Online data analysis of molecular dynamics simulations for exascale: An NWChemEx+CODAR collaboration

Hubertus van Dam (BNL), Jong Yol Choi (ORNL), Tashin Kurc (SBU), Abid Malik (BNL), Gyorgy Matyasfalvi (BNL), Sungsoo Ha (BNL), Klaus Mueller (SBU), Line Pouchard (BNL), Li Tang (LANL), Dingwen Tao (UA), Cong Xi (SBU), Wei Xu (SBU), Shinjae Yoo (BNL)

BROOKHAVEN NATIONAL LABORATORY

ECP EXASCALE COMPUTING PROJECT

CODAR

NWChemEx

U.S. DEPARTMENT OF ENERGY

CODAR: Center for Online Data Analysis and Reduction

System Attributes	NERSC Was	OLCF Going	ALCF	NERSC Now	OLCF Now	ALCF Now	NERSC Upgrade	OLCF Upgrade	ALCF Upgrade
Name	Edison	Titan	Mira	Cori	Summit	Theta	Perlmutter	Frontier	Aurora
Peak perf. (PF)	2.6	27	10	30	184	12	> 90	> 1,500	> 1,000
System memory (TB)	357	710	768	1,389	2,731	824	?	?	> 5,120
Filesystem	7.6 PB 168 GB/s	32 PB 1,000 GB/s	26 PB 300 GB/s	28 PB 744 GB/s	250 PB 2,560 GB/s	10 PB 210 GB/s	30 PB > 4,096 GB/s	> 500 PB > 5,000 GB/s	?

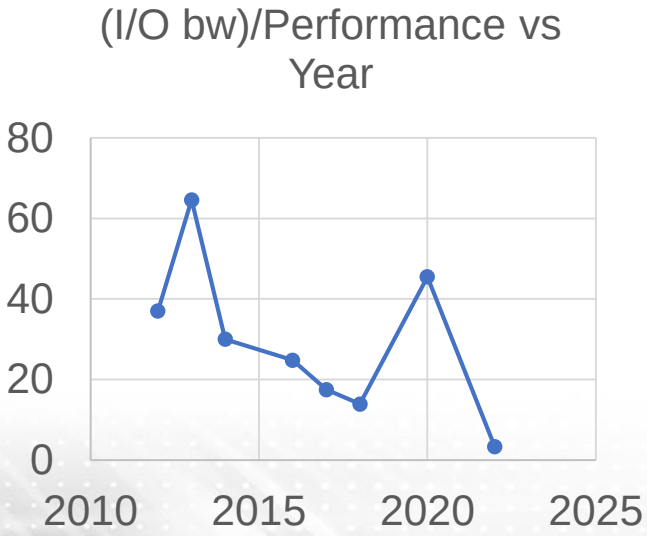

- Compute-Data gap is a challenge for exascale
- I/O bandwidth is not keeping up with peak performance
- **Solutions:**
  - **Data (de)compression**
    - Substitute cycles for I/O
  - **Online data analysis**
    - Extract the science
    - Do not store the data

U.S. DEPARTMENT OF ENERGY

ECP EXASCALE COMPUTING PROJECT

CODAR

NWChemEx

BROOKHAVEN NATIONAL LABORATORY

CODAR with NWChemEx

- Engagement 1
  - Adaptive Sampling of Molecular Dynamics Trajectories
  - Goal: Sampling important MD trajectories
- Engagement 2
  - Workflow Performance Analysis
  - Goal: Enabling workflow level performance analysis from performance traces

NWChem – Parallel Open Source Computational Chemistry

- NWChem started in the early 1990s [1,2].
- It supports a wide range of functionality:
  - Density Functional Theory (most popular)
    - For molecules
    - For materials
  - Coupled Cluster (very accurate, but expensive)
  - Molecular Dynamics for Biomolecules
  - And other methods and many properties
- NWChemEx is the successor under development as part of the exascale computing project
- NWChem availability:
  - Installed on all major DOE facilities (Oak Ridge, Argonne, NERSC)
  - Binary versions through Debian, and Ubuntu Linux distributions
  - Source code available from GitHub at https://github.com/nwchemgit/nwchem

NWCHEM
HIGH-PERFORMANCE COMPUTATIONAL CHEMISTRY SOFTWARE

[1] High Performance Computational Chemistry Group, Pacific Northwest Laboratory. Richland, Washington 99352. USA, NWChem. a computational chemistry package for parallel computers, version 1.0, 1994.
[2] http://www.nwchem-sw.org, [07/05/2018]

Use Case: NWChem MD

- MD simulates
  - Conformational changes of biomolecules
  - Transport processes
  - Try to find statistics and key points
- Processes involve
  - Many atoms (~1 million)
  - Long time scales (μs)
  - Many time step (~1 billion)
- Shown
  - Transmembrane Calcium channel
  - Regulated by pH
  - Plays a role in plant draught response
- Challenges
  - **~32PB Trajectories**

pH 8

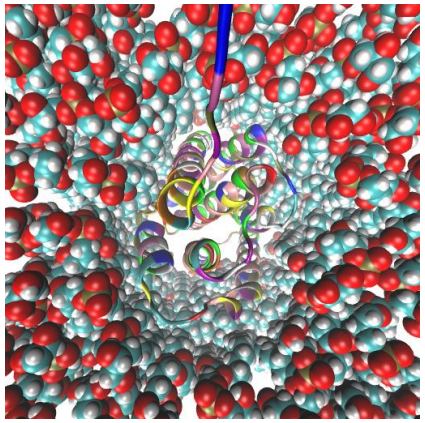

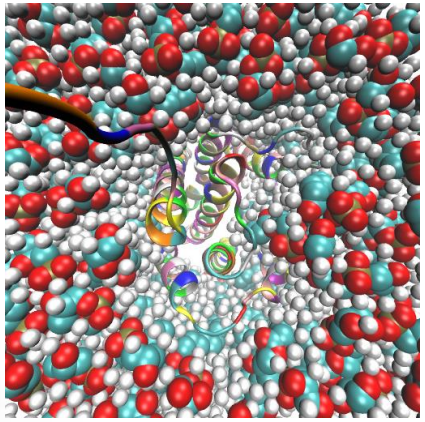

pH 6

Adaptive Sampling Motivation / Challenges

- Motivation
  - Not all parts of the trajectories are interesting
  - Sampling only significant phase change events
- Challenges
  - The atoms vibrate significantly
  - Difficult to compress

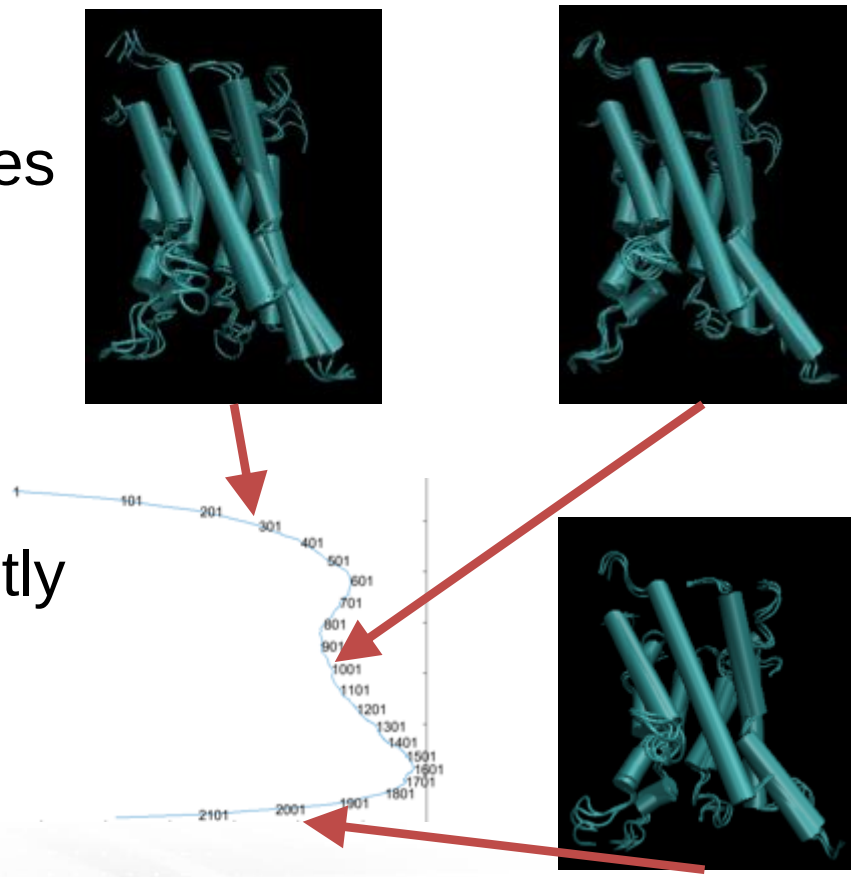

Low dimensional manifold projection of different state of MD trajectories

Batch Approach

- Manifold Learning for MD trajectories
  - Learn the intrinsic relations among time points
  - Less vibrational noise than original space
  - Effectively compress high dimensional coordinates
- Important / Interesting event detection
  - L2 normalized gradient changes
    - Not focus on how fast atoms moves
    - But focus on docking or conformational changes
  - Mean filter for smoothed trajectories
    - Gradient based changes can be sensitive local fluctuations
    - Smoothed manifold helps to stabilize trajectories

A Case Study

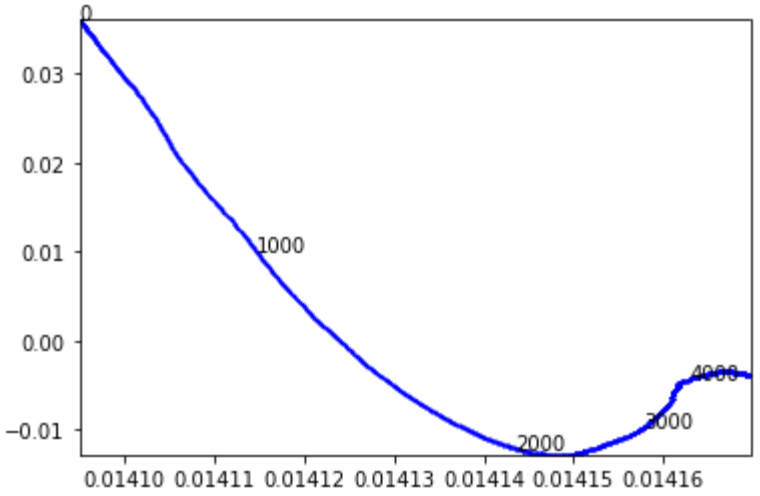

Manifold Trajectories

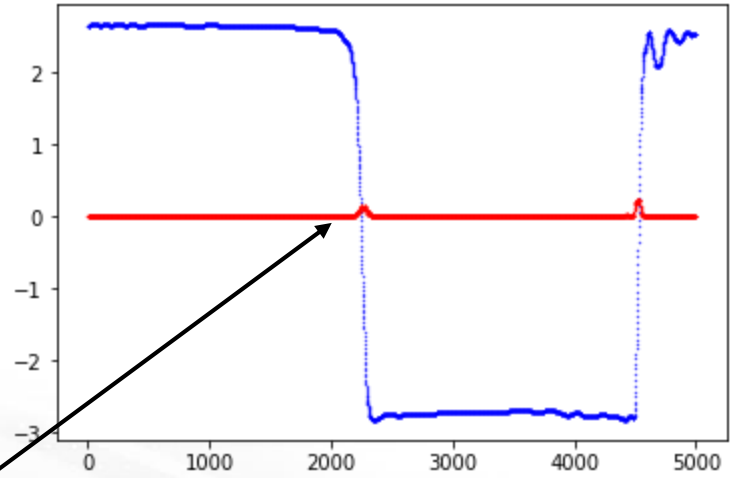

Smoothed Angles

Importance

A Case Study – 5 Samples (Batch)

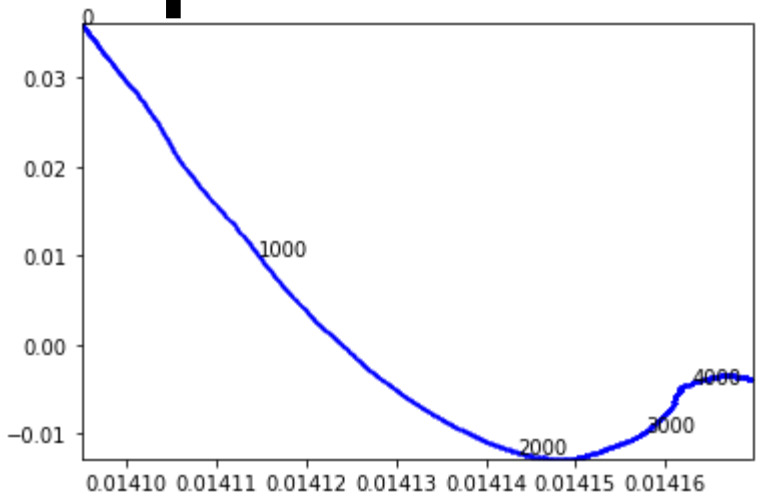

Manifold Trajectories

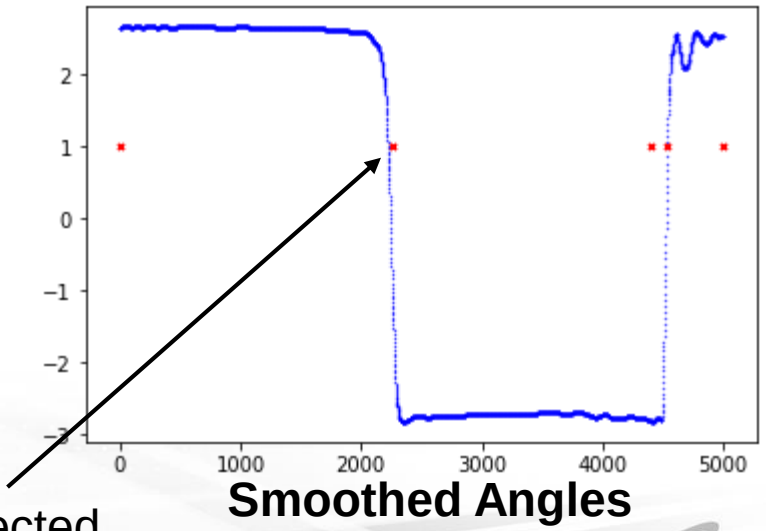

Smoothed Angles

Selected event

Online Adaptive Approach

- Manifold Learning for MD trajectories
  - Matrix Sketch of MD trajectories
- Important / Interesting event detection
  - Weighted Reservoir Sampling of gradient changes

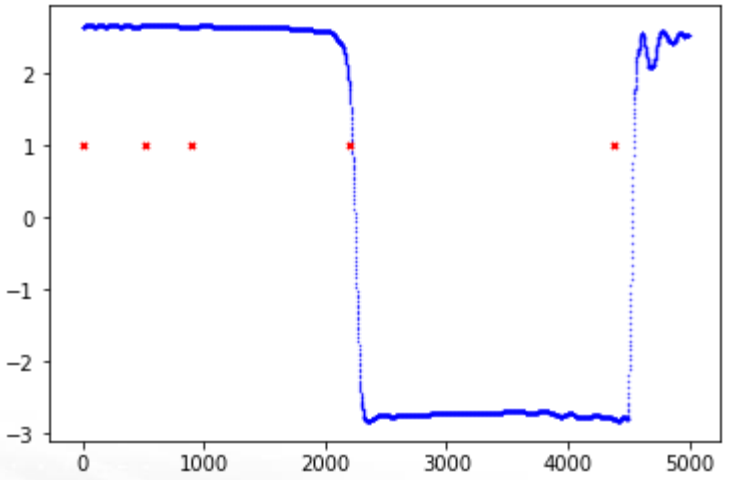

Chimbuko: Workflow Performance Analysis on NWChem

- Volume of performance data is too large
  - Difficult to do manual inspection
  - Not all data can be saved
- Performance data reduction is required
  - Focus on performance anomalous events
  - Store such events for provenance
  - Provide visual interface to performance anomaly events

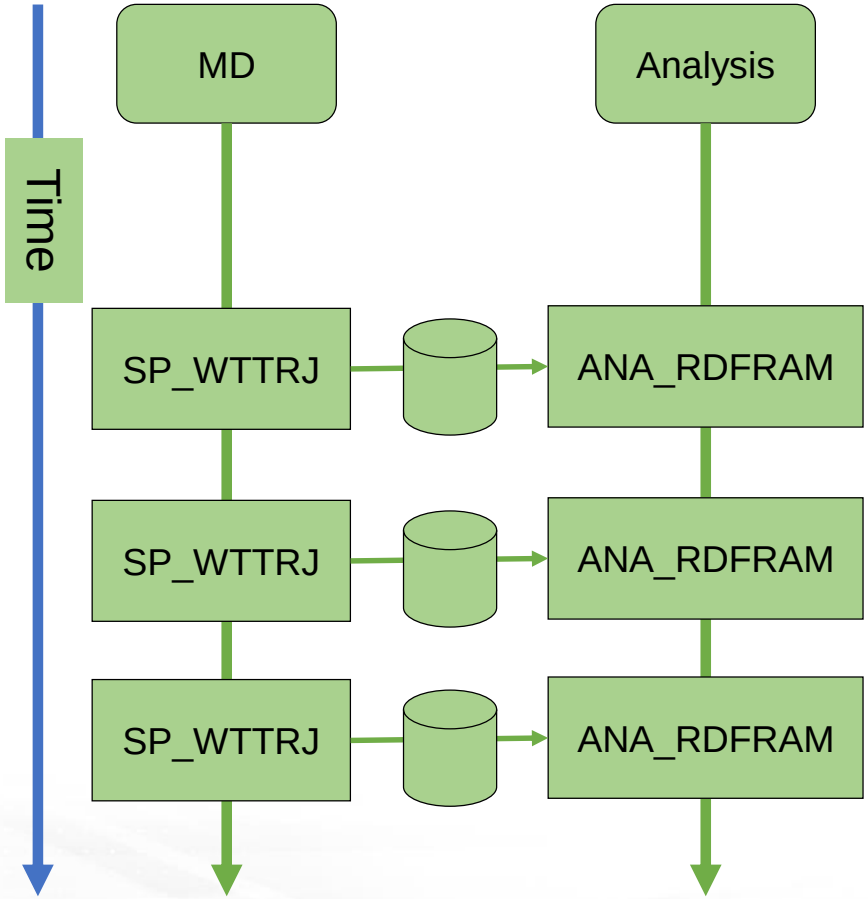

Chimbuko Architecture

- Chimbuko monitors workflow execution, extracts provenance and visualization data
- ADIOS provides data streaming among workflow components (blue line)
- TAU provides performance metrics for instrumented components
- SOSFlow aggregates the performance data (green lines)
- Trace data is dynamically analyzed to detect anomalies (solid red line)
- Selected metadata and trace data is stored and visualize (dashed red lines)

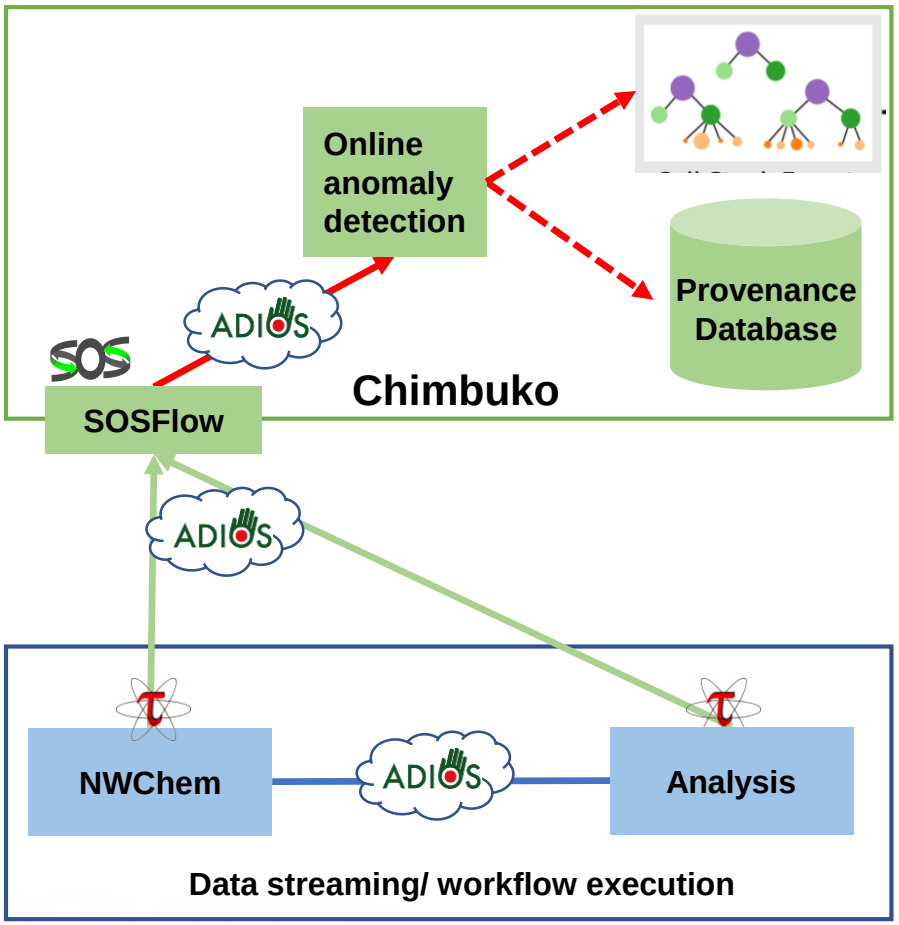

Anomaly Detection of Function Executions

- Detecting unexpected runtime behavior in HPC clusters
- Trace events: sequences of execution activities
  - Function entry/exit in each HPC node
  - Message passing between different HPC nodes

The offline use case used LAMMPS

```
event:entry, function:compute, timestamp:200, node-id:0
|  event:entry, function:reverse_comm_pair, timestamp:210, node-id:0
|  |  event:entry, function:MPI_Send, timestamp:215, node-id:0
|  |  |  event:send, destination-node:2, message: ...
|  |  event:exit, function:MPI_Send, timestamp:216, node-id:0
|  |  ... ...
|  event:exit, function:reverse_comm_pair, timestamp:280, node-id:0
|  ... ...
event:exit, function:compute, timestamp:400, node-id:0
```

Trace events inside one execution of function compute in an HPC node

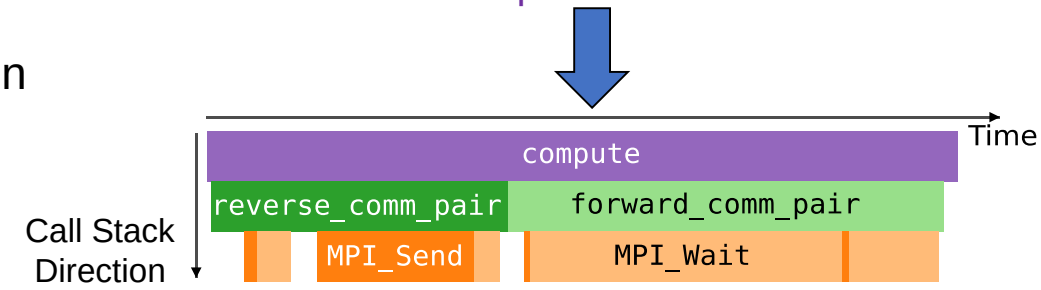

Functions invoked during that execution of compute

Call Stack Tree (CSTree)

- Call stack tree representation
  - Node: an execution
  - Link: invocation relations
  - Node weight: execution time
- Multiple executions => Call stack forest
- Anomalous behavior detection => Anomalous tree structure detection
  - A big node: long execution
  - Lots of children: large loop

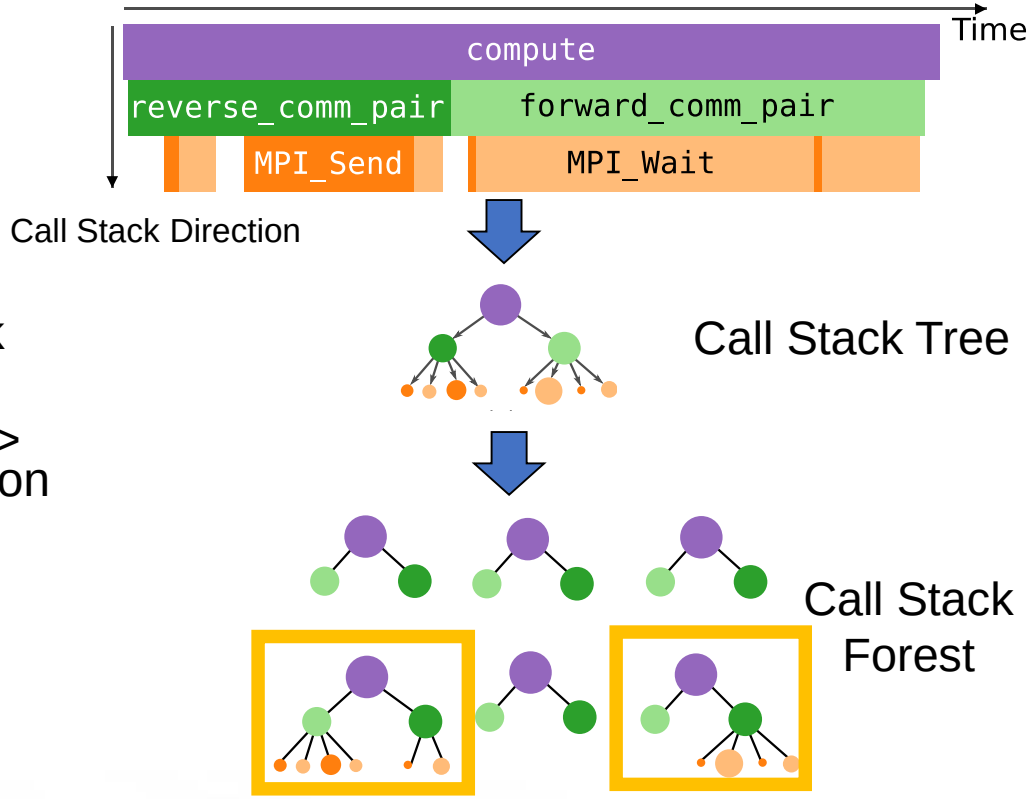

Online Architecture

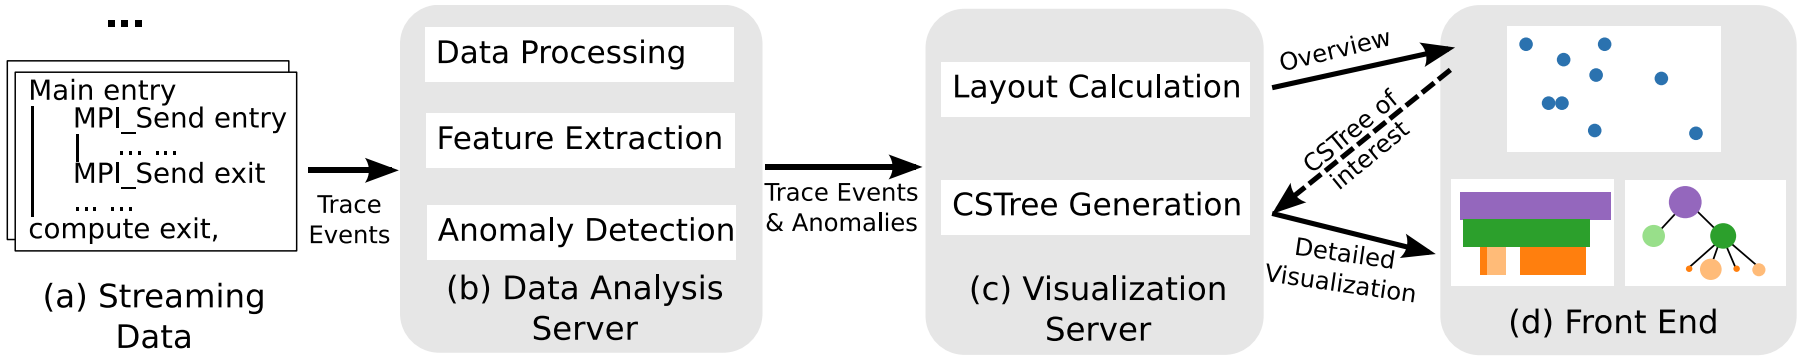

- Streaming trace events (a)
- Data analysis server for anomaly detection (b)
- Visualization server for CSTree exploration (c)
- Visual interface (d)

Chimbuko – Online Visualization

Acknowledgements

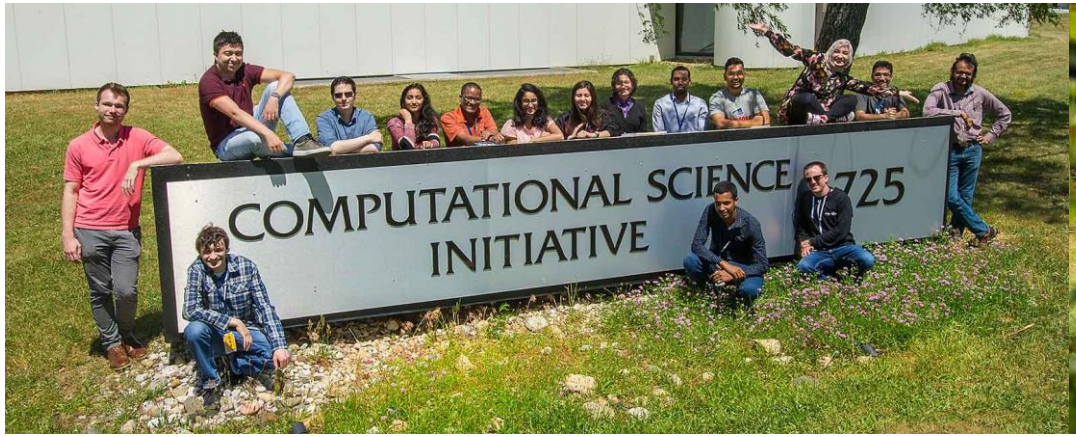

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

Thank You

ECP
EXASCALE COMPUTING PROJECT

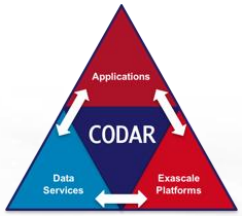

NWChemEx

U.S. DEPARTMENT OF ENERGY

BROOKHAVEN NATIONAL LABORATORY

Execution time without/with instrumentation

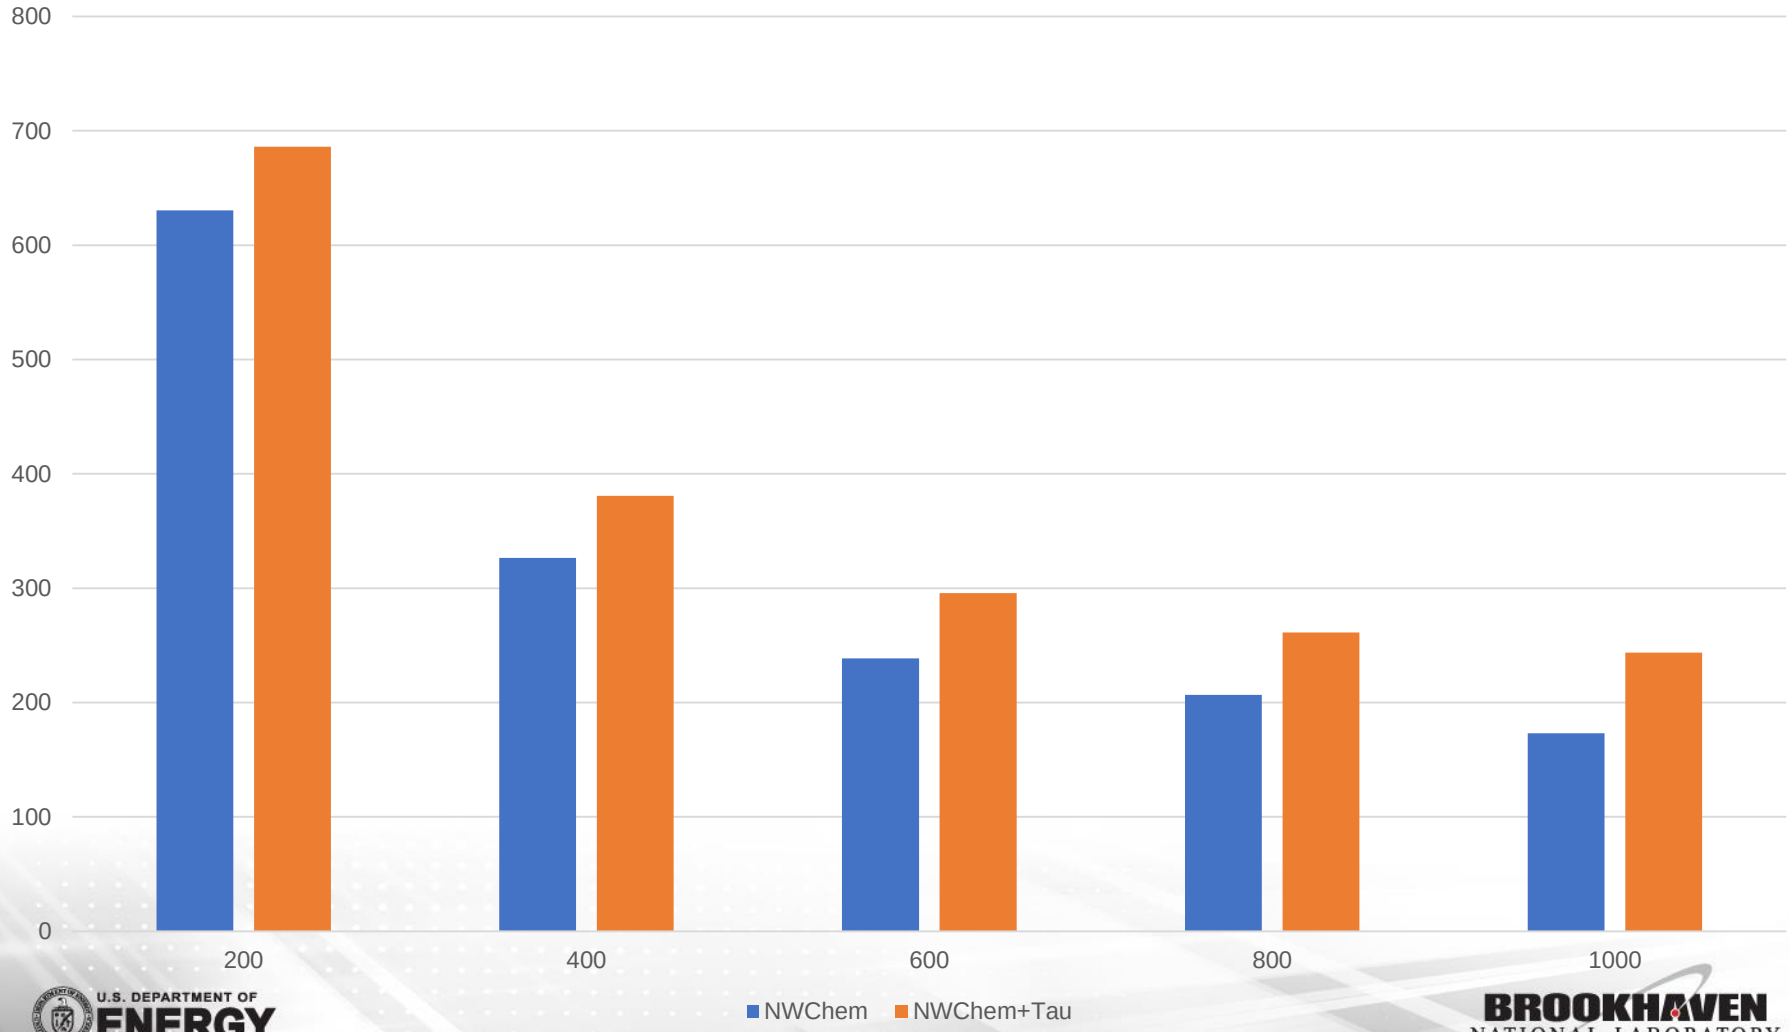

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov