WHPC Lightning Talks

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

Programme Committee

- Elsa Gonsiorowski, Lawrence Livermore National Laboratory, USA
- Raquell Holmes, Improvscience
- Elizabeth Bautista, NERSC, Lawrence Berkeley National Laboratory, USA
- Jo Adegbola, Amazon Web Services, USA
- Mozghan Kabiri, University of Sheffield, UK
- Karen Devine, Sandia National Laboratory, USA
- Zhiling Lan, Illinois Institute of Technology, USA
- Hadia Ahmed, Lawrence Berkeley National Laboratory, USA
- Lavanya Ramakrishnan, Lawrence Berkeley National Laboratory, USA
- Debbie Bard, Lawrence Berkeley National Laboratory, USA
- Rosa Filgueira, EPCC, UK
- Danielle Sikich, Intel, USA
- Mahwish Arif, CAM, UK
- Baiou Shi, PSU, USA
- Neelofer Banglawala, EPCC, UK
- Catherine Schumann, Oak Ridge National Laboratory, USA
- Shubbhi Taneja, Sonoma State University, USA

Thanks to the mentors for volunteering!

FIRST EVER WOMEN-IN-HPC SUMMIT!

In partnership with Simon Fraser University

Call For Participation (Papers, Tutorials & Posters) Open!

https://womeninhpc.org/events/summit-2020

Multi-Resource Scheduling in HPC

Yuping Fan | yfan22@hawk.iit.edu
Illinois Institute of Technology

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING
CHANGING THE FACE OF HPC

Job Scheduling in HPC

- Decides when and where to execute jobs
  - Policies: FCFS, SJF
  - Examples: Slurm, PBS, Mesos
- Traditionally, focuses on CPU utilization

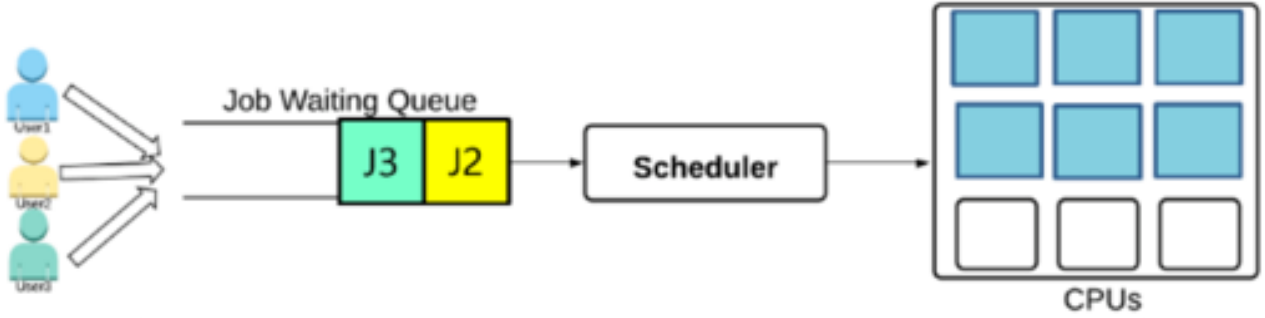

- Multiple Resources in HPC
  - Local resources: CPU, GPU, SSD
  - Shared resources: burst buffer, RAN
- **How can we efficiently use of multiple HPC resources?**

Existing Methods

- Naïve method: no optimization; first job in the queue can run when there are sufficient resources
- Optimize the utilization of one resource (e.g., CPU)
- Optimize the weighted sum of multiple resource utilization
- **Problem**: They are single-objective and can only provide one solution, but multiple optimal solution exist when scheduling multiple resources

BBSched: a Novel Multi-Resource Scheduling Framework

- Formulates the multi-resource scheduling problem in to multi-objective optimization problem
  - Maximize multiple independent objectives
    - Node utilization: f_1(x) = $\sum_{i=1}^{w} n_i \times x_i$
    - Burst Buffer utilization: f_2(x) = $\sum_{i=1}^{w} b_i \times x_i$
- Rapidly solves the problem by genetic algorithm
- Provides multiple scheduling options for system administrators
- Flexible to embrace new resources added to HPC systems

Moment Representation in the Lattice Boltzmann Method on Massively Parallel Hardware

Madhurima Vardhan| madhurima.vardhan@duke.edu
PhD Candidate | Biomedical Engineering
Duke University

Why is the lati… LBM algorithm an active area of research?

LBM is a widely used in CFD algorithm

- Explicit and second order accurate
- Highly scalable
- Handles complex geometries

However, LBM is

- Memory-expensive, 38 doubles per lattice site
- Imposes limit on the resolution

What do we propose to reduce memory requirement?

Adapt Regularized LBM

- First three order moments can accurately solve fluid dynamic equations[1]
- 10 doubles per lattice site (~74% reduction)

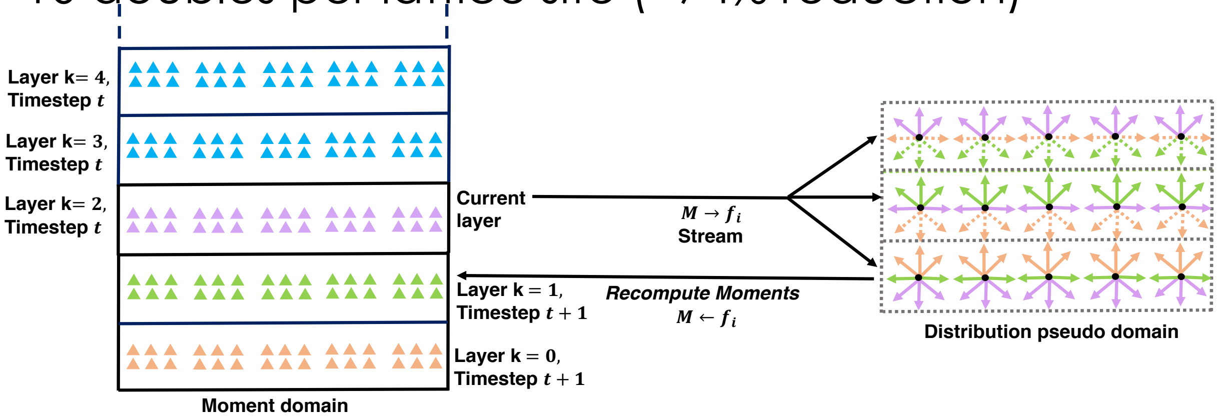

1. Latt and Chopard 2006, Latt 2007

What have our efforts yielded so far?

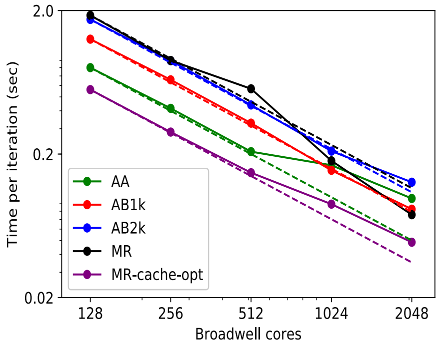

Strong scaling comparison on Intel Broadwell Cores

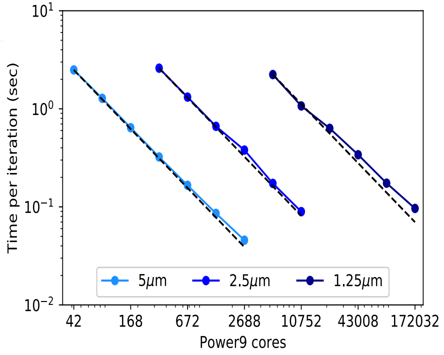

Strong scaling on Summit Power9 cores

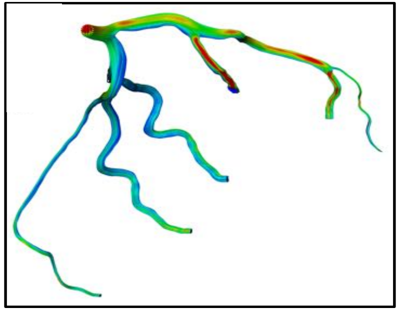
Patient-specific hemodynamics In left coronary arterial circulations

- Relative to other methods, best time to solution
- Near ideal scaling on Summit Power9 cores
- Applicable to real-world complex problems
- Challenges – Accelerators and GPUs

libCEED:
Lightweight High-Order Finite Elements Library

Valeria Barra | valeria.barra@colorado.edu
University of Colorado Boulder

libCEED

What is the point of having a **Ferrari**,
if you drive it stuck in second gear?

libCEED is a low-level API for achieving high-performance scientific computing on different architectures.

libCEED supports run-time selection of implementations tuned for a variety of computational device types (e.g., CPUs, GPUs, etc).

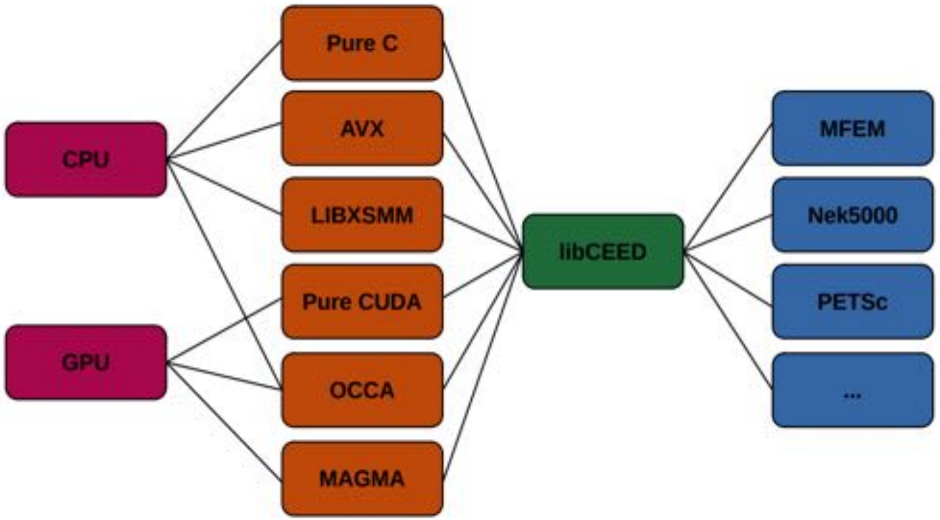

Why matrix-free?

A sparse matrix is no longer a good representation for high-order operators.
libCEED has a purely algebraic interface for matrix-free operator representation

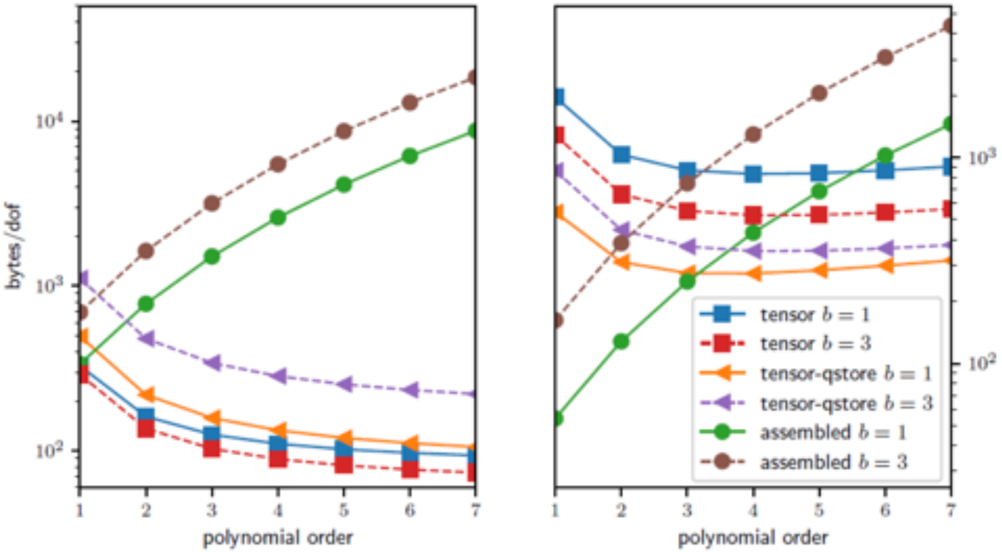

Memory bandwidth (left) and flops per dof (right) to apply a Jacobian matrix, obtained from discretization of a b-variable PDE system.

Performance

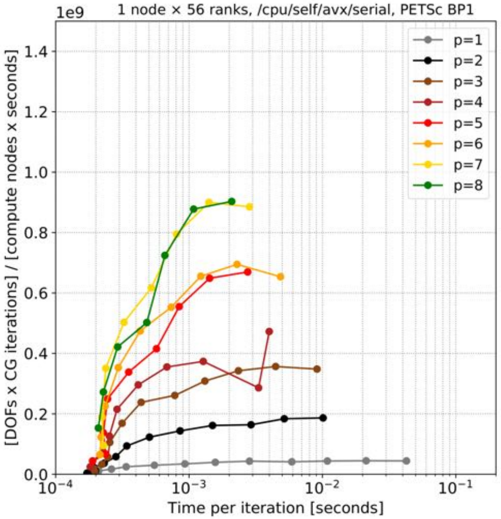

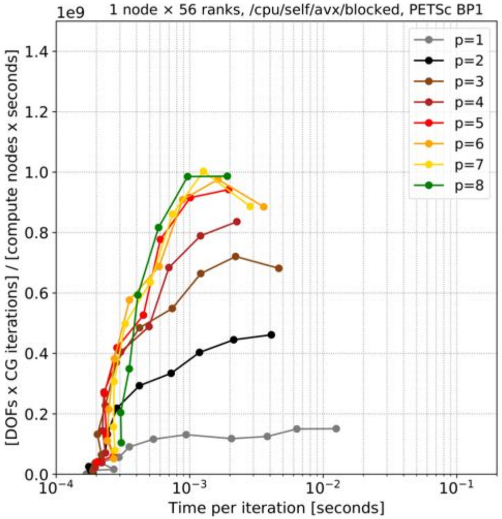

Application Example

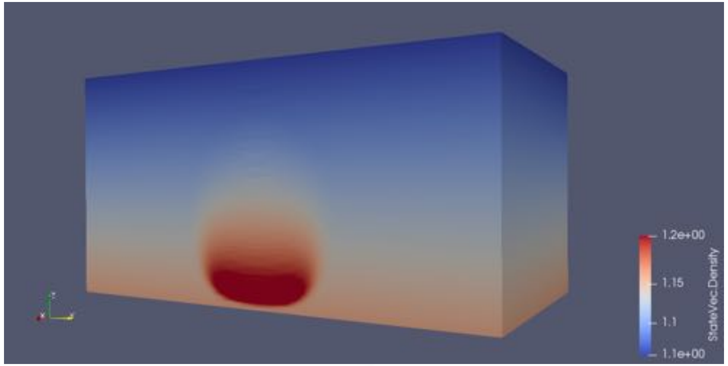

Not just toy problems: A fast and efficient **Navier-Stokes** solver.

Polynomial order of spectral elements: p=10

Computational domain:
[0,6000] x [0,6000] x [0,3000] m

Elem. Resolution: 500 m

893101 Nodes

Analysing Digital Historical Textual Data in HPC clusters and Cloud using Apache Spark and Jupyter Notebooks

Rosa Filgueira| rosa.filgueira@ed.ac.uk
EPCC, University of Edinburgh

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING
CHANGING THE FACE OF HPC

Context and Motivation

<u>Working with</u>

- Historians, Humanities and computational linguistics researchers
- Large digital collections been available for research

<u>Motivation – toolbox for Historians & Humanities comunities</u>

- Hunger for large scale text mining facilities
- Limited capacity and/or skills to use:
    - HPC/Cloud environments
    - analytic frameworks to create applications

<u>Challenges</u>

- Several large digital collections (semi-structured data)
- Different levels of quality of data – OCR
- Data with different physical representations and schemas

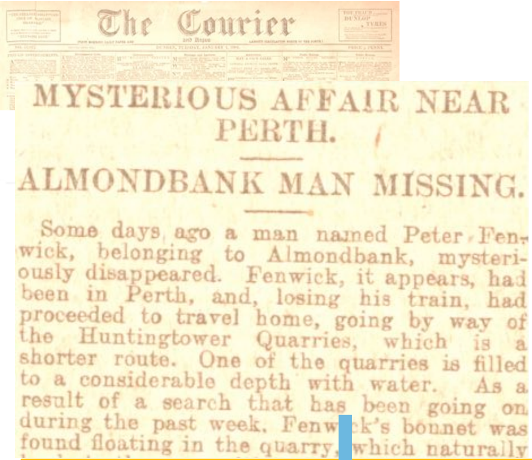

The Courier

MYSTERIOUS AFFAIR NEAR PERTH.

ALMONDBANK MAN MISSING.

Some days ago a man named Peter Fenwick, belonging to Almondbank, mysteriously disappeared. Fenwick, it appears, had been in Perth, and, losing his train, had proceeded to travel home, going by way of the Huntingtower Quarries, which is a shorter route. One of the quarries is filled to a considerable depth with water. As a result of a search that has been going on during the past week, Fenwick's bonnet was found floating in the quarry, which naturally

Digitalized Newspaper issue


```
...
<text.title>
  <pg pgref="5" clipref="1"
      pos="4069,3036,4949,3154"/>
  <p>
    <wd pos="4069,3036,4949,3154">MYSTERIOUS AFFAIR NEAR
PERTH.</wd>
  </p>
</text.title>
<text.cr>
  <pg pgref="5" clipref="1"
      pos="4039,3191,4987,4235"/>
  <p>
    <wd pos="4041,3192,4496,3241">ALMONDBANK</wd>
    <wd pos="4523,3200,4663,3246">MAN</wd>
    <wd pos="4696,3198,4976,3250">MISSING.</wd>
    <wd pos="4085,3290,4189,3323">Some</wd>
    <wd pos="4214,3290,4312,3329">days,</wd>
    ...
```


XML schema

defoe: new toolbox for historical research

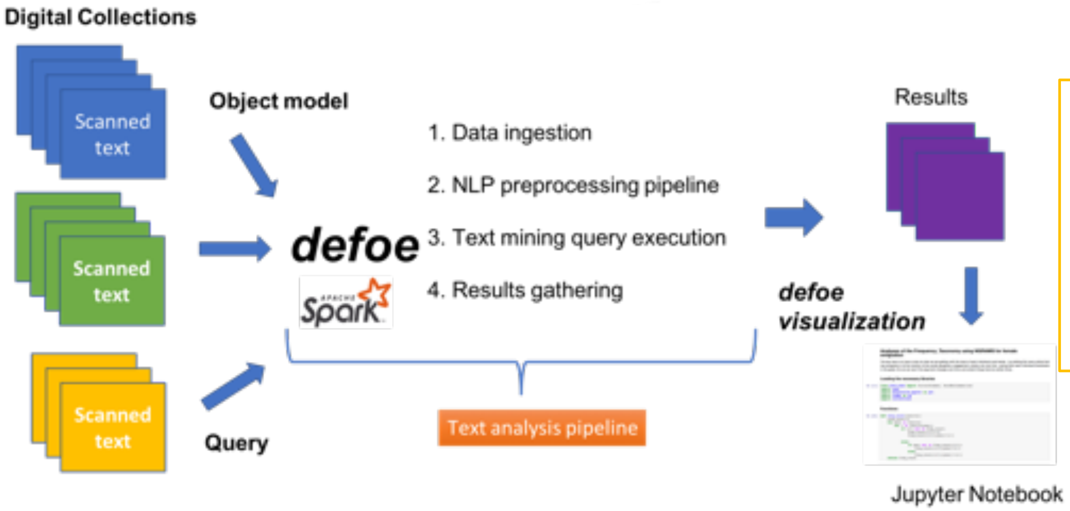

- Extracting knowledge from historical data.
- Running parallel text analyses across large collections.
- Rich set of text mining queries.
- Scalable and distributed analyses.
- NLP prepossessing techniques to mitigate OCR errors.
- Portability - computing environments and collections.

https://github.com/alan-turing-institute/defoe

https://github.com/alan-turing-institute/defoe_visualization

Developing the QuEST Library for Quantum Circuit Simulation

Lessons learnt concerning good software practice in HPC

Anna Brown | anna.brown@oerc.ox.ac.uk
University of Oxford; University of Southampton

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

QuEST

- Simulate quantum circuits on classical HPC
- Develop quantum algorithms that are robust to noise
- Memory requirements and run time scale exponentially in number of qubits
-

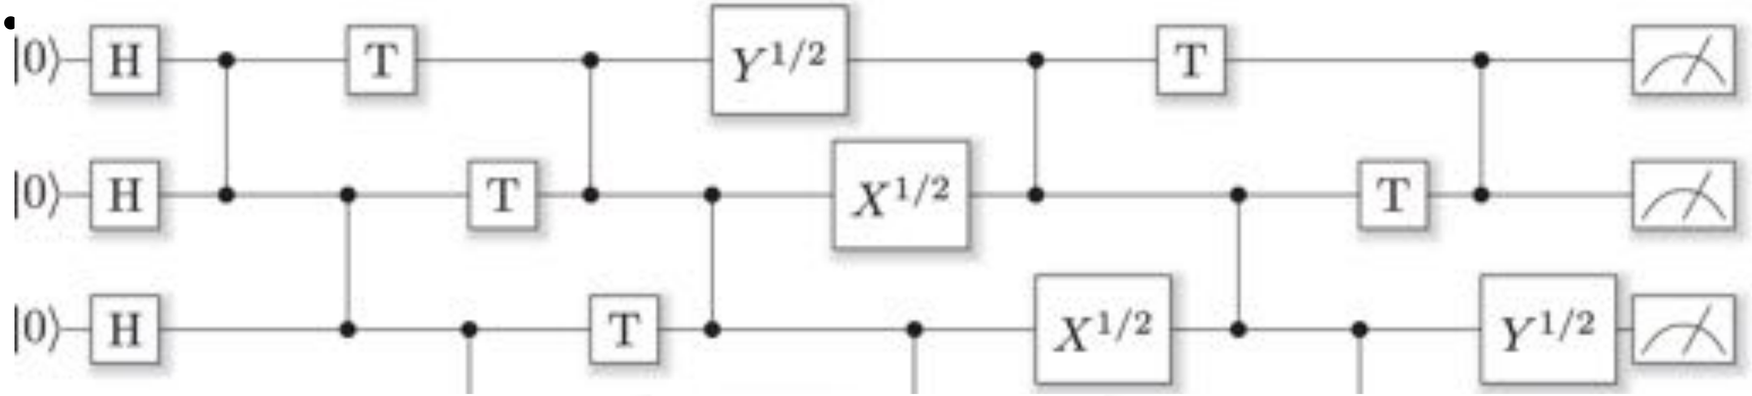

Good software practices

- Single CPU, MPI and GPU versions available through same API
- Complexity handled by build system
- Minimum dependencies – cmake, C99, Python3
- Github, doxygen, unit testing framework, CI

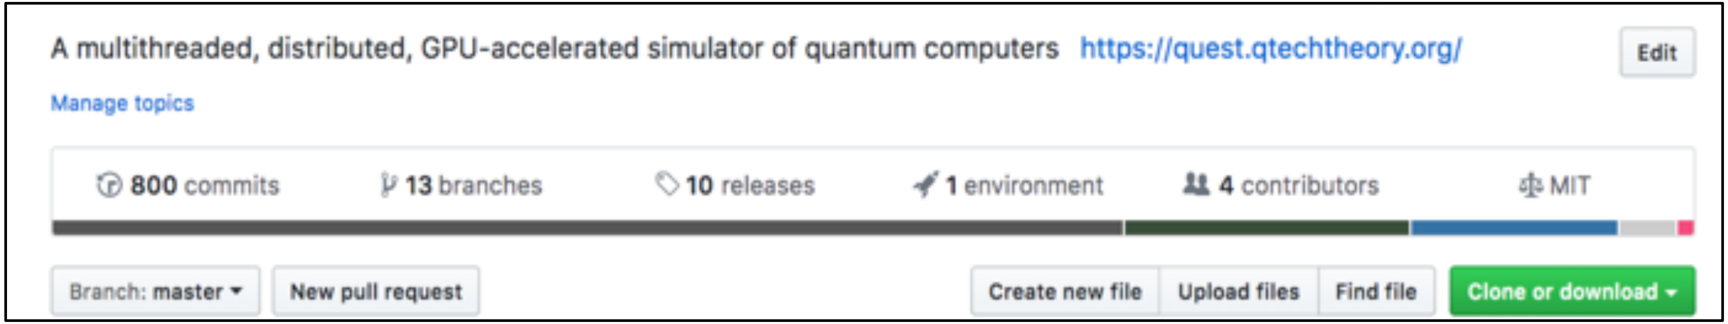

Visualization for Dynamic Fans Speed Control

Meagan Mak
meagan.mak@alumni.ubc.ca | meagan@lbl.gov
University of British Columbia | Lawrence Berkeley National Laboratory

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

Dynamic Fans Speed Control

- Put in place in April 2018
- Blower speeds are adjusted according to processor temperatures
- Reply on a variety of sensors (both internal and external)

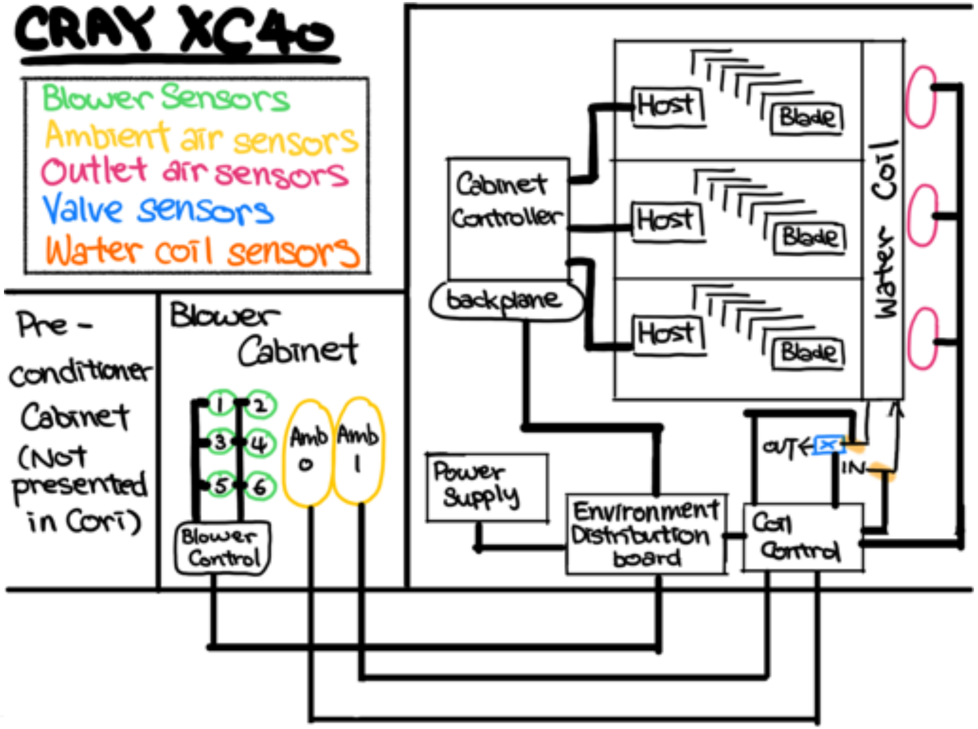

Real-time Monitoring

- Data collected from different sensors are processed through Elasticsearch and visualized with Grafana (Open-sourced)
- Provide a general picture of the conditions of the DFSC system to SREs
- Det
pur

Analysis and Optimization of DFSC

- Used at evaluating the efficiency of the system
- Provide data on the implementation of the dynamic temperature setpoint feature, which allows the system to be further optimized

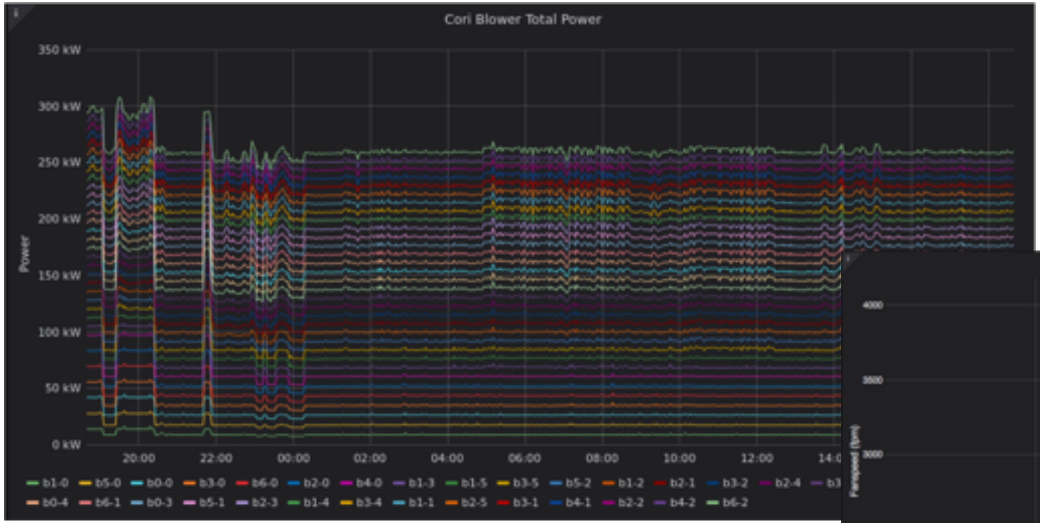

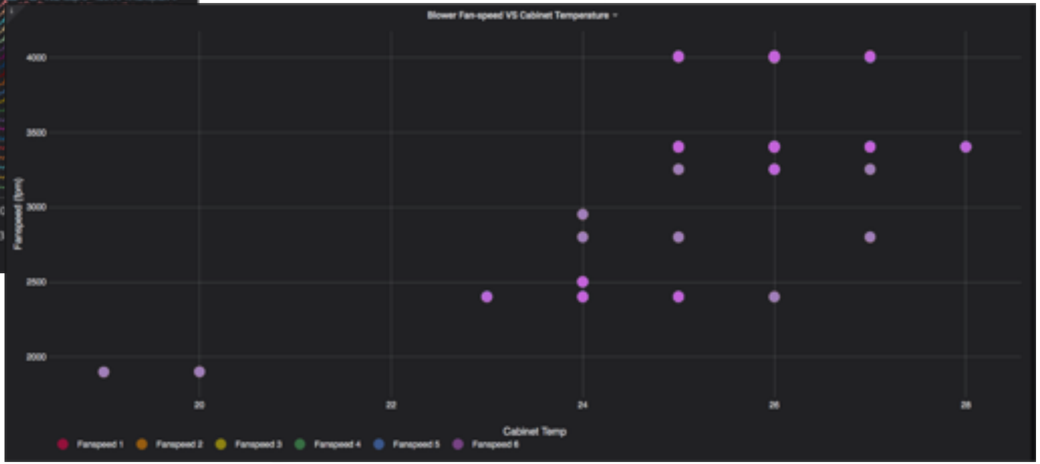

Backbone Network Design for a Sustainable Data Warehouse at NERSC

Meriam Gay Bautista | mgbautista@lbl.gov
Graduate Intern

Thomas Davis| tadavis@lbl.gov
System Architect Engineer

Operations Technology Group
NERSC
Lawrence Berkeley National Laboratory

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING CHANGING THE FACE OF HPC

National Energy Research Scientific Computing Center (NERSC)

- non-classified computational facility for Department of Energy

Instrumented a facility that gathers
from heterogeneous source;

- Syslogs
- I/O Disks
- Sflow network
- Cyber logs
- Building management data
  (power, temperature, humidity, etc.)

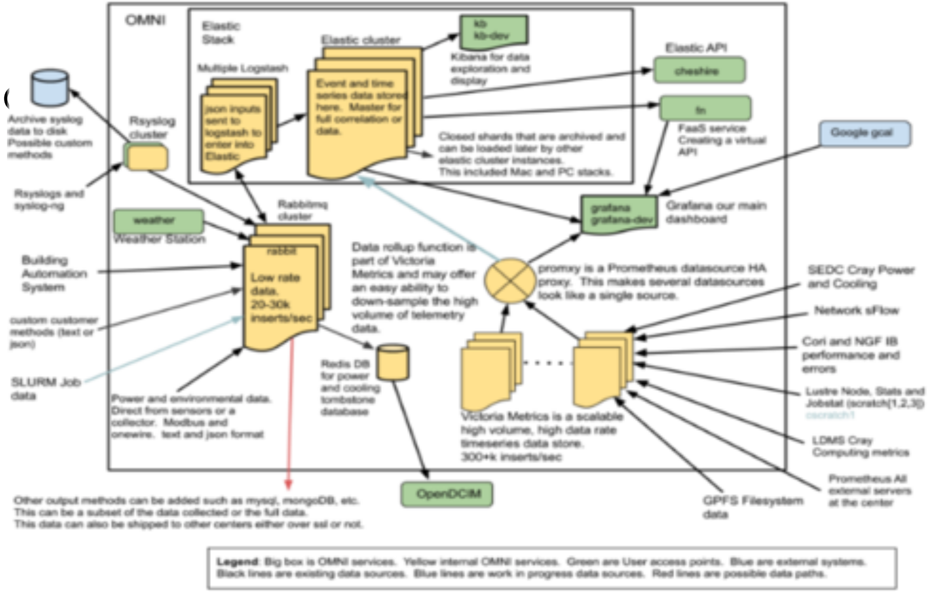

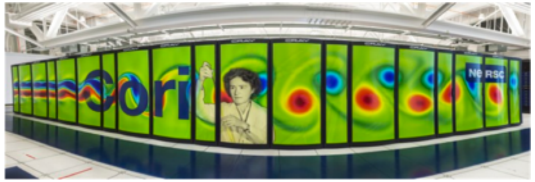

OMNI Architecture

(Operations Monitoring and Notification Infrastructure)

Current OMNI network

2 level Fat-Tree based network topology

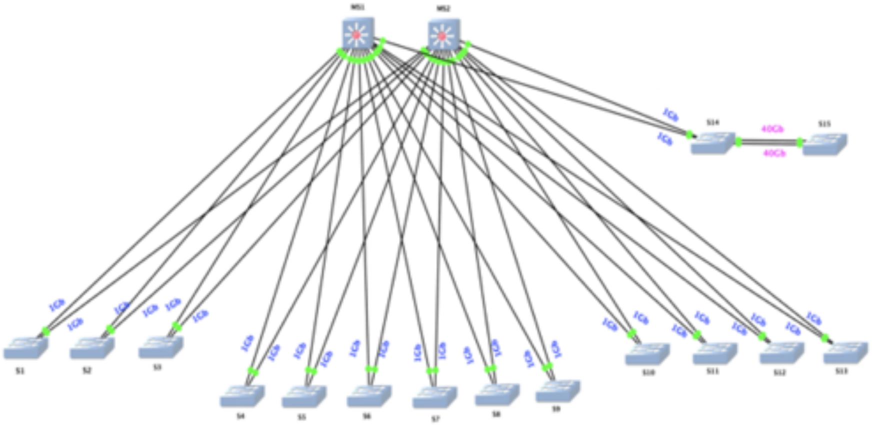

Upgrading OMNI Network

Spine-Leaf based Network Topology

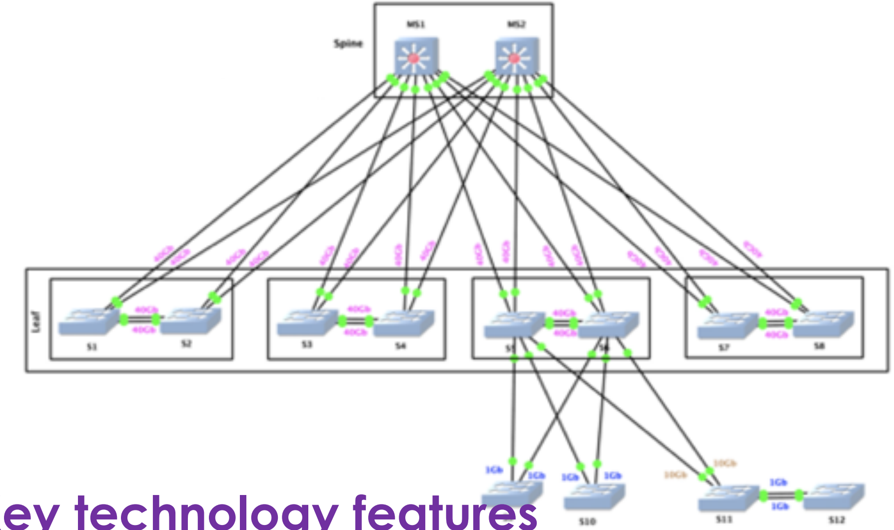

Why we need to upgrade ?

- The top of rack (TOR) switches are in End-of-life
- Cori - Compute node switch is also in End-of life
- No control plane for single management
- Leverage Ultra POE to implement new IoT technologies
- SPOF (single point of failure) configuration in network
- Reduce 3 switch OS to 1 single OS

Key technology features

- Use of VXLAN for SDN based networking
- Cumulux Linux: leverage automation, EVPN, MLAG, Ansible
- Use of Border Gateway Protocol (BGP): flexible, scalable routing, improves latency, minimize downtime
- Much more compatible to Perlmutter, collect more data

Modernizing Supercomputer Monitoring via Artificial Intelligence

Elisabeth (Lissa) Moore | lissa@lanl.gov
Ultrascale Systems Research Center
Los Alamos National Laboratory

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

LA-UR-19-30057

Anomaly Detection in Text Logs

Mix density estimation, user feedback, and explainable ML to finding interesting syslog messages.

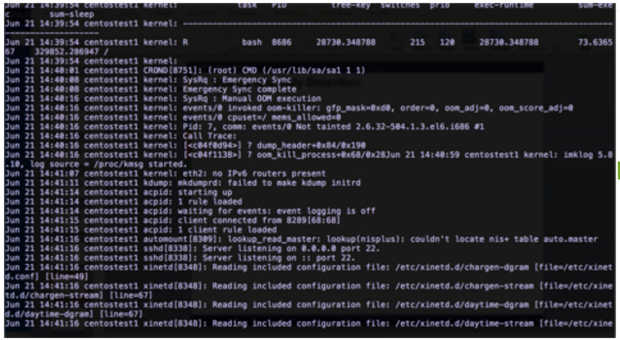

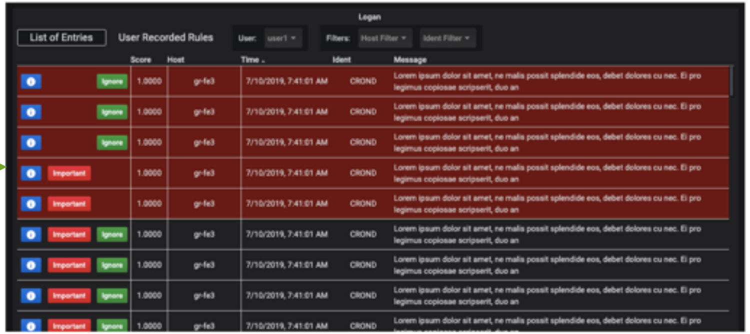

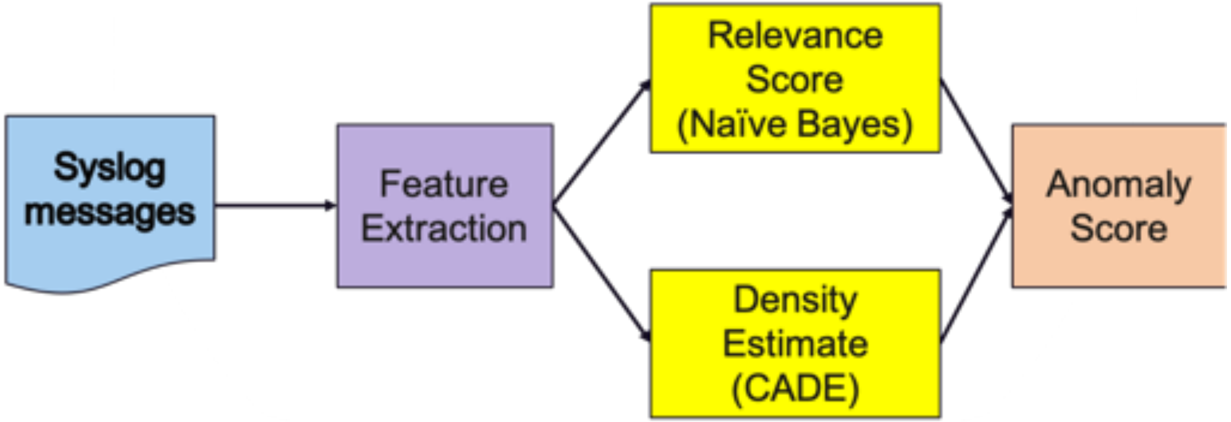

Job Outcome Prediction

Extract features from syslog messages for early detection of failing, timeout, and successful jobs.

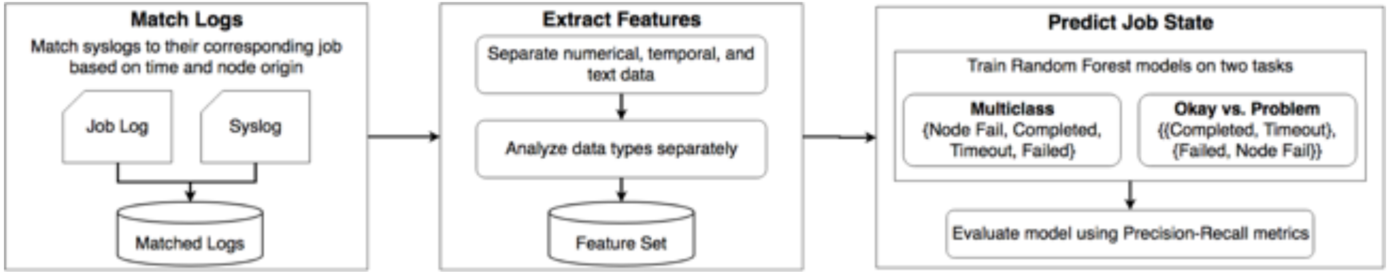

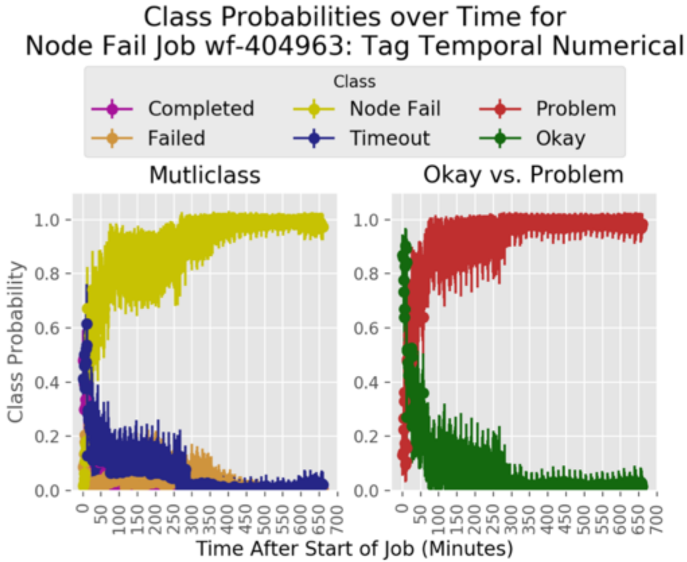

For more detail, also come to DAAC Workshop on Friday!

Telemetry Analysis

Characterize telemetry data from HPC systems to detect signals correlated with node failures.

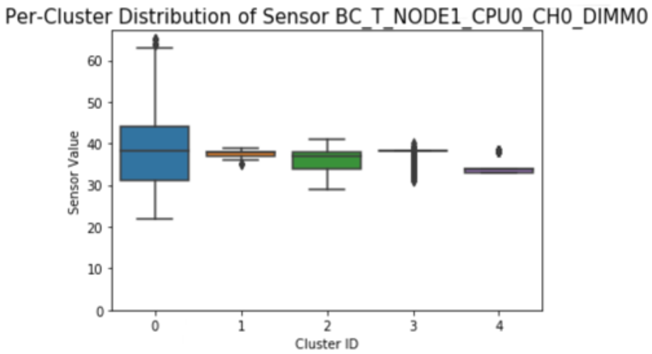

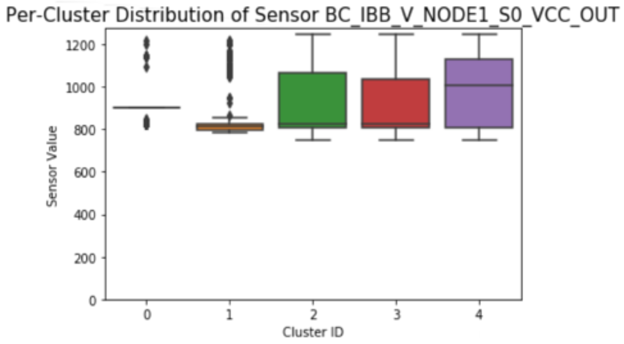

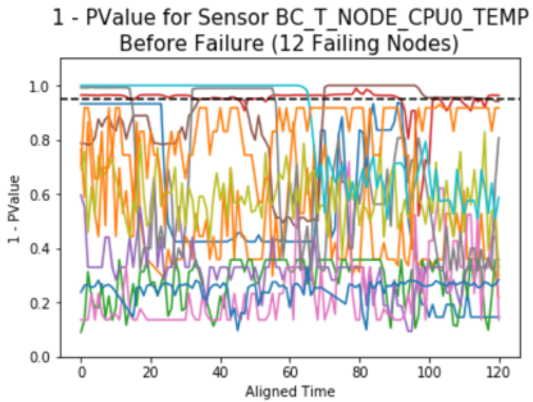

Workflow Visualization for Maintaining NERSC Data Center

Jameelah N. Mercer
JMercer@lbl.gov
Lawerence Berkeley National Laboratory

WHPC WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

Significance of Efficient Data Centers

Introduction:

This study has two main goals:

1. First, to troubleshoot NERSC supercomputing sensor functionality

   - Analyzing and debugging the temperature, humidity, and power sensor data to ensure complete workability of the NERSC supercomputing environment.

2. To create an instance go OMNI infrastructure using Kubernetes and Docker so that we can run a data set through the new version of Elastic Stack.

   - Specifically by updating the Elastic Stack by implementing K3's/Rancher and installing an ES-Operator.

Daemon Environment

1. Consistent troubleshooting of environmental sensors.
2. Periodic analysis of PDU/Power consumption.

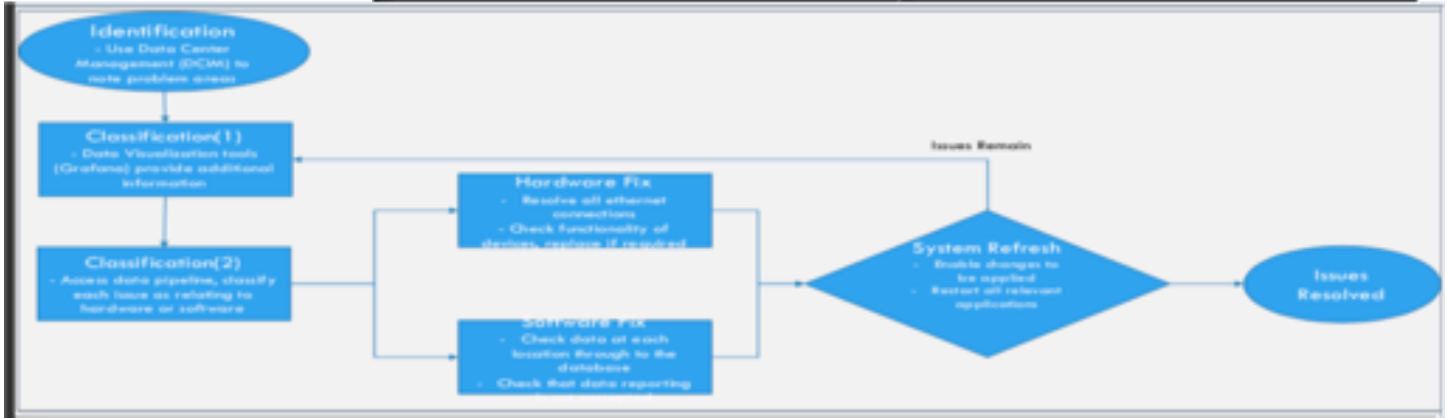

Operation Monitoring and Notification Infrastructure

1. OMNI is built using open source technologies.
2. OMNI contains over two years of operational data, accumulating over 125 of data.
3. An instance of ONMI is created to limit unforeseen problems and to increase reliability.

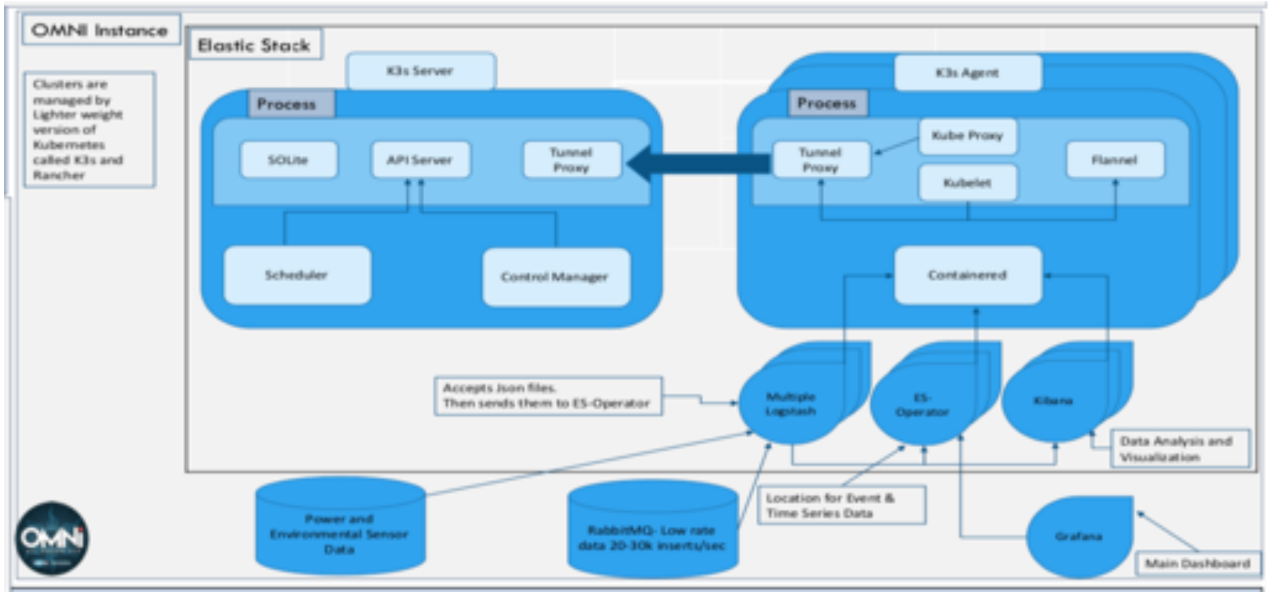

WHPC Fellows

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING
CHANGING THE FACE OF HPC

Scalable Assembly of Large Genomes

Priyanka Ghosh | Priyanka.ghosh@pnnl.gov
Pacific Northwest National Laboratory

WOMEN IN HIGH PERFORMANCE COMPUTING

CHANGING THE FACE OF HPC

Motivation

- Precision Medicine Initiative – research effort for disease prevention and treatment
  - ❑ Take into account individual differences in people's genes, environments, and lifestyles
  - ❑ Study and analyze genetic variants/mutations in diseased tissues (such as tumors) - facilitate development of targeted therapeutics
- Genome and Metagenome assembly recognized as one of the key applications in the DOE Exascale Project
  - ❑ Develop highly scalable algorithms/software to overcome high computational demands of assembling millions of (meta)genomes
  - ❑ Reduce assembly time by orders of magnitude and make feasible the assembly of larger complex genomes

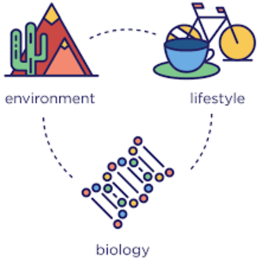

Figure courtesy: www.nih.gov/precisionmedicine

De Novo Genome Assembly: Problem Statement

<u>Goal:</u> Assemble the DNA sequence of an unknown target genome from numerous fragments (or 'reads') obtained from it

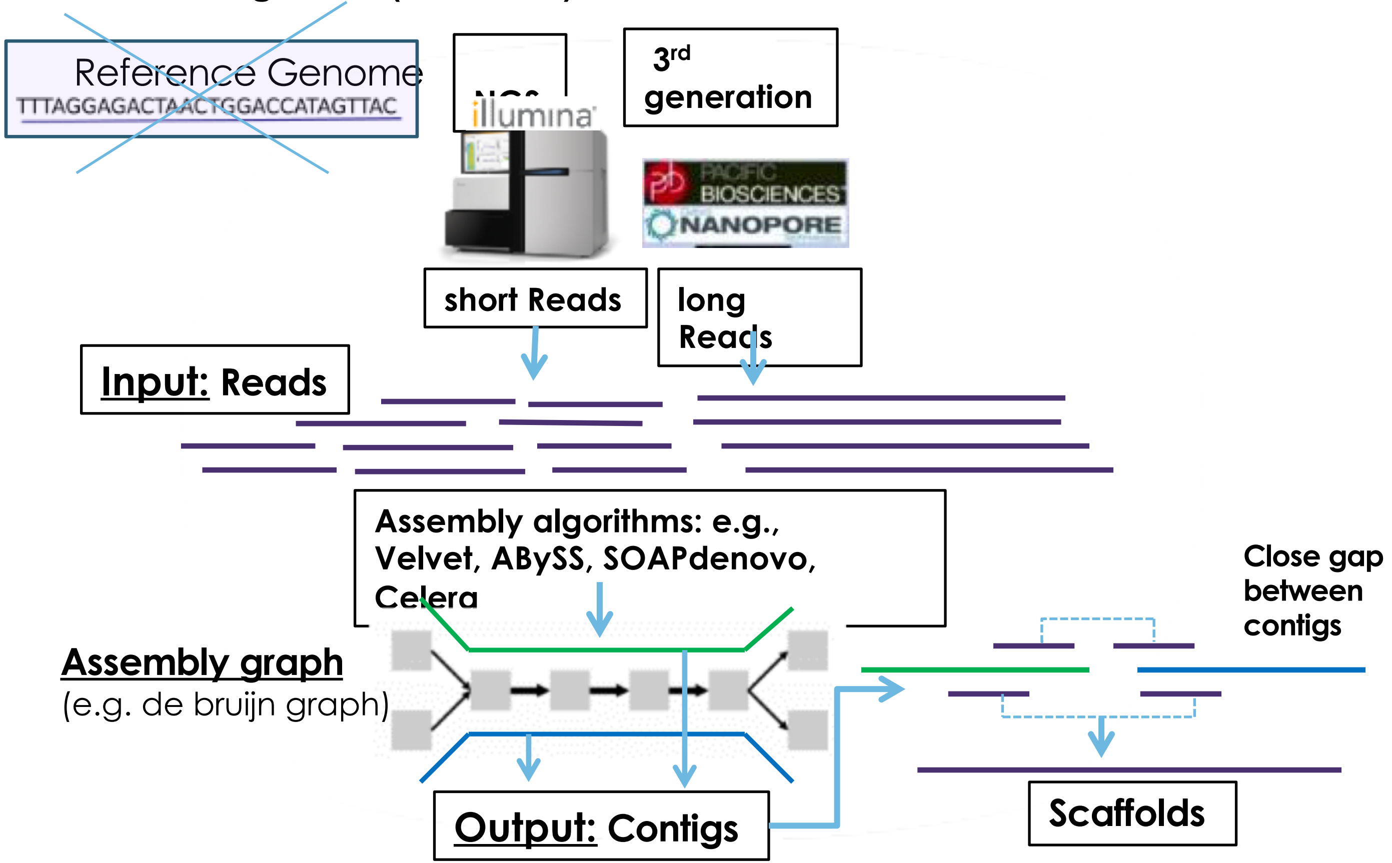

Distributed-memory approach (*PaKman*)

- Scalability Challenge: Typical read dataset comprises of billions of reads
  - Several hundred billions of vertices in the graph
  - Computationally demanding with respect to <u>memory</u> and <u>time</u>
- *PaKman*: scalable algorithm tackling assemblies of large genomes at extreme scale
  - novel distributed-memory graph data-structure (PaK-Graph) that enables minimal communication during contig enumeration
  - novel contig generation algorithm with simplified I/O and communication patterns

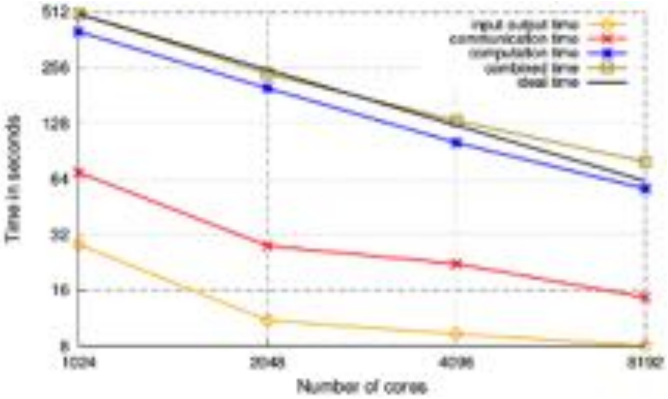

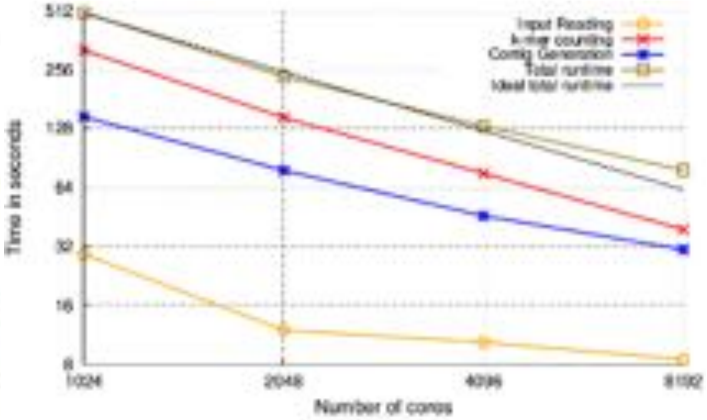

PaKman assembles a complete set of contigs for full human genome in 78.4 secs on 8k cores

Ghosh, Priyanka, Sriram Krishnamoorthy and Ananth Kalyanaraman. "PaKman: Scalable Assembly of Large Genomes on Distributed Memory Machines." In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. 2019.

Technical Support in Configuring HPC Systems

Raksha Roy | raksha.roy@icimod.org
ICIMOD, Nepal

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING
CHANGING THE FACE OF HPC

Objectives

- Get proper training on configuration and use of HPC Systems
- Install and configure scientific applications required for the Supercomputing facility
- High Performance Computing Benchmarks
- Train Students and Scientists on use of HPC Systems

Impact

Installation and configuration of OpenHPC, Lustre

Accomplishments

- A full fledged HPC Production System
- Introduction on HPC, Parallel Programming Techniques, Public Awareness on SuperComputing

Enabling HPC for neuroimaging science

- CT template creation using nonlinear image registration for TBI analysis

Zhe Bai | zhebai@lbl.gov
Computational Research Division
Lawrence Berkeley National Laboratory

WHPC
WOMEN IN HIGH PERFORMANCE COMPUTING
CHANGING THE FACE OF HPC

Background: Traumatic Brain Injury (TBI)

- Number of TBI cases occurring in the U.S. every year: ~1.7 million.
- Influences in life: physical, psychological, social, and spiritual.
- Complexities: *multi-modal data, large volume image, statistical varieties*.

data science + computer vision +
high-performance computing

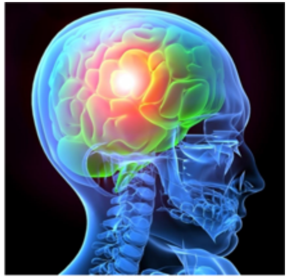

Centers for Disease Control and Prevention (CDC), Traumatic Brain Injury (TBI): Incidence and Distribution, 2004

CT template creation

- Subgroup study based on physiological features.
- Iterative algorithm: rigid + affine + nonlinear transformation.
- High performance: image similarities are optimized in parallel.

Created template based on 12 subjects (shown in MNI space)

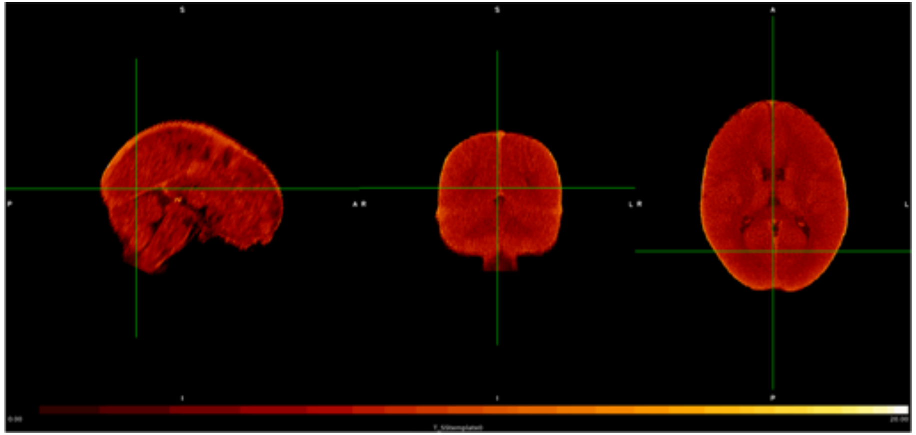

Computational time vs. # cores

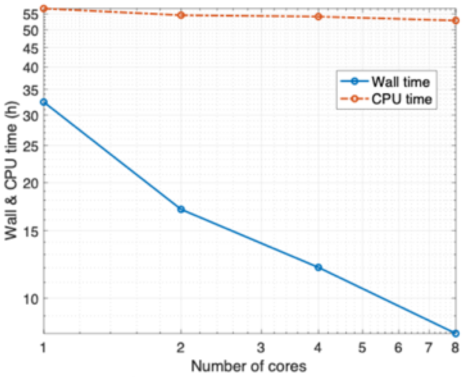

Image segmentation for TBI patients

- Segmented template & patients' CT scans.
- Automatic parcellation: 48 structural areas < *1 min*.

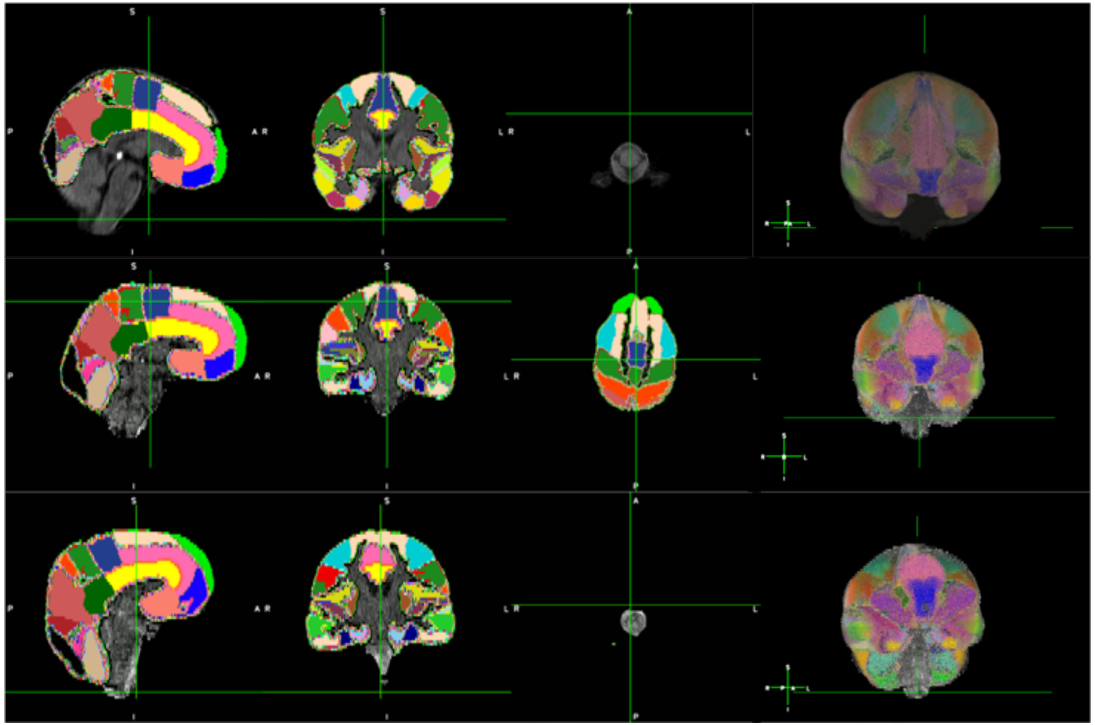

Atlas: HarvardOxford-Cortical
Bottom: Group skull-stripped CT template

Patient GCS: 15
Bottom: Patient's skull-stripped CT

Patient GCS: 4
Bottom: Patient's skull-stripped CT