ECP Update: Progress on Applications in Data Analytics and Optimization

Doug Kothe, ECP Project Director, Oak Ridge National Laboratory
Lori Diachin, ECP Deputy Director, Lawrence Livermore National Laboratory
Bill Hart, AD Leadership Team, The Sandia National Laboratories

September 30, 2022

ECP has successfully run two of our challenge problems on Frontier

WarpX
Accelerator Physics

Figure of Merit on 8576 Frontier nodes is **~500X** baseline run on NERSC Edison system

Science goal: First ever 3D simulation of a chain of tens of plasma accelerator stages for future colliders using multi-physics particle in cell codes.

FOM run: 1000 steps of plasma acceleration stage with a prefilled plasma. Physics included PIC dynamic maxwell equations solver using block AMR grids, field ionization of atomic levels, Coulomb collisions, macroscopic materials.

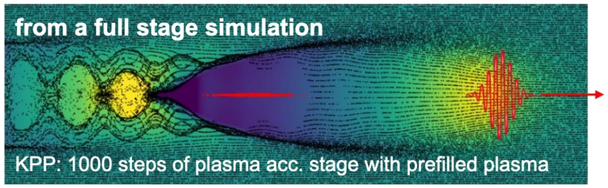

ExaSky
Cosmology

Figure of Merit on 8192 Frontier nodes is **~230X** baseline run on ANL Theta system

Science goal: Enable extraction of science from upcoming cosmological surveys.

FOM run: Gravity-only and hydro simulations using HACC at the scale of galaxies, groups and clusters. Physics included gravity, gas dynamics, heating/cooling, star formation, wind models, etc. Code design focuses on a small number of kernels that are optimized for each system; 95% of the code remains unchanged across systems.

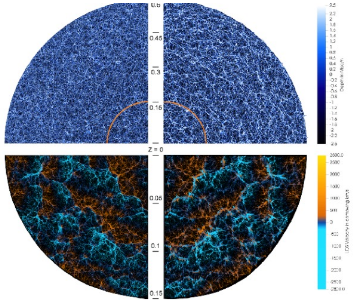

Particle light-cone visualizations from the 'Farpoint' run (Frontiere et al., ApJS 2022), a might mass resolution, large-volume cosmological simulation run with HACC.

ECP EXASCALE COMPUTING PROJECT

ECP teams are ready and eager to have access to exascale systems

Frontier

- ECP expected to get access late October 2022
- Most teams are ready for a medium to large-scale run; If additional pre-acceptance time slots become available, 5-7 teams (CANDLE, ExaSMR, EQSIM, Subsurface, LatticeQCD, WDMApp) to be prioritized
- Teams make extensive use of access to Crusher, hackathons, vendor office hours, OLCF user support

Aurora

- Sunspot TDS system ready for ECP teams in early November (2 racks of final Aurora hardware)

KPP-1 App	Aurora EAS (Intel Proprietary)	Frontier TDS	Frontier
LatticeQCD	Verified	Ready	
NWChemEx	Full Build/Test	Initial Build/Test	
EXAALT	Verified	Improving Perf.	
QMCPACK	Initial Build/Test	Improving Perf.	
ExaSMR	Improving Perf.	Ready	
WDMApp	Improving Perf.	Improving Perf.	
WarpX	Verified	Ready	Run
ExaSky	Improving Perf.	Ready	Run
EQSIM	Initial Build/Test	Ready	
E3SM-MMF	Improving Perf.	Improving Perf.	
CANDLE	Ready	Ready	

KPP-2 App	Aurora EAS (Intel Proprietary)	Frontier TDS	Frontier
GAMESS	Improving Perf.	Improving Perf.	
ExaAM	Initial Build/Test	Improving Perf.	
ExaWind	Verified	Improving Perf.	
Combustion-PELE	Initial Build/Test	Improving Perf.	
MFIX-Exa	Verified	Improving Perf.	
ExaStar	Full Build/Test	Improving Perf.	
Subsurface	Stretch	Ready	
ExaSGD	Stretch	Improving Perf.	
ExaBiome	Stretch	Improving Perf.	
ExaFEL	Full Build/Test	Blocked	

EXASCALE COMPUTING PROJECT

The first cohort of the ECP Broadening Participation Initiative Group was a big success!

SRP-HPC students and faculty, summer 2022

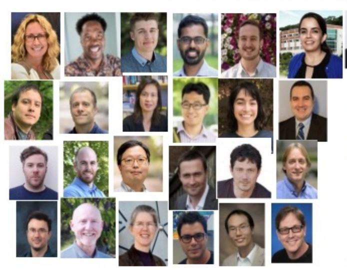

SRP-HPC mentors/co-mentors

Dan Martin, LBL
SRP thrust lead
for ECP

Keisha Moore, SHI
SRP Program
Coordinator

- A multi-pronged initiative to expand the pipeline and workforce for DOE HPC led by Dan Martin (LBNL) and Lois McInnes (ANL)
- Partnership with Sustainable Horizons Institute
- Cohort 1
  - 61 participants: 13 student track, 16 faculty track (+29 students), 3 self-funded students
  - 82% of overall participants represent at least 1 element of diversity
  - Mentors/hosts through ECP and Facilities community
  - Matches at all 9 participating labs
  - All participants to present posters at the 2023 ECP Annual Meeting
- Cohort 2
  - Funded for summer 2023
  - Kickoff meeting was Sept 29; matching workshop January 11-13

ECP is very active in agency outreach with many conversations around use of E4S

NOAA

- NOAA deep dive meeting on July 20 was very successful. Discussed NOAA goals and shared lessons learned
- NOAA experimenting with Spack build caches to significantly reduce compile times and, using E4S, build their code AM4 for the first time on AWS cloud.
- Working on ideas for collaboration projects post-ECP

NSF

- Planning an exascale system; very interested in E4S software stack. Exploring deployment of E4S on NFS commodity clusters
- Joint NSF-DOE workshops on E4s
- Shared lessons learned in ECP project management for portfolios of applications and software technologies.
- Inviting them to give a plenary panel at the 2023 ECP Annual Meeting.

NASA

- ST presentation at the NASA Science Mission Directorate Open Source Science Initiative Data and Compute Architecture study
- Planning technical deep dives; collected first round of topics of interest.

DoD

- Deployment and evaluation of E4S on DoD Narwal HPC system planned (Navy)
- Planning technical deep dive; requested topics of interest.

E4S

https://e4s.io

E4S lead: Sameer Shende (U Oregon)

Spack

https://spack.io

Spack lead: Todd Gamblin (LLNL)

All IAC agencies invited to cross-cut workshops on Cloud Computing in an Exascale World and Fortran planned for October and November, respectively

ECP EXASCALE COMPUTING PROJECT

ECP leadership helping to drive conversations around post-ECP sustainability

Application Development

- Summary slides of post-ECP status for funding and opportunity
- Ongoing interactions with stakeholder offices
- White paper on application role in broader sustainability efforts

Software Technology

- Monthly meetings with ASCR task force on software stewardship
- Response to ASCR RFI; summary information of dependencies
- ST / Co-design product summary slides for post ECP funding and opportunity

Hardware and Integration

- Active discussions with facilities on staff retention post ECP; particularly for application integration staff
- Several activities integrated with ST sustainability vision (productivity, continuous integration, software deployment)

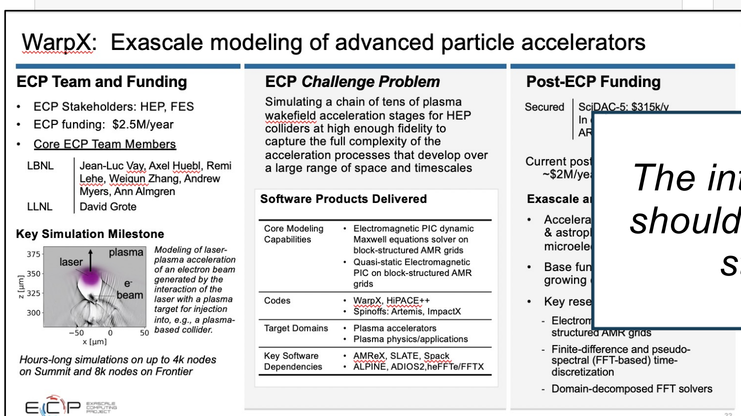

WarpX: Exascale modeling of advanced particle accelerators

ECP Team and Funding

- ECP Stakeholders: HEP, FES
- ECP funding: $2.5M/year
- Core ECP Team Members

LBNL	Jean-Luc Vay, Axel Huebl, Remi Lehe, Weiqun Zhang, Andrew Myers, Ann Almgren
LLNL	David Grote

Key Simulation Milestone

Modeling of laser-plasma acceleration of an electron beam generated by the interaction of the laser with a plasma target for injection into, e.g., a plasma-based collider.

Hours-long simulations on up to 4k nodes on Summit and 8k nodes on Frontier

ECP *Challenge Problem*

Simulating a chain of tens of plasma wakefield acceleration stages for HEP colliders at high enough fidelity to capture the full complexity of the acceleration processes that develop over a large range of space and timescales

Software Products Delivered

Core Modeling Capabilities	• Electromagnetic PIC dynamic Maxwell equations solver on block-structured AMR grids • Quasi-static Electromagnetic PIC on block-structured AMR grids
Codes	• WarpX, HiPACE++ • Spinoffs: Artemis, ImpactX
Target Domains	• Plasma accelerators • Plasma physics/applications
Key Software Dependencies	• AMReX, SLATE, Spack • ALPINE, ADIOS2,heFFTe/FFTX

Post-ECP Funding

Secured | SciDAC-5: $315k/y
In
AR

Current post
~$2M/yea

Exascale a

- Accelera & astrop microele
- Base fun growing
- Key rese
  - Electron structured AMR grids
  - Finite-difference and pseudo-spectral (FFT-based) time-discretization
  - Domain-decomposed FFT solvers

22

The integrated AD/ST/HI interactions initiated by ECP should be continued – the complex systems, software stacks, etc. necessitates this going forward

EXASCALE COMPUTING PROJECT

A Response to the "Stewardship of Software for Scientific and High-Performance Computing" Request for Information

Michael A. Heroux, Sandia National Laboratories
Lois Curfman McInnes, Argonne National Laboratory
Rajeev Thakur, Argonne National Laboratory
Jeffrey S. Vetter, Oak Ridge National Laboratory
Sherry Li, Lawrence Berkeley National Laboratory
James Ahrens, Los Alamos National Laboratory
Todd Munson, Argonne National Laboratory
Kathryn Mohror, Lawrence Livermore National Laboratory
Timothy Germann, Los Alamos National Laboratory
Ryan Adamson, Oak Ridge National Laboratory
Todd Gamblin, Lawrence Livermore National Laboratory
Sameer Shende, University of Oregon
James Willenbring, Sandia National Laboratories

December 13, 2021

U.S. DEPARTMENT OF ENERGY | Office of Science
NNSA National Nuclear Security Administration

ECP is entering a very exciting time!

- Executing on KPP challenge problems and integration goals
- Engaging stakeholders on the new capabilities developed
- Engaging industry and other agencies with outreach and lessons learned to broaden the community of exascale-ready applications and technologies

Date	ECP Events
October 26-27, 2022	IAC Meeting at ORNL
October 31, 2022	Cloud Computing Workshop (IAC)(Virtual)
October 2022	Access to Frontier
November 2022	Access to Aurora TDS
November 2022	Fortran Workshop (IAC)(Virtual)
November 14-19, 2022	SC22 in Dallas
January 11-13, 2023	Broadening participation cohort 2 matching workshop
January 17-20, 2023	2023 ECP Annual Meeting

AD stakeholder engagement

Office	POC	Briefing Date
FECM	Jennifer Wilcox	December 8, 2021
FES	James Van Dam	December 23, 2021, next update Oct 13, 2022
HEP	Harriet Kung	June 10, 2022
BES	Linda Horton	June 29, 2022
WETO	Benjamin Hallissy	September 26, 2022
NE	Katie Huff	Setting date
BER	Gary Geernaert	**TBD**
NP	Tim Hallman	**TBD**
EERE	Mike Anderson	**TBD**
OE	Gil Bindenwald	**TBD**
CESER	Puesh Kumar	**TBD**

Will invite program managers from these offices to the 2023 ECP Annual Meeting

ECP Data Analytics and Optimization Applications

William Hart, SNL

Overview of ECP Data Analytics and Optimization Applications

Focus: Applications that employ modern data analysis and machine learning techniques as a fundamental component of understanding and predictability

Project	PI	Description
CANDLE	Rick Stevens, ANL	Accelerate and Translate Cancer Research
ExaFEL	Amedeo Perazzo, Stanford	Light Source-Enabled Analysis of Molecular Structure
ExaSGD	Chris Oehmen, PNNL	Reliable and Efficient Planning of the Power Grid
ExaBiome	Kathy Yelick, LBNL	Improve Understanding of the Microbiome

DAO applications are new challenges for HPC systems

- Large-scale data-driven computations
- Kernels for sparse, irregular computations, etc.

➡ DAO applications are high-risk investments for ECP

CANDLE

ECP-CANDLE: CANcer Distributed Learning Environment

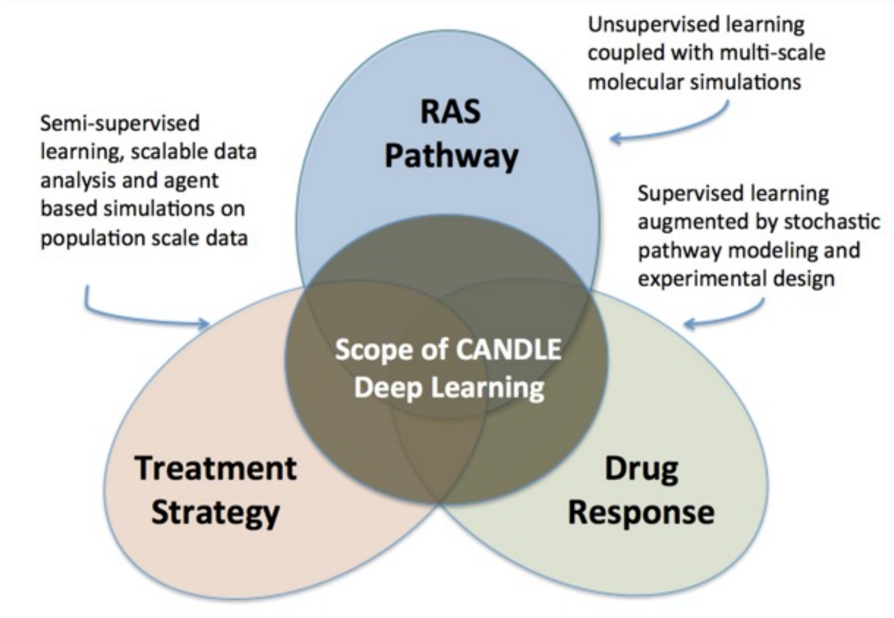

CANDLE Goals

Develop an exscale deep learning environment for cancer

Build on open source deep learning frameworks

Optimize for CORAL and exascale platforms

Support all three pilot project needs for deep learning

Collaborate with DOE computing centers, HPC vendors and ECP co-design and software technology projects

ECP EXASCALE COMPUTING PROJECT

Candle Functional Goals

- Enable high productivity for deep learning centric workflows
- Support Key DL frameworks on DOE supercomputers (Keras, TF, Mxnet, CNTK)
- Support multiple paths to concurrency (Ensembles, Data and Model Parallel)
- Manage training data, model search, scoring, optimization, production training and inference (End-to-End Workflow)
- CANDLE runtime/supervisor (interface with batch schedulers)
- CANDLE Python library for improving model development (UQ, HPO, CV, MV)
- Well documented open examples and tutorials on Github
- Leverage as much open source as possible (build only what we need to add to existing frameworks)

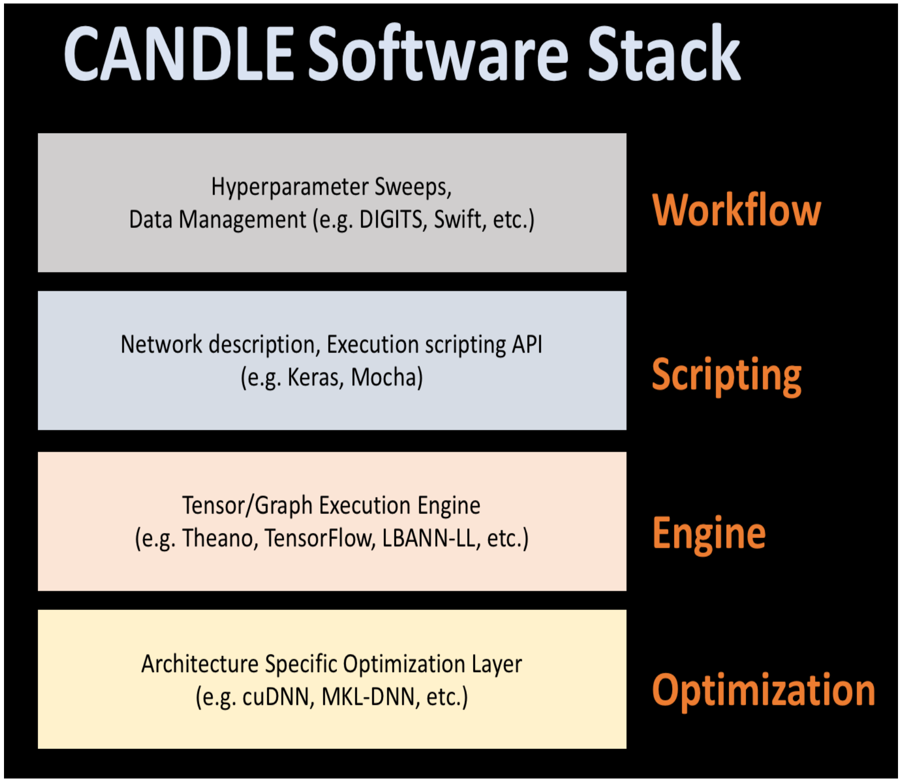

Deep Learning in Cancer ⟹ many Methods

- **AutoEncoders** – learning data representations for classification and prediction of drug response, molecular trajectories
- **VAEs and GANs** – generating data to support methods development, data augmentation and feature space algebra, drug candidate generation
- **CNNs** – type classification, drug response, outcomes prediction, drug resistance
- **RNNs** – sequence, text and molecular trajectories analysis

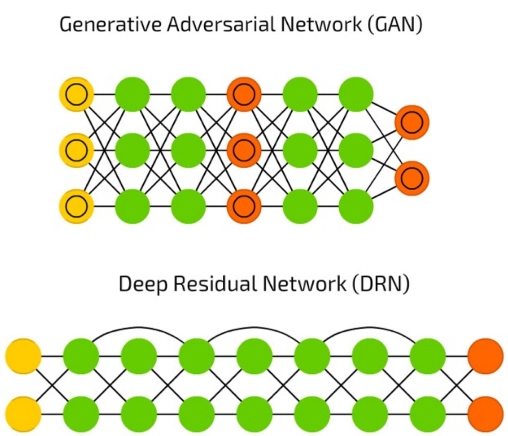

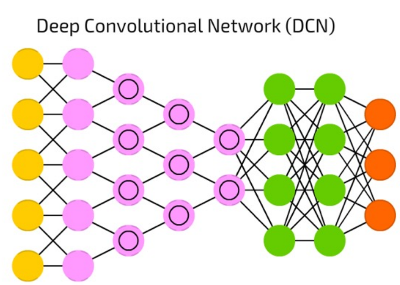

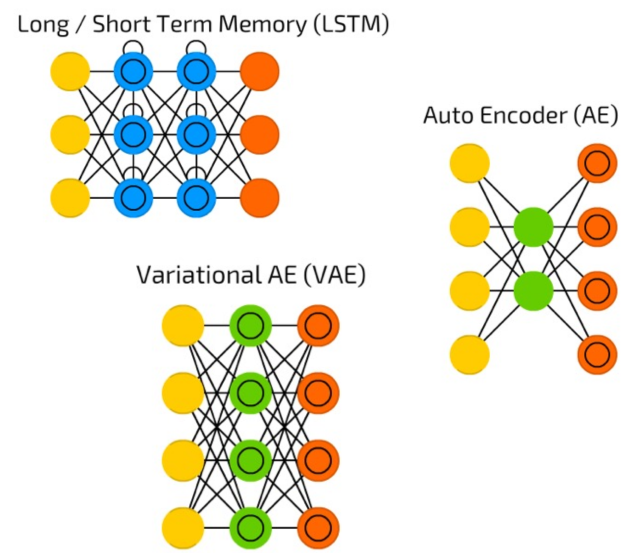

EXASCALE COMPUTING PROJECT

CANDLE KPP and FOM

Achieved an estimated KPP of 50 in Apr 2020

CANDLE FOM is throughput rate measured in "models trained per hour"

Improvements continue with new optimizations to deep learning stacks

Improvements continue with new I/O strategies

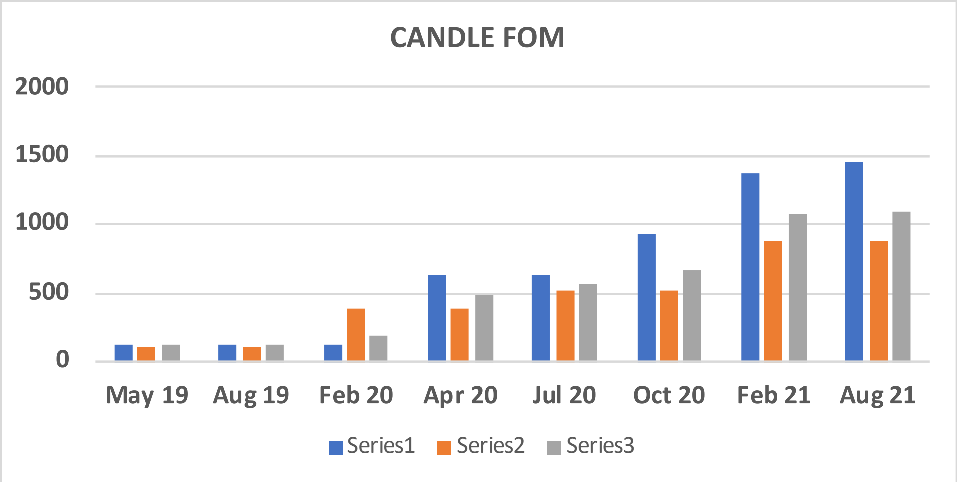

	Pilot 1	Pilot 3	CANDLE FOM	CANDLE KPP
May 19	128.00	114.00	120.60	14.98
Aug 19	128.00	112.60	119.98	14.90
Feb 20	128.00	392.90	193.09	23.99
Apr 20	631.67	392.9	484.46	60.18
Jul 20	631.67	511.20	565.09	70.20
Oct 20	931.04	511.20	660.01	81.99
Feb 21	1376.56	874.20	1069.32	132.83
Aug 21	1449.83	874.2	1090.73	135.49

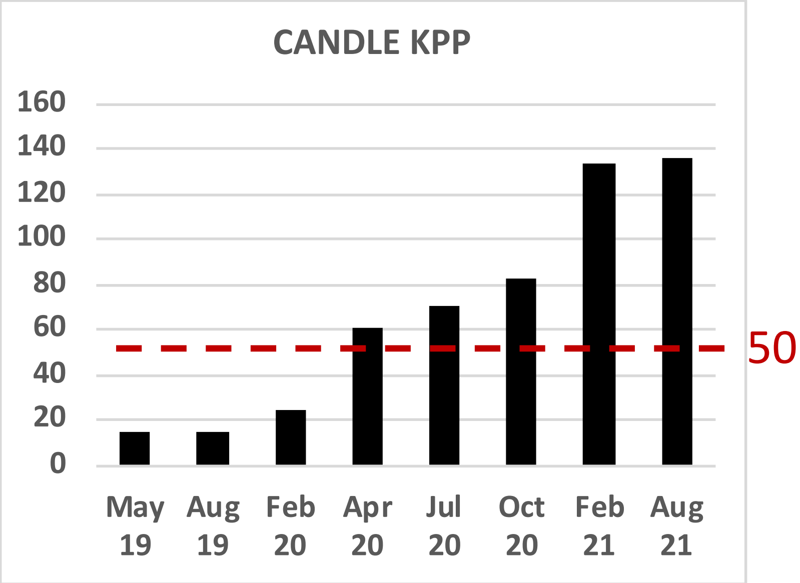

EXASCALE COMPUTING PROJECT

ECP KPP: Readiness for Frontier

- CANDLE Benchmark P3B4 from MOSSAIC
- Focus on single device performance
- Treating Frontier GPU as two logical devices ('GCDs') and benchmarking single GCD vs V100 from Summit
- Seeing 31-33% decrease in runtime on MOSSAIC datasets from different state cancer registries
- Expect to report Frontier KPP well in excess of ECP's 50x standard during FY23Q1

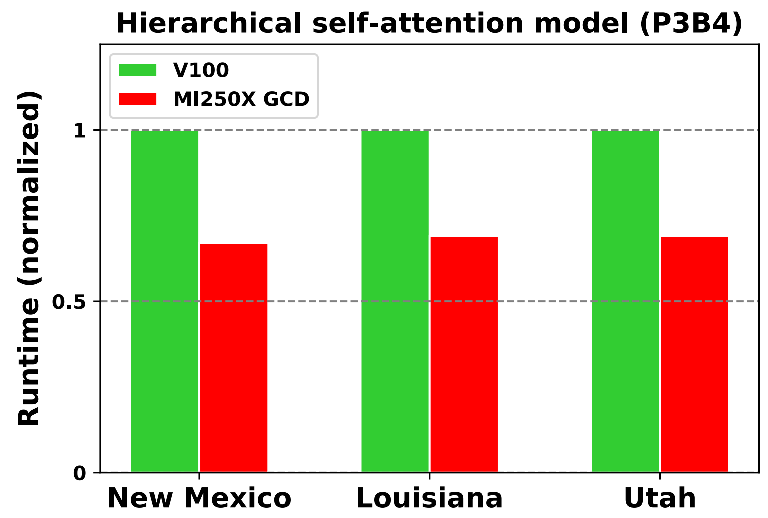

KPP Risk Assessment	KPP Timeline
Near Certainty	FY23/Q2

ExaFEL

ExaFEL:
Data Analytics for High Repetition Rate Free Electron Lasers

FEL Data Challenge:

- Ultrafast X-ray pulses from LCLS are used like flashes from a high-speed strobe light, producing stop-action movies of atoms and molecules
- Both data processing and scientific interpretation demand intensive computational analysis

LCLS-II will increase data throughput by three orders of magnitude by 2025, creating an exceptional scientific computing challenge

Challenge Problem:

- Serial Femtosecond Crystallography (SFX): using x-ray tracing in nanocrystallography reconstruction

Stretch Goals:

- Single Particle Imaging (SPI): simultaneously determine conformational states, orientations, intensity, and phase from single particle diffraction images
- Real Time End-to-end Workflows: automate the coordination of resources to execute end-to-end workflows from SLAC to NERSC

Broad ExaFEL Goal: Near real-time analysis of FEL data

Resource orchestration

- A key challenge for all ExaFEL applications
- This requires making reservations for both communication and computational resources

Filtering LCLS Data

- Critical to enable high quality, near real-time analysis
- Impacts requirements for communication bandwidths and HPC nodes

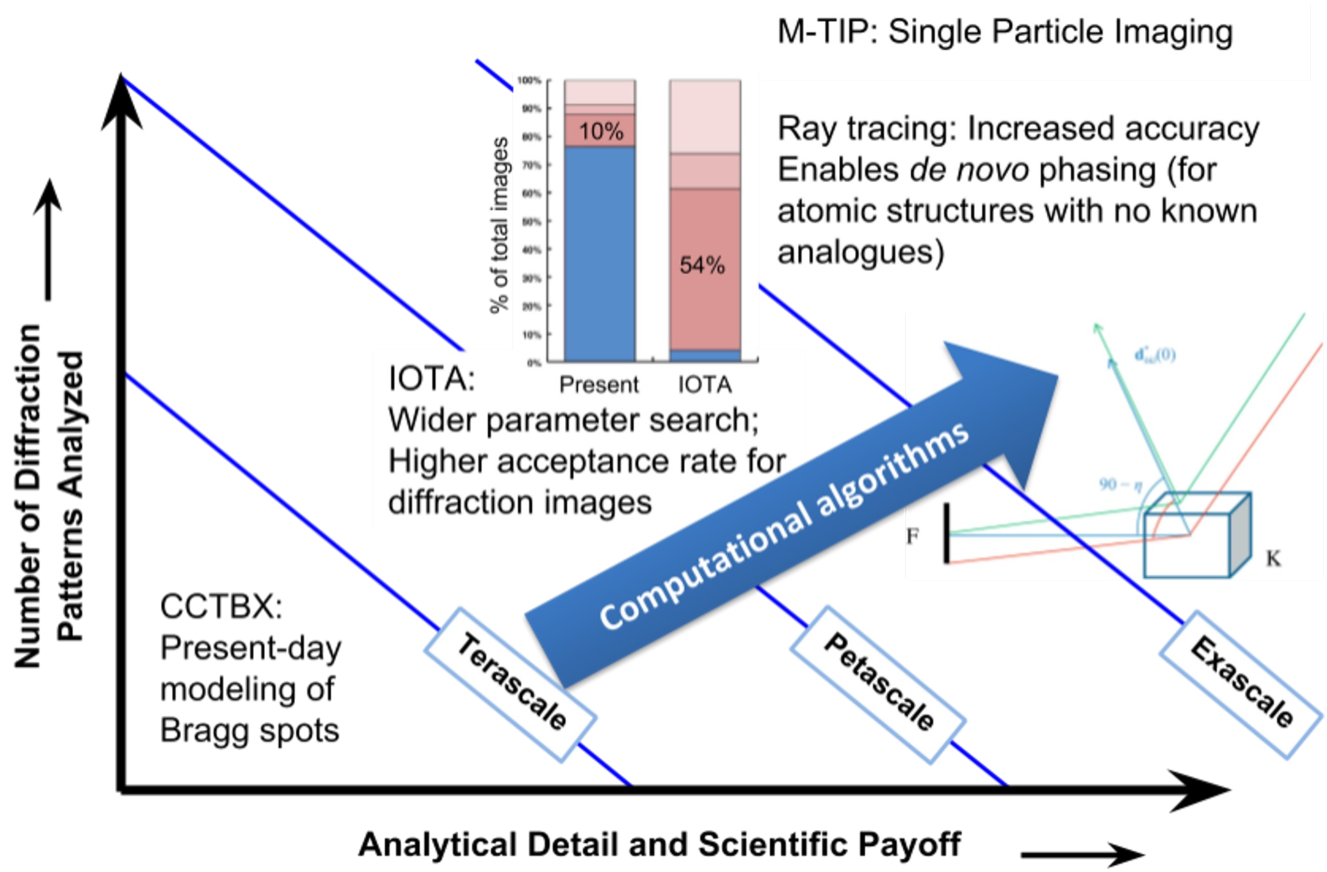

EXASCALE COMPUTING PROJECT

ExaFEL Challenge Problem: X-ray Crystallography (SFX)

Science Goal: enable the time-domain "movie" of an enzymatic reaction

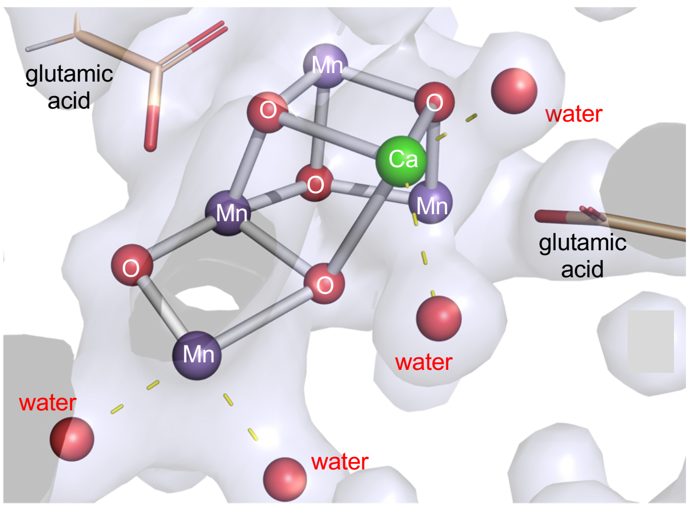

Photosystem II protein complex:

$2H_2O \rightarrow O_2 + 4H^+ + 4e^-$

Performance Goal: analyze LCLS-II data in near real time: 1 TB/s

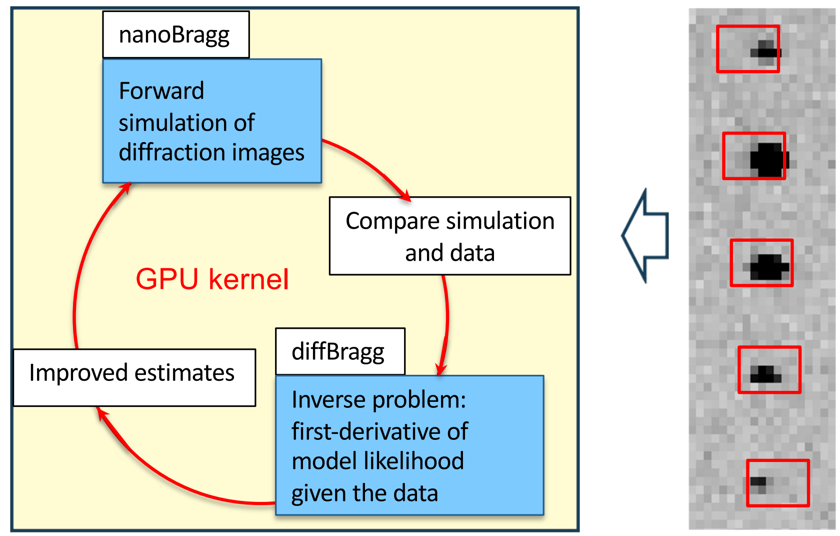

SFX Approach: apply L-BFGS to estimate structural parameters from diffraction data

Progress Towards KPP Demonstration

KPP Challenge Problem

- Analysis will simulate data processing in real time
- Planned Computation: ⅓ of Frontier for 1 hour
  - SFX analysis of 500k diffraction patterns
  - LCLS-II-HE: 5000 Hz (after data reduction)
  - Transfer 20 datasets to flash on Frontier
  - Each representing 100 seconds of data acquisition

KPP Risk Assessment	KPP Timeline
Very Likely	FY23/Q4

Status on Crusher

- Finalizing port of diffBragg kernel to Kokkos
- Currently no blockers

KPP demonstration planned for FY23/Q4 with new data collected in FY23/Q3

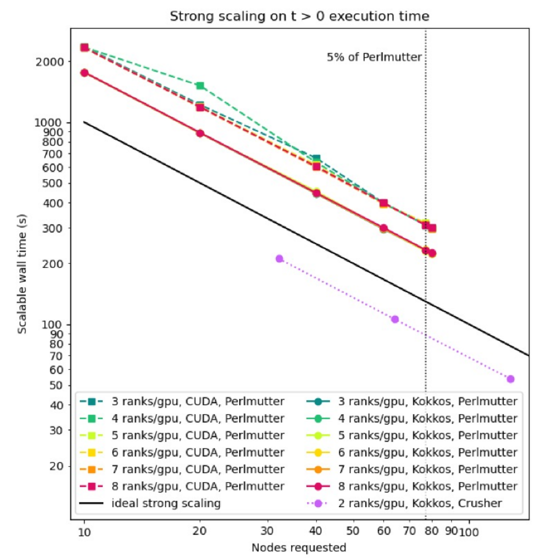

Science Impact of CCTBX Investments

- ExaFEL is extending the Computational Crystallography Toolbox (CCTBX) to support massive data sets
- Example: new data merging capability in CCTBX can process ~500GB of crystallography data in 1000 seconds

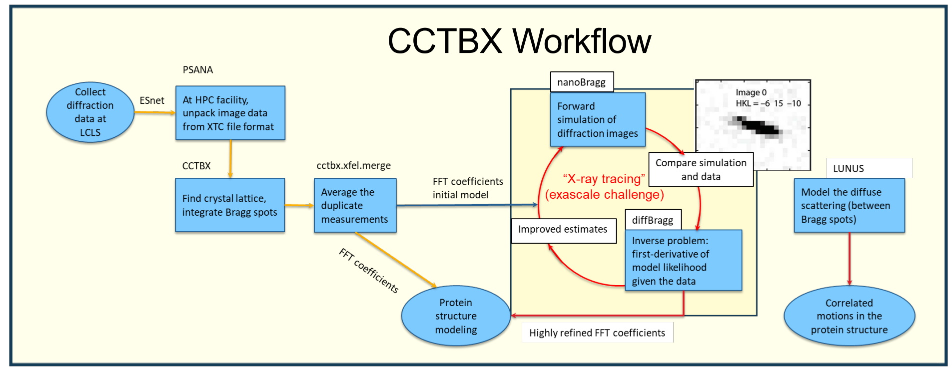

Using XFEL analysis for small molecules to design advanced materials

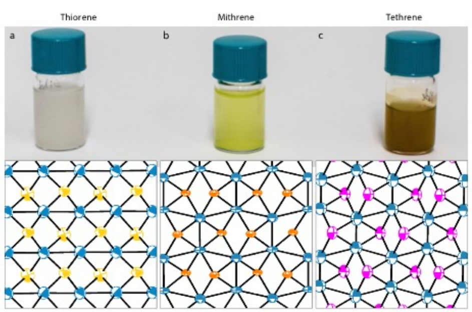

"Chemical crystallography by serial femtosecond X-ray diffraction", Nature, https://doi.org/10.1038/s41586-021-04218-3

Discovering enzyme mechanisms for photosystem II

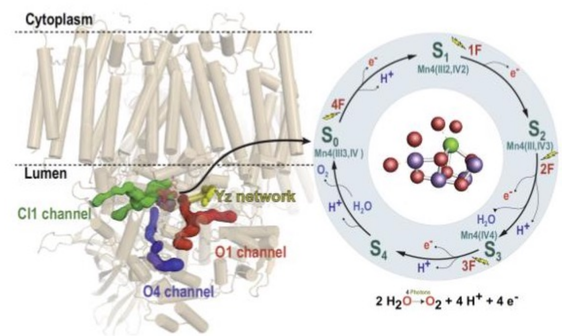

"Structural dynamics in the water and proton channels of photosystem II during the S2 to S3 transition", Nature Communications, https://doi.org/10.1038/s41467-021-26781-z

Coupling Experimental and Computational Facilities

- ExaFEL has been developing a data management infrastructure capable of
  - streaming data to remote computational resources
  - launching remote workflows automatically
  - returning results to experimenters in a web browser
- This infrastructure relies on the superfacility API to connect to NERSC

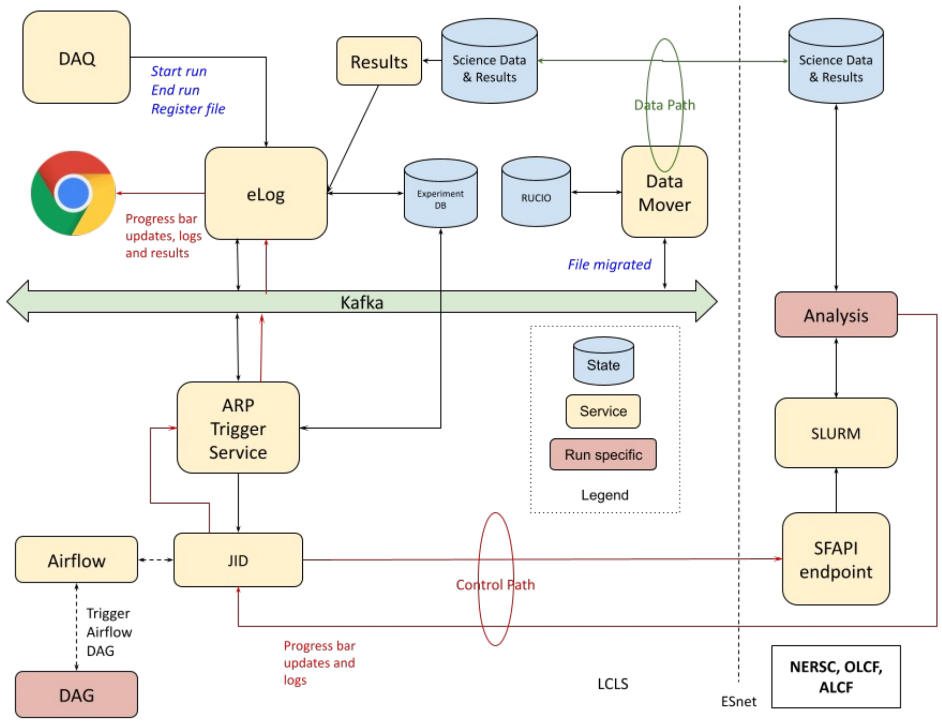

ExaSGD

ExaSGD:
Optimizing Stochastic Grid Dynamics at Exascale

Scalability Challenges for Power Grid Planning and Operation

- Goal: Computing optimal power flow based on forecasts and contingencies
- Objective: Reduce cost of power generation and power supply
- Constraints: Demand and supply balance, security, and stability
- Challenges:
  - Large number of forecast scenarios
  - Large number of contingencies
  - Multiperiod analysis for recovery/restoration

Project Base Goal:

- Security Constrained AC Optimal Power Flow (SC ACOPF): large-scale OPF calculation with many forecasts and contingencies (*challenge problem*)

Project Stretch Goals:

- Stochastic Multiperiod SC ACOPF: Multiperiod analysis of SC ACOPF (5-10 periods) to account for ramping of dispatchable power generators
- Frequency Restoration: Multiperiod analysis of SC ACOPF including frequency dynamics

ExaSGD Base Goal: Large Scale SC ACOPF

Posed as a 2-stage optimization problem:

Minimize

$\sum_t F_t(x_t)$	generator fuel cost
$+ \sum_{tsk} G_{tsk}(x_t, y_{tsk})$	wind curtailment, load shedding, power imbalance, etc.

Subject to:

$H_{tsk}(x_t, y_{tsk}) = 0$	flow definitions, power balance
$Q_{tsk}(x_t, y_{tsk}) \leq 0$	bounds: generator power, voltage, branch flow
$R_t(x_t, x_{t+1}) \leq 0$	generator ramping limit

The optimization problem structure is apparent in the underlying linear system:

$$
\begin{bmatrix}
K_1 & & & & & B_1 \\
& K_2 & & & & B_2 \\
& & K_3 & & & B_3 \\
& & & \dots & & \dots \\
& & & & K_N & B_N \\
B_1^T & B_2^T & B_3^T & \dots & B_N^T & K_0
\end{bmatrix}
\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \dots \\ y_N \\ x \end{bmatrix}
=
\begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ \dots \\ r_N \\ r_0 \end{bmatrix}
$$

Key Idea: The block-arrow coupling structure can be exploited to decompose the optimization problem.

Progress Towards KPP Demonstration

KPP Challenge Problem

- Security constrained optimal power flow
  - 10,000-bus US Western Interconnect
  - 10^5 contingency scenarios
- Planned Computation: 25% Frontier for 1 hours

KPP Risk Assessment	KPP Timeline
Likely	FY23/Q2

- Software stack deployed and tested on Crusher.
- Nightly tests running.
- Preliminary scaling results show promising performance
- Several bugs identified, debugging in progress.
- Peak performance ~10 PFLOPS on 64 Crusher nodes

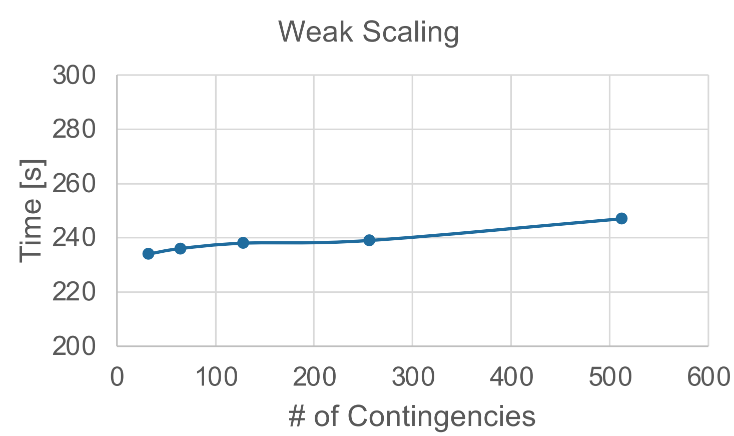

Optimal Power Dispatch for Texas Grid

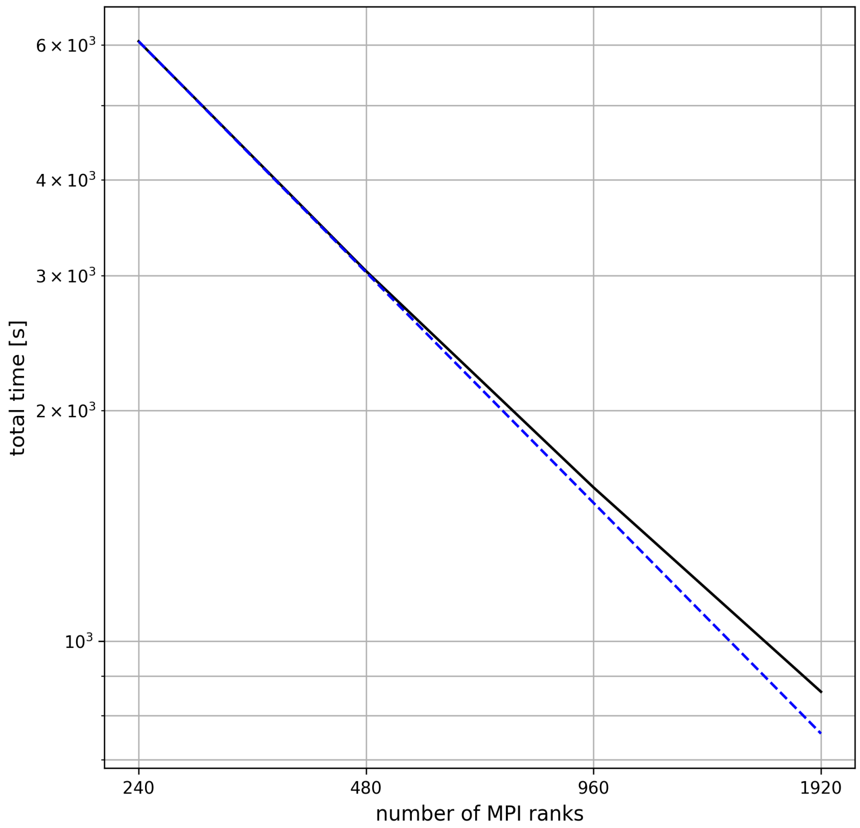

Strong scaling on Summit for the synthetic Texas-2000 bus grid with 1000 contingencies and 10 scenarios

Goal: 10 PF with all optimization loops running on GPU

Status: Ran formulation 3 on SUMMIT with HiOp primal decomposition solver (all 2nd stage model and optimization solver code running on GPU)

7TF/GPU * 1920 GPU * 0.7 = **9.4 PF**
(70% utilization was the max observed when using non-pivoting factorization in Magma)

Optimal power dispatch computed within 20 min for a 2000-bus grid with 1000 contingencies and 10 different weather scenarios.

ECP EXASCALE COMPUTING PROJECT

Key Challenge: Tailoring Second-Stage Problem for GPUs

$$
\begin{aligned}
R_1(x) = \min_{y_1} \quad & f_1(x, y_1) \\
\text{s.t.} \quad & g_1(x, y_1) = b_1, \\
& y_1 \geq 0.
\end{aligned}
$$

$$
\begin{bmatrix} i_\alpha \\ 0 \end{bmatrix} = \begin{bmatrix} Y^{\text{bus}}_{\alpha,\alpha} & Y^{\text{bus}}_{\alpha,\beta} \\ Y^{\text{bus}}_{\beta,\alpha} & Y^{\text{bus}}_{\beta,\beta} \end{bmatrix} \begin{bmatrix} v_\alpha \\ v_\beta \end{bmatrix}
$$

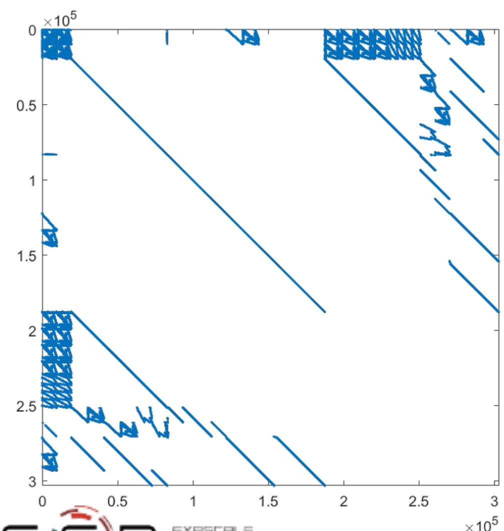

Challenge: Suitable GPU-accelerated sparse direct solvers are not anticipated for Frontier or Aurora*

Idea: Reformulate ACOPF subproblems to create dense formulations suitable for dense linear solvers

Impact:

- Dense 10k bus ACOPF formulations can fit on target GPUs
- 20-70% of peak GPU performance, depending on stability requirements
- Efficiency degrades (compared to sparse) as problem size increases

Kron Reduction (CPU; 1/solve)

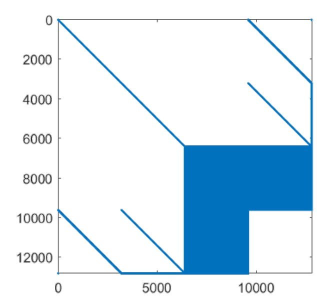

Schur complement Reduction (GPU; 1/iteration)

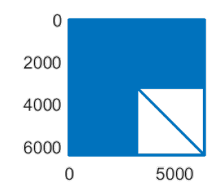

Challenge: We need sparse linear solvers for constrained optimization that run efficiently on GPUs

- Linear systems arising in optimal power flow analysis are symmetric, indefinite, very sparse and often ill-conditioned.

- ExaSGD Engagement (starting in FY22)
  - Sparse matrix factorization within Ginkgo framework
    - Collaboration with LBNL and KIT within ECP
    - Solver developed with CUDA and HIP backends and interfaced with HiOp.
  - Hybrid direct-iterative method for optimal power flow analysis
    - Collaboration with Stanford to develop HyKKT solver
    - Preliminary results show 2x speedup over CPU-only MA57 software
  - Developed cuSolver-based solver with help from NVIDIA
    - Extended the ExaSGD HiOp solver.
    - Latest profiling results show 3x speedup in numerical factorization.

ExaSGD will have a broad legacy for power grid optimization and related applications

Problem Formulation and Solution Methods	ExaSGD Capabilities
Application formulation	Stochastic Multiperiod SC ACOPF, Frequency Restoration
Optimization equations	Scalable optimal power flow formulations
Computational models*	ExaGO, ExaTron, Powerscenarios
Optimization solvers	Hiop
Numerical methods	Ginkgo, HyKKT

*ExaSGD is exploring both C++ and Julia approaches for modeling and solvers

ExaBiome

ExaBiome:
Exascale Solutions to Microbiome Analysis

How do microbes affect disease and growth of switchgrass for biofuels?

What happens to microbes after a wildfire?

Challenge Problem: Metagenomic Assembly
- Find species, genes and relative abundance in microbial communities

Stretch Goal: Metagenome Analysis
- Improve understanding of tree of life for microbes; aid in identifying gene function
- Track microbiome over time or space, changes in environment, etc.

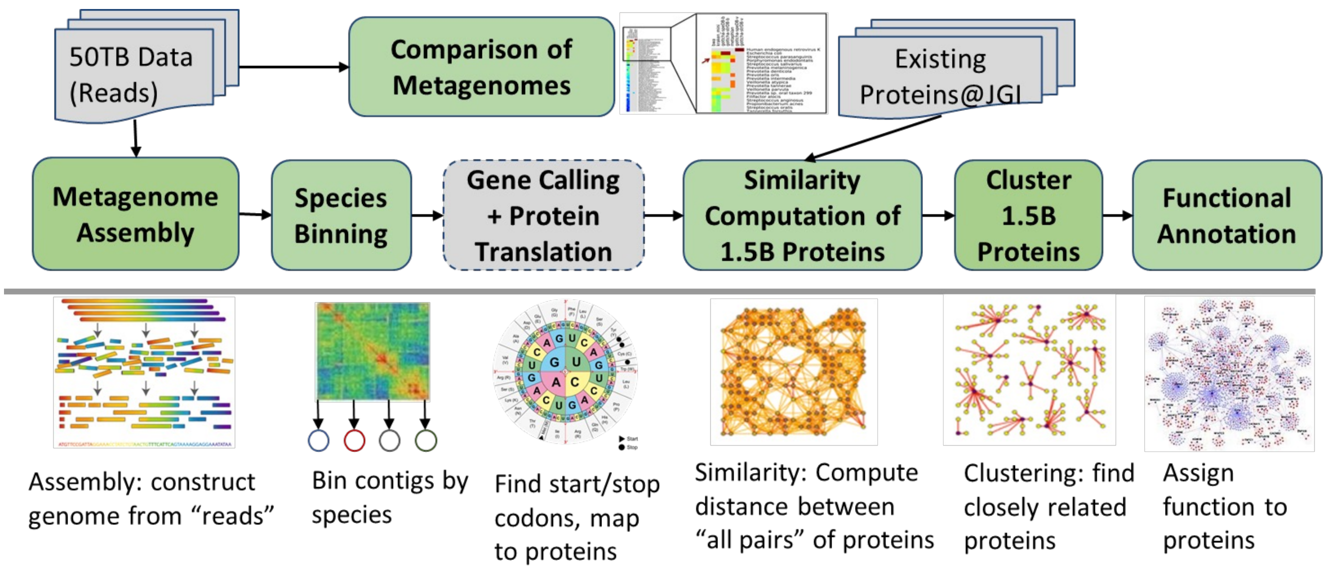

MetaHipMer: Metagenomic Sequence Assembly

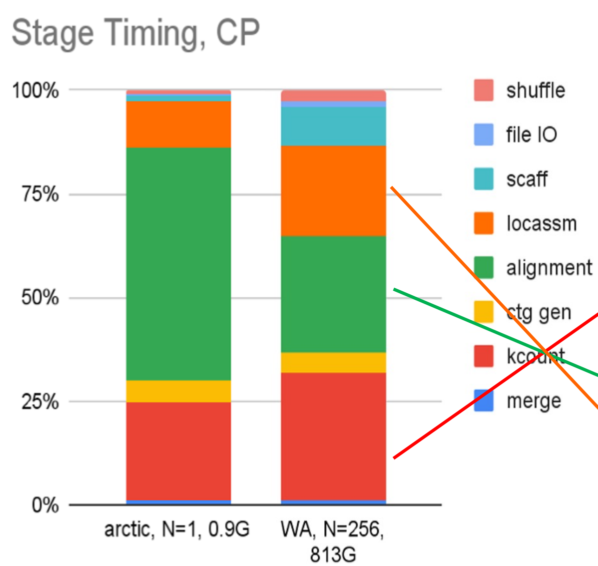

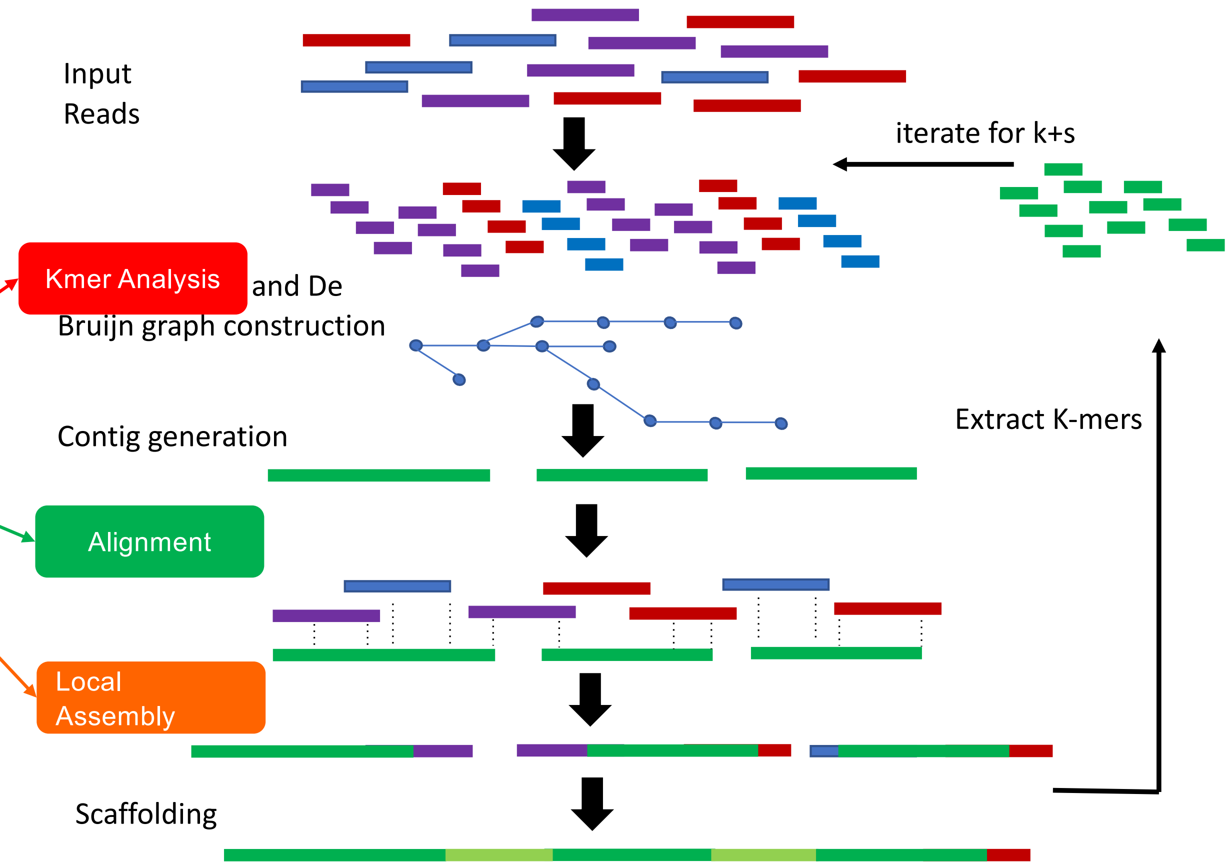

On CPU systems, 3 computations dominate runtime

Progress Towards KPP Demonstration

KPP Challenge Problem

- Tara Oceans Assembly
  - Microbial data from all oceans, collected from 2009-13
  - 84 Terabytes, never before co-assembled
- Planned Computations:
  - 50TB: ½ of Frontier for 1 hour
  - 84TB: ¾ of Frontier for 1 hour

KPP Risk Assessment	KPP Timeline
Likely	FY23/Q3

Created miniapps and ported these to HIP using hipify script

- Alignment kernel - ported and integrated
- Kmer Analysis kernel - ported and integrated
- Local Assembly - uses NVIDIA-specific intrinsics; have a workaround and will integrate while looking for better approach

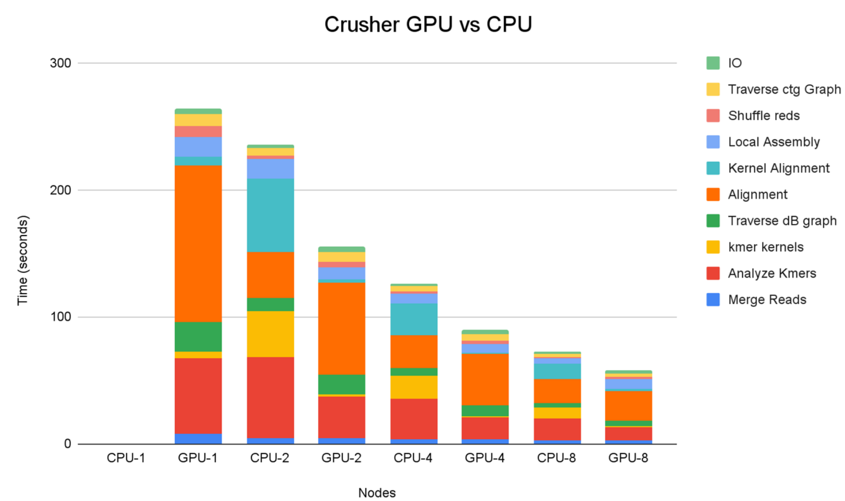

Scientific Impact of Large-Scale Assembly Computations

Enabling a sequence of largest-ever metagenomic assemblies

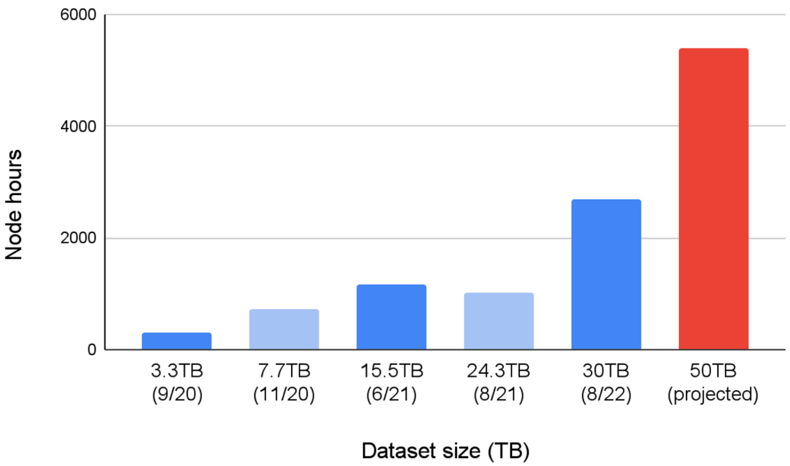

Big Data, Big Compute → Better Science

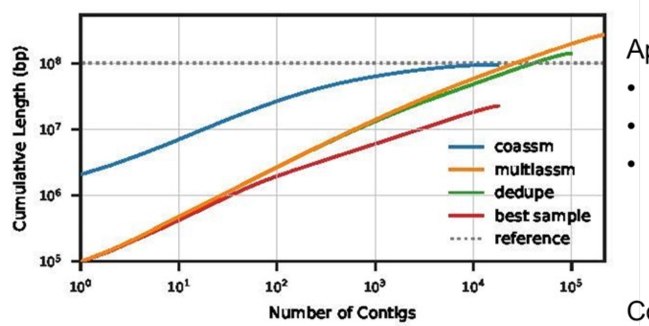

800 GB of soil (Western Arctic, 12) data plus synthetic data from 64 reference genomes

Approaches:

- Co-assembly: entire data set
- Multi-assembly: lane at a time
- Dedup: remove duplicates from multi-assembly

Co-assembly gives longer, less redundant assemblies, and is only possible with HPC

Microbial metabolic dependency and its impact on the soil carbon (3.3 TB)

GRE: Microbial carbon transformations in Wet Tropical Soils (7.7TB)

Lake Mendota time series: temporal dynamics of microbial carbon processing

Tara oceans: worldwide expedition to sample microbes from across the oceans (image: ©G.Bounaud/ C.Sardet/ Soixanteseize/ Tara Expéditions)

EXASCALE COMPUTING PROJECT

Terabase-scale metagenome assembly drives discovery at JGI

Challenge: Scaling metagenome assembly to increasingly large datasets

Terabase-scale metagenome assemblies would not be possible without MetaHipMer

Peatland ecosystems (**0.6 TB**)

CSP #1445
Erik Lilleskov

Ectomycorrhizal fungi (**1.5 TB**)

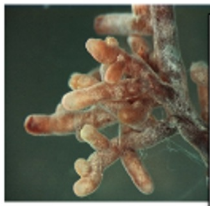

FICUS #503579
Jenny Bhatnagar

Subsiding subtropical soils (**1.5 TB**)

CSP #503337
Willm Martens-Habbena

Great Redox Experiment (GRE) in wet tropical soils (**8TB**)

CSP #502924
Jennifer Pett-Ridge

Alicia Clum

Robert Riley

Rob Egan

ExaBiome

Prescribed fire (**1.5 TB**)

CSP #503203
Nicholas Dove

Dark matter fungi (**1.7 TB**)

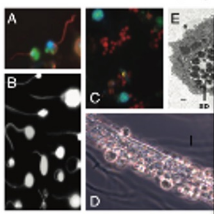

CSP #1626
Tim James

Coastal mangroves (**2.1 TB**)

CSP #503542
Brandi Reese

Largest ever metagenome assembly enabled by Summit

Mountainous watershed (**2.7 TB**)

FICUS #503568
Eoin Brodie

Soil carbon cycle (**3.3 TB**)

CSP #503502
Dan Buckley

Lake Mendota time series (**25 TB**)

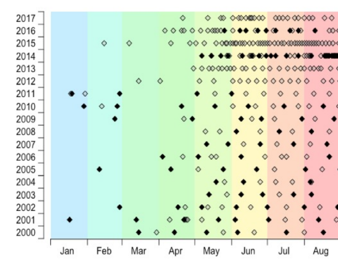

CSP #504350
Trina McMahon

Under development
Largest assembly of ~500 metagenomes from 20-year time series to evaluate long-term ecological and evolutionary microbial processes

Transformative contributions of ExaBiome

First-of-kind HPC tools, now exascale capable

1. MetaHipMer: metagenome short read assembler
2. PASTIS/HipMCL: protein alignment & clustering
3. KmerProf: tool for comparing metagenomes
4. ELBA: assembler for long read assembly

Worked closely with and relied heavily on tools from the PAGODA (1,3) and ExaGraph (2,4) teams

SC18 Dallas, TX | hpc inspires. — Best Paper Finalist

SC21 St. Louis, MO | science & beyond. — Best Paper Finalist

SC22 Dallas, TX | hpc accelerates. — Gordon Bell Finalist

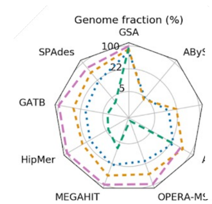

2022 Critical Assessment of Metagenome Interpretation

"...best ranking method across metrics and all data sets..."

Future: production support, plus new and augmented analyses using AI methods and hardware

ECP EXASCALE COMPUTING PROJECT

Summary

ECP Impact Extends Far Beyond Core KPP Demonstrations

Project	Stretch Goals
CANDLE	• Large Language Models (e.g. Transformers) • COVID-19 Docking Models
ExaFEL	• M-TIP: Single Particle Imaging • Cross-Facility Resource Management
ExaSGD	• Multi-Period Stochastic Optimal Power Flow • Frequency Restoration • Sparse Linear Solvers
ExaBiome	• Protein Clustering • Long-Read Metagenomic Assembly

Next Steps

- Frontier
  - All of the projects seem likely to meet their KPP objectives on Frontier
- Aurora
  - CANDLE is already well-positioned for a KPP demonstration
  - ExaBiome and ExaFEL are good candidates for FY24 investments
  - Possible deployment of ExaSGD using Ginkgo sparse linear solver
- Post-ECP Funding
  - CANDLE has a strong post-ECP funding plan (ASCR/NIH)
  - Funding plans are uncertain for other projects

Observations and Lessons Learned

- Performance portability abstractions
  - Existing performance portability abstraction techniques appear well-suited for DAO applications
- GPU parallelism
  - Half of the DAO applications have seen non-trivial challenges developing GPU kernels
- Key dependencies
  - Most of the DAO applications have critical dependencies that are shared with few other ECP applications
- Performance bottlenecks
  - Data management and movement is a key challenge, and on-node memory can be constraining
- Continuous Integration
  - Critical to catalyze the development of new capabilities

Comparing DAO Projects

- CANDLE vs ExaFEL
  - Both rely on similar numerical algorithms for continuous, unconstrained parameter estimation
  - CANDLE was able to leverage significant prior investment in AI/ML toolkits
- ExaFEL vs ExaSGD
  - Both are focused on continuous optimization
  - ExaSGD requires *constrained* optimization methods, which generally require sparse linear solvers
  - The lack of suitable linear solvers significantly inhibited the development of ExaSGD capabilities
- ExaBiome: MetaHipMer vs HipMCL
  - ExaBiome's sequence assembly algorithms in MetaHipMer involve discrete algorithms
    - These are qualitatively distinct from methods used by other ECP applications
    - Parallelization of sequence assembly was quite complex (e.g. GPU assembly kernels)
  - ExaBiome's protein clustering algorithm in HipMCL leverage distributed matrix algebra
    - Parallelization of matrix algebra is relatively well-understood, which accelerated HipMCL development