

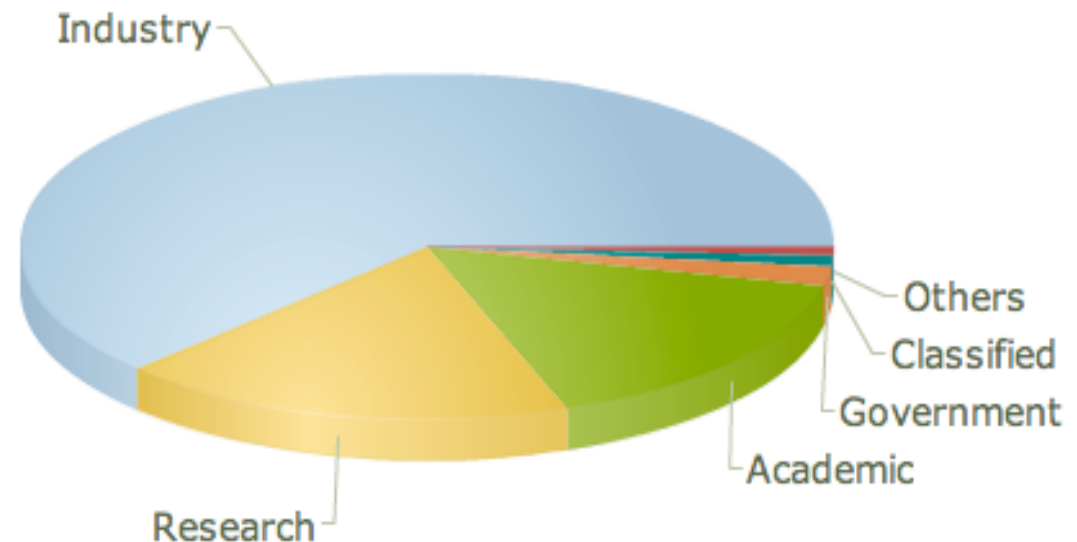
High Performance Computing

Didem Unat
University of California, San Diego

High Performance Computing

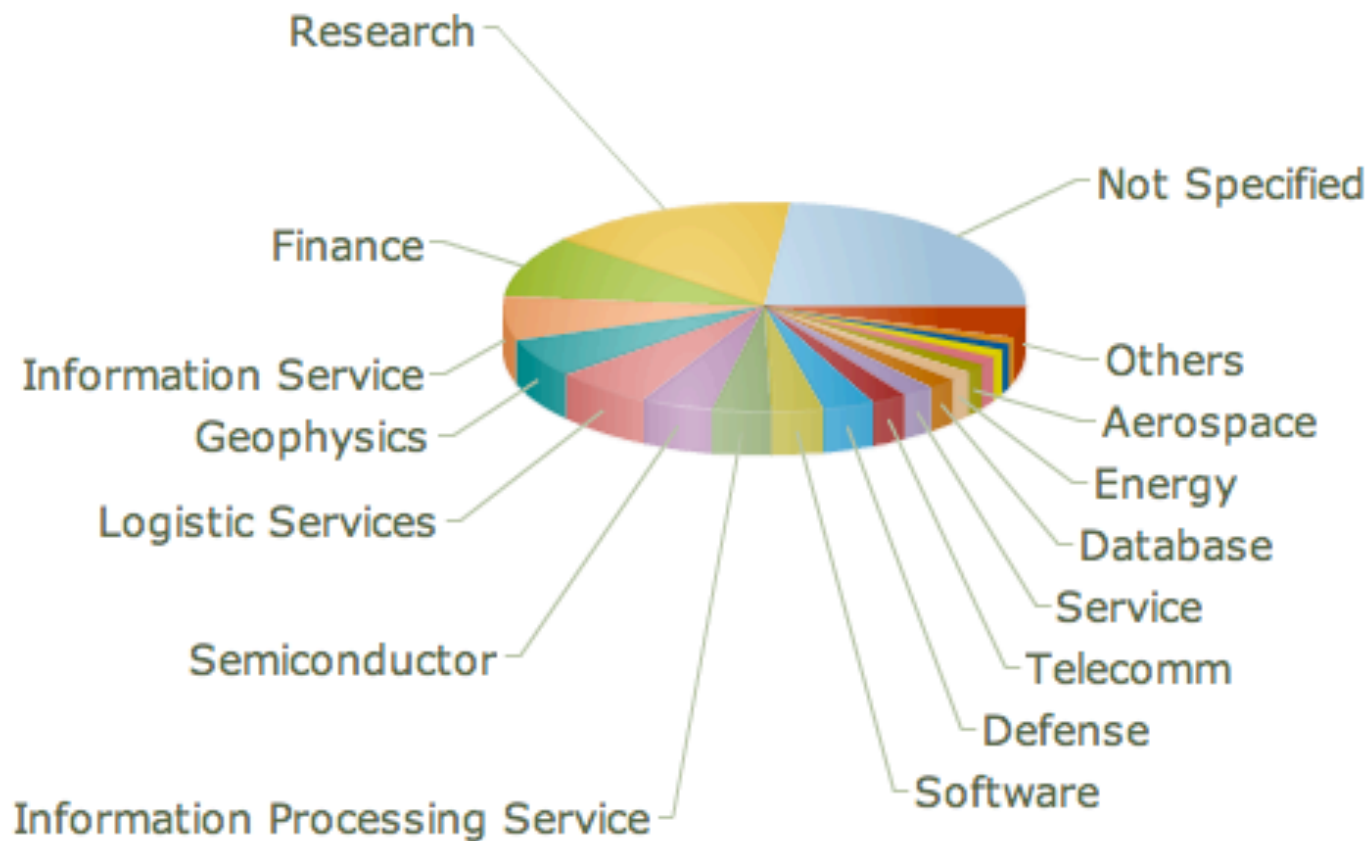
- Performance and Parallelism
- Works on large problems:
 - Long execution time OR large data sets
- Computing for scientific research (our focus)
 - Our clients are scientists
 - It is fun to work with them

Distribution of
Top 500 supercomputers



Application Areas are diverse

**Application Area / Systems
November 2009**



Then, research areas are diverse

- Scientific Simulation Optimizations
 - Dense & Sparse Linear Algebra
 - Monte Carlo Simulations
 - Structured & Unstructured Grid
 - FFT
 - N-body Simulations
- Developer Tools
 - Compilers and runtime systems
 - Parallel Programming Models
 - Numerical libraries
 - Profilers, debuggers
- Performance Prediction and Analysis
- Visualization & Large-scale Data Analysis/DB & Storage
- Energy Utilization and Cooling Systems



UCSD & SDSC

- My group, led by Scott Baden
 - Develops parallel programming models for emerging technologies
 - Domain specific languages
 - Numerical Method Optimization (Heart Simulations, Circuit Simulations)
 - Computational Databases for Computational Fluid Dynamics
 - Collaborates with Bio-medical and Mechanical Engineers.
- PMaC at SDSC, led by Alan Snavely
 - Performance Modeling and Characterization
- DASH and Flash Gordon (Steven Swanson and others)
Low cost and low power solution compared to traditional disk based approaches.
- Energy Efficiency in Multi-scale computing
 - Rajesh Gupta, Tajana Rosing, Amin Vahdat
- Earthquake Simulations, Geosciences, Biological Sciences
- Visualization and many more...

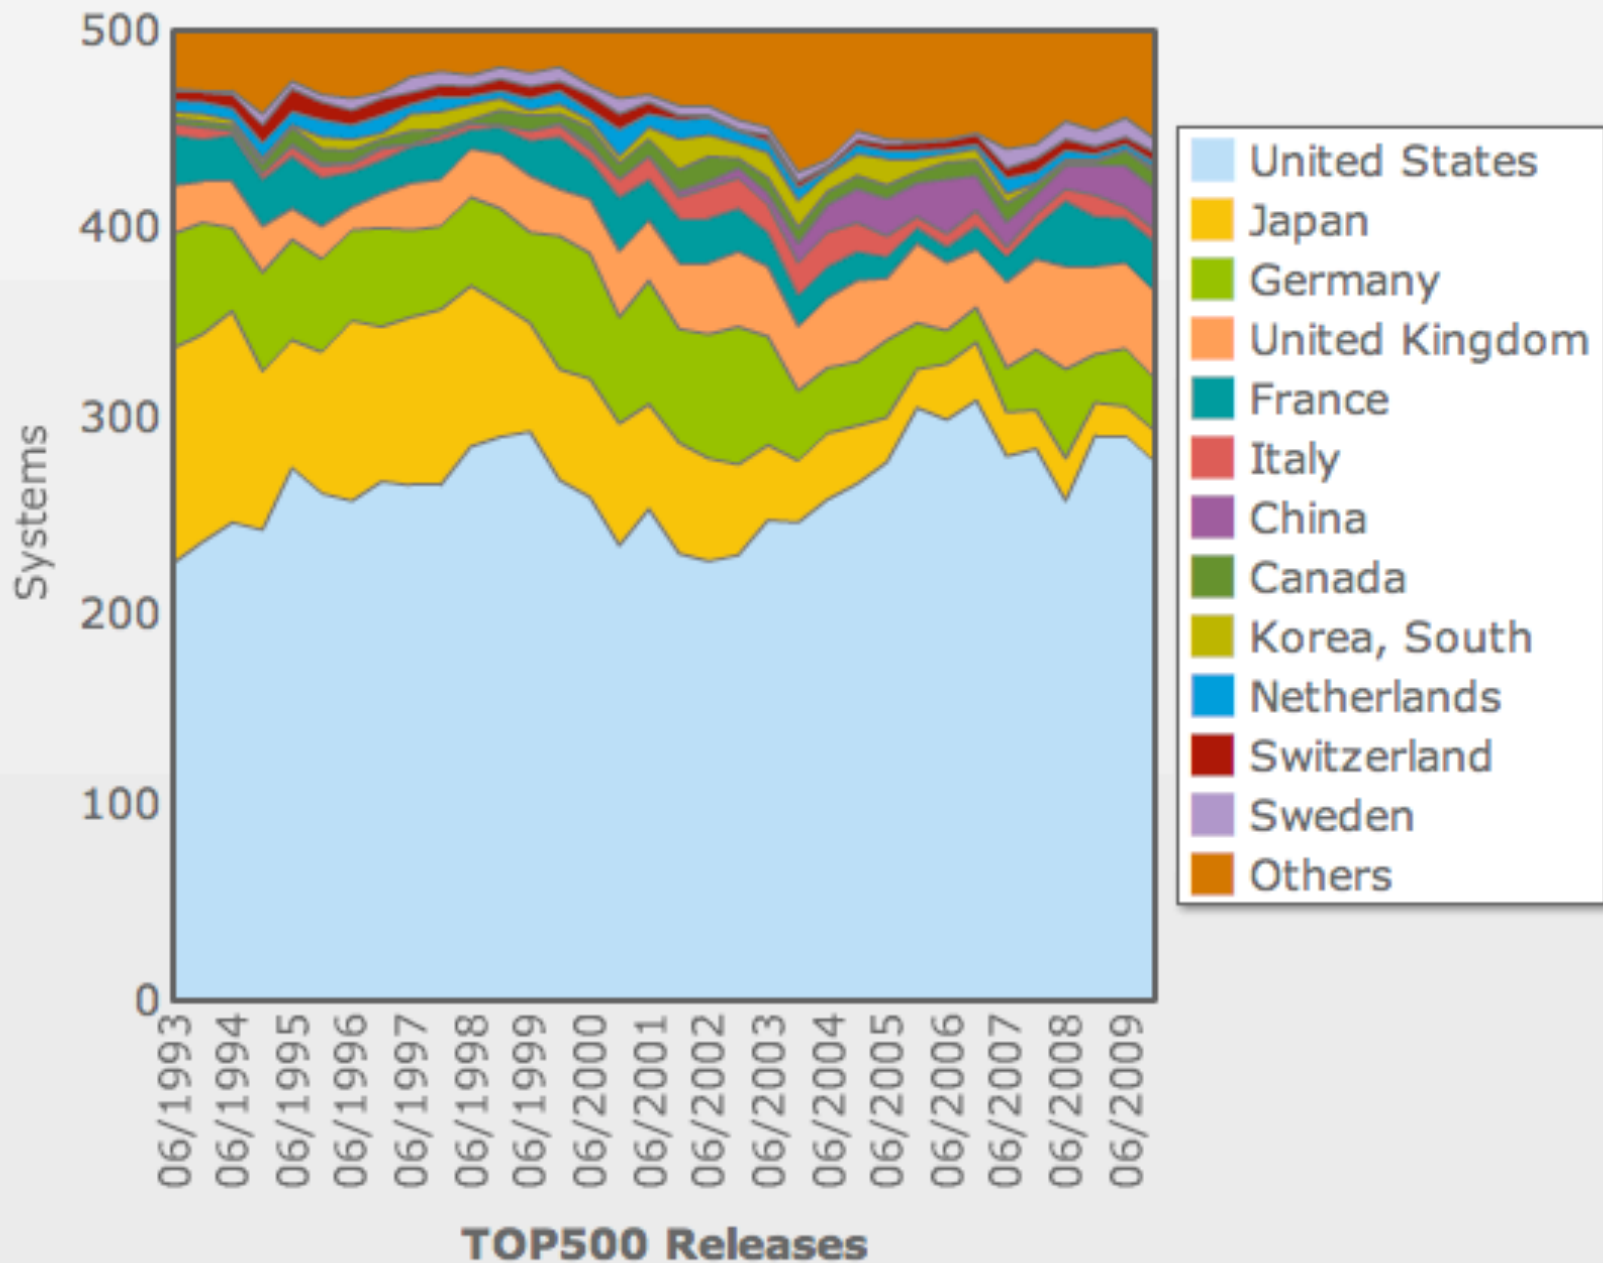
What is HOT?

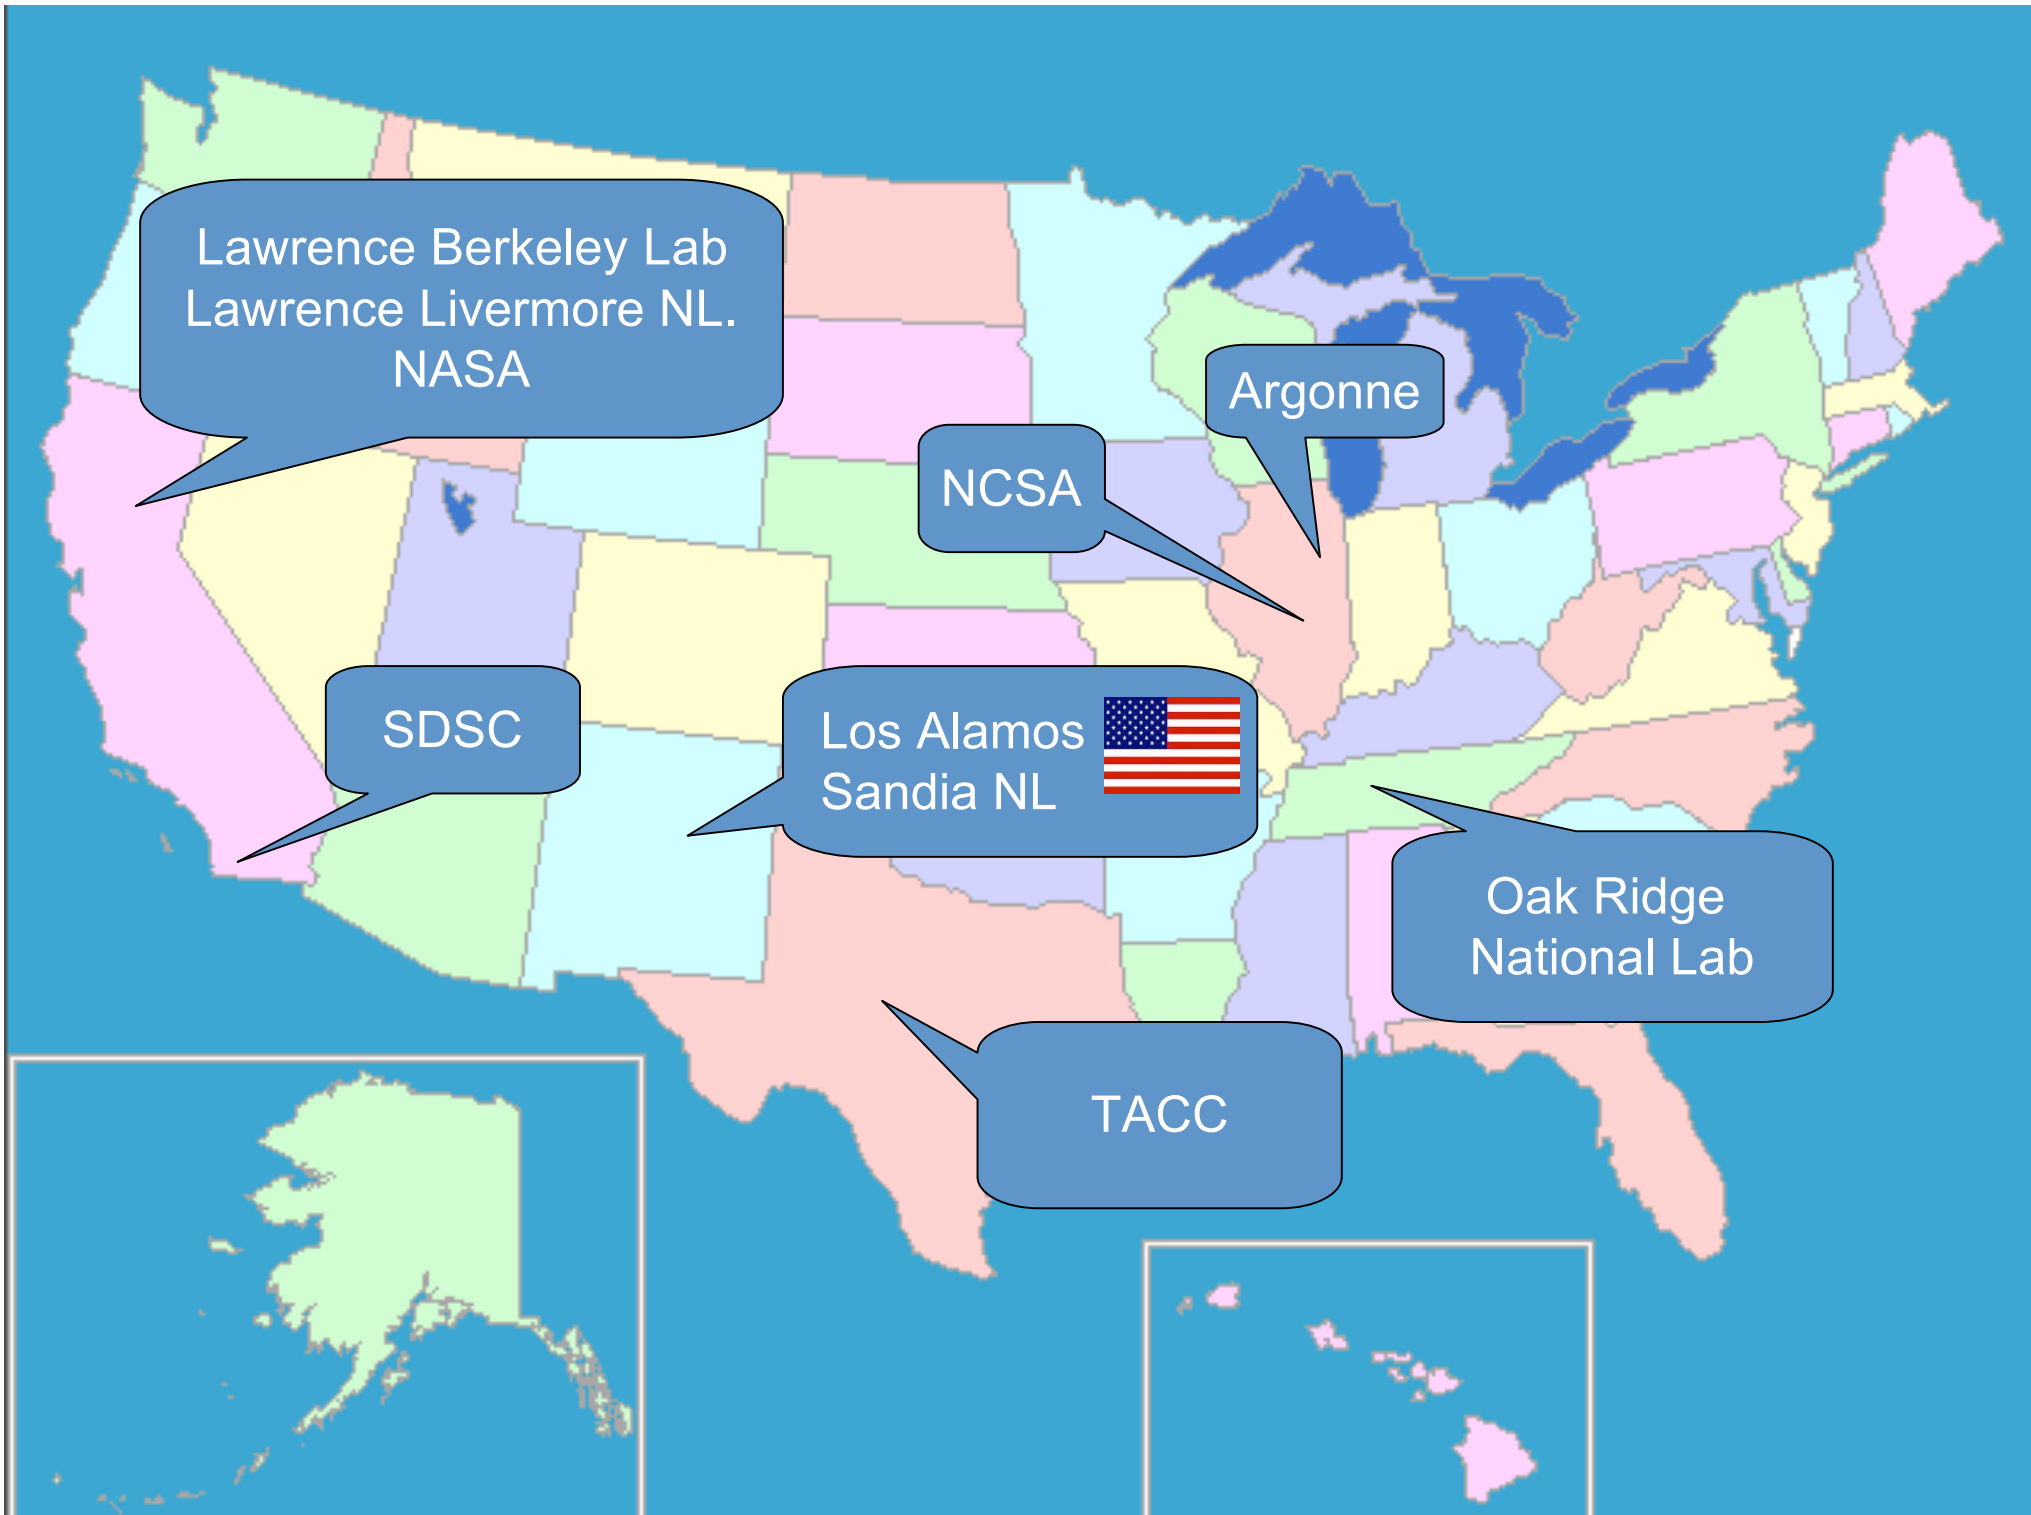
- Heterogeneous Computing with Accelerators
 - Many think accelerated HPC is the primary path to Exascale computing.
 - The Road Runner supercomputer and Lincoln machine with the GPU accelerator
 - Focus on software infrastructure for making effective use of accelerator
- Data Intensive Computing
 - Climate, Earthquake Simulations and Bio-computing
 - Creates huge data requirements from both their simulation and specific experimental activities.
 - How data is shared and communicated among participating scientists
- Energy Savings
 - Applications Architecture Power Puzzle
 - Low Cost and low power consumption
 - Cooling technologies
 - Power-aware computing/scheduling

Jobs?

- Faculty Positions
 - University of CA, Berkeley
 - University of IL, Urbana Champaign, Purdue University, University of Tennessee, University of Texas, Austin
- Supercomputer Centers and National Laboratories
 - Research Scientists or Postdoc positions
 - Well-paid
 - Flexible working environment
 - If you have a major in a life sciences, that helps but not a requirement

Countries Share Over Time 1993-2009





Lawrence Berkeley Lab
Lawrence Livermore NL.
NASA

SDSC

Los Alamos
Sandia NL

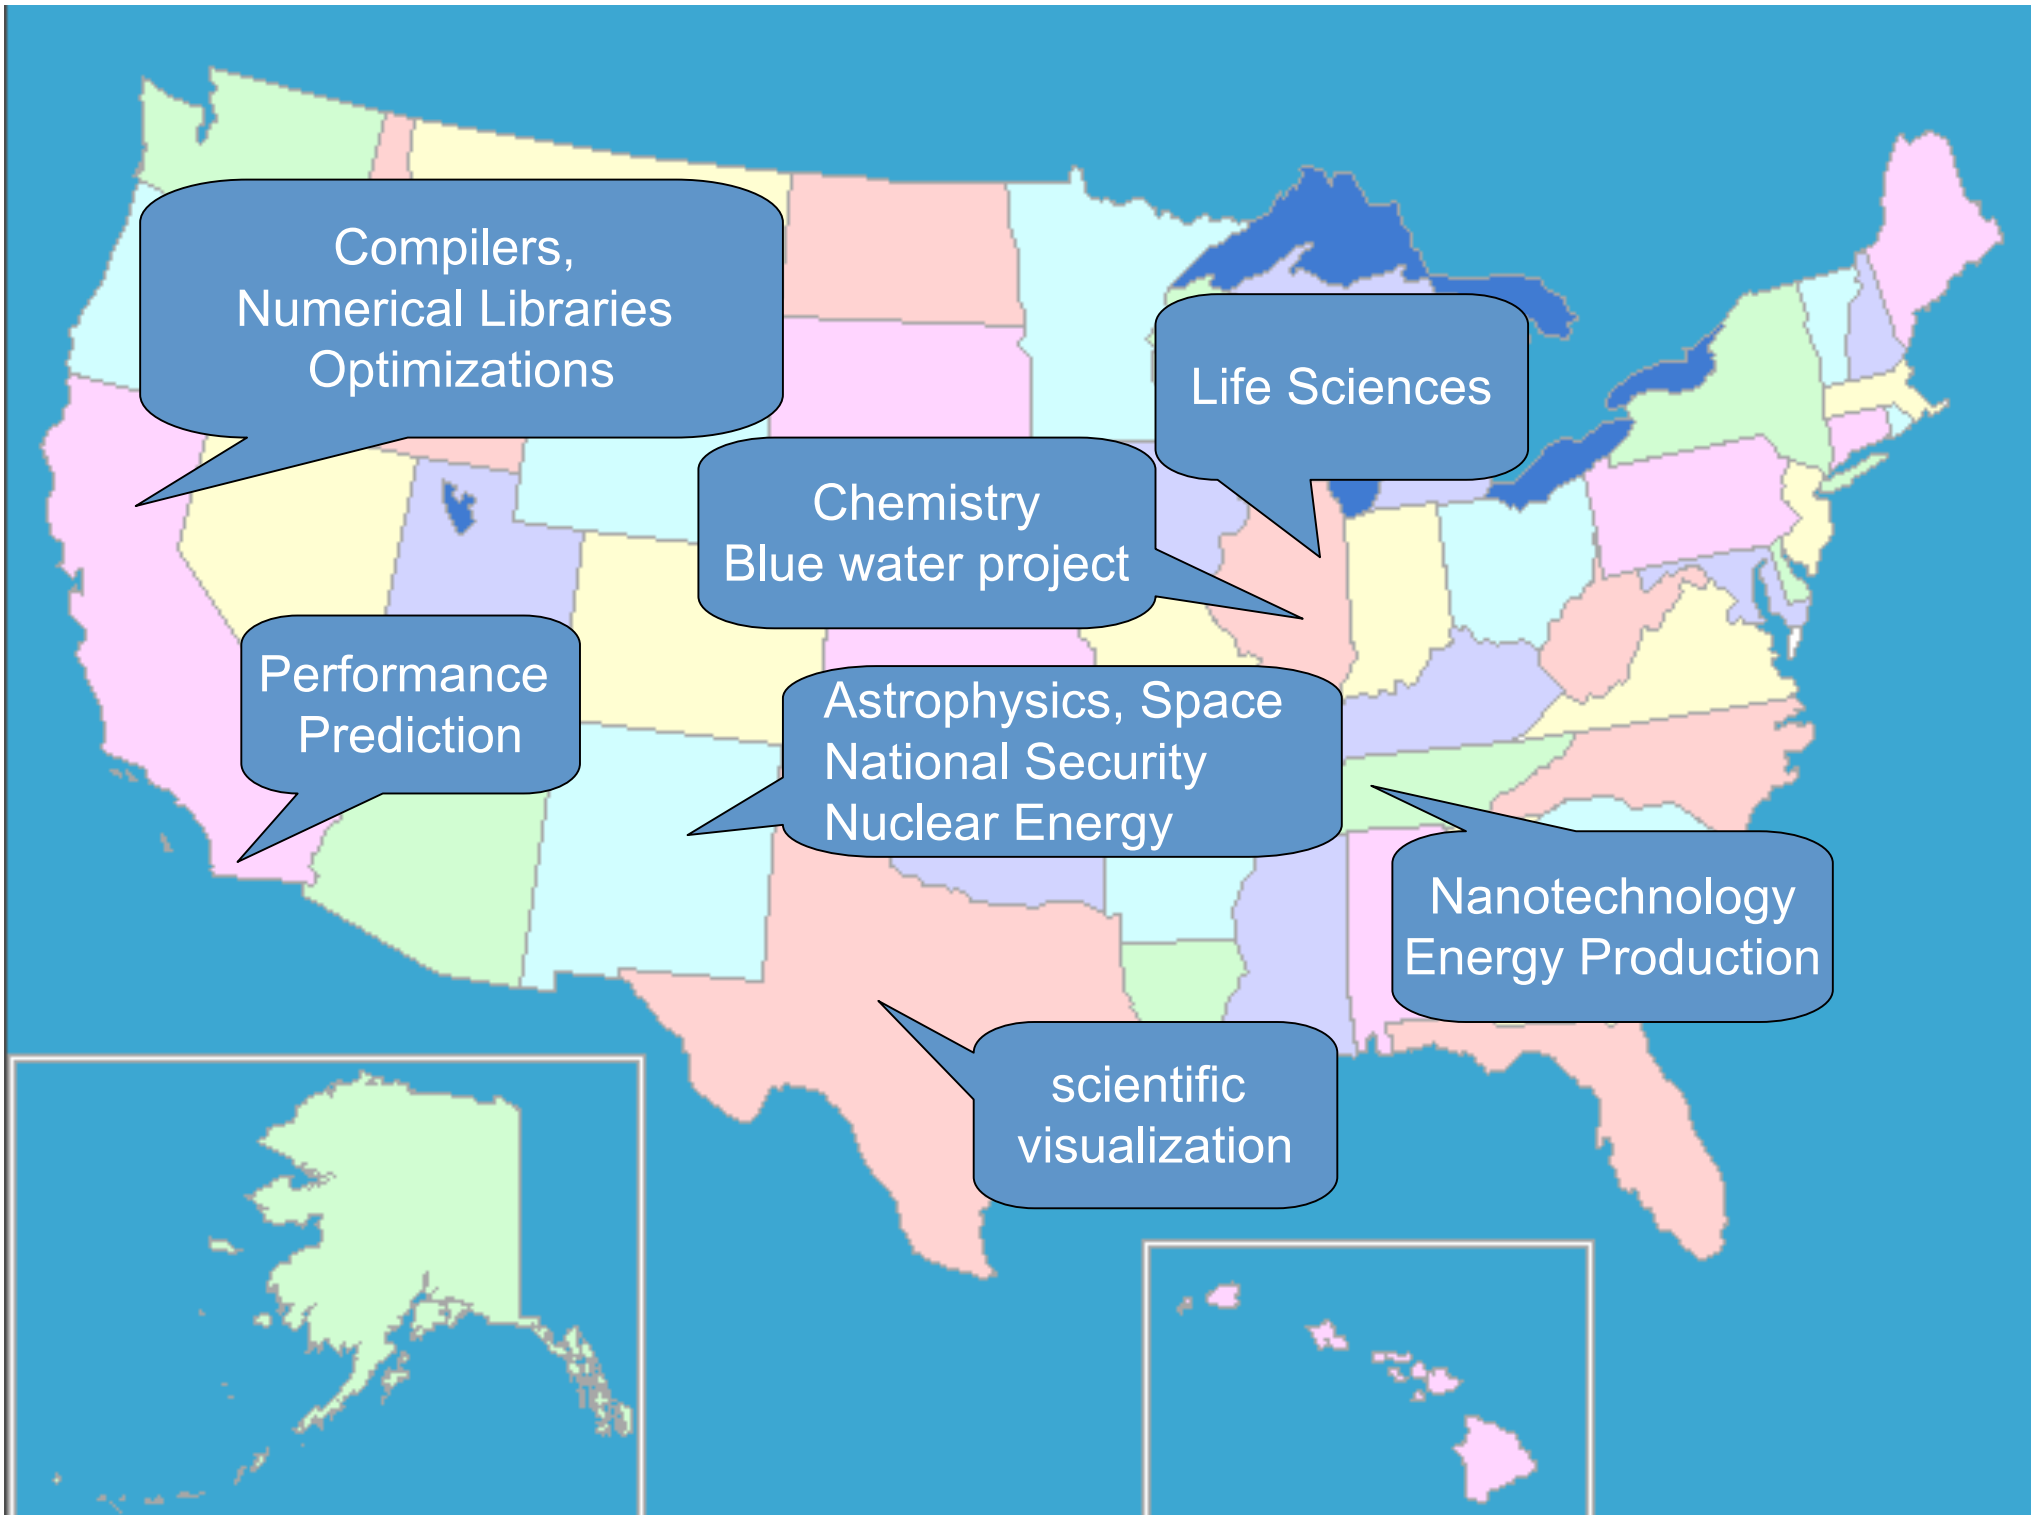


NCSA

Argonne

Oak Ridge
National Lab

TACC



Compilers,
Numerical Libraries
Optimizations

Life Sciences

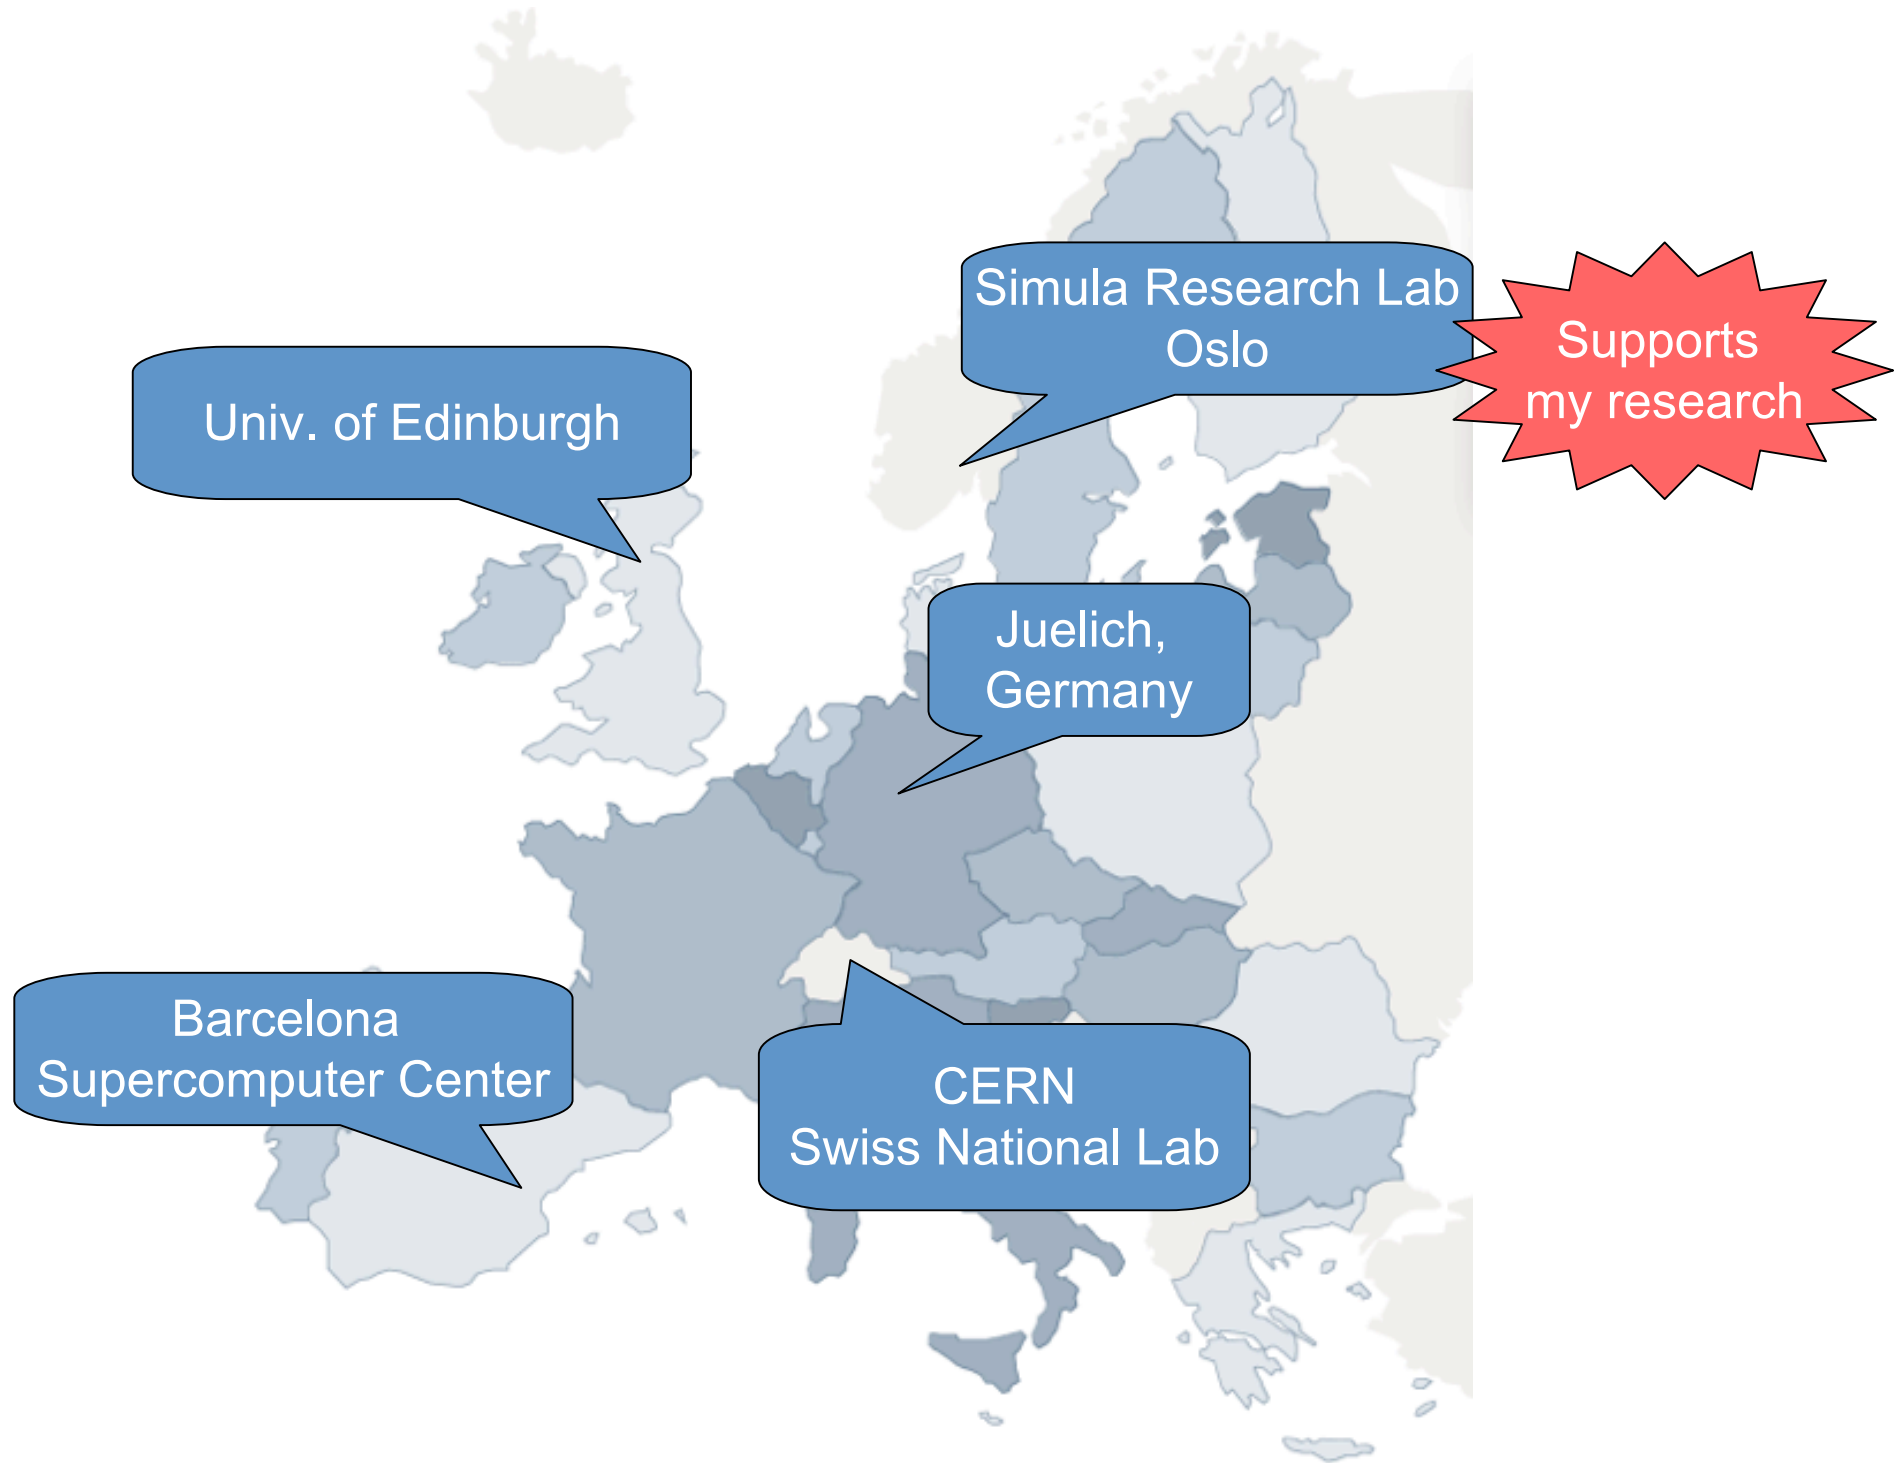
Chemistry
Blue water project

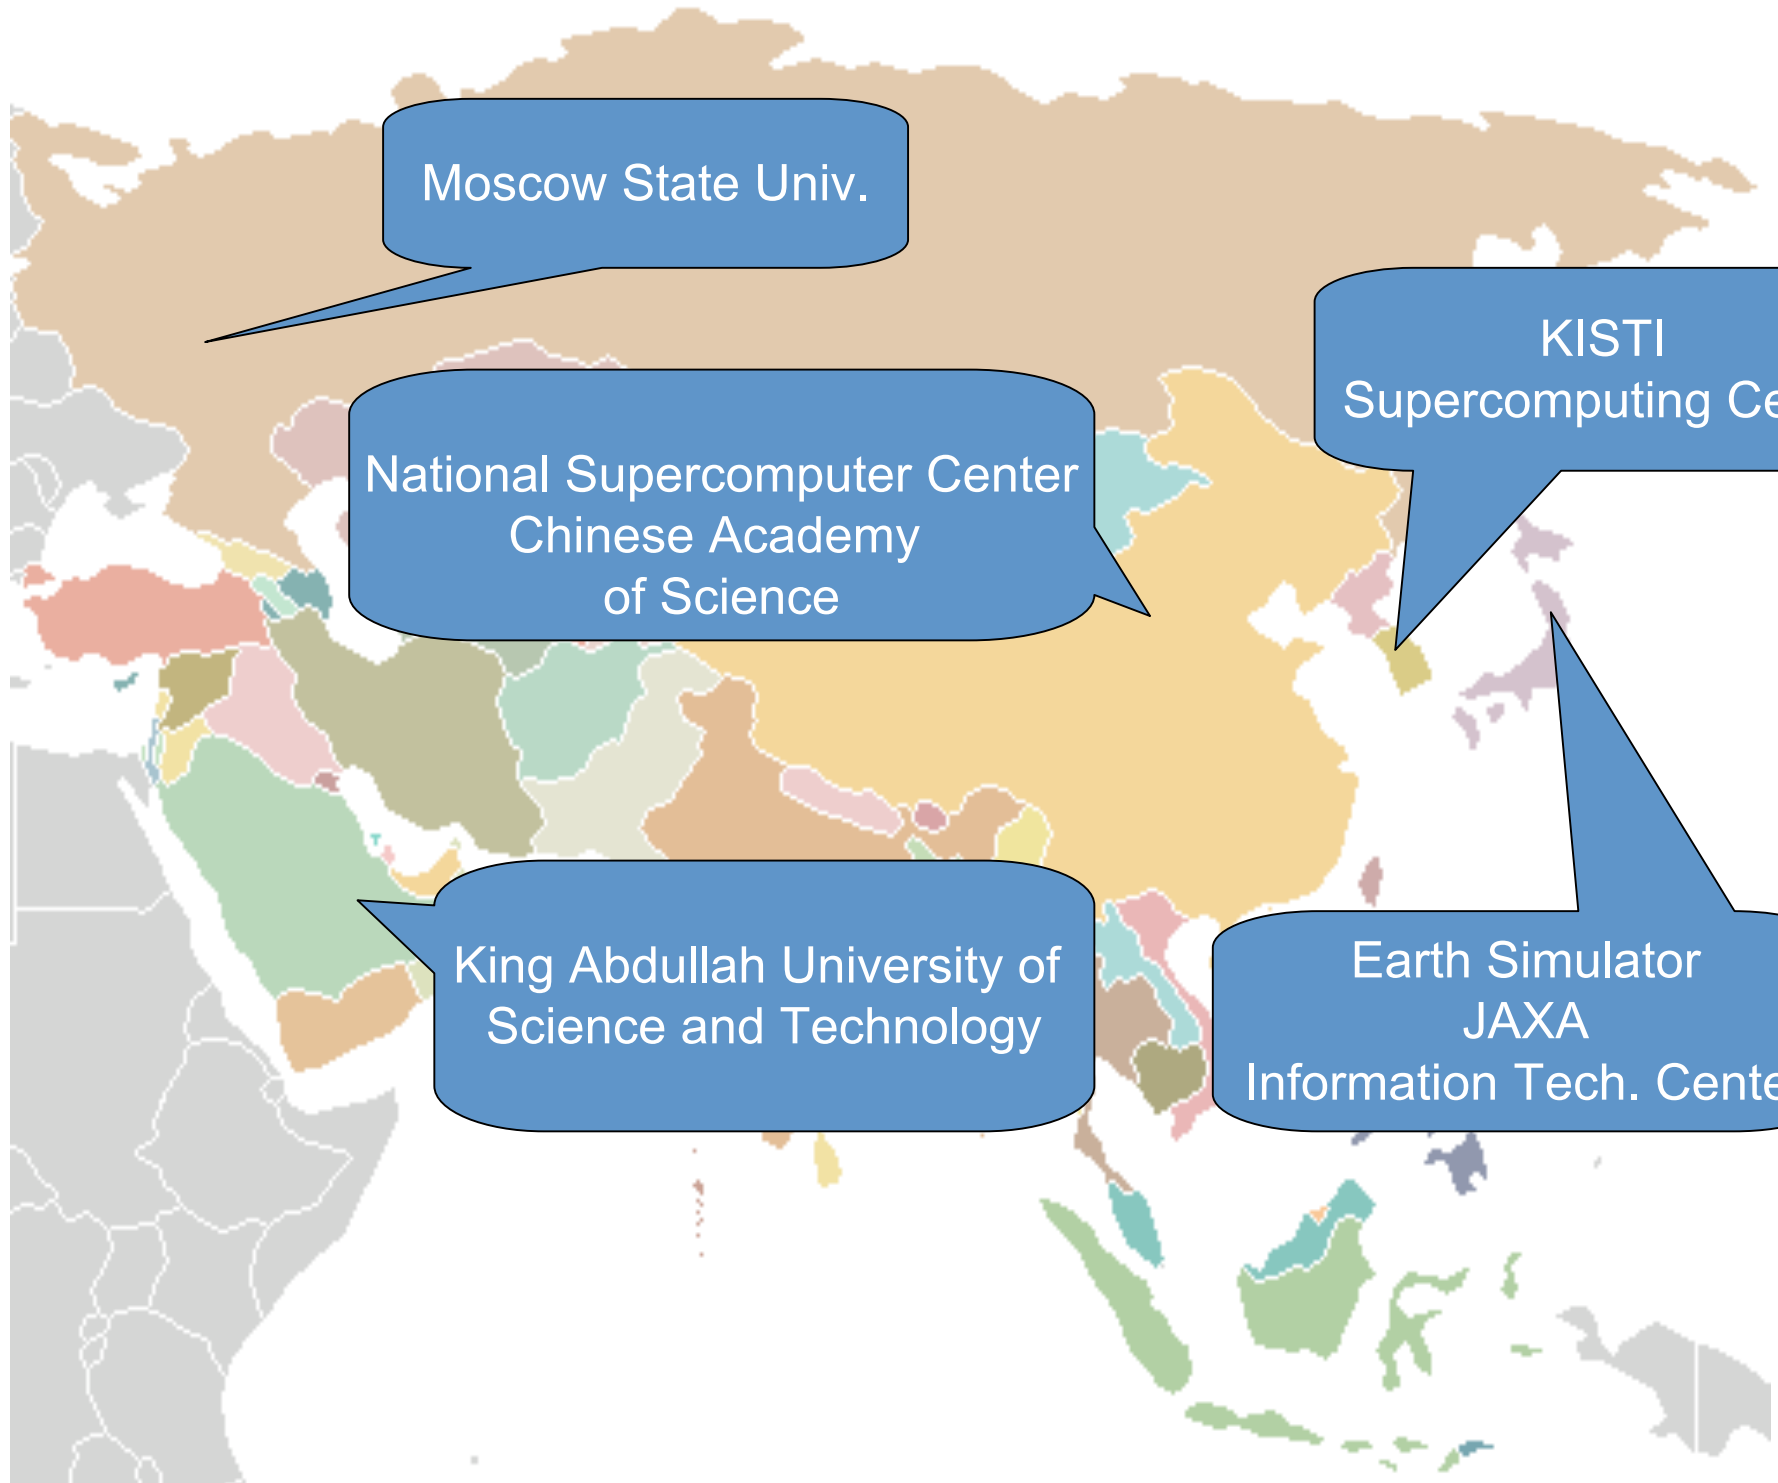
Performance
Prediction

Astrophysics, Space
National Security
Nuclear Energy

Nanotechnology
Energy Production

scientific
visualization





Moscow State Univ.

National Supercomputer Center
Chinese Academy
of Science

King Abdullah University of
Science and Technology

KISTI
Supercomputing Center

Earth Simulator
JAXA
Information Tech. Center

Links

Supercomputing Conference

<http://sc10.supercomputing.org/>

HPC wire Magazine

<http://www.hpcwire.com/>

Top 500 supercomputers:

<http://www.top500.org/>