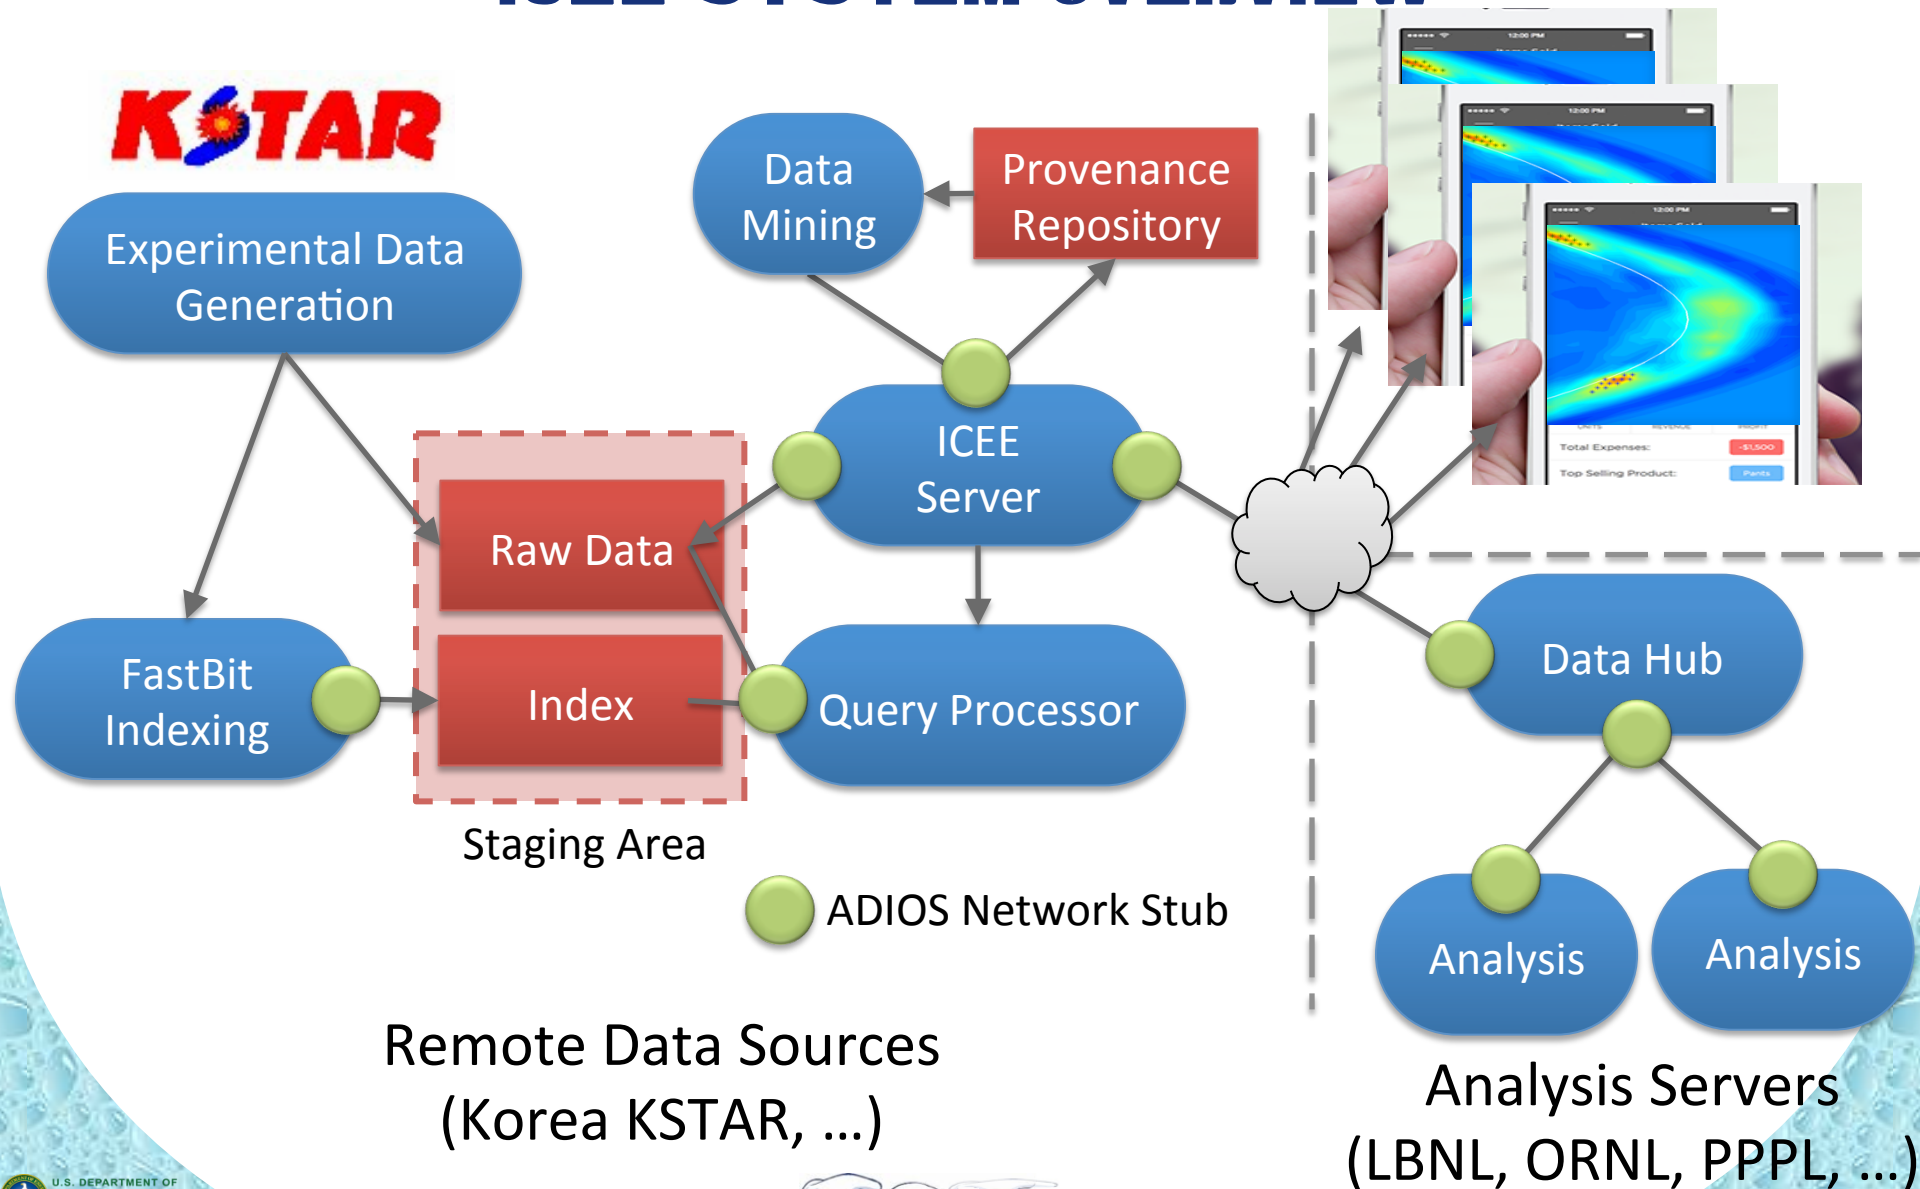


PROJECT OVERVIEW

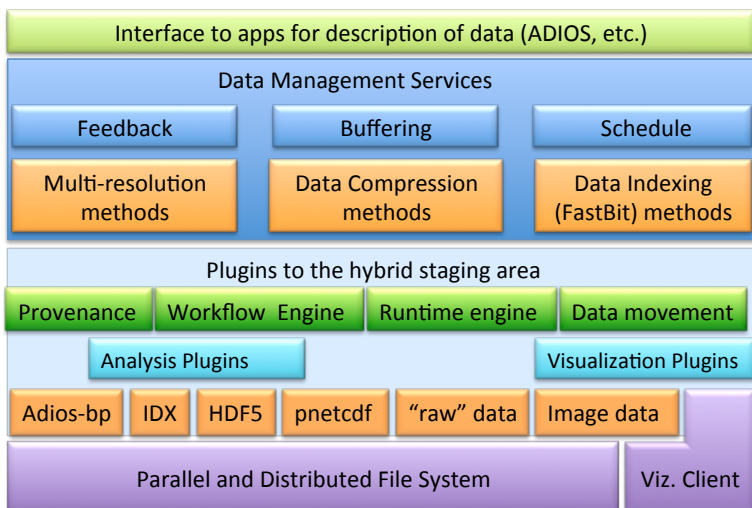
- **Challenge:**
 - Many large experiments produce massive amounts of data
 - Require immediate feedback to control the execution
 - But, the compute resources or experienced analysts are far away
- **Question:**
 - Can a distributed workflow respond in near real-time?
- **Approach**
 - Create a framework to enable efficient distributed analyses on extreme scale data
 - Select the most relevant data records through indexing (FastBit)
 - Reduce response time by keeping transient data in memory (ADIOS)
 - Increase usability with mobile device front-ends



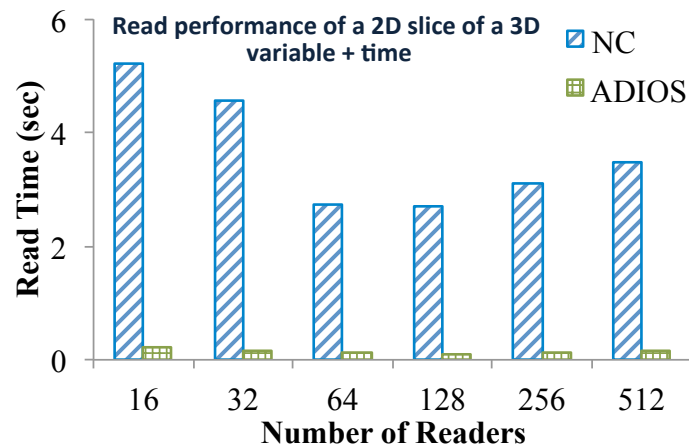
ICEE SYSTEM OVERVIEW



BUILDING BLOCK: AN I/O ABSTRACTION

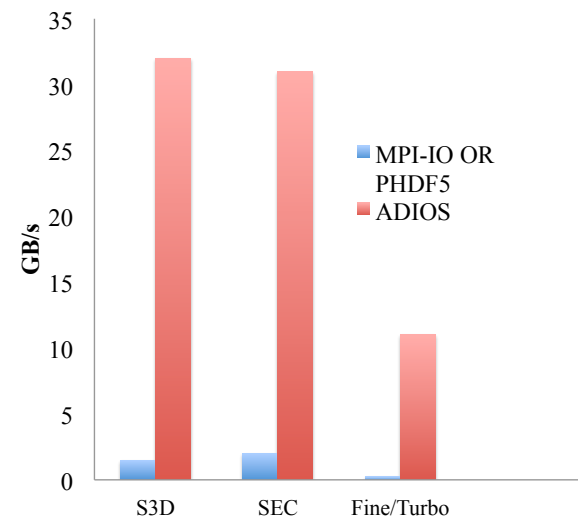


- An I/O abstraction framework
- Provides portable, fast, scalable, easy-to-use, metadata rich output
- Change I/O method on-the-fly
- Abstracts the API from the method
<http://www.nccs.gov/user-support/center-projects/adios/>



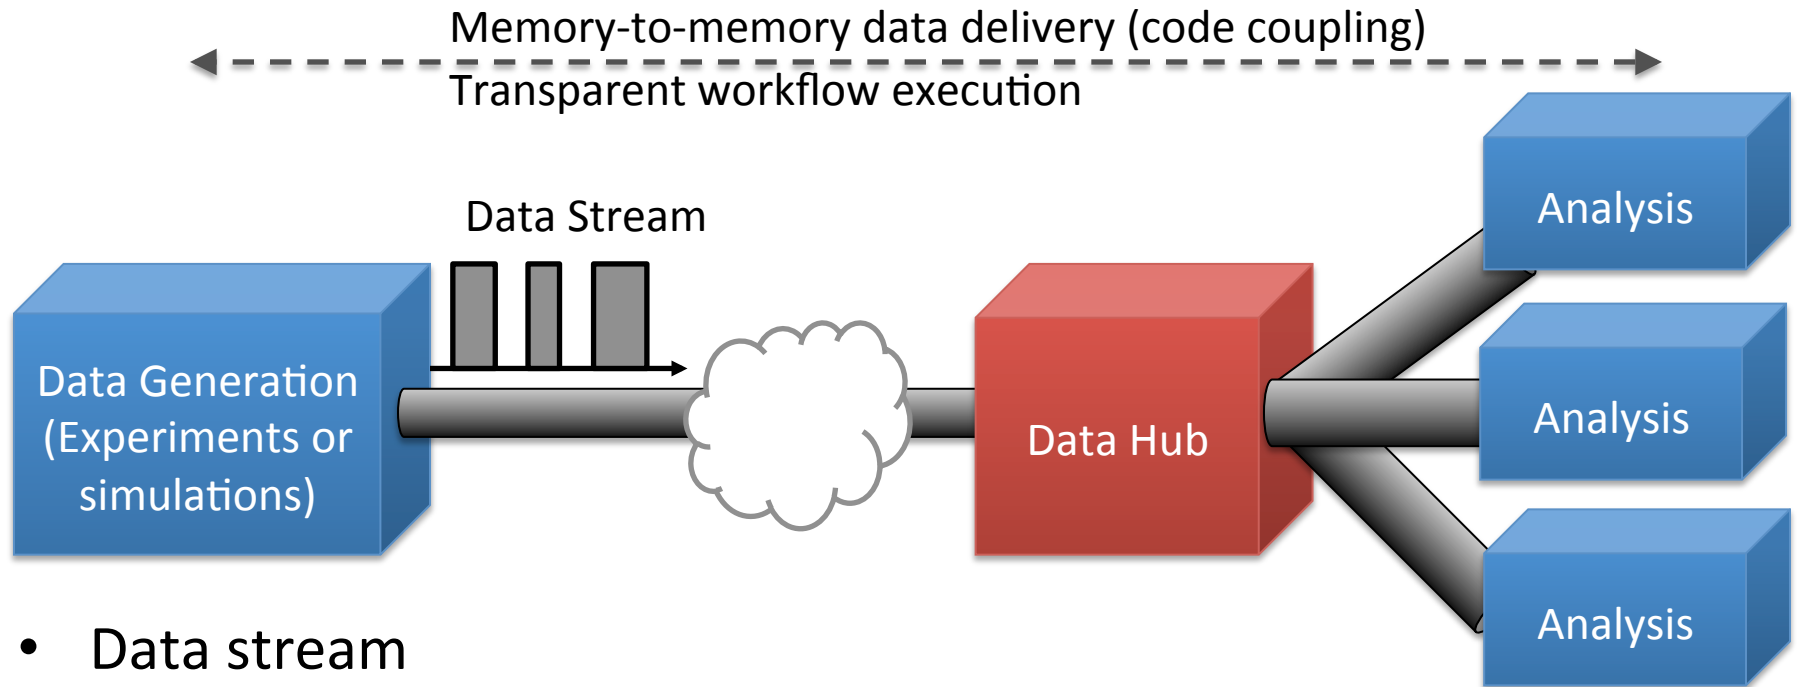
- Provides typical 10X performance improvement for synchronous I/O

I/O performance of the Combustion S3D code (96K cores), the SCEC PCML3D (30K cores), and the Fine/Turbo (4K cores) codes.



Simulations with and without ADIOS

BUILDING BLOCK: WIDE-AREA DATA STREAMING



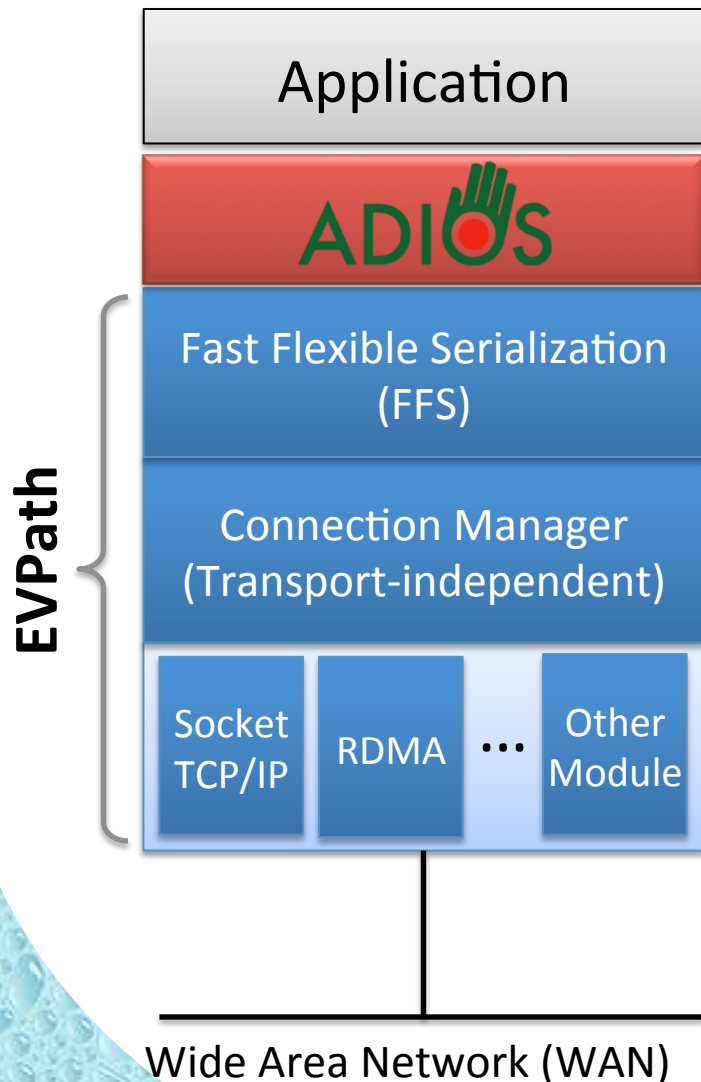
- Data stream
 - A limited windows of data is available for analysis
 - Stream could come from data in memory, which is much faster than reading disk
- Data from files
 - Any part of the data in a file is available for analysis
 - Reading data from file is much slower than reading data in memory

BUILDING BLOCK: ADIOS

- Integrating I/O system with workflow concept
 - Provide end-to-end data access pipeline
 - Support seamless data processing
 - Support stream-based scientific data process
- Driver applications
 - Fusion experiment data from KSTAR
 - Burst data generation (hundreds GBs per shot)
 - Need near real time support for running analysis codes
 - EPSI fusion simulation (PI: C.S. Chang)
 - Integrating large number of coupling analysis codes is challenging
 - Require flexible methods for various coupling execution scenarios



ICEE ARCHITECTURE




- Built on top of ADIOS
 - Designed for remote staging through WAN communication
 - Using EVPPath package developed by collaborators in Georgia Tech
 - Support different transportation methods
- Advantages
 - FFS (Fast Flexible Serialization): Data transformations for high performance applications
 - Easy to build an overlay network
 - CoD (C-on-Demand) features to process data while moving




EXAMPLE WORKFLOW: BLOBS IN FUSION

- Blobs are intermittent bursts of particles near the edge of the confined plasma
 - ➔ Driven by turbulence
- Blobs are bad for fusion performance because they:
 - Transport heat and particles away from the confined plasma
 - May damage the main chamber wall
 - Lead to increased levels of neutrals and impurities, bypassing control mechanisms



EXAMPLE WORKFLOW: BLOBS IN FUSION

- Need to understand blob physics to predict ITER performance
- Experimental diagnostics measure blobs on $\sim\mu\text{s}$ timescales, with significant radial resolution
 - Typical measured blob lifetimes 10's μs
- Diagnostics measuring blobs generate large amounts of data in short periods of time (GB's/s)
- Near real-time understanding of blobs at remote analysis site is preferred for efficient collaboration

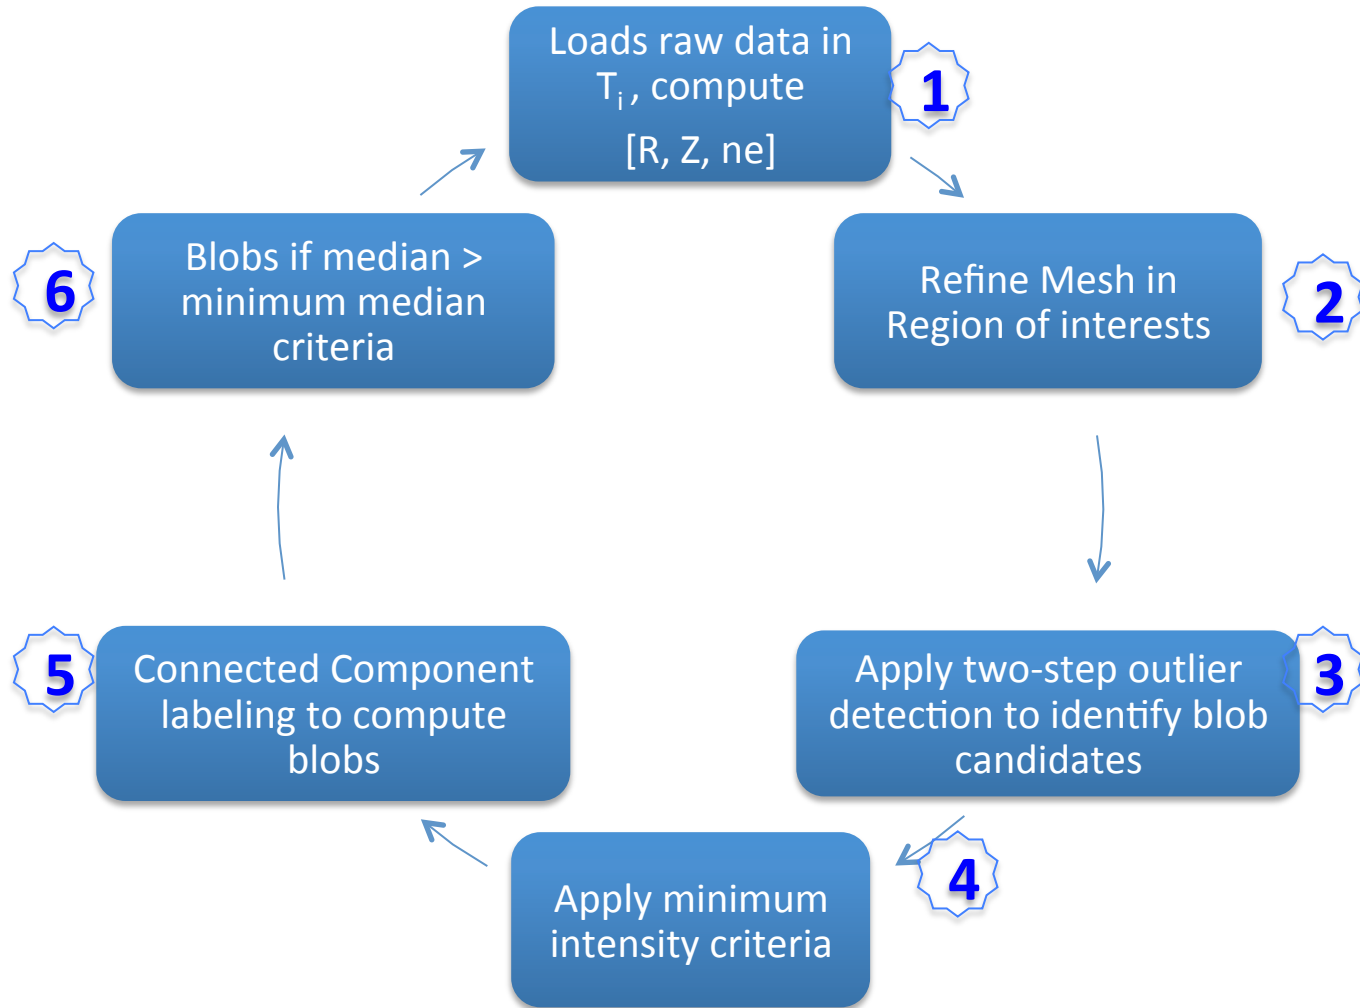


BLOB DETECTION ALGORITHM

- We present an outlier detection algorithm for efficiently finding blobs in numerical simulations and fusion experiments
 - Two-step outlier detection with various criteria after normalizing the local intensity
 - Leverage a fast connected component labeling method to find blob components based on refined triangular mesh
- Contributions:
 - A new method will not miss detection of blobs in the edge of the region of interests compared to recently developed contouring method
 - This is a first research work presenting a real-time blob detection approach to achieve blob detection in a few milliseconds
 - Demonstrate the possibility of supporting in-shot-analysis and between-shot-analysis for fusion experiments



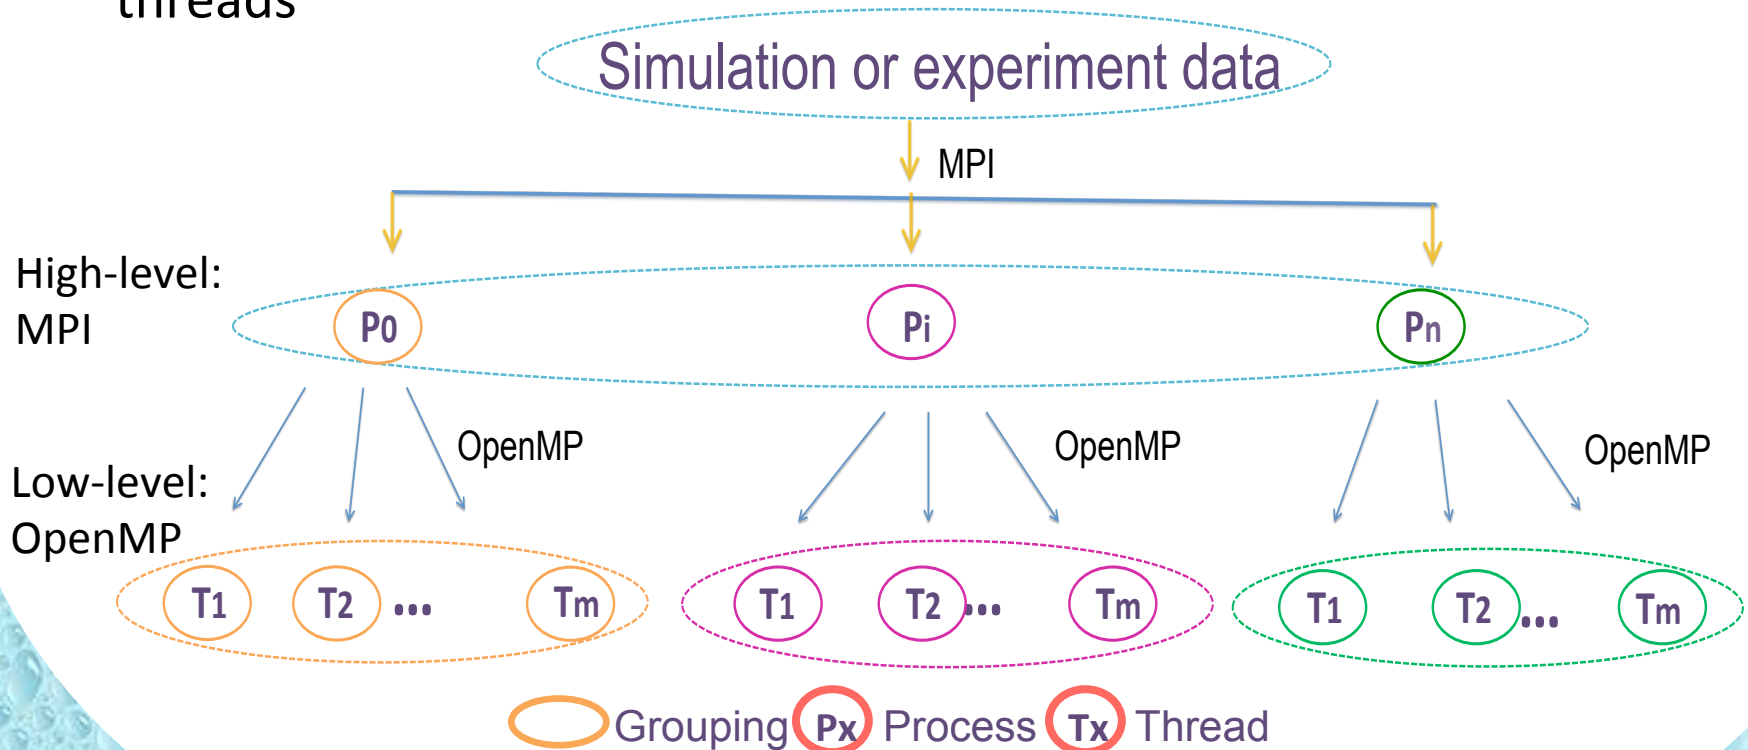
AN OUTLIER DETECTION ALGORITHM FOR FINDING BLOBS IN PLASMA





PARALLELIZATION OF BLOB DETECTION APPROACH


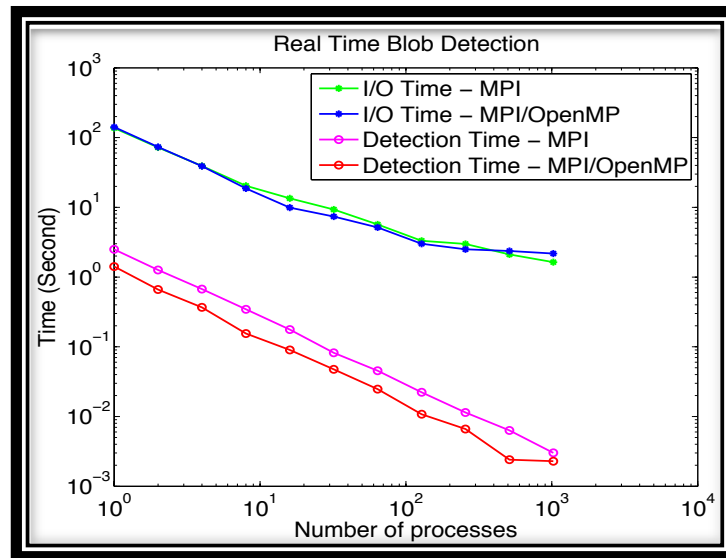
A hybrid MPI/OpenMP parallelization on many-core processor architecture:

- High-level: use MPI to allocate n processes to process each time frame
- Low-level: use OpenMP to accelerate the computations with m threads





An example of the blob detection in two time frames



Blob detection time, I/O time and speedup varying number of processes






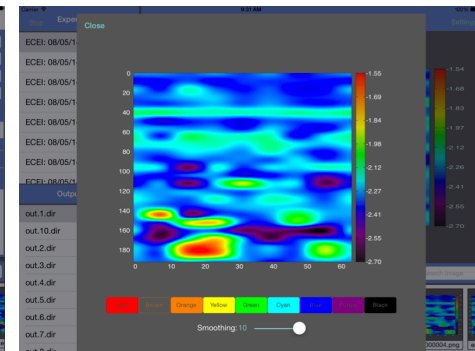
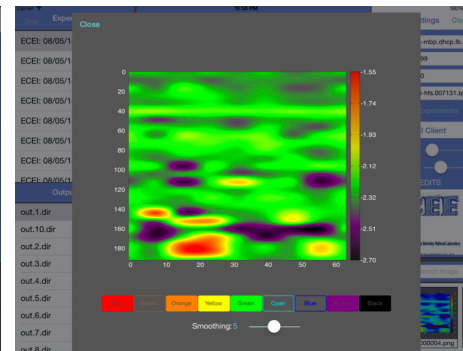
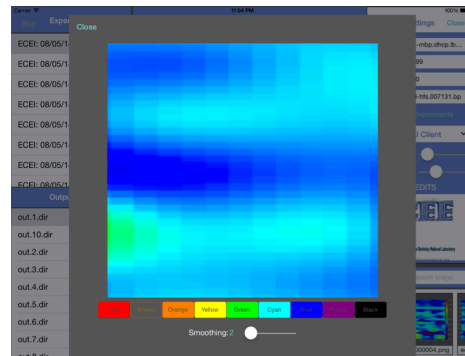
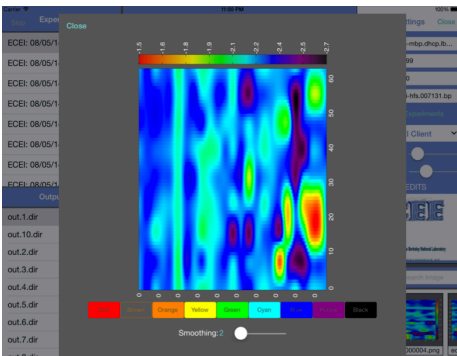

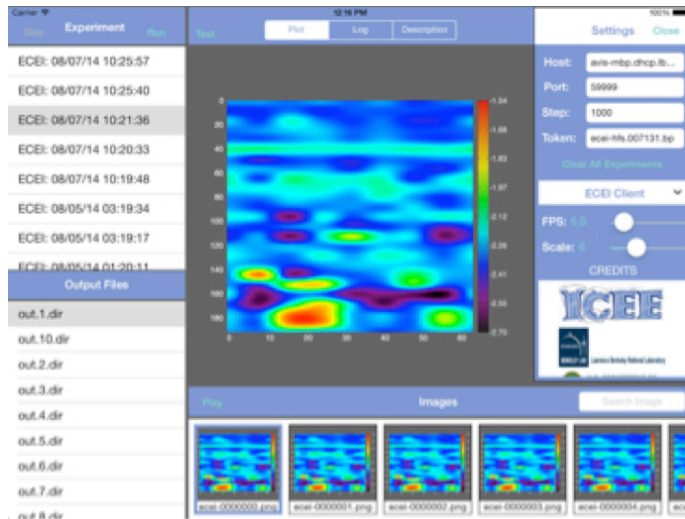
MOBILE ACCESS TO REMOTE WORKFLOW MONITORING

- To explore tablet mobile computing support in collaborative environment
 - Remote monitoring of data analysis over ICEE framework
 - Mobile distributed access to the monitoring for the data analysis, through tablets with iOS (i.e. iPad)
 - Tablet access to monitor the workflow
 - Tablet access to analysis results
 - Small computation and/or image manipulation on the tablet
 - Plan to study exploration of mobile tablet computing in science
 - Exploring mobile tablet computing as a computing resource for data analysis
 - Exploring mobile tablet computing as a collaborative tool



Main screen showing the output of experiment runs and a settings popup menu.

Edit Image screen allows image rotation, zooming, and panning by finger gesture, as well as color coding optimization by selection and graphical smoothing with a sliding scale.



GNU plot

Our plot on iPad

Before and after smoothing

SUMMARY AND PLANS

- ICEE demonstrated the feasibility of supporting near real-time analysis workflows
 - Takes advantage of a new method in I/O abstraction for WAN data movement
- Hybrid parallel data analysis engine for feature detection
 - SC14 paper in preparation
- Plans for next year
 - Connecting to multiple data sources
 - Optimize data movement
- Next questions
 - Could we support multiple real-time workflows simultaneously?
 - Could we automatically adapt to the varying network and compute resources?

