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

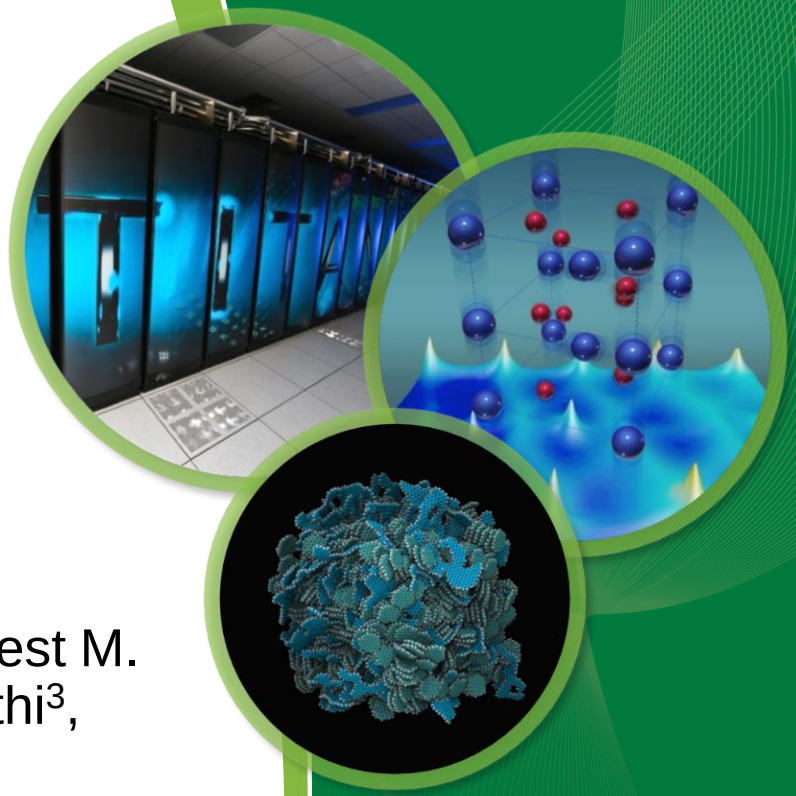

Sarat Sreepathi[1], Jitendra Kumar[1], Forrest M. Hoffman[1], Richard T. Mills[2], Vamsi Sripathi[3], William W. Hargrove[4]

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

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

Applications

- Vegetation mapping and characterization
- Development of ecoregions
- Species distribution

- Climate zone classification
- Understand climate regime changes in future
  - Under various predicted climate change scenarios

Datasets

Great Smoky Mountains National Park (GSMNP)

- Airborne multiple return Light Detection and Ranging (LiDAR) data
  - Vertical canopy structure of the vegetation
  - 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
- 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)

Datasets

DESCRIPTION OF DATA SETS USED IN THE CURRENT STUDY

Description	Dimensions	Size
GSMNP LiDAR	3,186,679 × 74	900 MB
CMIP3 Climate States	123,471,198 × 17	7.9 GB

Preprocessing

- Standardized the data set along each dimension
  - A mean of zero and standard deviation of one
- Allowing every dimension to be equally and fairly represented in the clustering algorithm

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 $T_{avg} \geq 5°C$	°C
Integer number of consecutive months where $T_{avg} \geq 5°C$	–
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

Parallel k-means (Baseline)

- 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)

Squared Euclidean Distance: $\mathbf{dist}_{i,j} = \left\|\mathbf{obs}_{i,*} - \mathbf{cent}_{i,*}\right\|^2$

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

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

xGER

xGEMM

$A := alpha * x * y' + A$

$C := alpha * op(A) * op(B) + beta * C$

BLAS Subroutines

Triangular acceleration (Application Phase 2)

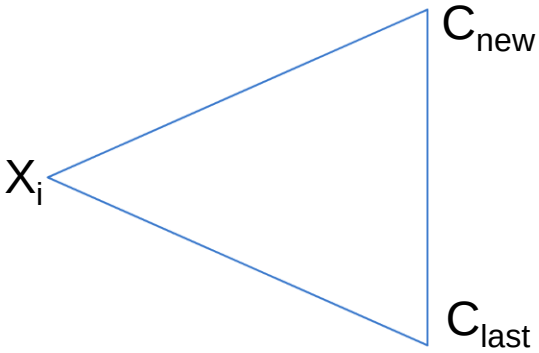

- Triangle inequality states :
  $d(C_{last}, C_{new}) \leq d(X_i, C_{last} + d(X_i, C_{new})$
- If $d(C_{last}, C_{new}) \geq 2d(X_i, C_{last})$,
  $\Rightarrow d(X_i, C_{new}) \geq d(X_i, C_{last})$ without computing
- Distance computations can be further reduced by sorting the inter-centroid distances, $d(C_{last}, C_{new})$
- New candidate centroids are evaluated as per sorted distance order
- Once the critical distance, $2d(X_i, C_{last})$ is surpassed all subsequent candidate centroids can be safely discarded

HPC Platforms

Summit

Summit Architecture

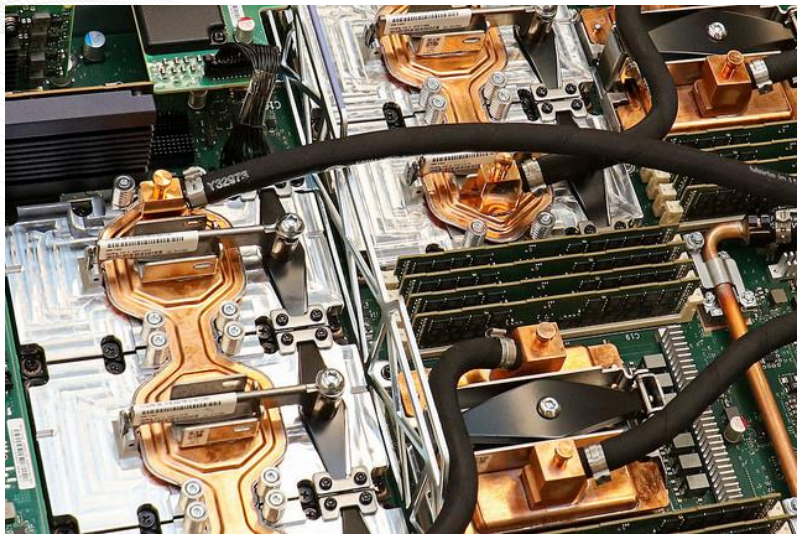

- ~200 PF (143 PF Linpack)
- 4608 Compute nodes

Each node:

- 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 Comparison

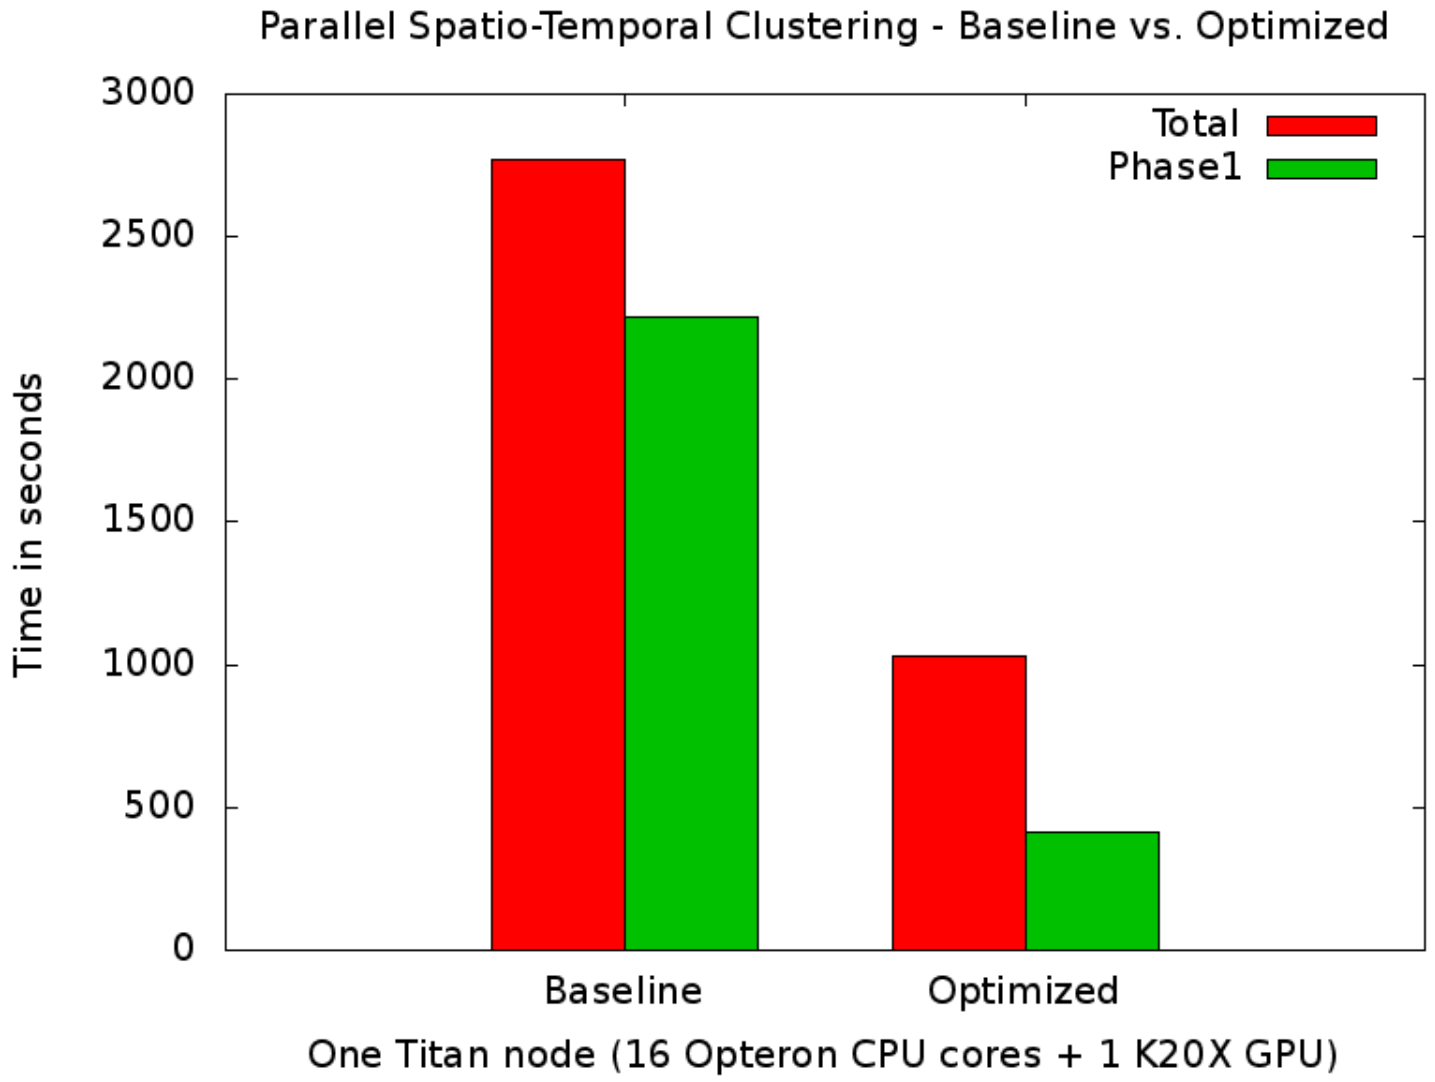

Performance: Varying Number of Clusters (k)

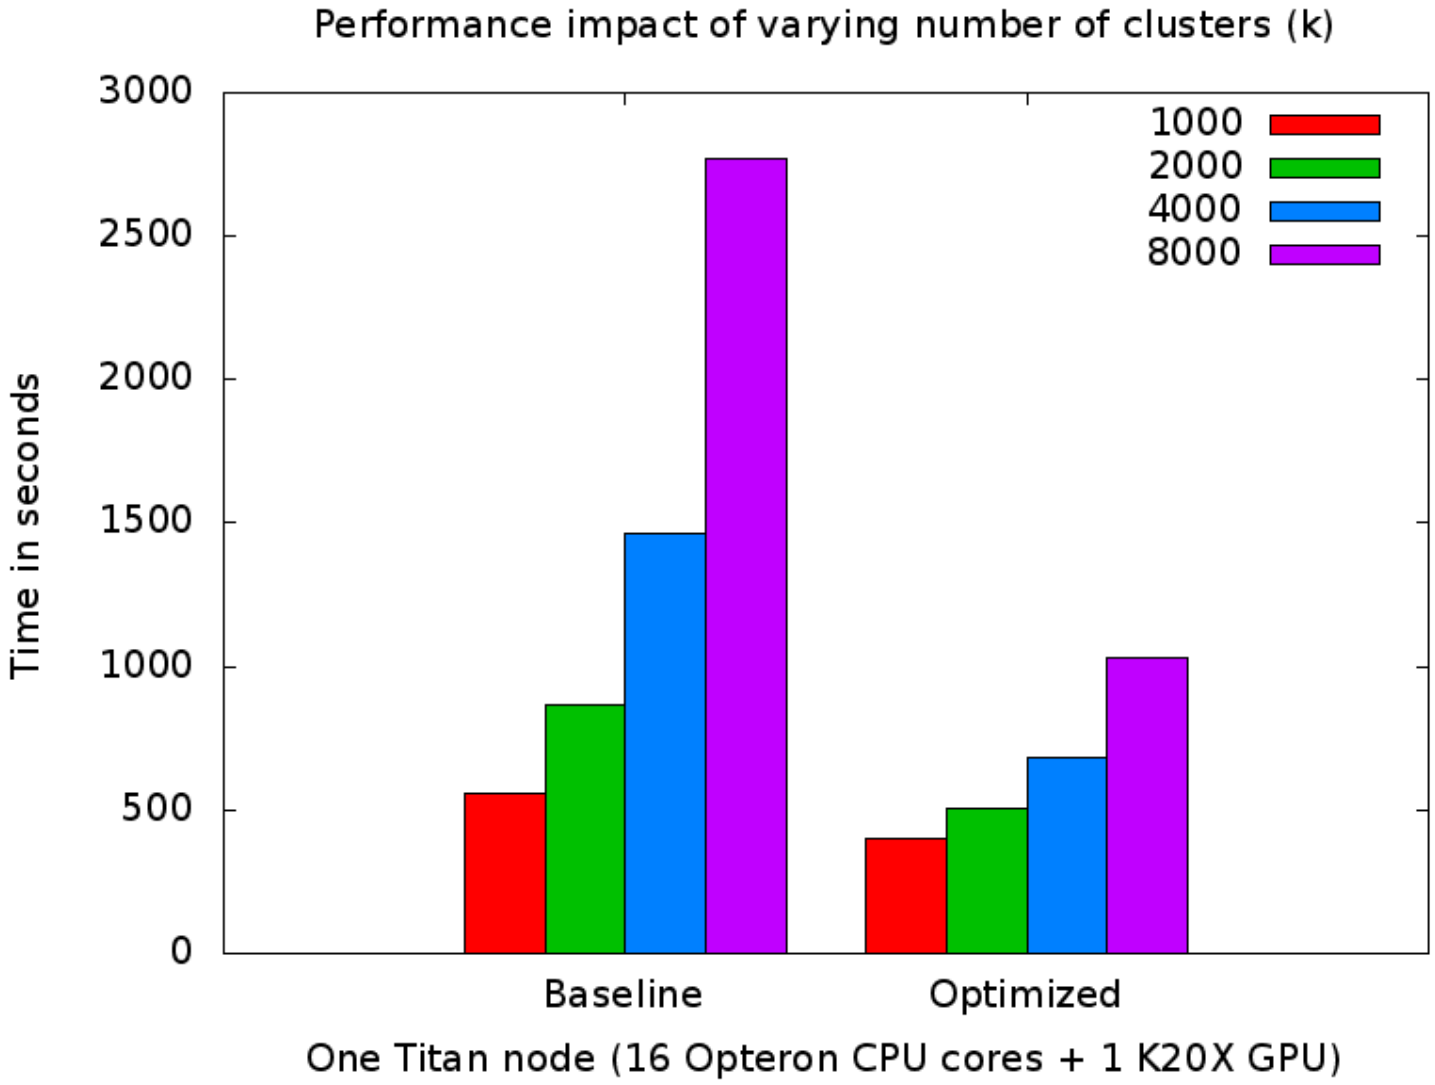

Summit: Early Results

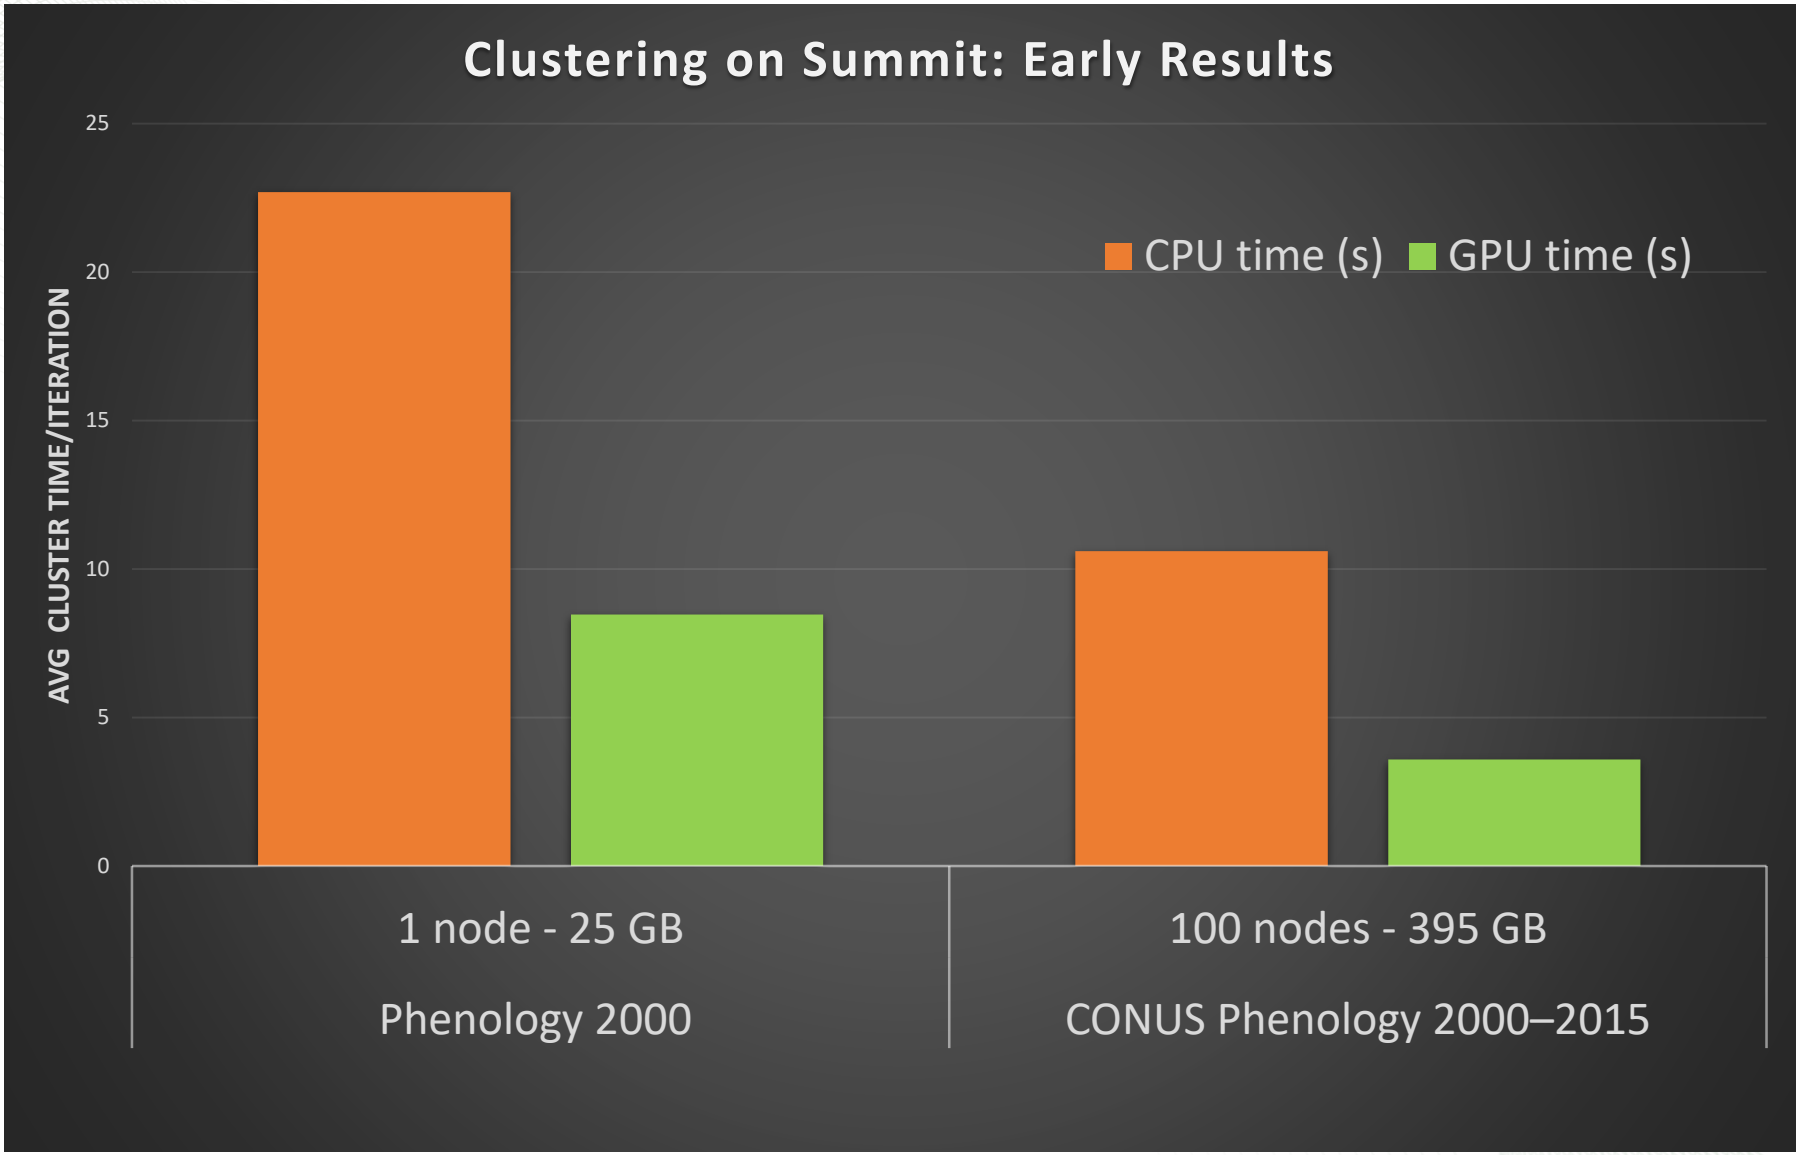

Summit : Future Plans

Build upon our hybrid implementation

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

Applications

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

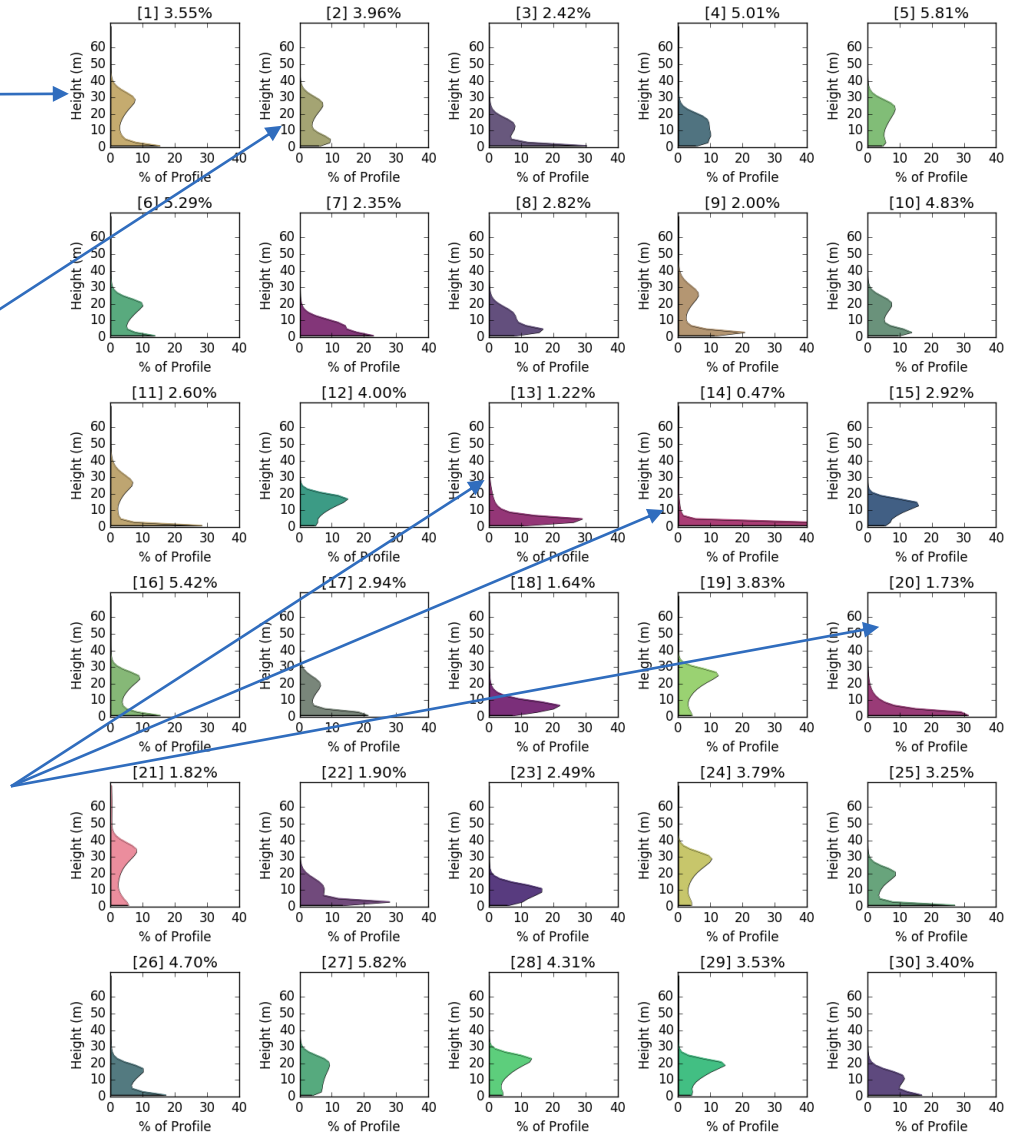

GSMNP: Spatial distribution of the 30 vegetation clusters across the national park

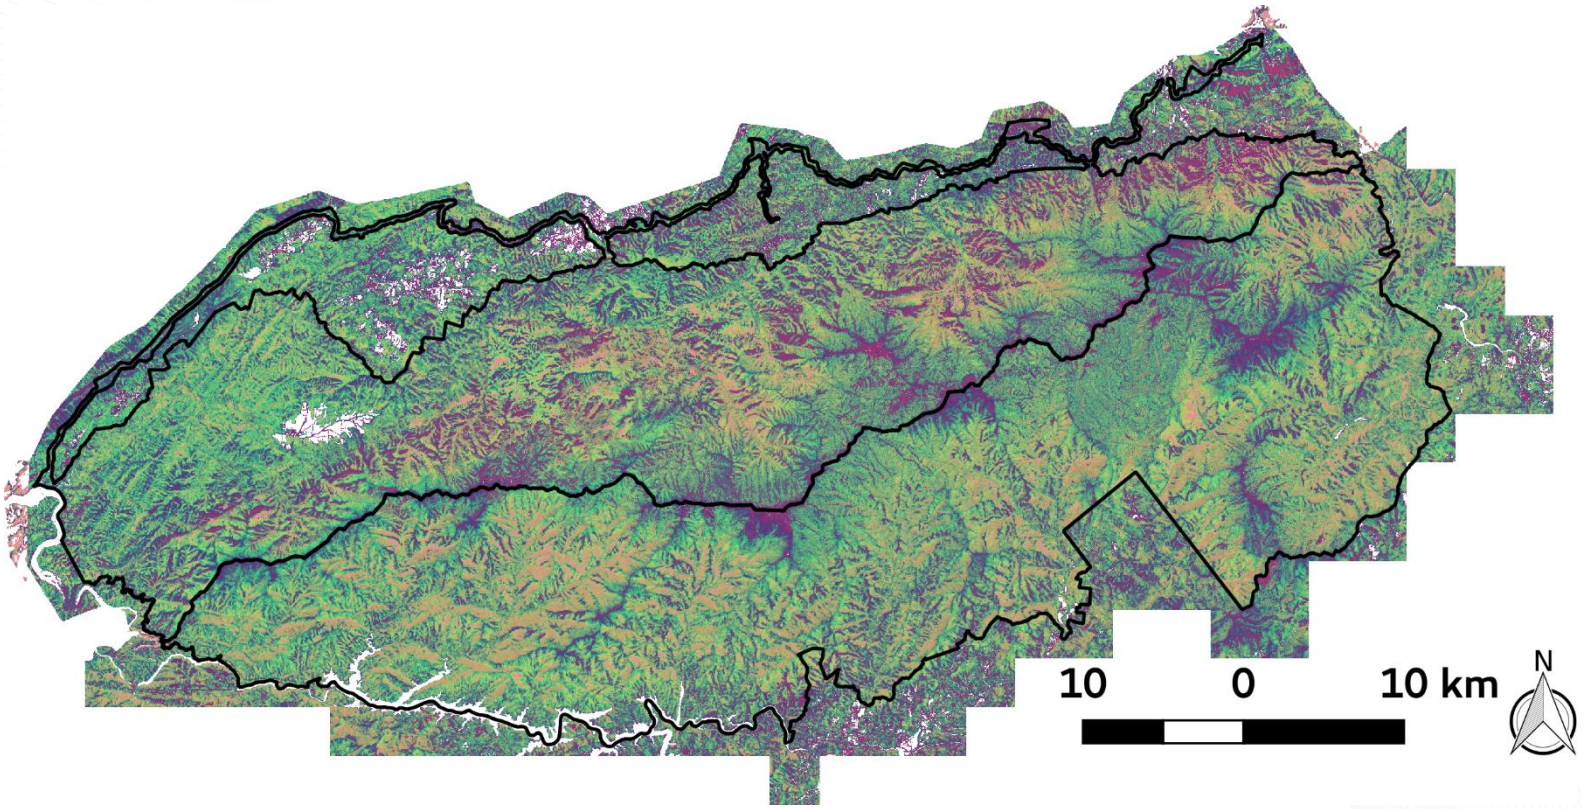

Global Climate Regimes: 1000 clusters Contemporary using Similarity color scheme

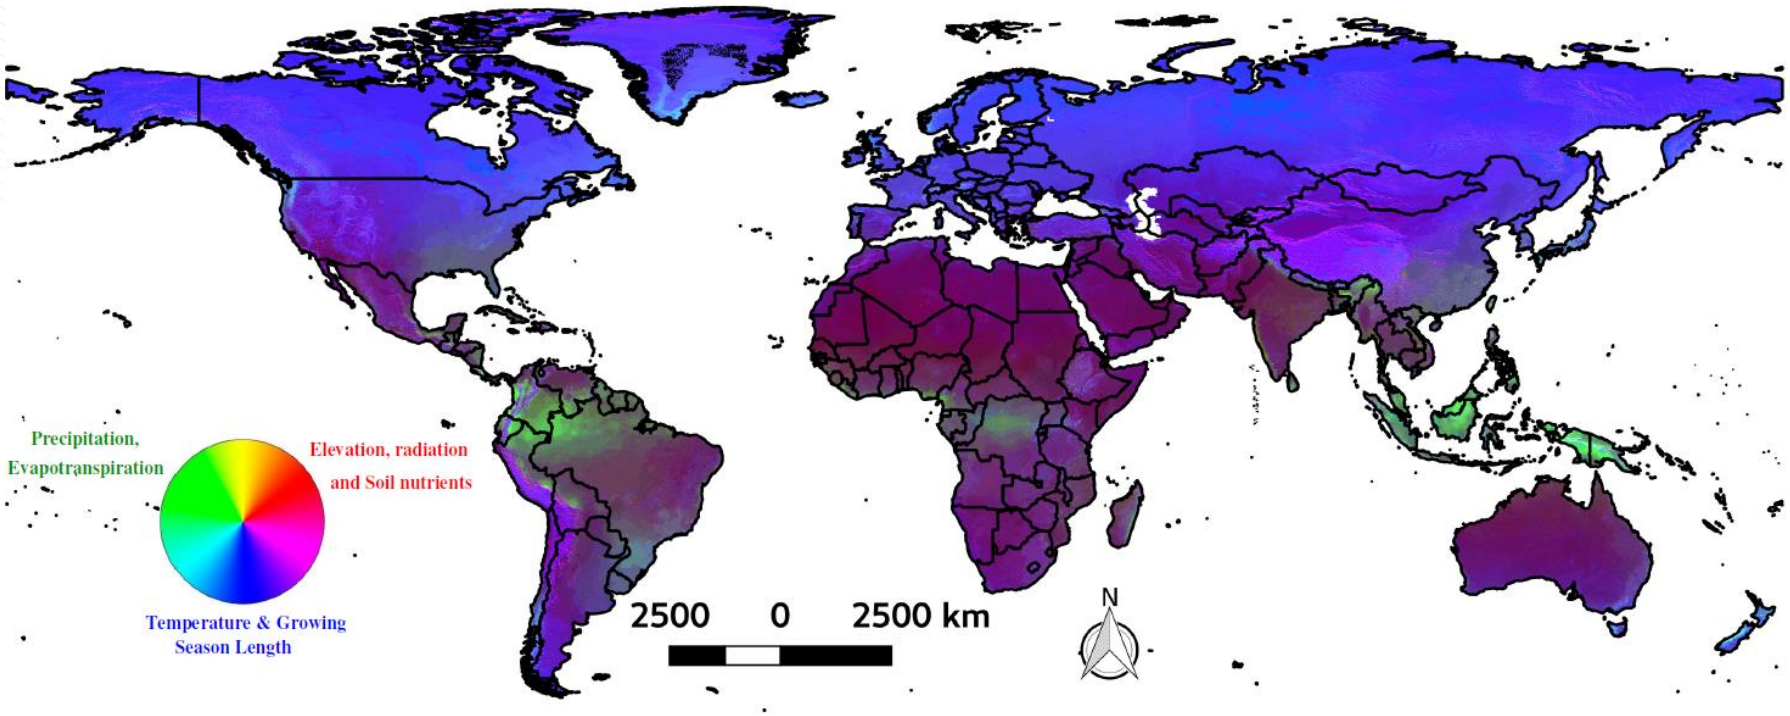

Global Climate Regimes: 1000 clusters 2100 using Similarity color scheme

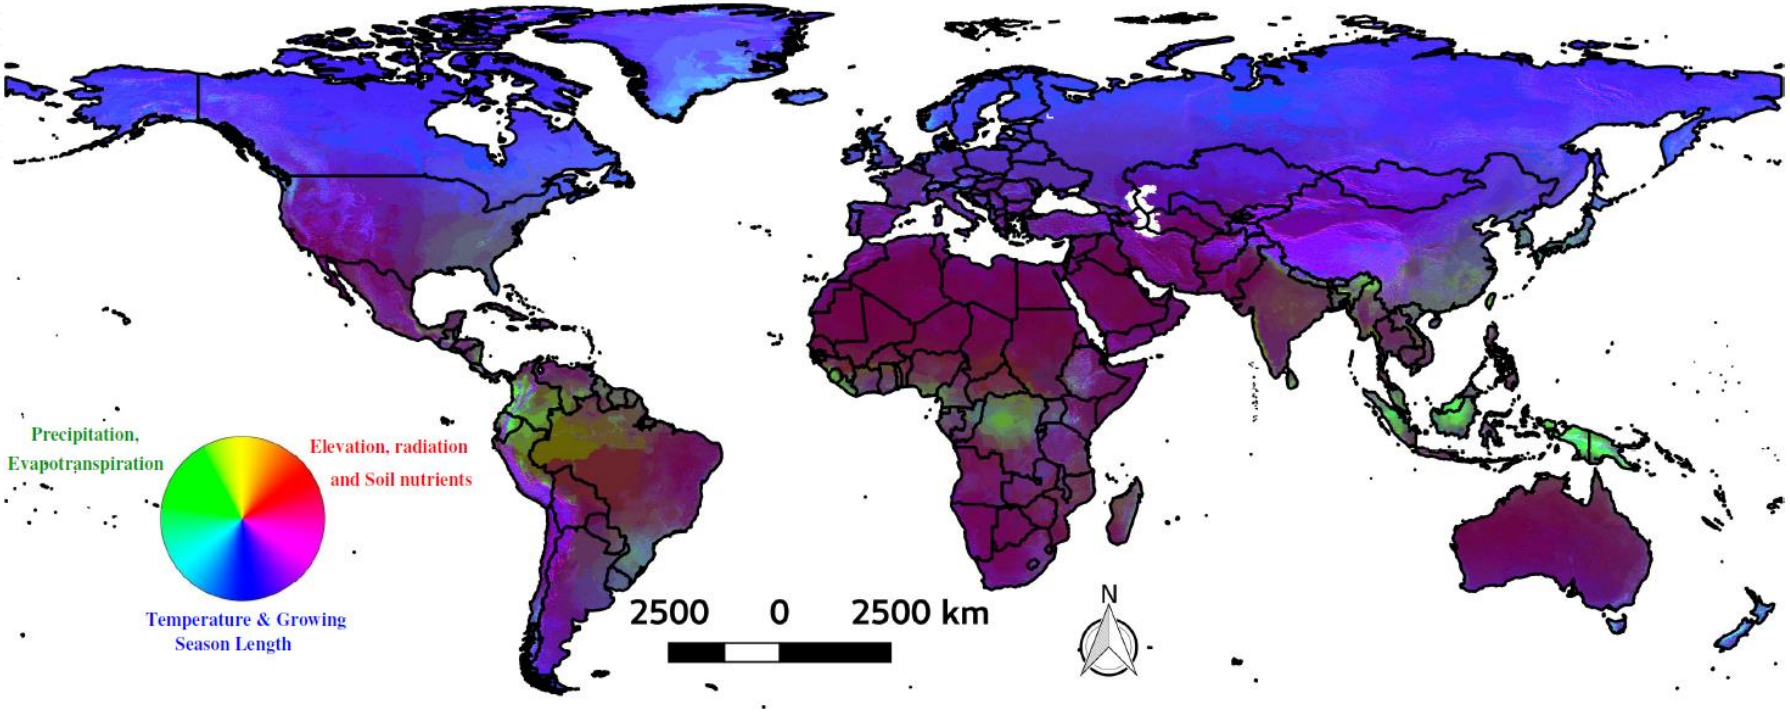

Conclusions

- 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
- Two Earth science applications
  - Great Smoky Mountains National Park: identification of vegetation structure
  - Global Climate Regimes: understanding global patterns of climate, vegetation and terrestrial ecology

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

USDA

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.

Global Climate Regimes: 1000 clusters Contemporary using Random color scheme

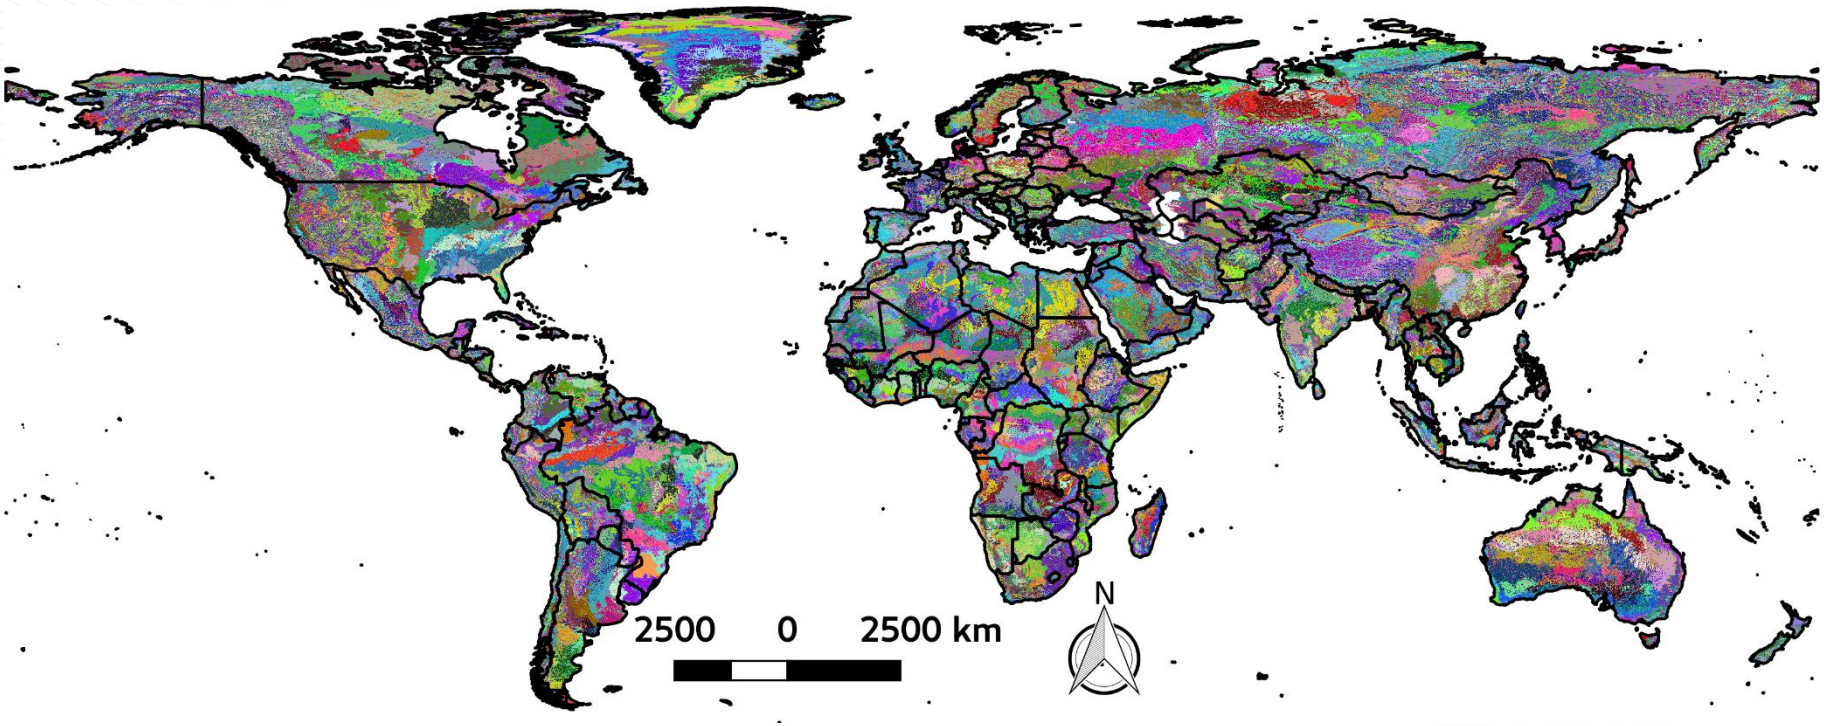

Backup

Limitations and Future Work

- Centralized master: inherent scaling limits
  - Decentralized version in development
- Larger datasets: Exceeds available on-node memory
  - Cluster assignment table and intermediate data structures
  - Short term: Decentralized version should partially address
  - Long term: Looking into NVRAM
- Application phases
  - Heuristic for switching
  - Combination
- Ported to KNL

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

Exascale Roadmap

Home » Secretary of Energy Rick Perry Announces $1.8 Billion Initiative for New Supercomputers

Systems Will Solidify U.S. Leadership in the "Exascale" Computing Era

CORAL
COLLABORATION
OAK RIDGE • ARGONNE • LIVERMORE

CORAL-2 ACQUISITION

CORAL-2 RFP No. 6400015092

RFP Components | CORAL-2 Benchmark Codes | Q&A

Proposal Due Date: May 24, 2018 by 5:00 pm Eastern Time

Frontier (ORNL) : 2021/22

El Capitan (LLNL) : 2022/23

Potential System (ANL)

First US Exascale System : ANL A21

- Planned Acceptance : 2021
- Outside this RFP

Application Characterization: Baseline k-means

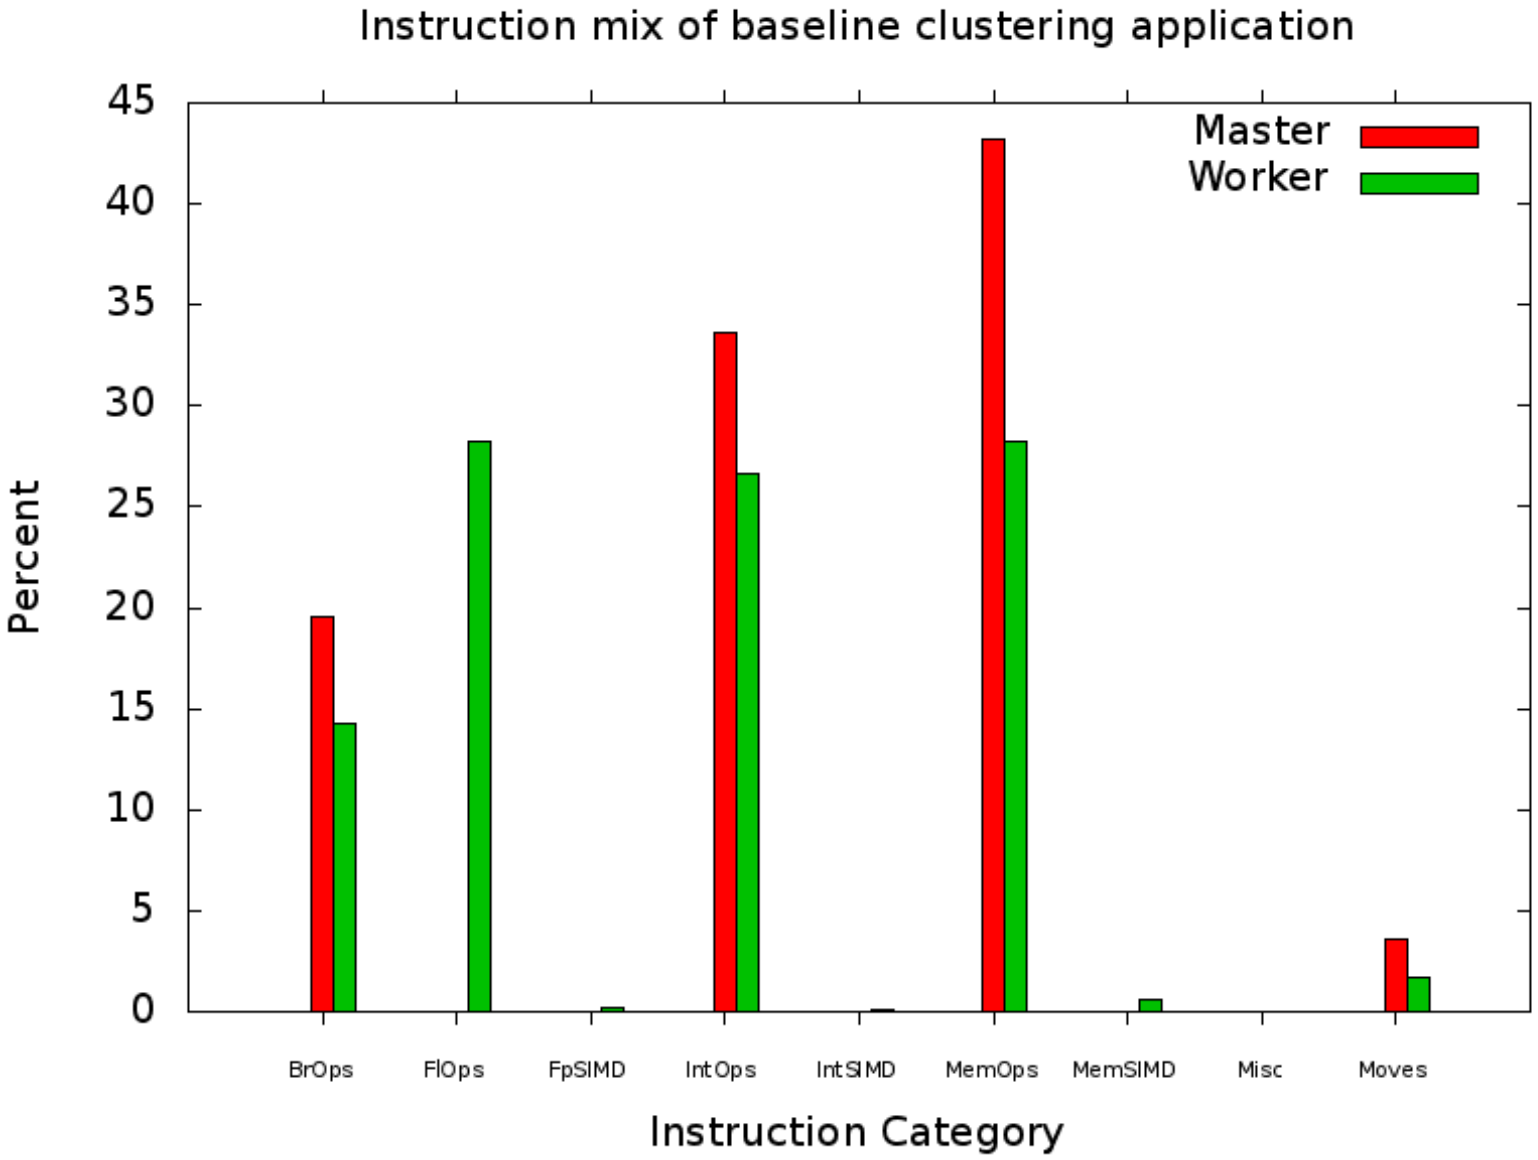

Performance: Impact of no. of work chunks

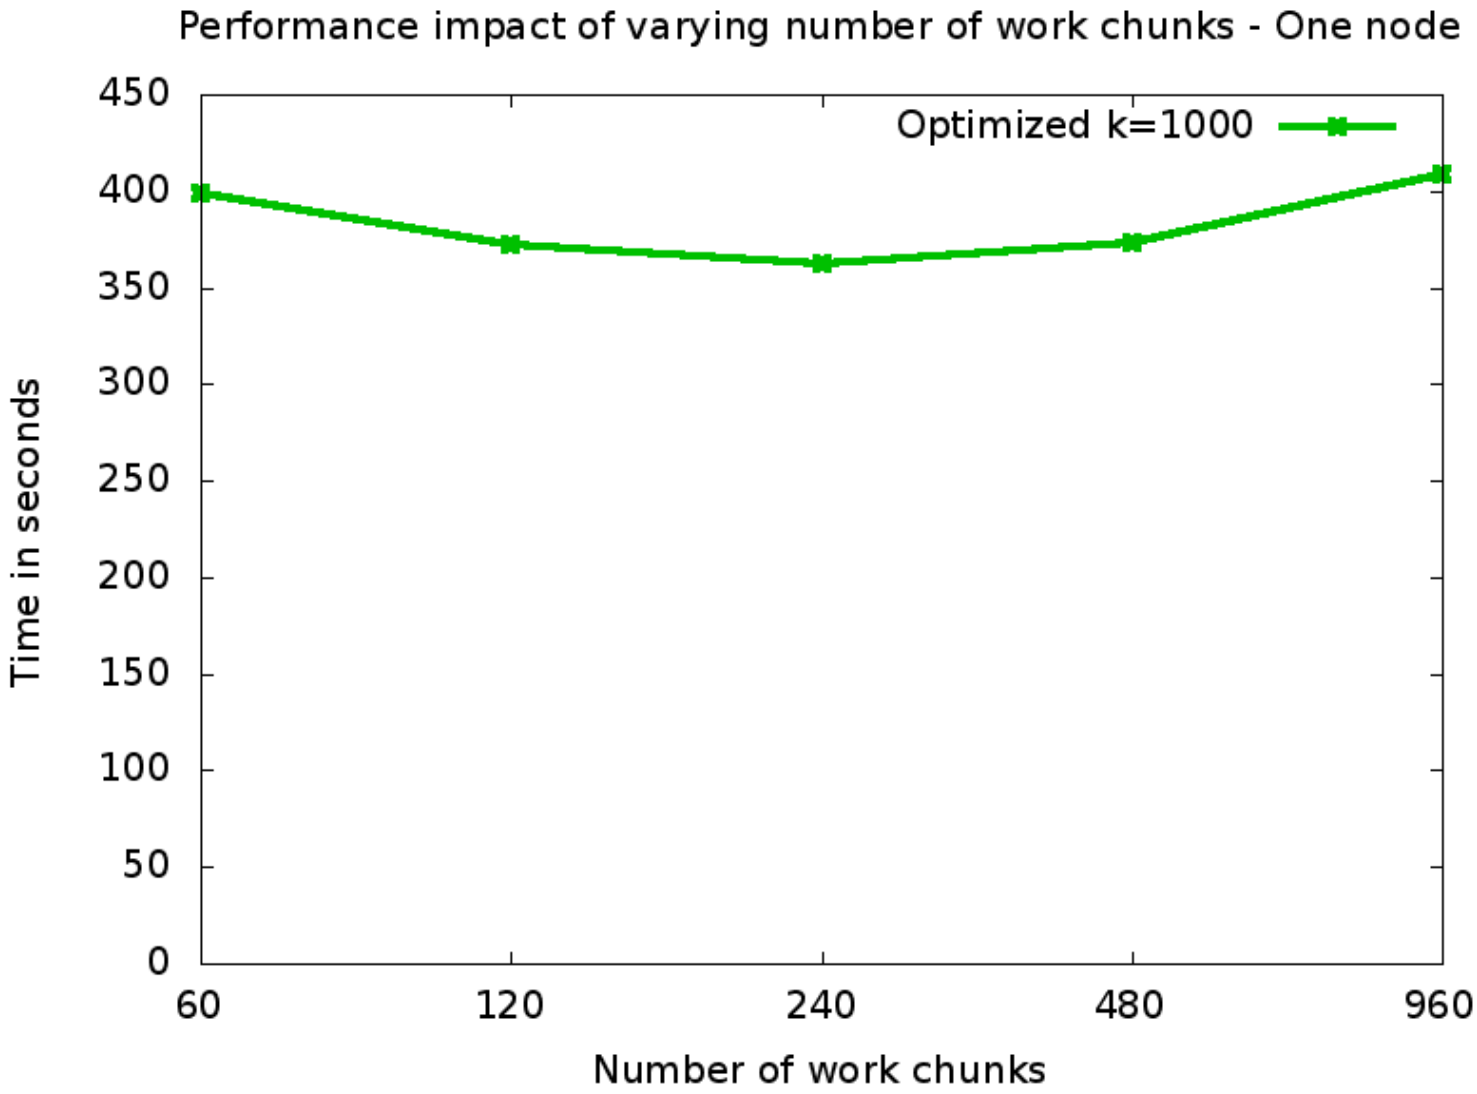

Performance: Strong Scaling