Workflows are the New Applications. So What?

Shantenu Jha
Rutgers University & Brookhaven National Laboratory

http://radical.rutgers.edu

HPC Workflows: Motivation

- Unsustainable trends for performance & scale
  - **Think beyond single task performance!**
  - "..room at the bottom, even more room sideways!"
- Scientific applications are increasingly
  - **Assembled** using multiple components (e.g., modsim, post-processing), into **higher-level parallel patterns and motifs** (e.g., ensembles, steer)
  - Consist of **multiple stages**, leveraging multiple compute resources and facilities.

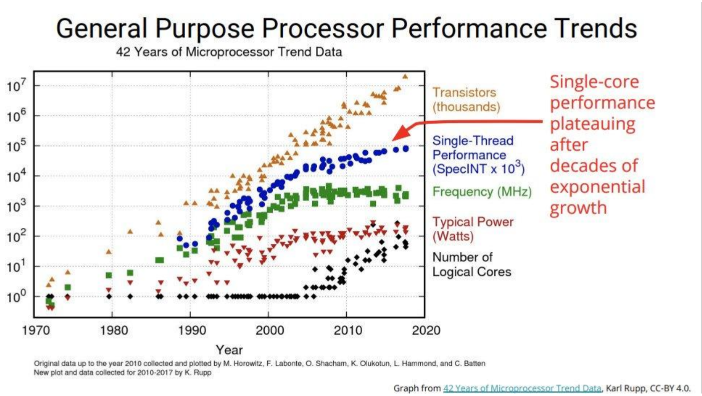

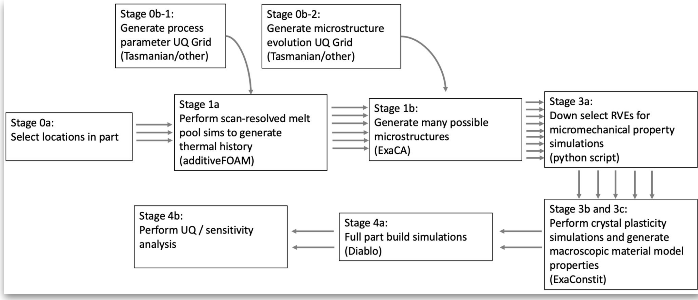

Workflow Motivation: Major Changes at Extreme Scales.

- Rate of computing capability, measured in **peak** FLOPs, will not increase as rapidly as in the past
  - Scientific advances to be driven by innovative algorithms and software
- Supply and demand: Not all ExaFlops are equal!
  - Microsoft #3 on Top 100
  - EF-days vs EF-hours vs EF-second
- How & when, not if convergence of HPC & Cloud
  - Cloudify HPC vs HPCfy Clouds ?
  - Currently economics on-premises HPC deployment; may change?
  - Will clouds define new architecture for modeling & simulation as they have for AI?

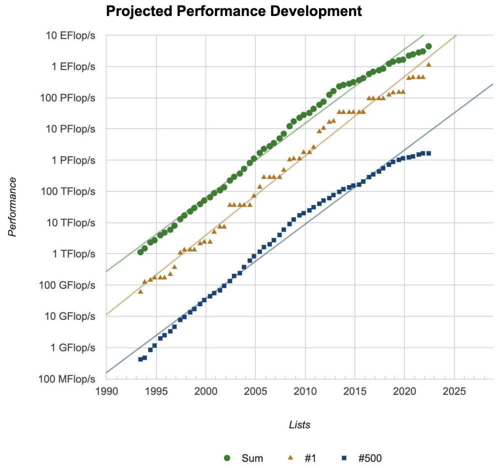

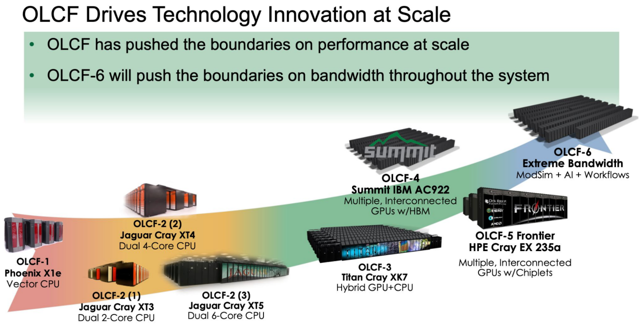

Two illustrative Examples

National Virtual Biotechnology Lab (NVBL)

- National Virtual Biotechnology Lab (NVBL)
    - https://science.osti.gov/nvbl
- Aid U.S. policymakers in responding to the COVID-19 pandemic with epidemiological information for decision making
- Accelerate production of critical medical supplies across the nation
- **Supercomputing and artificial intelligence for design of targeted therapeutics**
- Leverage chemical testing & analysis to facilitate new antigen and antibody testing

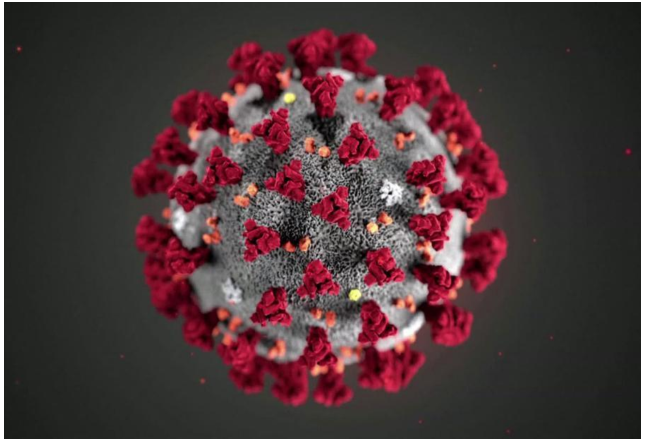

BROOKHAVEN
NATIONAL LABORATORY

National Virtual Biotechnology Lab (NVBL)

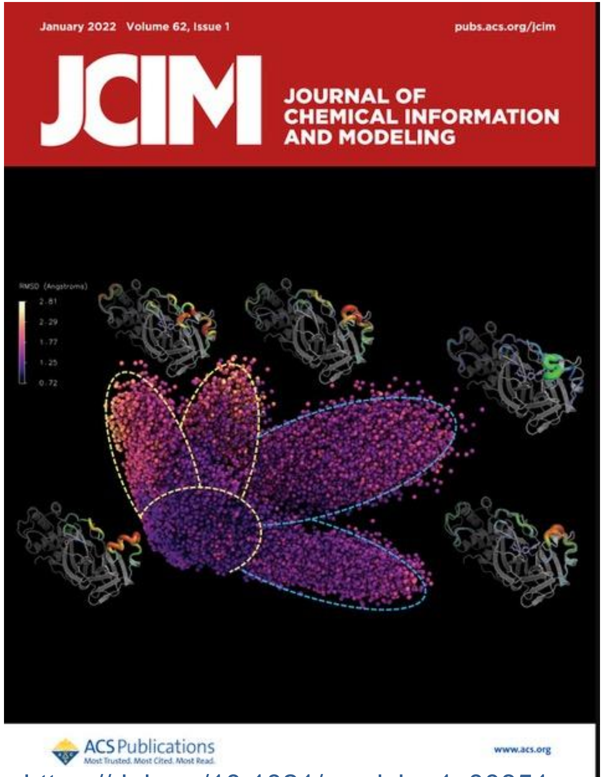

After analyzing 100B candidates, we proposed: MCULE-5948770040 interacts with the SARS-CoV-2 main protease observed in the crystal structure 7LTJ.

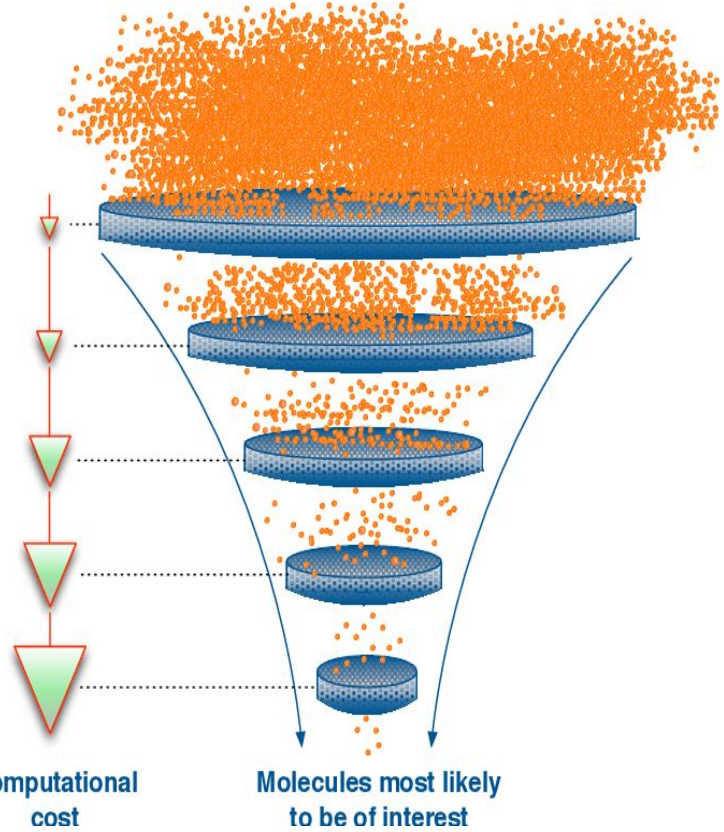

https://doi.org/10.1021/acs.jcim.1c00851

NVBL given US Secretary of Energy Honour Award (2021)

Impacting SARS-CoV-2 Medical Therapeutics

- **Scale of Operation:**
  - **~10^{11}** Docking calculations
  - ~10^{3} ML-driven MD calculations
  - ~5 x 10^{4} Binding Free Energy Calculations
  - **~2.5 x 10^{6} node-hours**
- Peak Performance
  - ~ **8000** nodes on Frontera
  - ~ **4000** nodes on Summit
- ***O(100B) possible due to both methodological advances (e.g., AI-coupled HPC methods) at multiple stages and scalable execution***

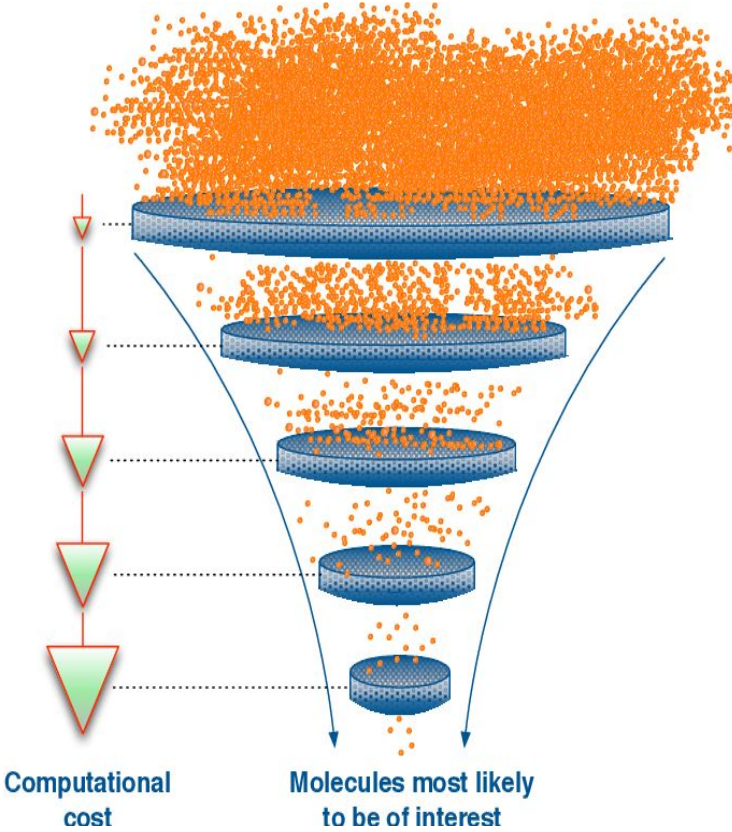

BROOKHAVEN NATIONAL LABORATORY

ExaAM UQ Campaign: Large Ensemble-based Workflow

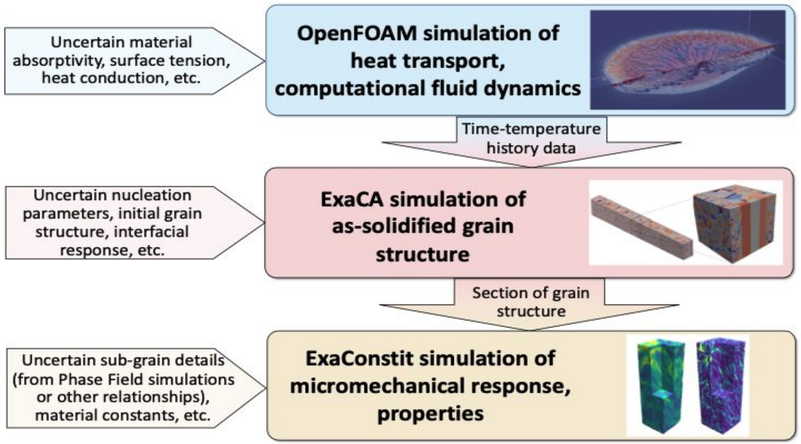

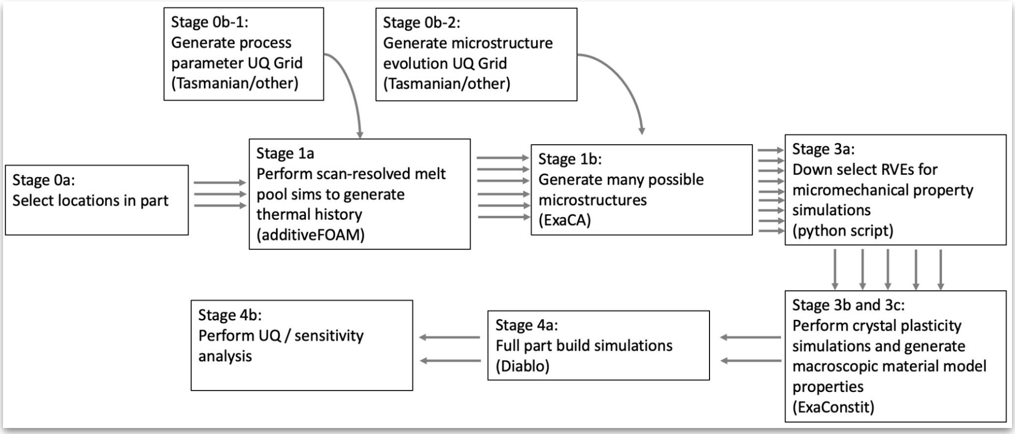

- Ability to coordinate a campaign of ensemble of pipelines
  - Parallelism at multiple levels; pipelines progress asynchronously
- Each stage of a pipeline is a different workflow, with different workloads
  - Each workflow has diverse spatio-temporal properties, resource management, and scheduling requirements
- Fault-tolerance of the tools and executing processes
  - Manage and recover from failures at large-scale

Stage 3 (ExaConstit) on Frontier

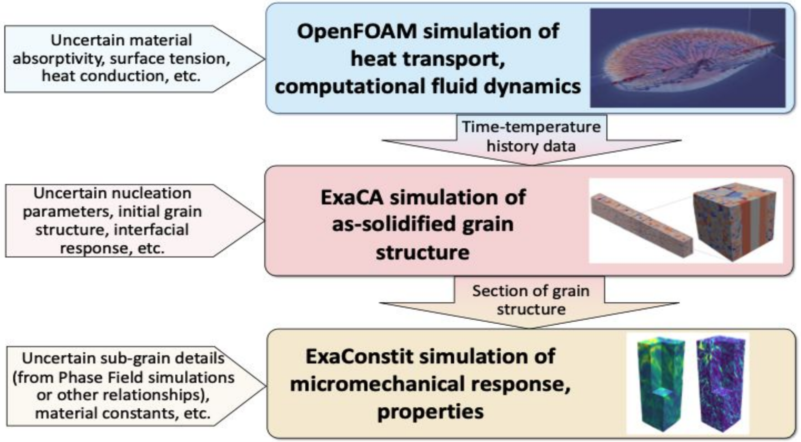

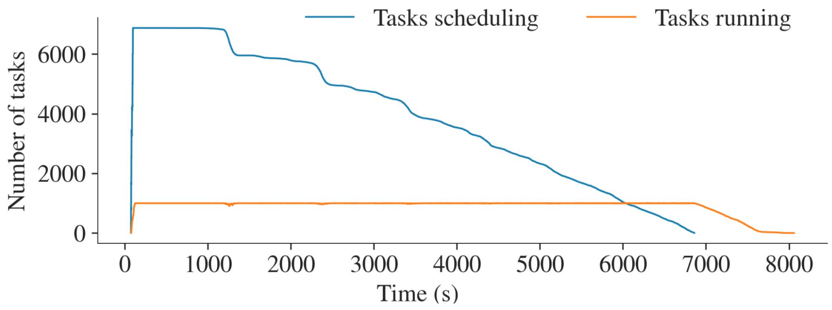

7875 concurrent task in scheduling and running states.

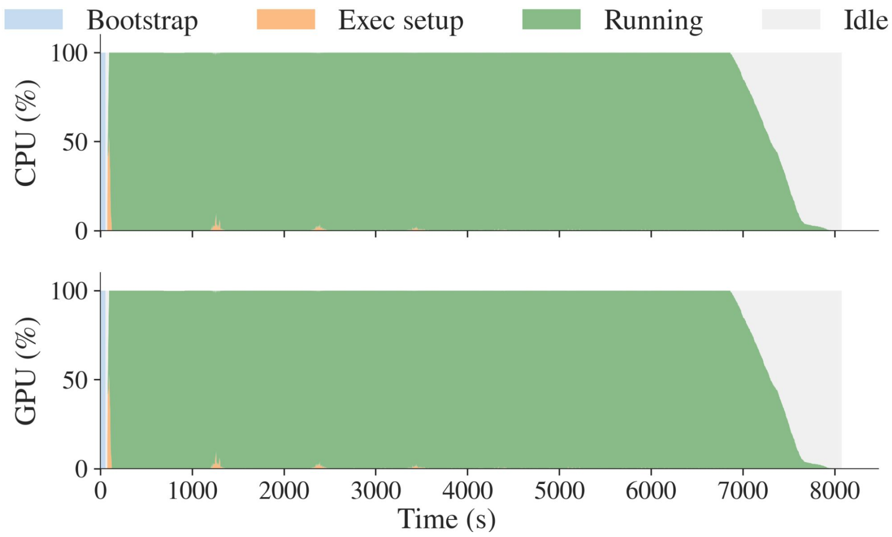

Resource utilization: 100% corresponds to 448,000 CPU cores and 64,000 GPUs.

2. AI-coupled HPC Workflows

AI-coupled HPC Workflows

- Online concurrent coupling of AI and HPC, **not offline** training and inference
  - Heterogeneous mix of simulations, training, inference, retraining tasks
  - Compute time partition: Building and using AI **vs** generate data and insight

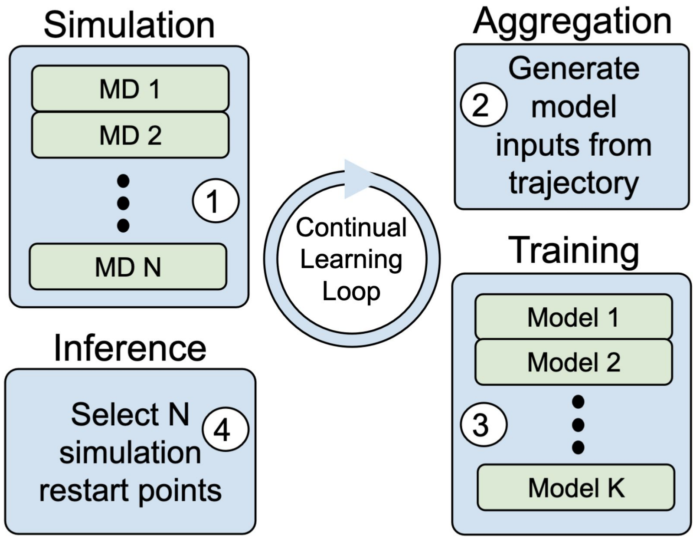

AI-coupled HPC Workflows: Motivation and Modes

- **Effective performance** measured by *"science for a given amount of computing"*
- Algorithmic Moore's Law: **Learning Everywhere** [1,2] Effective performance of 10^N by coupling AI/ML with HPC at multiple levels
- **AI-x-HPC** modes not mutually exclusive
  - AI-**in**-HPC is learning and replacing traditional simulations
  - AI-**out**-HPC is "in charge" of simulation and optimizing execution
  - AI-**about**-HPC is AI running synergistically with simulations

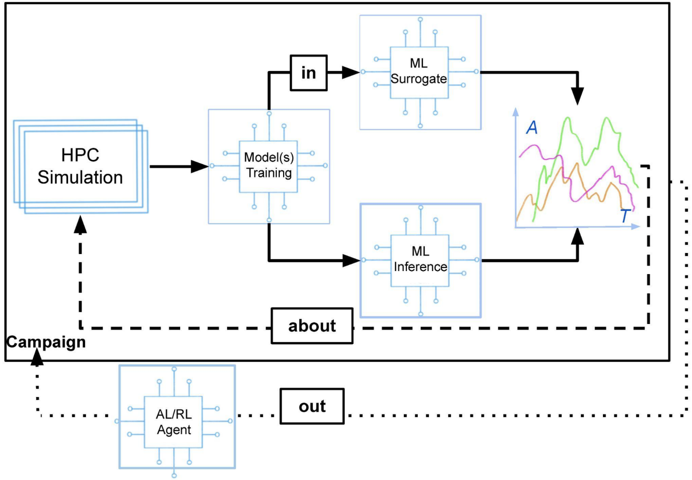

Review

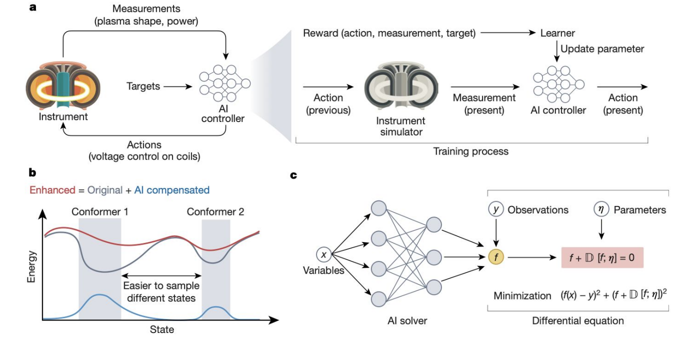

Fig. 4 | Integration of AI with scientific experiments and simulation. potential energy, allowing the system to escape local minima (in grey) and

[1,2] https://arxiv.org/abs/2208.11745 and https://arxiv.org/abs/1909.13340

https://doi.org/10.1038/s41586-023-06221-2

Surrogates for Simulations

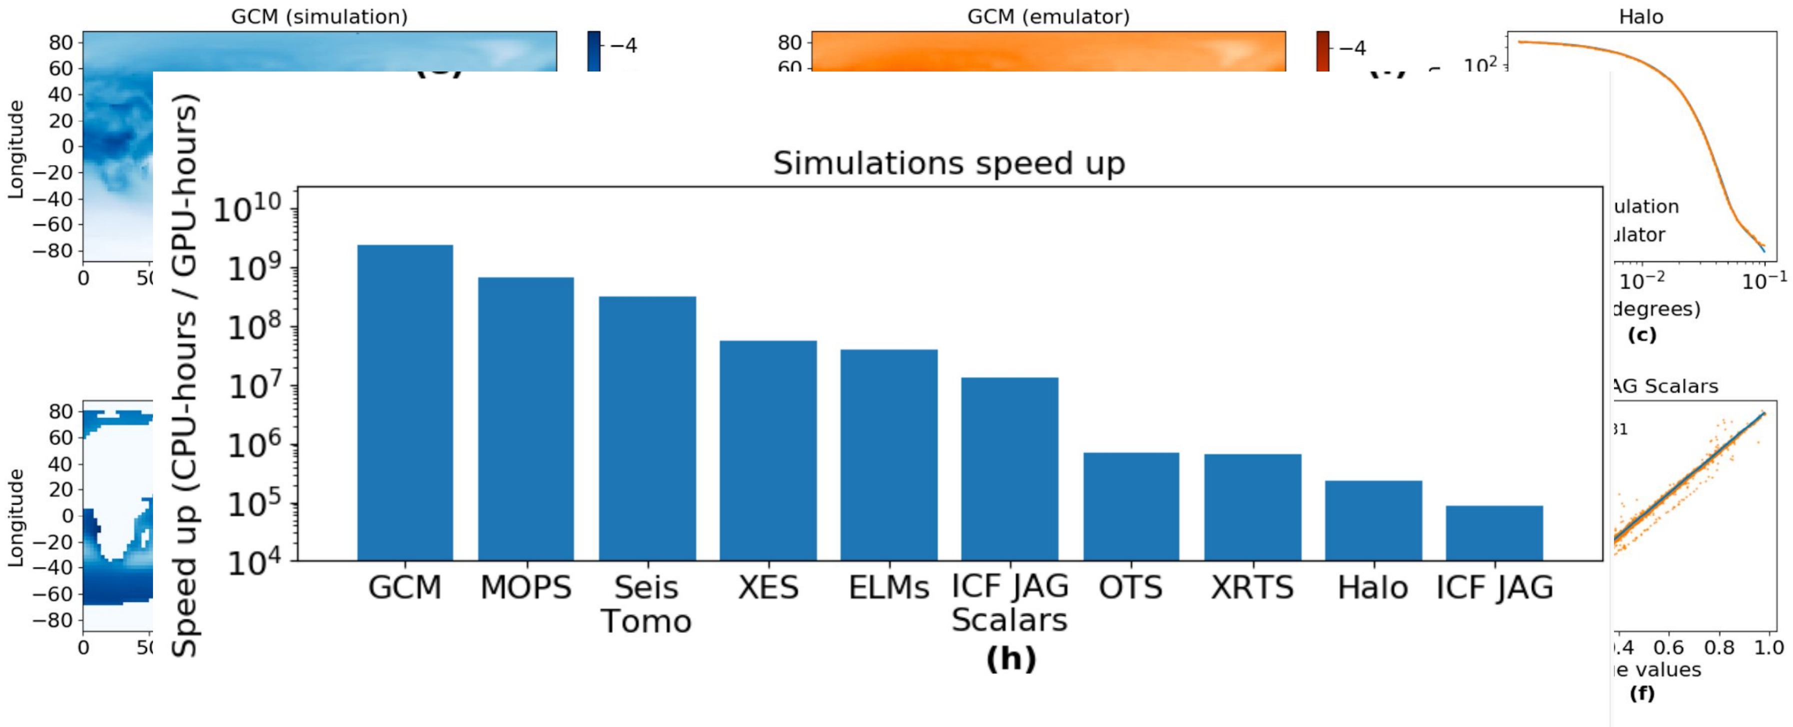

https://arxiv.org/abs/2001.08055

Surrogates enhance performance without loss of accuracy

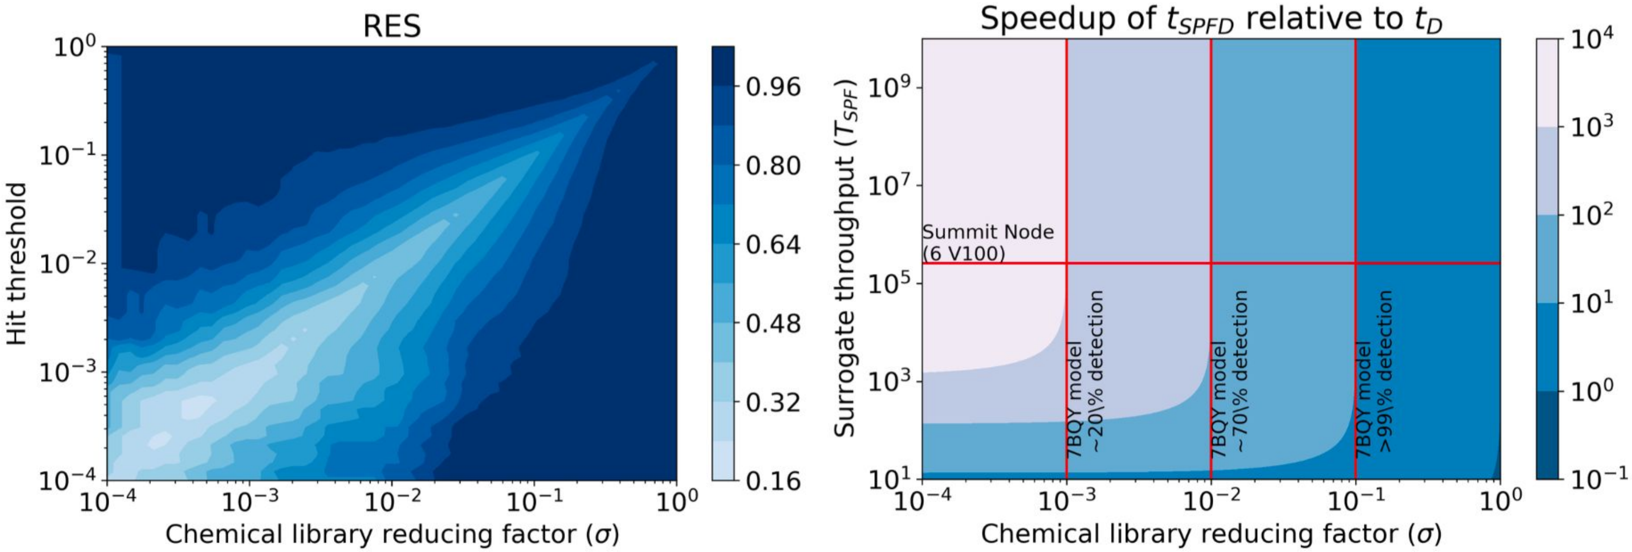

AI-accelerated Protein-Ligand Docking for SARS-CoV-2 is 100-fold Faster With No Significant Change in Detection (i.e., accuracy)

https://doi.org/10.1038/s41598-023-28785-9

AI-about-HPC: 100x protein folding with the physics

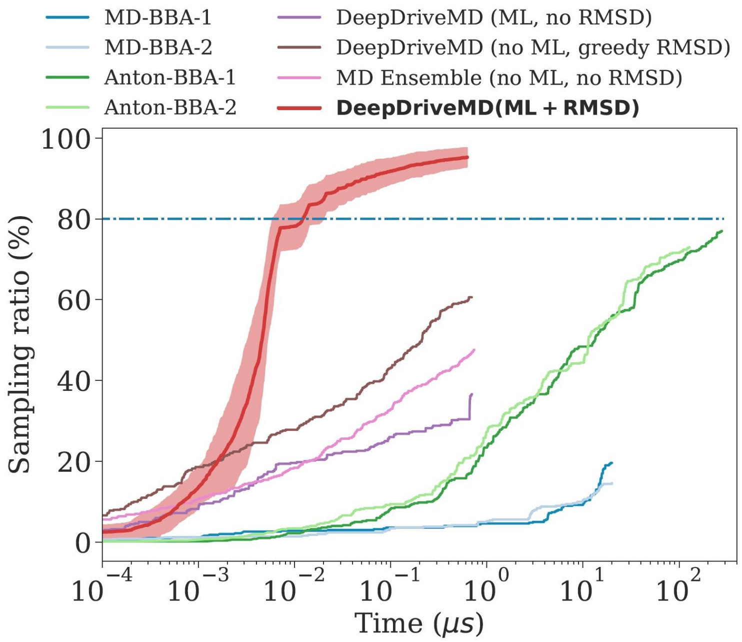

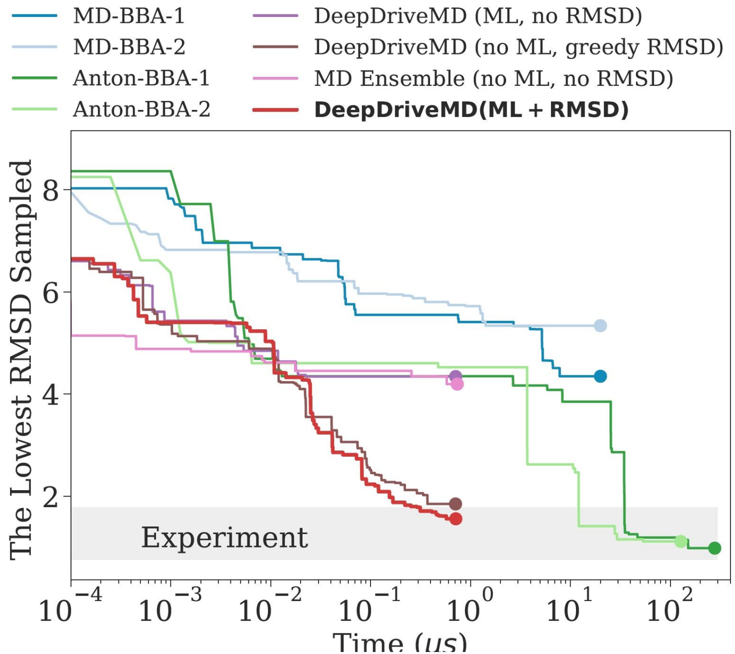

Achieving 100X faster simulations of complex biological phenomena by coupling ML to HPC ensembles. Alexander Brace, Hyungro Lee, Heng Ma, Anda Trifan, Matteo Turilli, Igor Yakushin, Todd Munson, Ian Foster, Shantenu Jha, Arvind Ramanathan: https://arxiv.org/pdf/2104.04797.pdf

AI-coupled HPC Workflows

- HPC Workflows to train and fine-tune AI models
  - Few can/will develop LLM/FM from scratch, but **many** will fine-tune them locally
  - RADICAL Optimal Surrogate Explorer (ROSE) For UQ-based O[ptimal model]
- HPC Workflows using multiple AI models
  - Concurrent and parallel workflows within AL loop
  - Scalable Active-Learning for Material Structure Determination with Neutron Diffraction

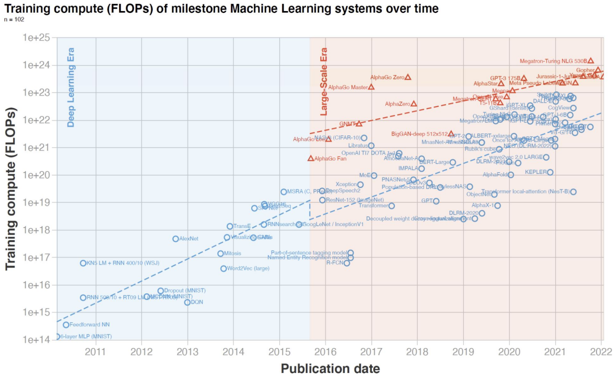

EF day

EF second

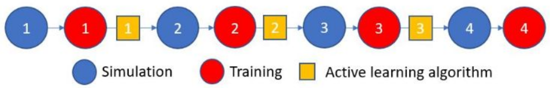

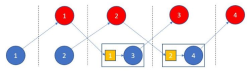

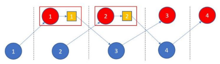

Learning Everywhere: #CovidIsAirborne

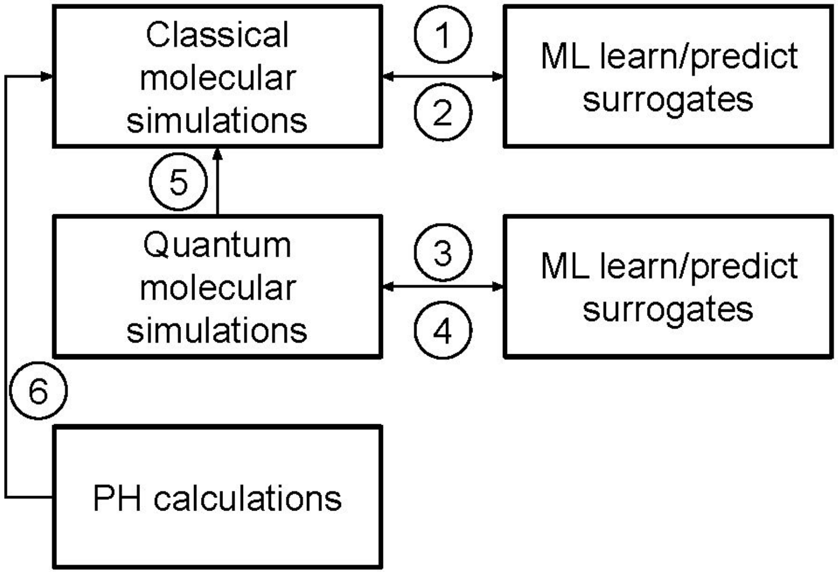

(b)

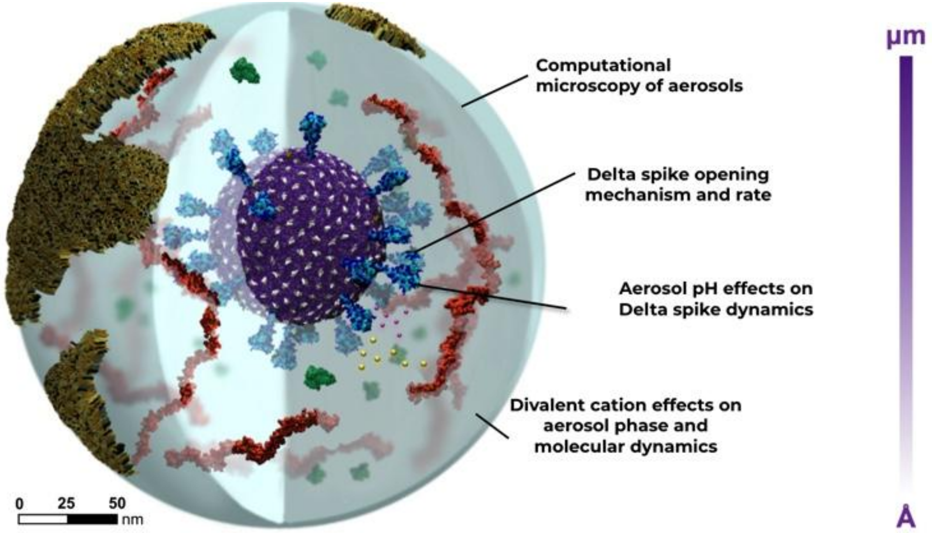

Collaboration led by Rommie Amaro (UCSD)
https://doi.org/10.1177/10943420221128233

Improving molecular dynamics accuracy using OrbNet

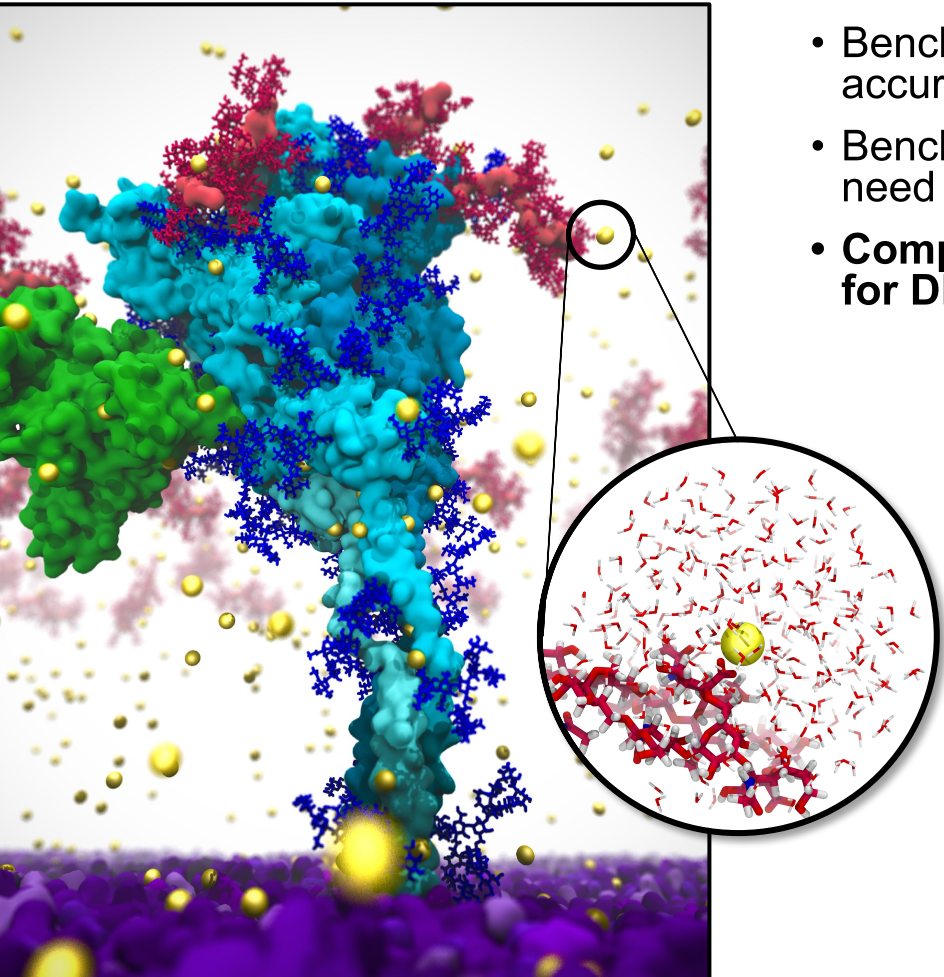

- Benchmarking OrbNet vs. DFT confirms quantum accuracy for Ca2+ interactions
- Benchmarking CHARMM36M vs. OrbNet reveals the need for electrostatic polarization
- **Computational cost reduced from 115M core-hours for DFT to 100k core-hours for OrbNet**

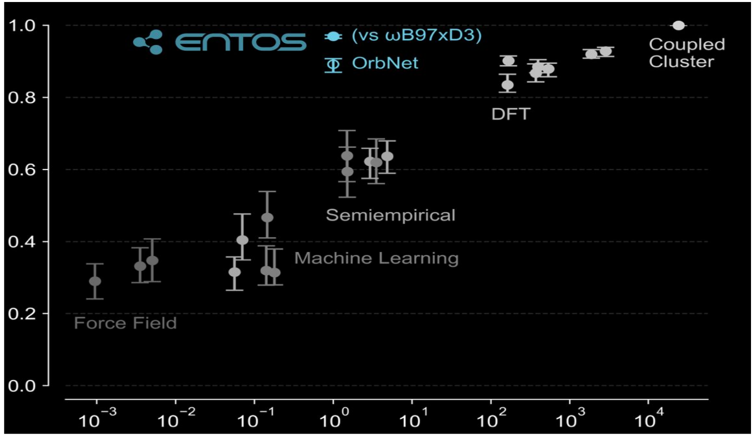

Learning Everywhere: IMPECCABLE 2.0

IMPECCABLE 2.0 Workflow - Accelerating COVID-19 drug discovery

ML components

- **Docking surrogate:** faster predictions of docking scores for huge chemical libraries.
- **Pose optimizer:** a model to rank binding poses to select the best starting pose for binding affinity calculations.
- **Structure generator:** A generative model to create new "optimized" compounds with desired properties such as low binding affinities.

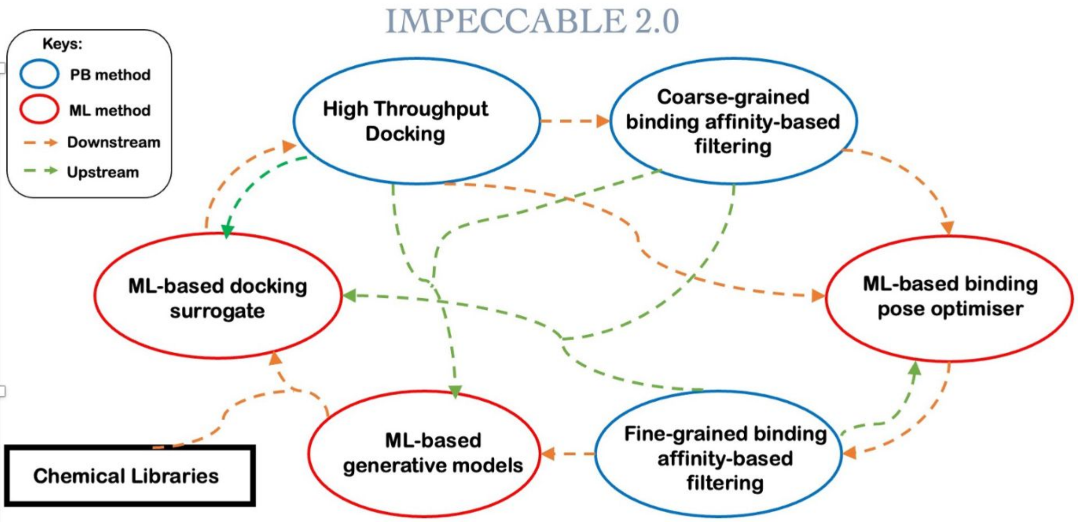

PB components

- High-throughput docking
- Reproducible and precise binding affinity calculations: ESMACS

3. So What? Systems and Software

Workflows @ Extreme Scale: Why is this challenging?

- **Workload comprised of Heterogeneous tasks** at multiple levels
  - Coupled AI-HPC
  - High-throughput function calls
  - Ensembles of MPI tasks
  - Spatio-temporal variation within each
- **Collective** versus single-task performance
  - Campaigns are ”integrated” workflows: WF1 and WF4 differ by 10^7x in computational cost
  - Producers of data (WF1) and consumers (ML1)
- **Adaptive Execution** at multiple levels
  - Workload: Task mix varies over campaign
  - Tasks: Run for varying duration

1000x variation in workflow throughput

Table 3: Throughput and performance measured as peak flop per second (mixed precision, measured over short but time interval) per Summit node (6 NVIDIA V100 GPU).

Comp.	#GPUs	Tflop/s	Throughput
ML1	1536	753.9	319674 ligands/s
S1	6000	112.5	14252 ligands/s
S3–CG	6000	277.9	2000 ligand/s
S3–FG	6000	732.4	200 ligand/s

10^7x variation in cost across workflows

Table 2: Normalized computational costs on Summit.

Method	Nodes per ligand	Hours per ligand (approx)	Node-hours per ligand
Docking (S1)	1/6	0.0001	~0.0001
BFE-CG (S3-CG)	1	0.5	0.5
Ad. Sampling (S2)	2	2	4
BFE-FG (S3-FG)	4	1.25	5
BFE-TI (not integrated)	64	10	640

BROOKHAVEN NATIONAL LABORATORY

- Master Task

One or more **Master tasks** are scheduled, placed and launched.

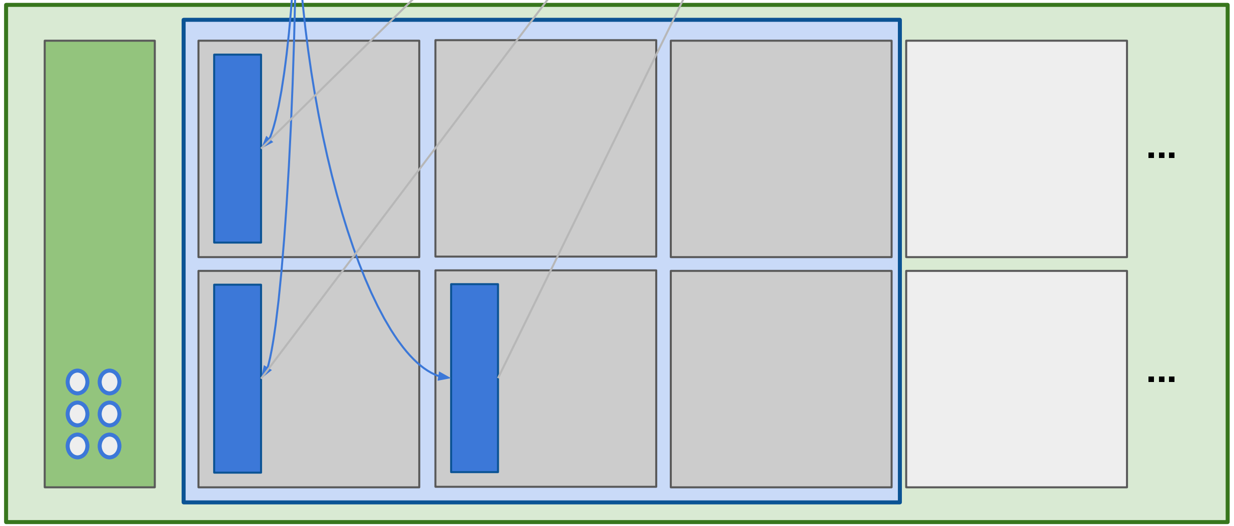

- Master Task
- Worker Task

The Masters each define one or more **Worker tasks** which are scheduled, placed and launched on the compute nodes.

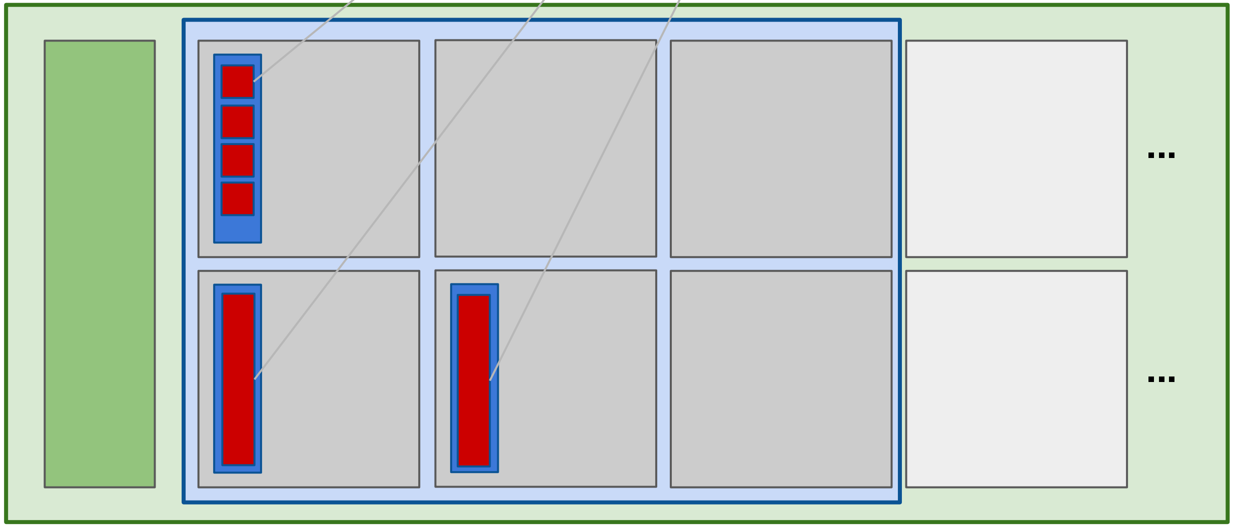

- Master Task
- Worker Task

Once the **Master** and **Worker** tasks are successfully bootstrapped, each Master directly coordinates its pool of **Workers**.

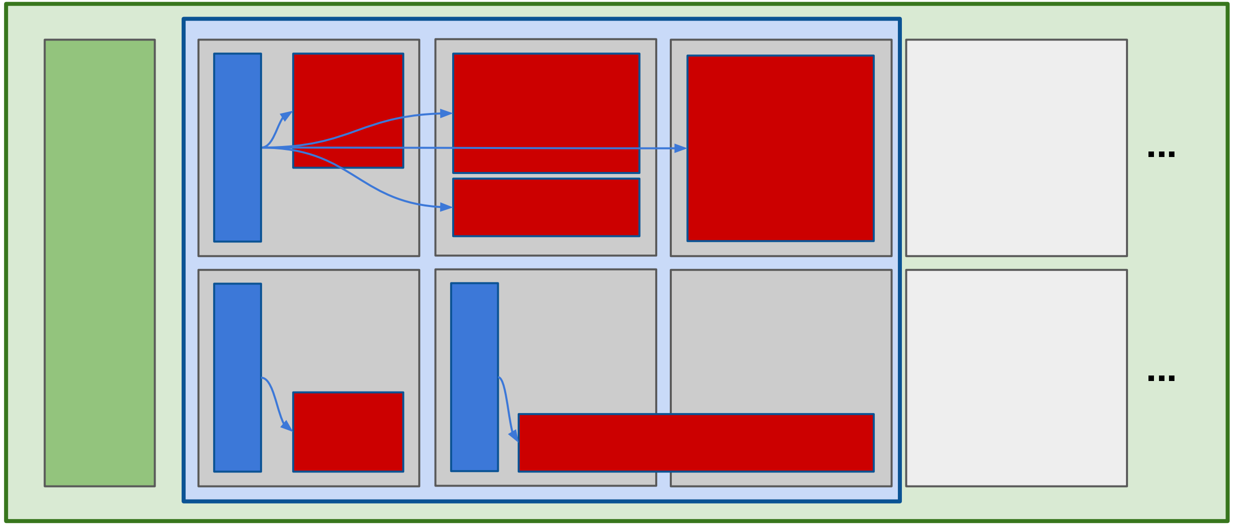

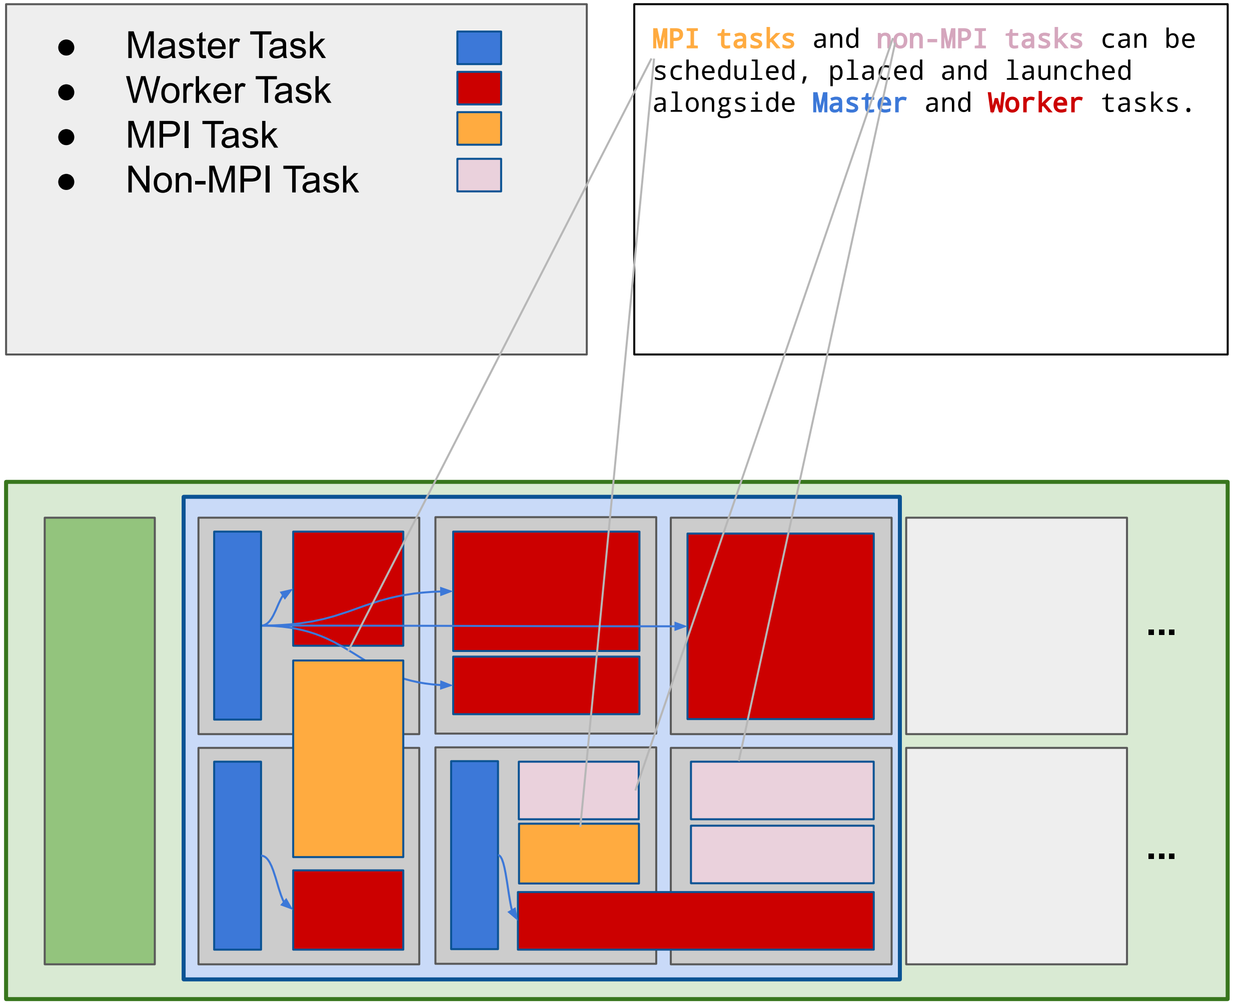

Master Task
Worker Task
MPI Task
Non-MPI Task
MPI tasks and non-MPI tasks can be scheduled, placed and launched alongside Master and Worker tasks.
...
...

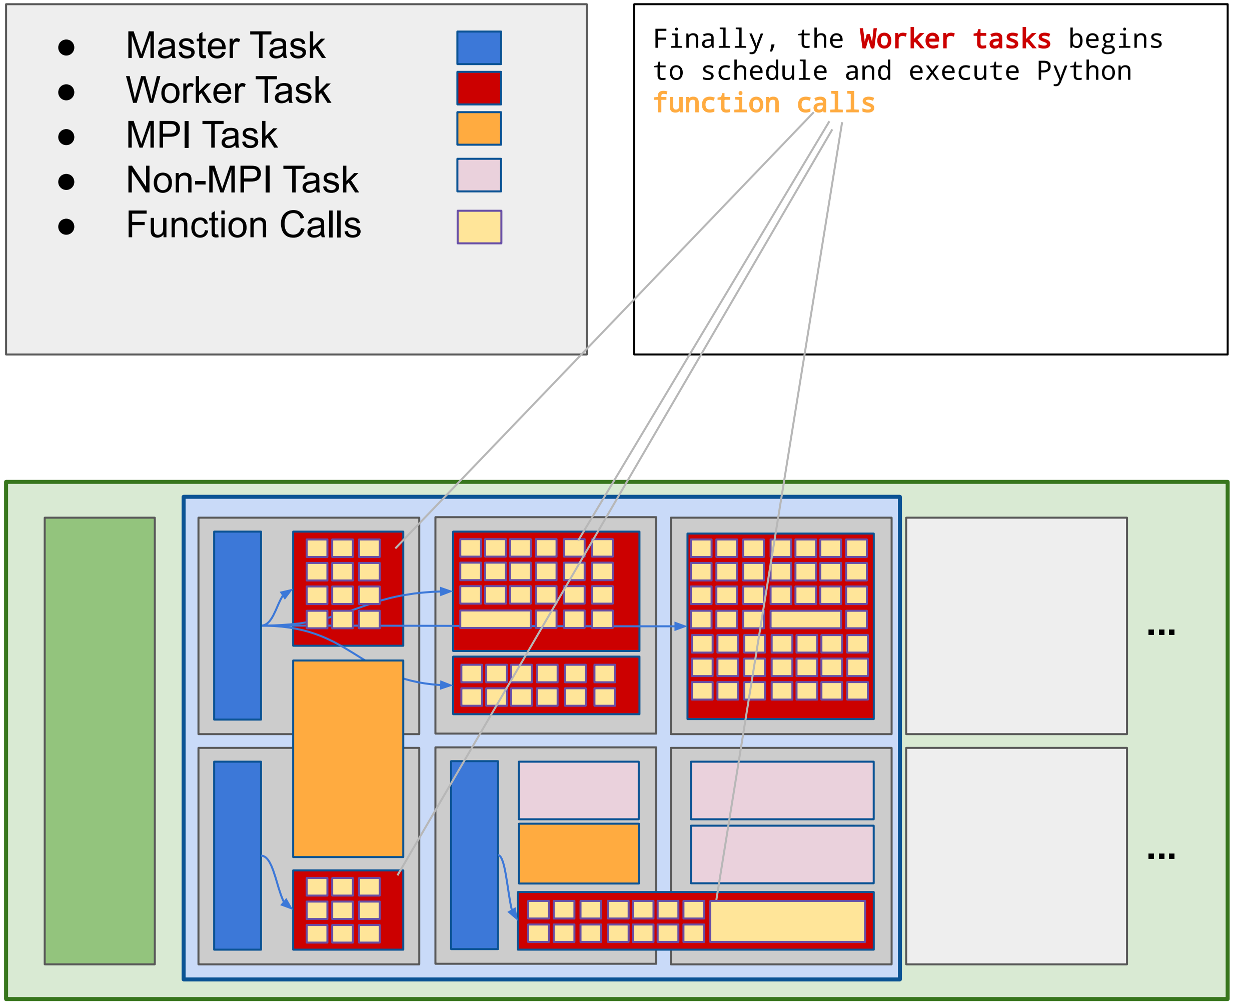
Master Task
Worker Task
MPI Task
Non-MPI Task
Function Calls
Finally, the Worker tasks begins to schedule and execute Python function calls
...
...

RAPTOR: High-Throughput Function Calls

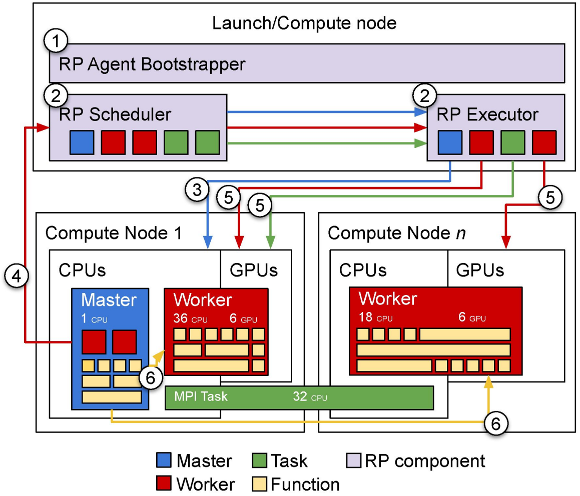

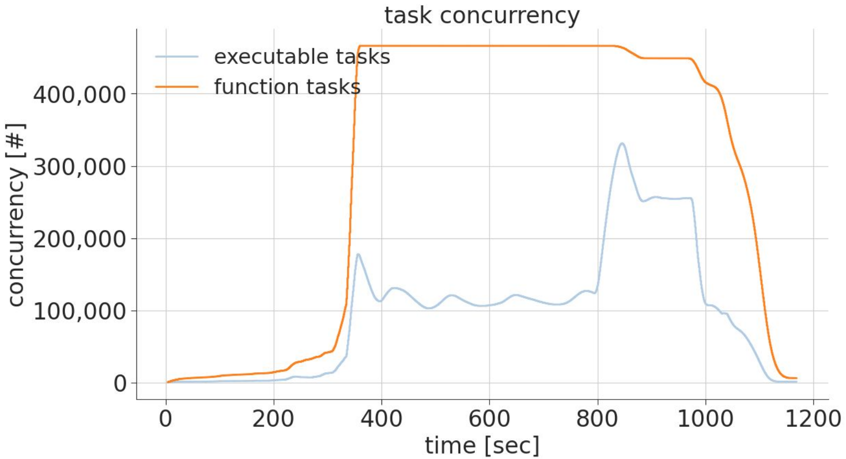

ID	Platform	Application	Nodes	Pilots	Tasks [$\times 10^6$]	Startup [sec]	Utilization avg / steady	Task Time [sec] max	Task Time [sec] mean	Rate [$\times 10^6/h$] max	Rate [$\times 10^6/h$] mean
1	Frontera	OpenEye	128	31	205	129	90% / 93%	3582.6	28.8	17.4	5.0
2	Frontera	OpenEye	7600	1	126	81	90% / 98%	14958.8	10.1	144.0	126.0
3	Frontera	OpenEye	8336	1	13	451	63% / 98%	219.0	25.3	91.8	11.0
4	Summit	AutoDock	1000	1	57	107	95% / 95%	263.9	36.2	11.3	11.1

Workflow Performance ≠ Aggregate of Task Performance

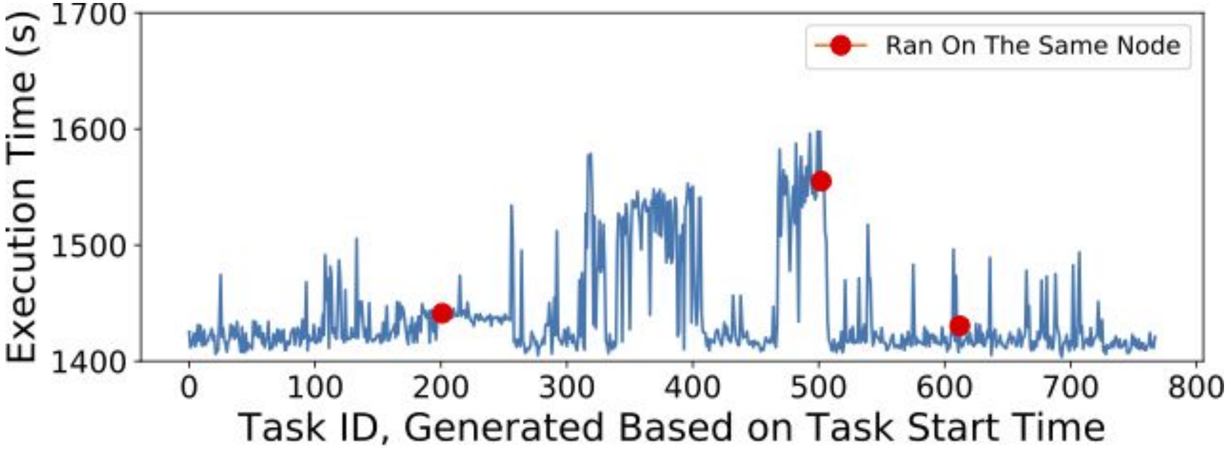

- Performance portability: a more complex consideration
- Measure & optimize **collective performance**, not single task performance
- Makespan optimization under performance uncertainty
  - More than makespan optimization; non-traditional trade-offs

BROOKHAVEN NATIONAL LABORATORY

Workflows: Software Challenge

- **Proliferation** of workflow solutions:
  - Mostly bespoke, often “local solutions”
  - Fragile 4P: Performance, Portability, Provenance, and Productivity
- Need ecosystem to manage proliferation and development models to facilitate
  - **Building Blocks** - Functional components with *common* interfaces, yet *differentiated* assurance, resilience and performance
- Open Building Blocks enable reuse, yet permit specific solutions
  - Build upon robust capabilities, while reducing development, testing, maintenance, etc. costs and enhance portability
  - Focus on innovation at higher levels of the workflows stack

ExaWorks (SWAS) Workflows SDK

- Providing a production-grade Software Development Kit (SDK)
- Implemented via **progressively integrated levels**
    - Level 0: Technologies packaged together
    - Level 1: Component interfaces or pairwise integration
    - Level 2: Community developed and supported API
- **SDK democratizes access** to hardened, scalable, and interoperable workflow technologies and components, for both developers and users
    - E4S-based community policies for software quality
    - Open community-based design and implementation process
    - Scalability of components on **Exascale Systems**
    - Standard packaging and testing

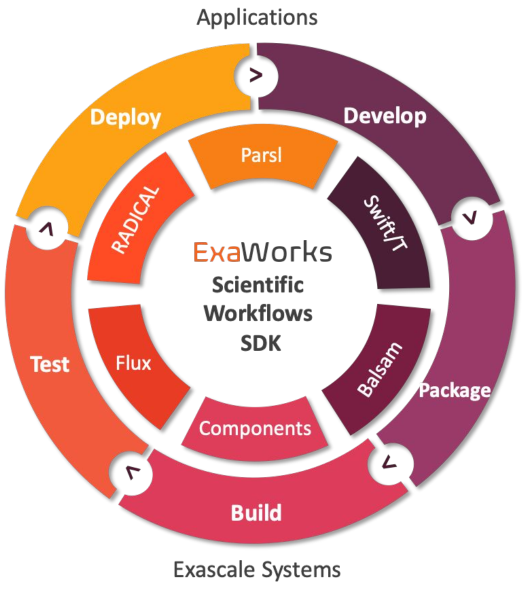

PSI/J: Portable Submission Interface for Jobs

- A set of interfaces that allow the specification and management of “jobs”
- Support for Slurm, LSF, Cobalt, Flux, PBS
- Open document to define a language-independent specification

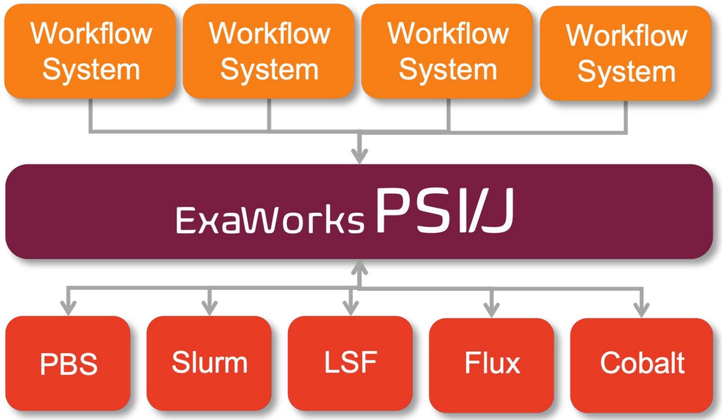

- Community specification http://exaworks.org/job-api-spec/specification.html
- PSI/J: https://arxiv.org/abs/2307.07895
- ExaWorks: https://arxiv.org/abs/2108.13521
  - https://www.hpcwire.com/off-the-wire/exaworks-provides-access-to-community-sustained-hardened-and-tested-components-to-create-award-winning-hpc-workflows/
  - https://www.exascaleproject.org/workflow-technologies-impact-sc20-gordon-bell-covid-19-award-winner-and-two-of-the-three-finalists

Summary

- Workflows are increasingly important as the basis of future scientific discovery
  - Technology trend and application requirements
  - Workflows are the highest-level of execution, parallelism, and programming

- Workflows require a fundamental rethink of performance
  - Thinking collectively, beyond single task performance
  - Advances in software systems engineering

- Next generation HPC platforms privilege workflows as first class application
  - Arguably, the most significant change in supercomputer software ecosystem
  - Design and evaluation (e.g., workflow benchmarks)

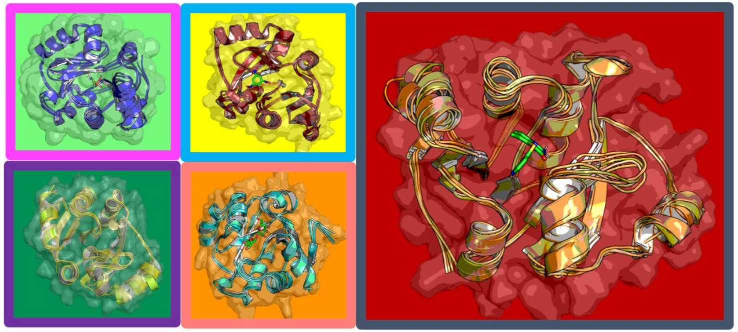

Thank you!

Funding acknowledgement:

- DOE ECP ExaWorks
- DOE ECP CANDLE and ECP ExaLearn
- DOE ASCR Surrogates Benchmarking Initiative
- NSF MolSSI, NSF SmartSurrogates, NSF RCT

Acknowledgement and thanks

- RADICAL team: Matteo Turilli, Andre Merzky, Mikhail Titov, Ozgur Kilic
- ExaWorks team
- Arjun Shankar (OLCF/ORNL)
- Geoffrey Fox
- Srinivas Rangan (HP Labs)
- Many domain scientists/collaborators