data in various domain sciences
- Earth Science
  - Advanced sensors – high fidelity data
  - Remote Sensing Platforms
    - Satellites
    - Unmanned Aircraft Systems (UAS)
    - Airborne systems
  - Observational Facilities
- Critical need for High Performance Big Data Analytics

History : Scalable algorithms for Clustering

One of the first Beowulf clusters (ca '97)

Data Analytics/Machine Learning

Originally developed in 1996-97 for use on the Stone Soupercomputer, a very early Beowulf cluster (see "The Do-It-Yourself Supercomputer", Scientific American, 265 (2), pp. 72-79, 2001.)

Applications

- Vegetation mapping and characterization
- Development of ecoregions
- Species distribution
- ForWarn: Real-time Operational Forest health monitoring
- Global Fire Regimes
- Climate zone classification
- Understand climate regime changes in future
  - Under various predicted climate change scenarios

Some Datasets

Problem Config	Data Size	No. of Clusters
GSMNP LiDAR	900 MB	30
CMIP3 Climate States	8 GB	1000
Phenology 2000	25 GB	1000
CONUS Phenology 2000–2017	500 GB	1000

Preprocessing

- Standardized the data set along each dimension
- Fairly represent every dimension in the clustering algorithm for high-dimensional datasets.

Parallel k-means Outline

Goal: Divide observations into k clusters

- Centralized Master-Worker paradigm
- Pick initial centroids
- Iterative method
- Workers
  - Compute distances
  - Update centroids and cluster assignments
  - Repeat till convergence is achieved
- Typical target convergence: $< 0.5\%$ changes

BLAS Formulation (Application Phase 1)

Binomial expansion:

$$\mathbf{dist}_{i,j} = \|\mathbf{obs}_{i,*}\|^2 + \|\mathbf{cent}_{i,*}\|^2 - 2 \cdot \mathbf{obs}_{i,*} \cdot \mathbf{cent}_{j,*}$$

$$\mathbf{dist} = \overline{\mathbf{obs}} \cdot \mathbf{1}^T + \mathbf{1} \cdot \overline{\mathbf{cent}}^T - 2 \cdot \mathbf{obs} \cdot \mathbf{cent}^T$$

xGER

xGEMM

$$A := alpha * x * y' + A \qquad C := alpha * op(A) * op(B) + beta * C$$

BLAS Subroutines

Triangular acceleration (Application Phase 2)

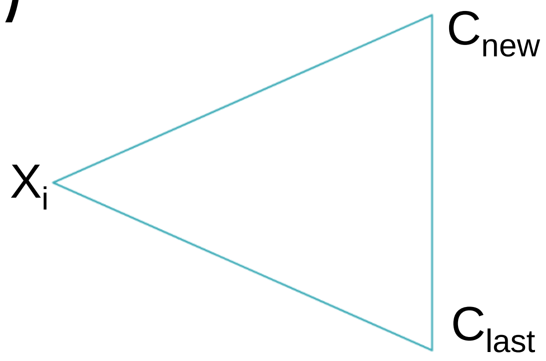

- Classical k-means performs far more distance calculations than required!
- Use the triangle inequality to eliminate unnecessary distance computations
- If $d(C_{last}, C_{new}) \geq 2d(X_i, C_{last})$
  $\Rightarrow d(X_i, C_{new}) \geq d(X_i, C_{last})$ without computing
- Distance computations can be further reduced by sorting the inter-centroid distances, $d(C_{last}, C_{new})$
  - Evaluate candidate centroids in sorted distance order
  - Skip computations once the critical distance, $2d(X_i, C_{last})$ is surpassed

HPC Platforms

Summit Architecture

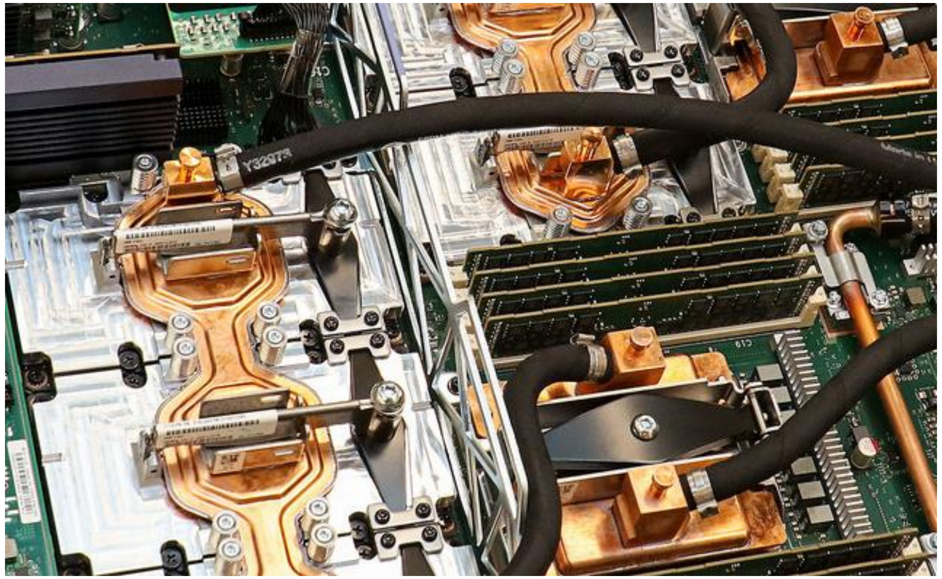

- ~200 PF (143 PF Linpack)
- 4608 Compute nodes

Each node: 42 TF

- Compute
  - 2 x Power 9 (22 cores)
    0.5 DP TF/s
  - 6 x Volta V100 GPU (80 SMs – 32 FP64 cores/SM)
    7.8 DP TF/s
- Memory/node
  - 512 GiB DDR4 memory
  - 96 (6x16) GiB High-bandwidth memory (GPU)
  - 1.6 TB NVMe

Titan

- Cray XK7 system
- Each node
  - 16-core AMD Opteron CPUs
  - NVIDIA Kepler K20X GPUs
  - 32 GB memory
- Total of 18,688 nodes
  - 299,008 CPU cores and 18,688 GPUs.

Software

- CPU (MKL + OpenMP)
- GPU (cuBLAS + OpenACC)
- MPI for communication

Performance and Scaling
Varying Number of Clusters (k)

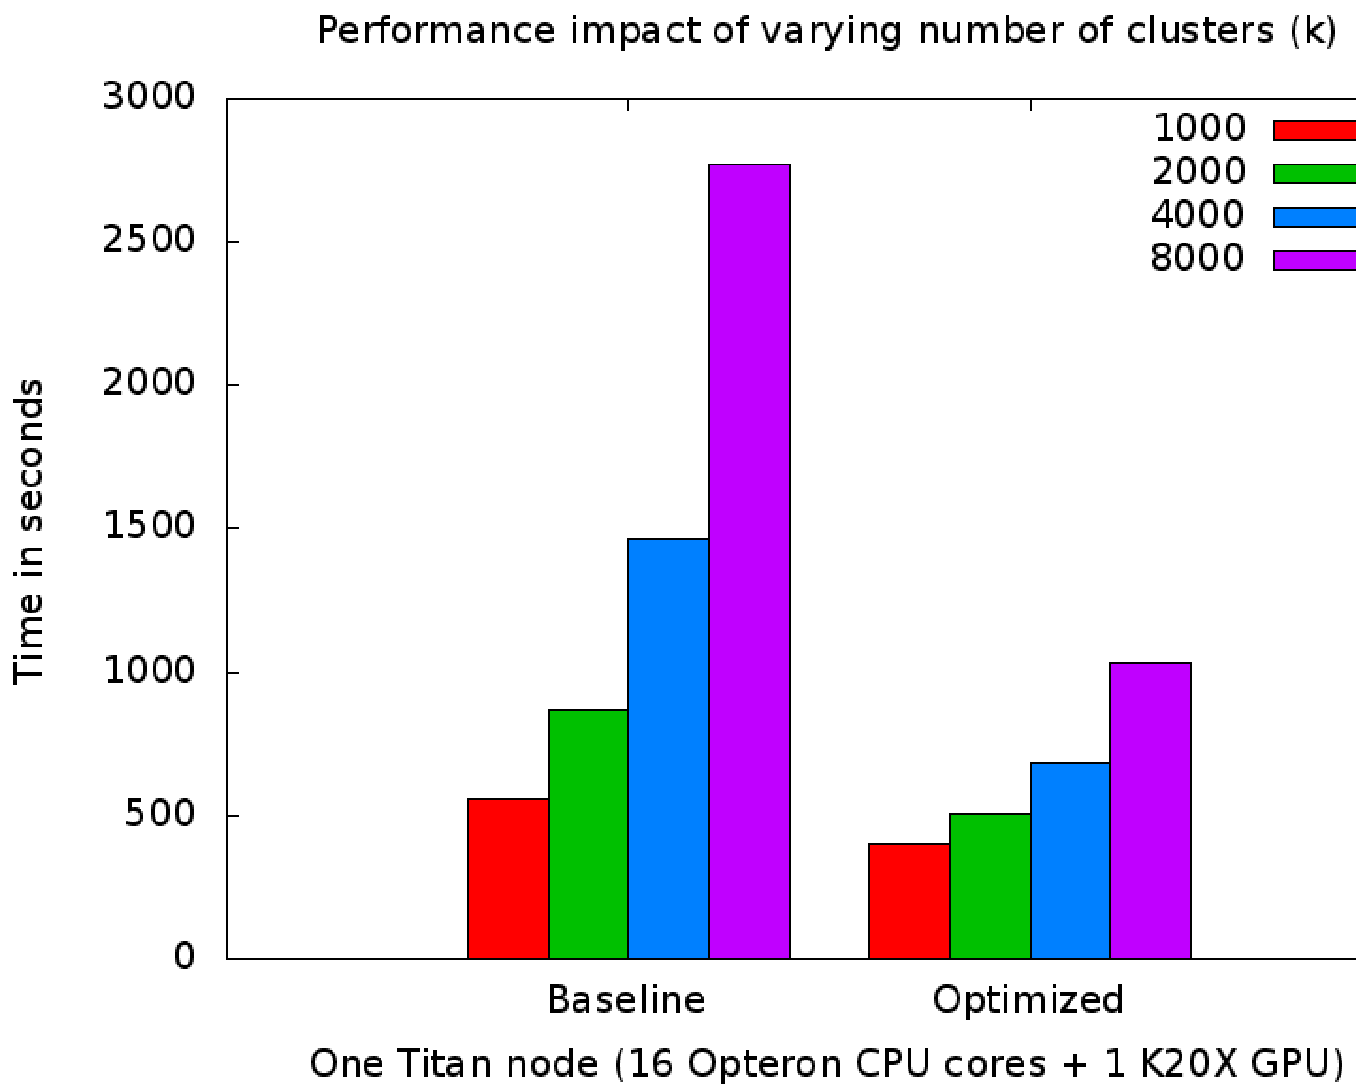

Summit: Early Results

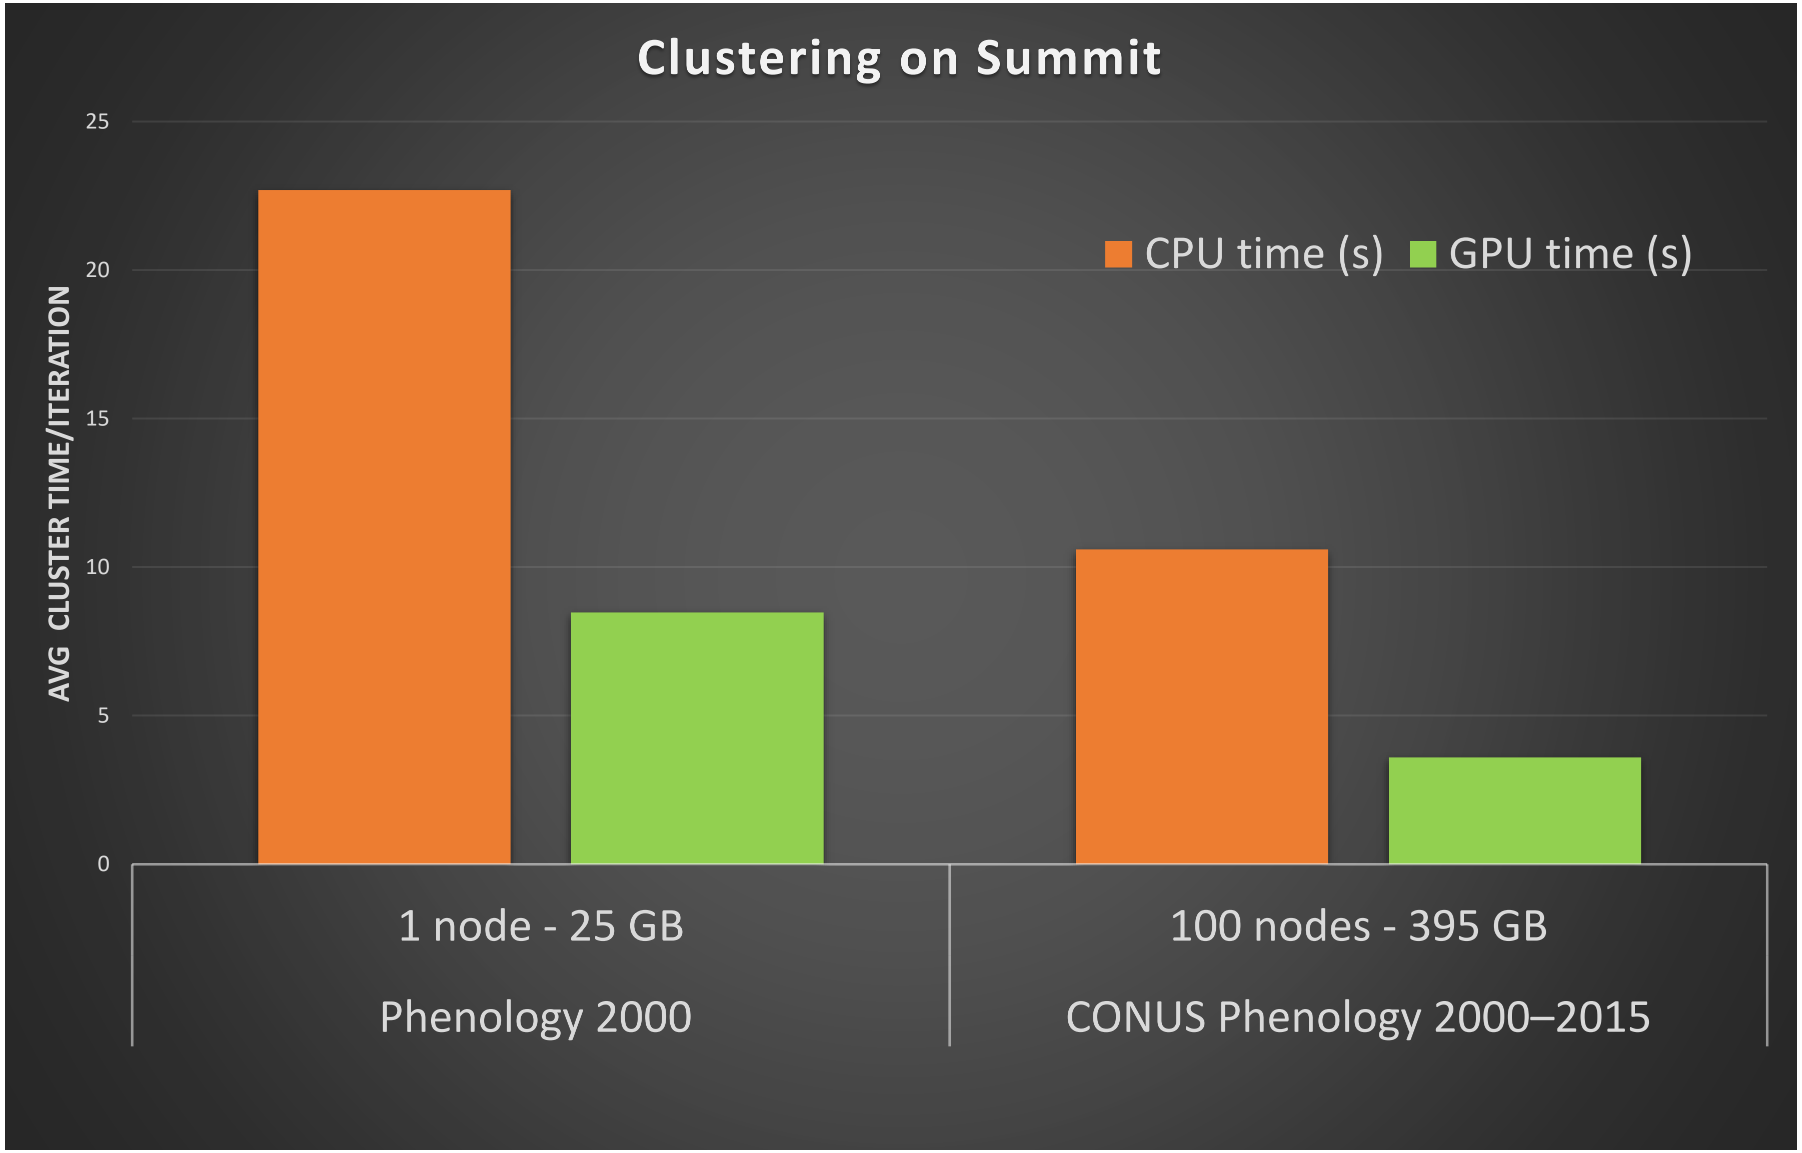

Applications

GSMNP: Spatial distribution of the 30 vegetation clusters across the national park

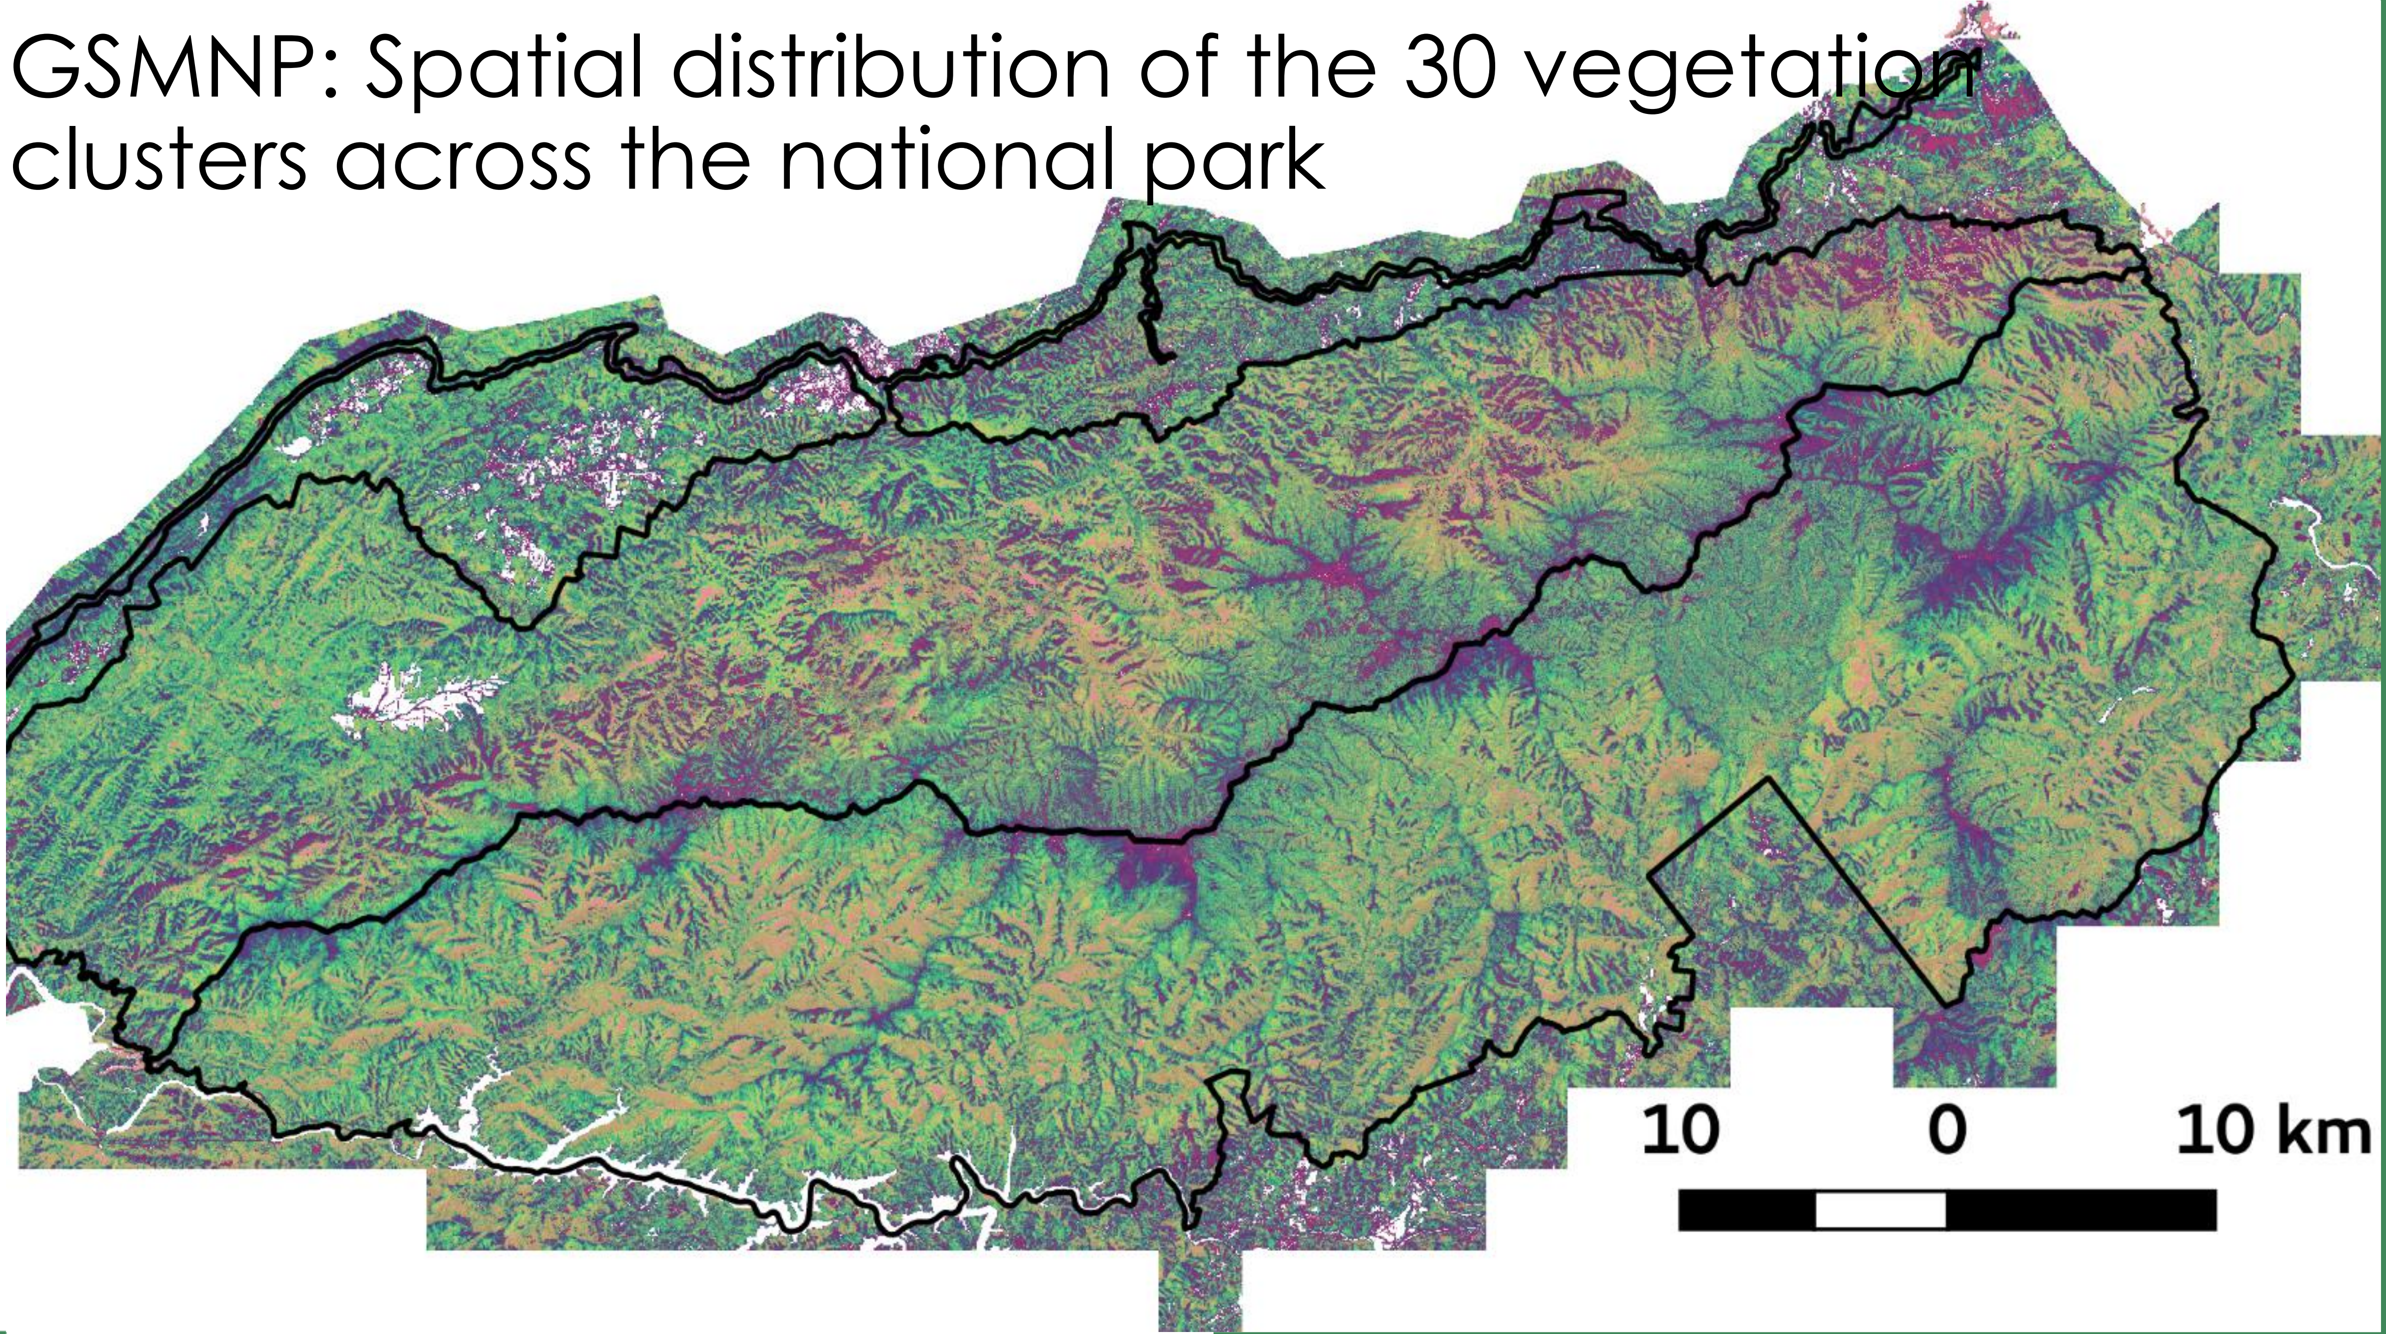

OAK RIDGE National Laboratory

GSMNP: 30 representative vertical structures (cluster centroids) identified

tall forests with low understory vegetation

forests with slightly lower mean height with dense understory vegetation

low height grasslands and heath balds that are small in area but distinct landscape type

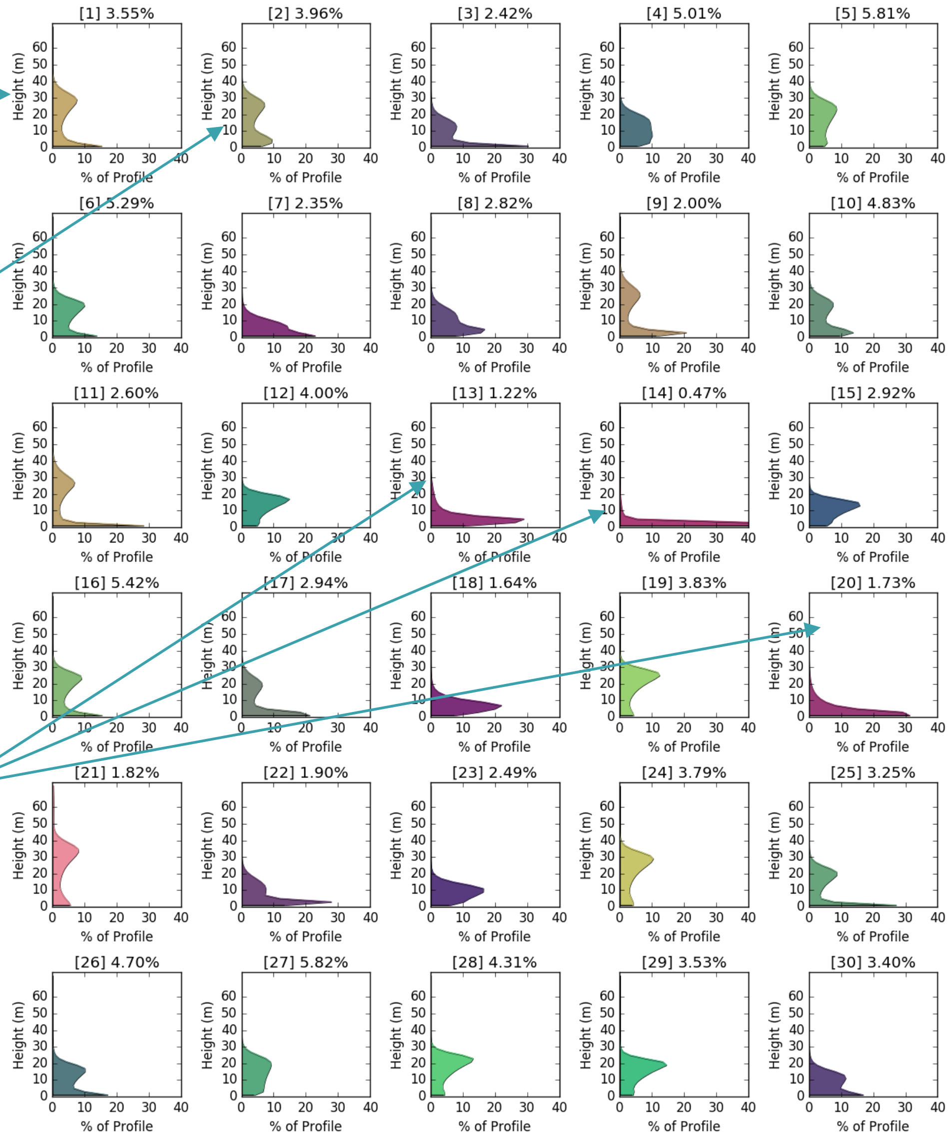

Global Climate Regimes: 1000 clusters Contemporary using Similarity color scheme

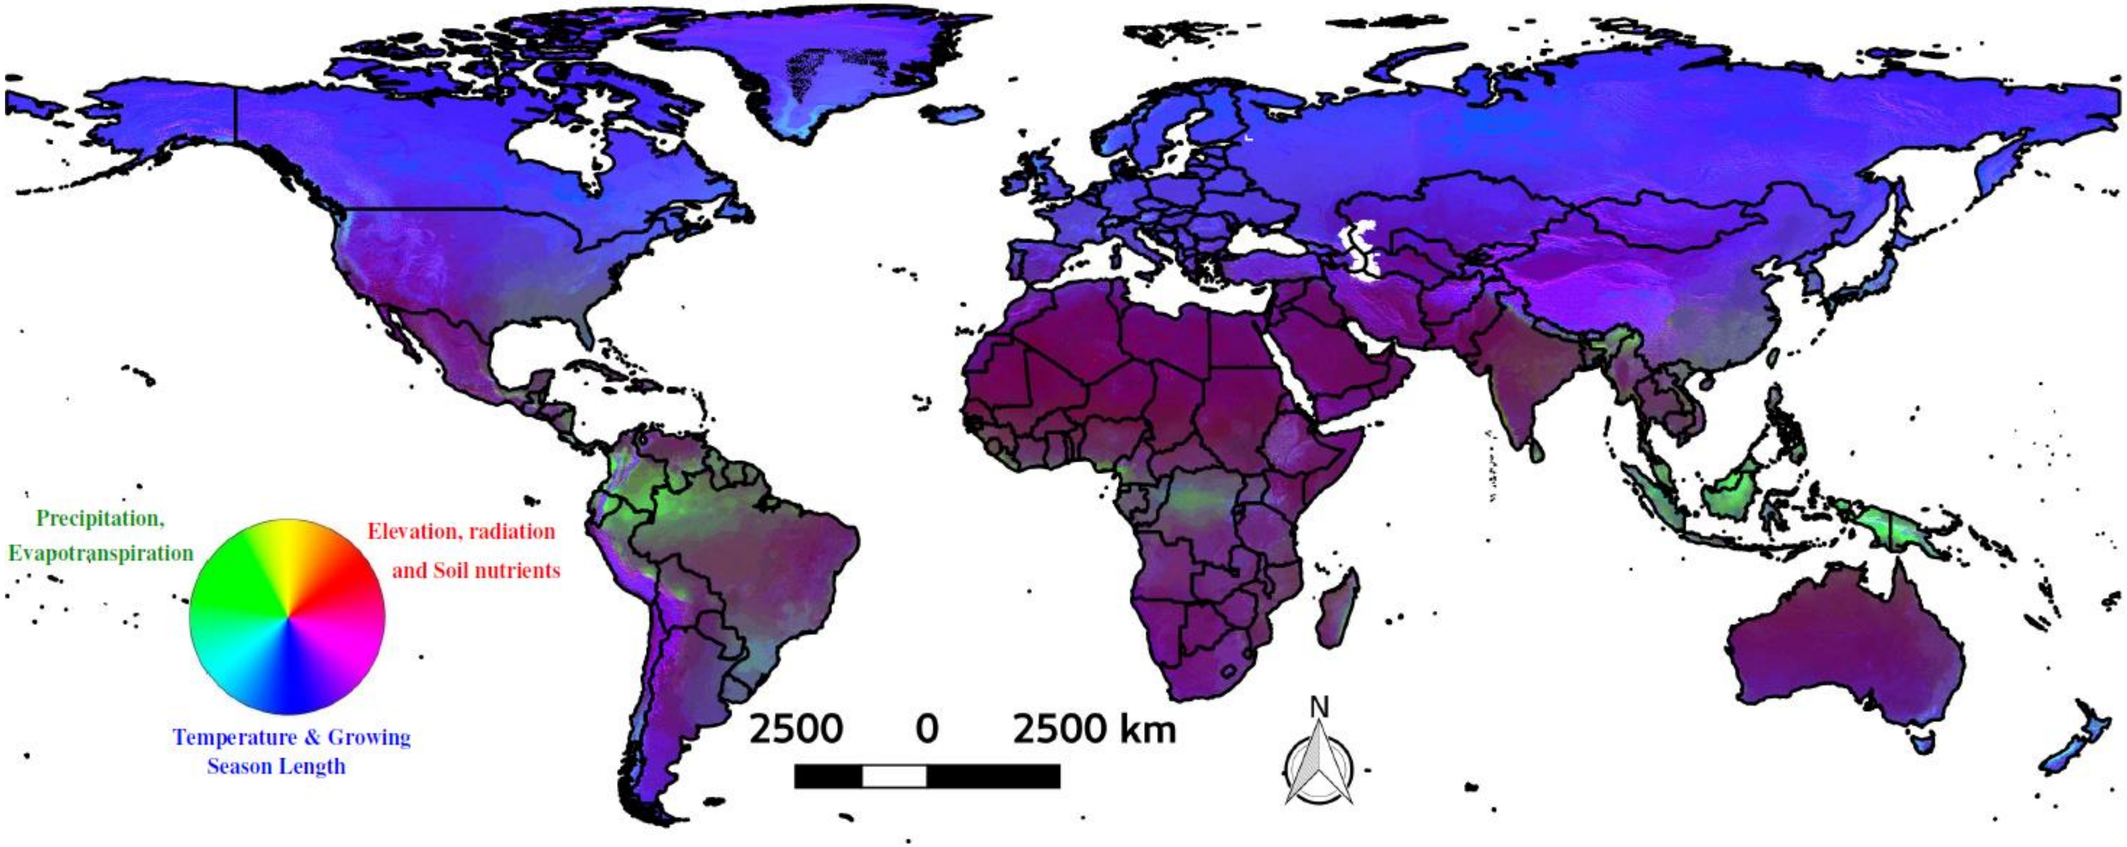

Global Climate Regimes: 1000 clusters 2100 using Similarity color scheme

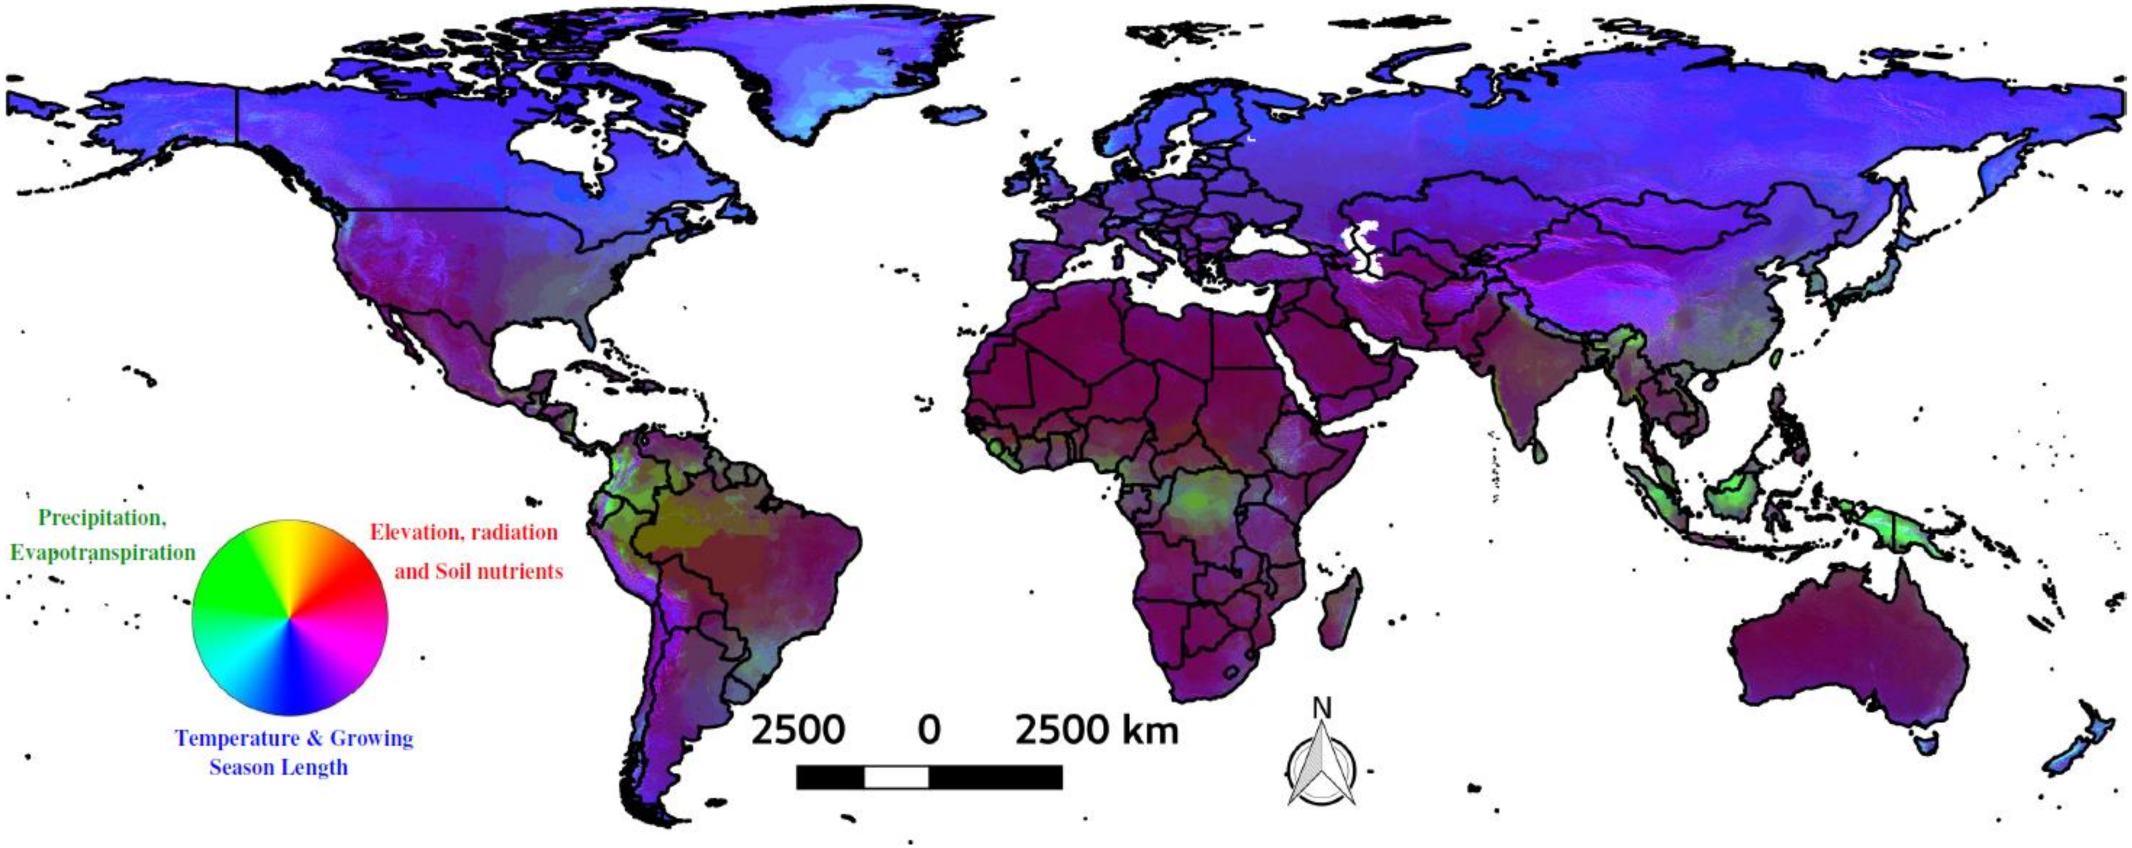

OAK RIDGE
National Laboratory

CONUS dynamic phenoregions

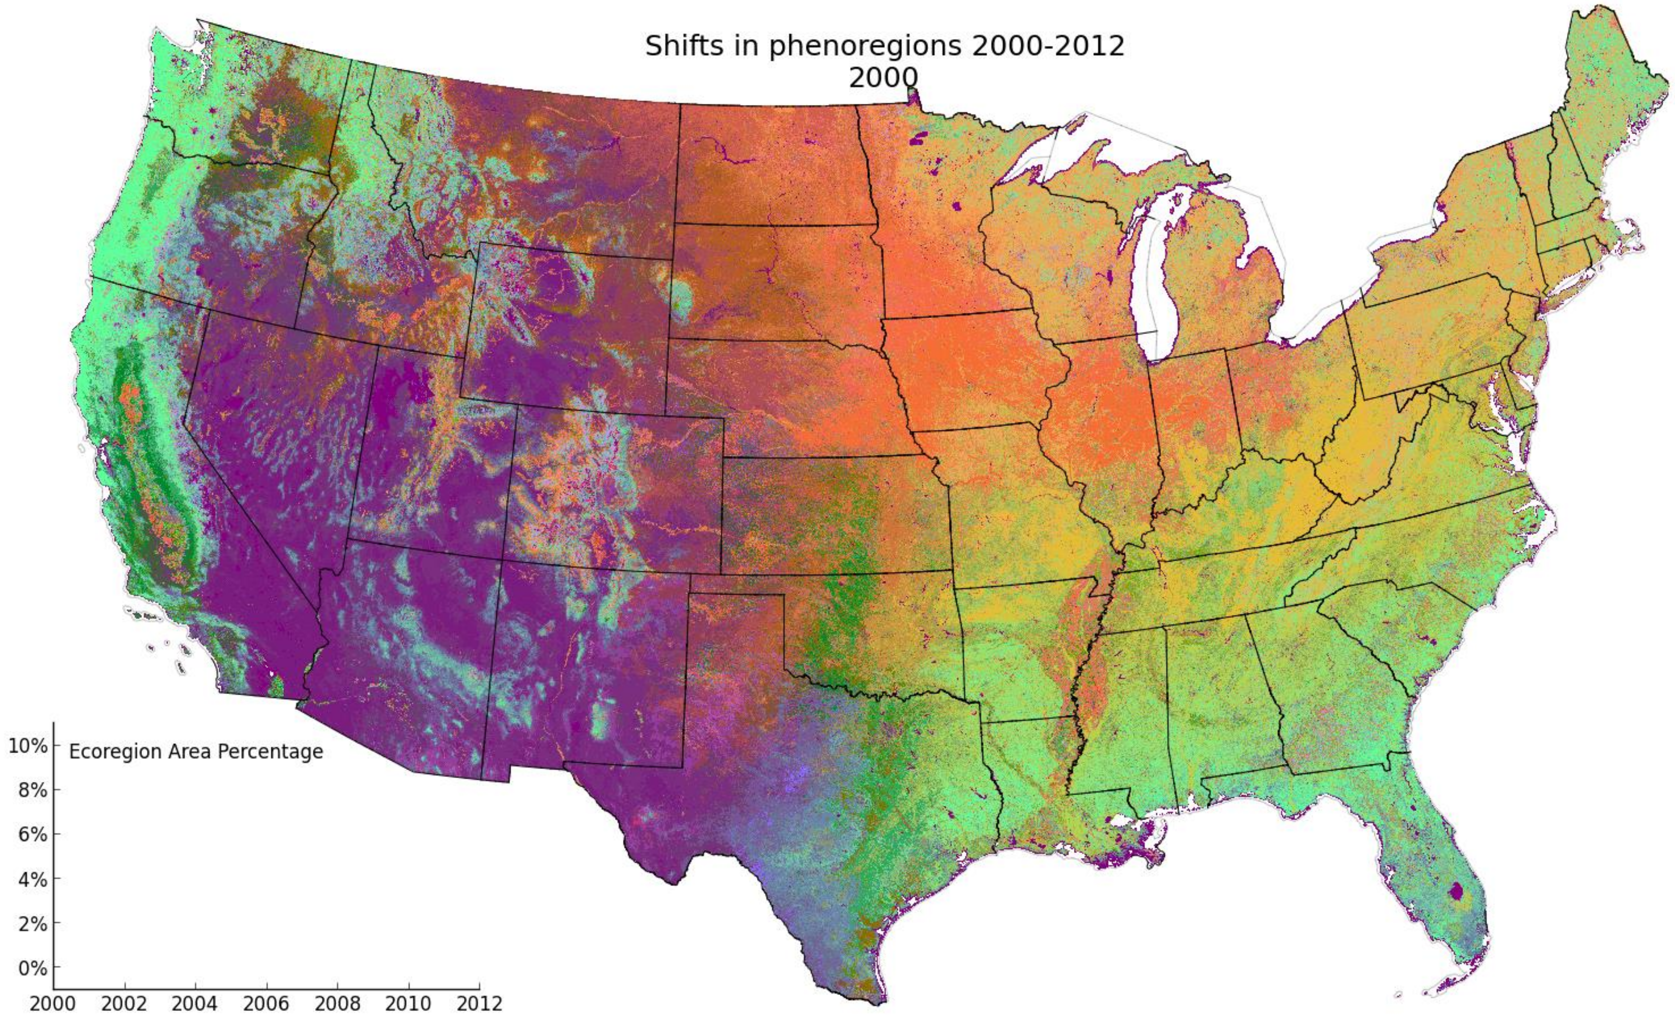

Global Fire Regimes

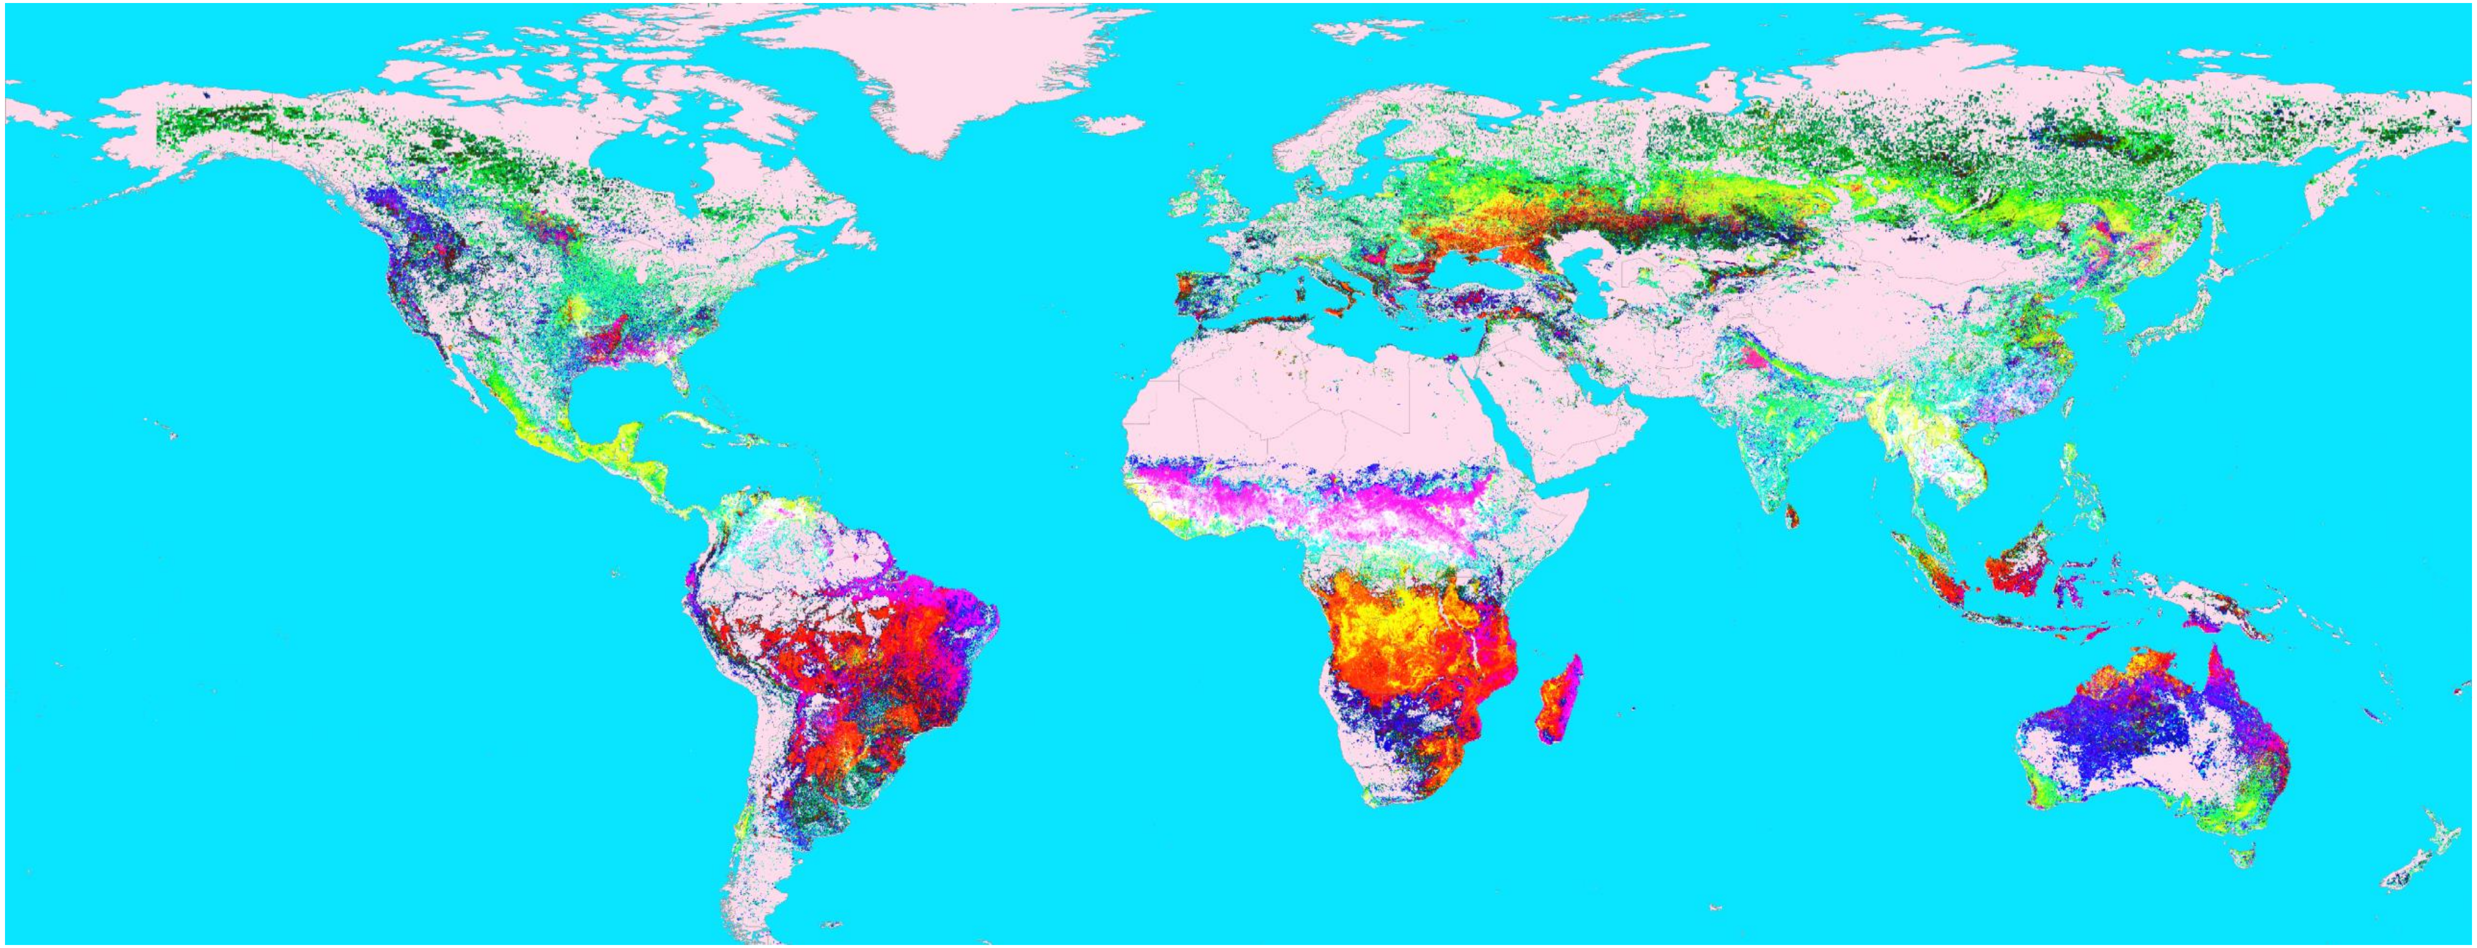

Arctic: High-resolution vegetation mapping

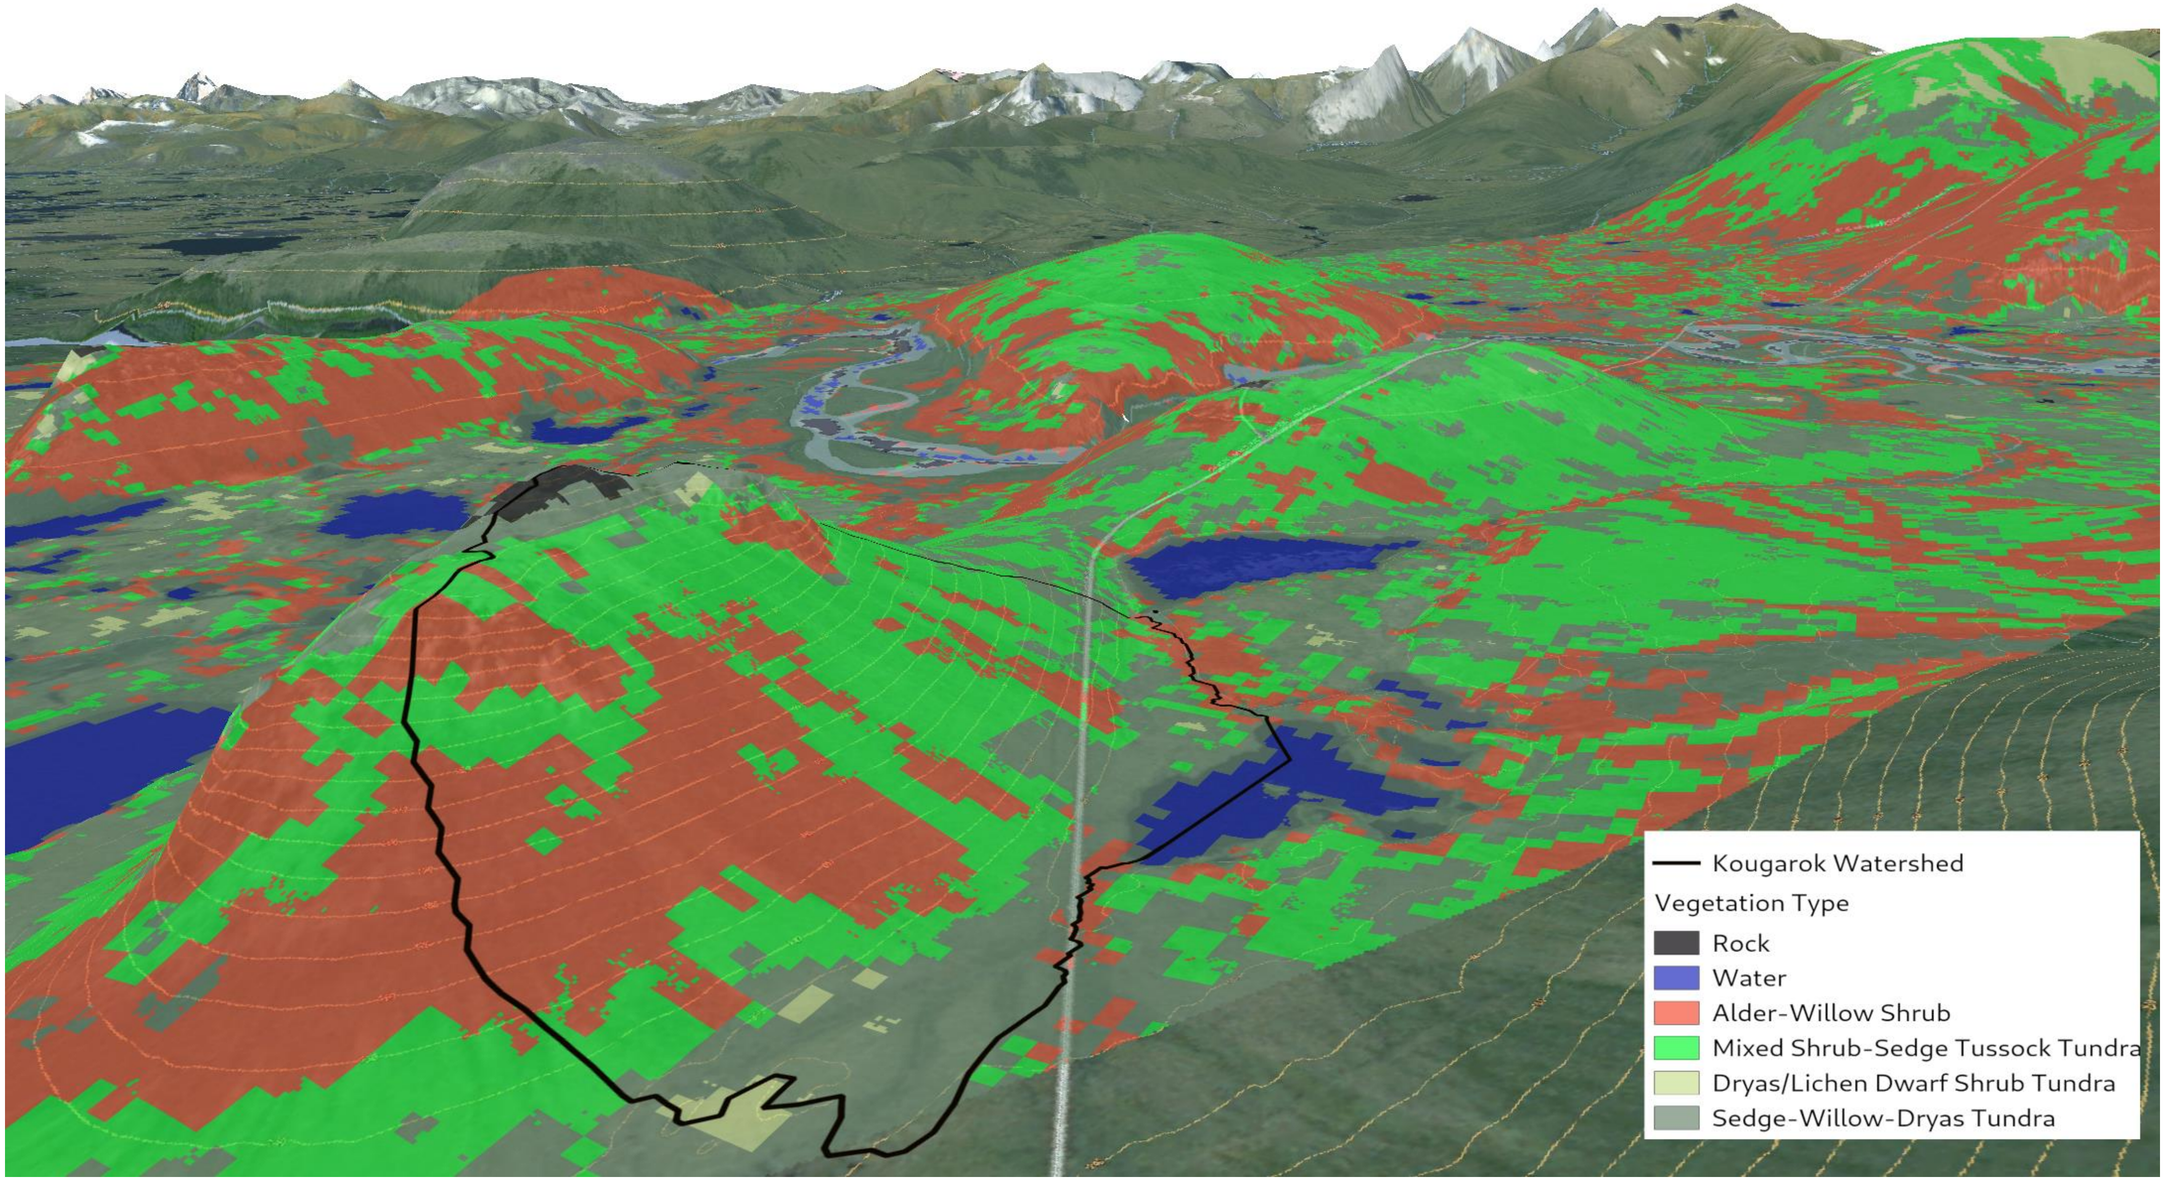

Summit : Future Plans

Build upon our hybrid implementation

- Utilize non-volatile memory (NVM) technologies
  - staging inputs
  - storing intermediate data structures
- Improve load balancing and utilization of GPUs
- Design a decentralized version
  - Handle very large data sets - O(10) TB+
- Improve integration of different algorithm formulations for better load balancing

E3SM

- Global Earth System Model
- Atmosphere, Land, Ocean and Ice component models
- 8 DOE labs, 12 university subcontracts, 53 FTEs spread over 87 individuals
- Development driven by DOE-SC mission interests: Energy/water issues looking out 40 years
- **Key computational goal: Ensure E3SM effectively utilizes DOE pre-exascale and exascale supercomputers**
- E3SM is open source / open development
  - Website: www.e3sm.org
  - Github: https://github.com/E3SM-Project
  - DOE Science youtube channel: https://www.youtube.com/channel/UC_rhpi0lBeD1U-6nD2zvlBA

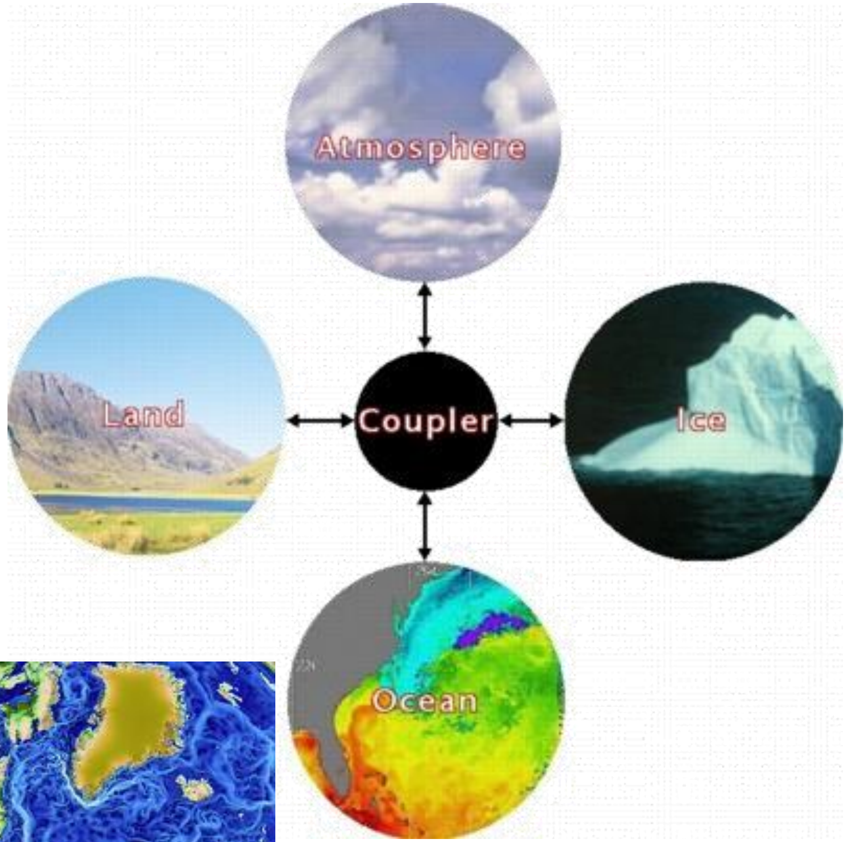

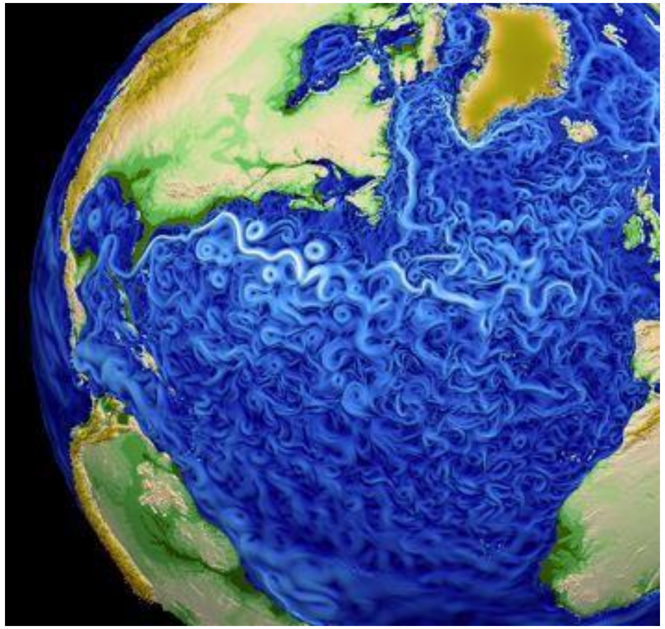

E3SM-MMF Cloud Resolving Climate Model

Goal: Develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.

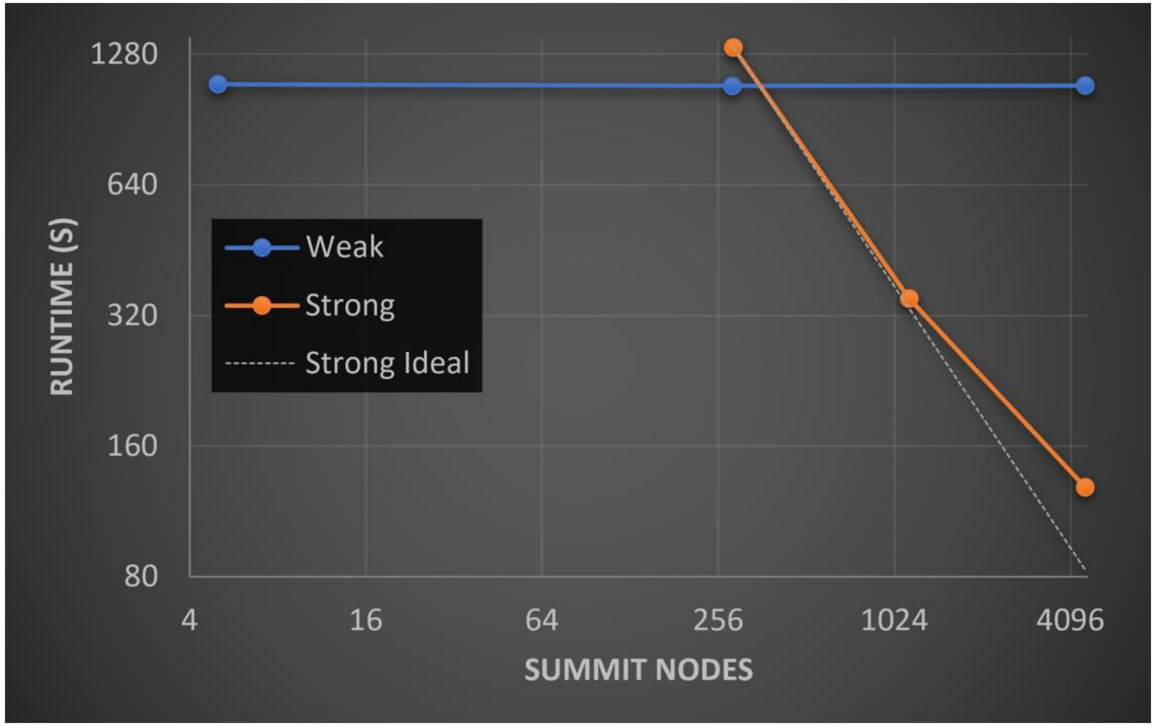

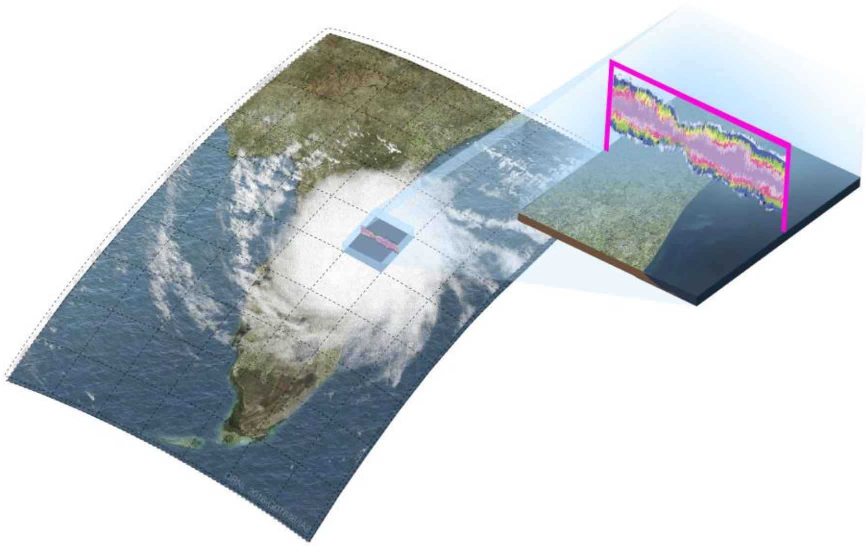

- Multiscale Modeling Framework / Super-Parameterization
- Replaces traditional parameterizations with cloud resolving model within each grid cell of global climate model

Co-design

- Feedback loop between applications, system s/w and computer architecture
- Application requirements inform (influence?) hardware design
- Technology choices and constraints guide problem formulation and design of algorithms.

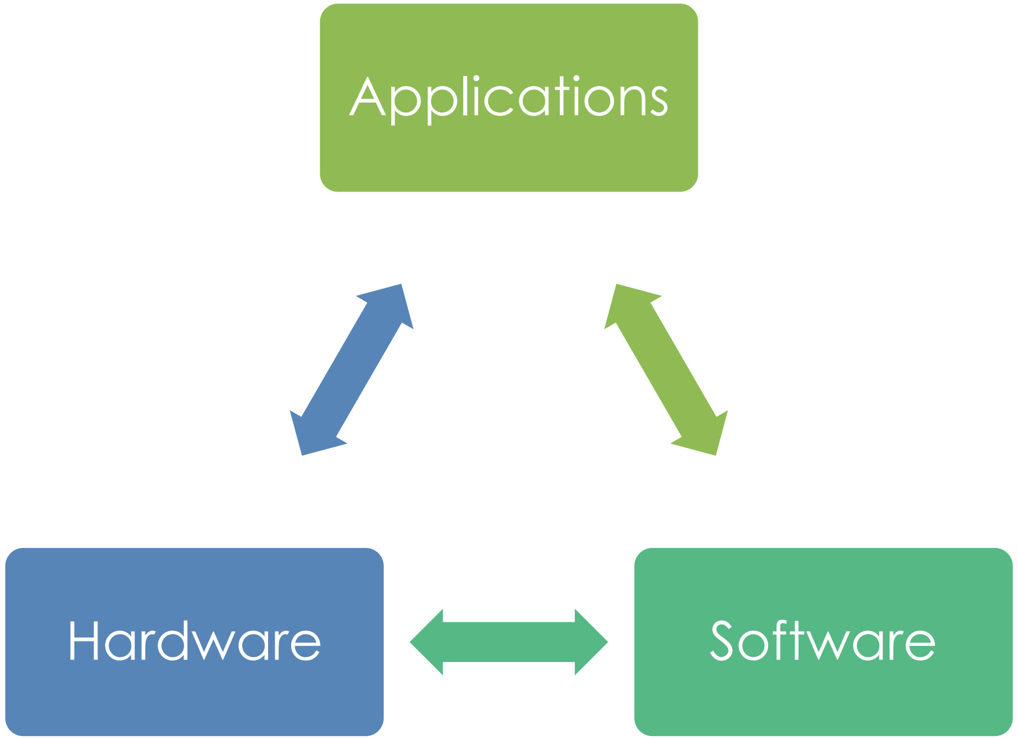

Acknowledgments

CLIMATE CHANGE SCIENCE INSTITUTE
OAK RIDGE NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY INCITE LEADERSHIP COMPUTING

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

SciDAC Scientific Discovery through Advanced Computing

U.S. Department of Agriculture, U.S. Forest Service,
Eastern Forest Environmental Threat Assessment Center.

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Backup

E3SM-MMF Results

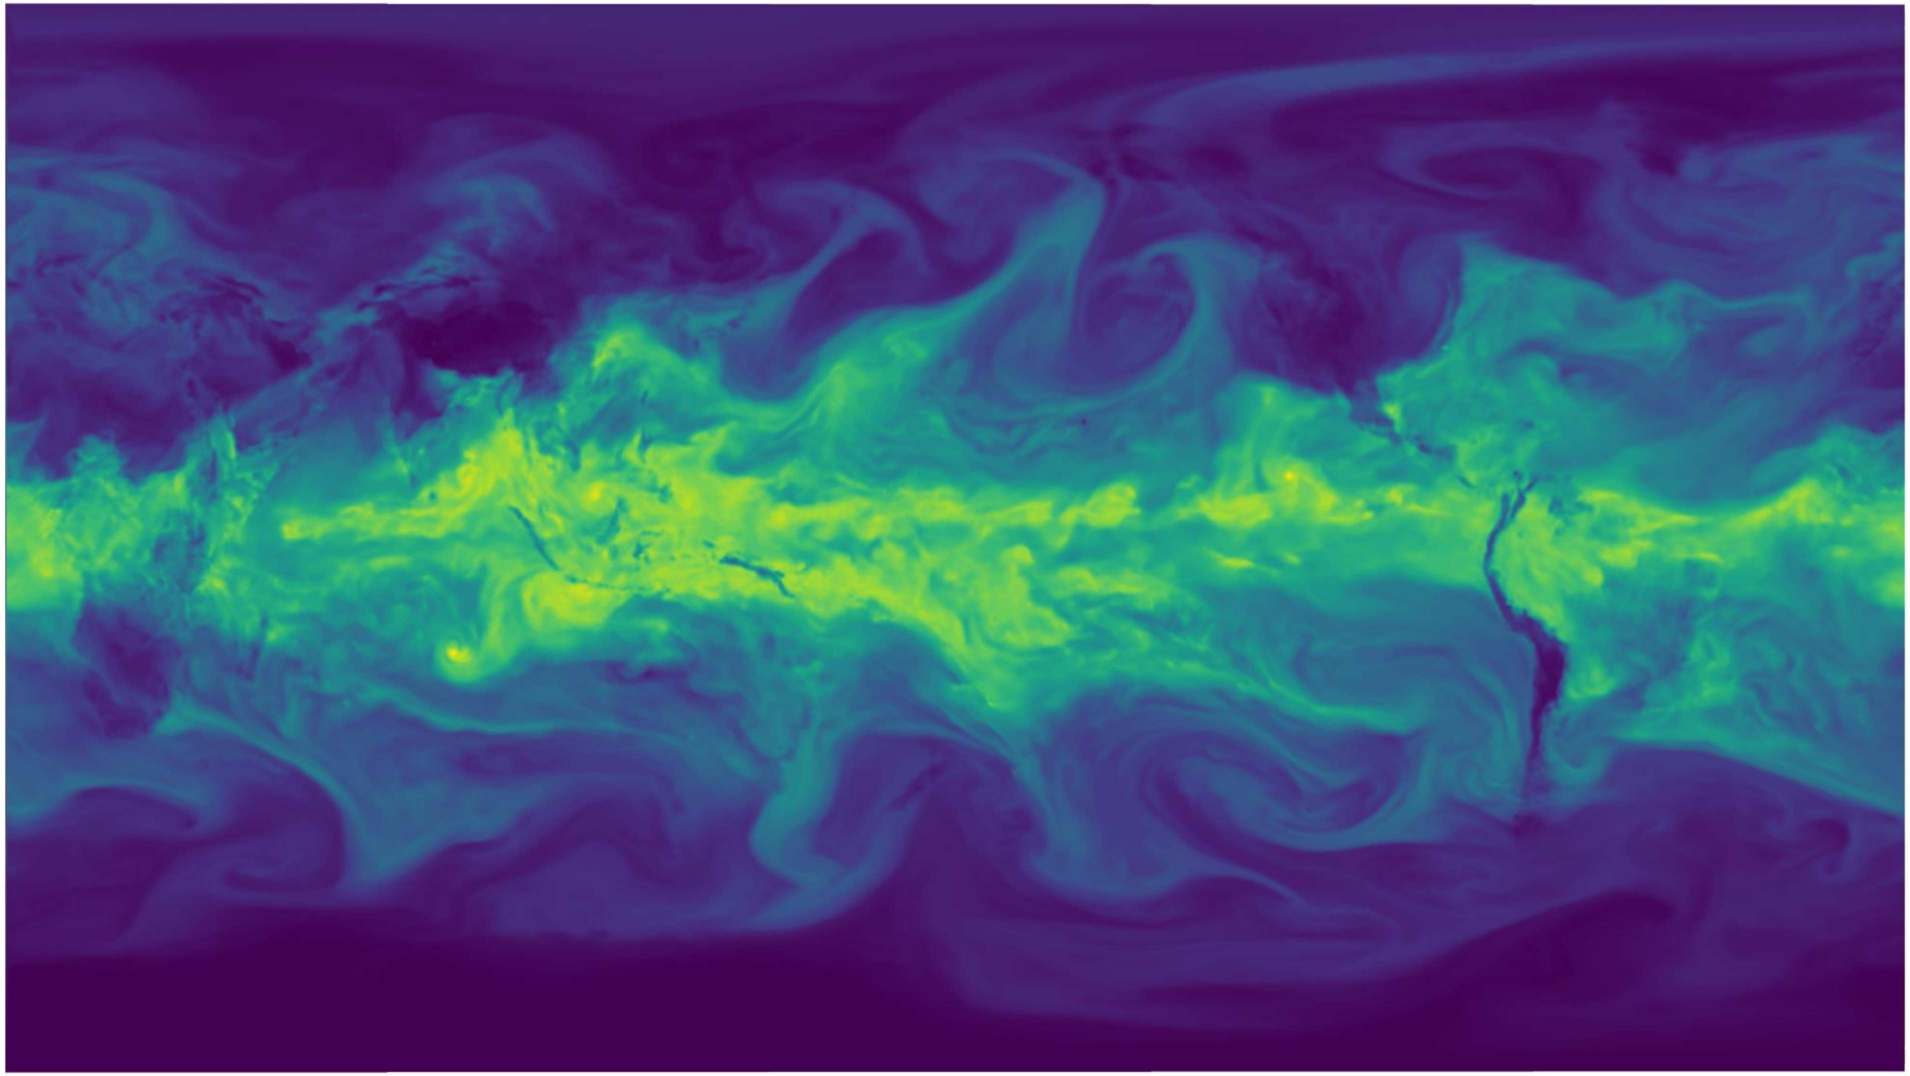

Water vapor contours after 40 days of simulation with E3SM-MMF using a 28km global grid and CRMs with 64 internal columns on 1,024 Summit nodes.
Credit: Kyle Pressel

Conclusions

- Parallel k-means clustering implementation for hybrid supercomputers
- BLAS formulation to accelerate Euclidean distance calculations
- Demonstrated up to 2.7x and 2.95x speedup over baseline CPU version in specific problem configurations on Titan and Summit
- Demonstrated capability to process large datasets
- Several Earth science applications
  - Great Smoky Mountains National Park: identification of vegetation structure
  - Global Climate Regimes: understanding global patterns of climate, vegetation and terrestrial ecology

Datasets

Great Smoky Mountains National Park (GSMNP)

- Airborne multiple return Light Detection and Ranging (LiDAR) data
  - Vertical canopy structure of the vegetation (74 dimensions)
  - 30 *m* × 30 *m* spatial resolution horizontal grid
  - 1 *m* vertical resolution to identify vegetation height from the ground surface

Global Climate Regimes

- Bioclimatic (BioClim) data for the contemporary period (17 dimensions)
- Climate models from IPCC Third Assessment Report (CMIP3) – Parallel Climate Model (PCM) and HadCM3 model
- Two different emissions scenarios:
  - B1 (lower emissions), A1FI (high emissions)

Global Climate Regimes: Variables

TABLE II
VARIABLES USED FOR DELINEATION OF GLOBAL CLIMATE REGIMES.

Variable Description	Units
Bioclimatic Variables	
Precipitation during the hottest quarter	mm
Precipitation during the coldest quarter	mm
Precipitation during the driest quarter	mm
Precipitation during the wettest quarter	mm
Ratio of precipitation to potential evapotranspiration	–
Temperature during the coldest quarter	°C
Temperature during the hottest quarter	°C
Day/night diurnal temperature difference	°C
Sum of monthly T_{avg} where $T_{\text{avg}} \geq 5°\text{C}$	°C
Integer number of consecutive months where $T_{\text{avg}} \geq 5°\text{C}$	–
Edaphic Variables	
Available water holding capacity of soil	mm
Bulk density of soil	g/cm^3
Carbon content of soil	g/cm^2
Nitrogen content of soil	g/cm^2
Topographic Variables	
Compound topographic index (relative wetness)	–
Solar interception	(kW/m^2)
Elevation	m

Summit Early Science Results

- Initial porting to Summit
- Performance Analysis and optimization
- Process large datasets – infeasible on Titan

Problem Config	Data Size	No. of Clusters	Nodes	CPU (avg cluster time/iter)	GPU (avg cluster time /iter)	Speedup
Phenology 2000	25 GB	1000	1	22.69 s	8.47 s	2.67
CONUS Phenology 2000–2015	395 GB	1000	100	10.60 s	3.59 s	2.95

Performance: Strong Scaling

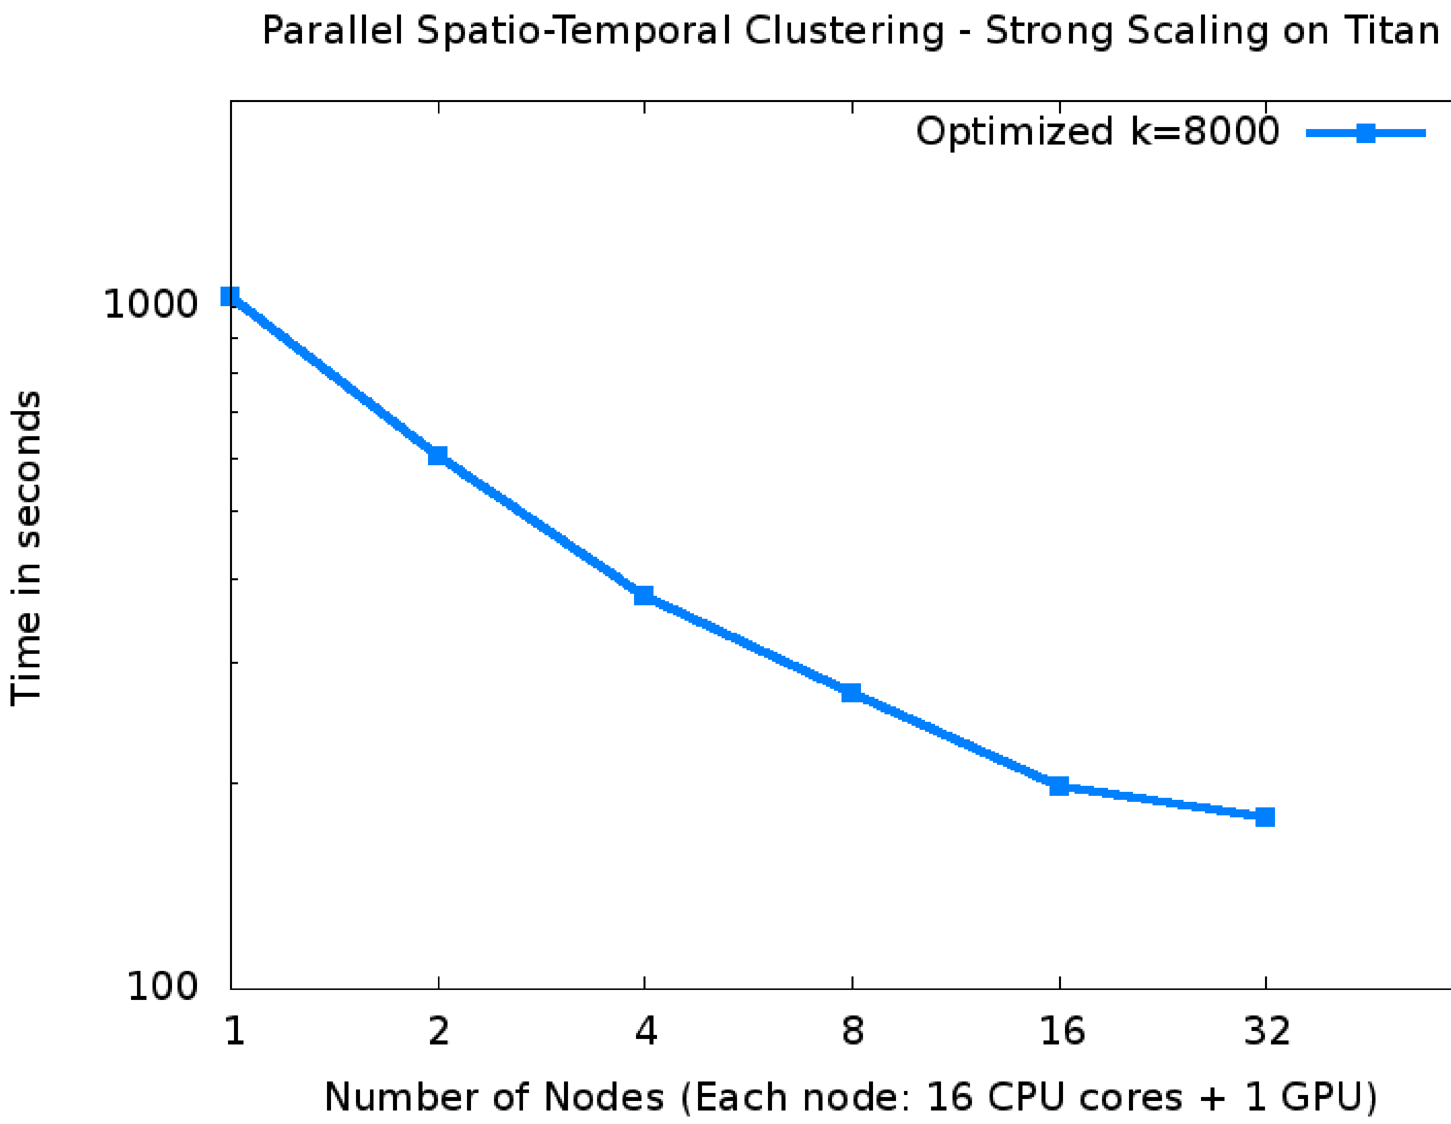

Performance Comparison