

Colmena: Seamless Computational Campaigns across Multiple Computing Clusters with Parsl/FuncX and Object Proxies

Cleared for public release



Logan Ward

Assistant Computational Scientist
Data Science and Learning Division
Argonne National Laboratory

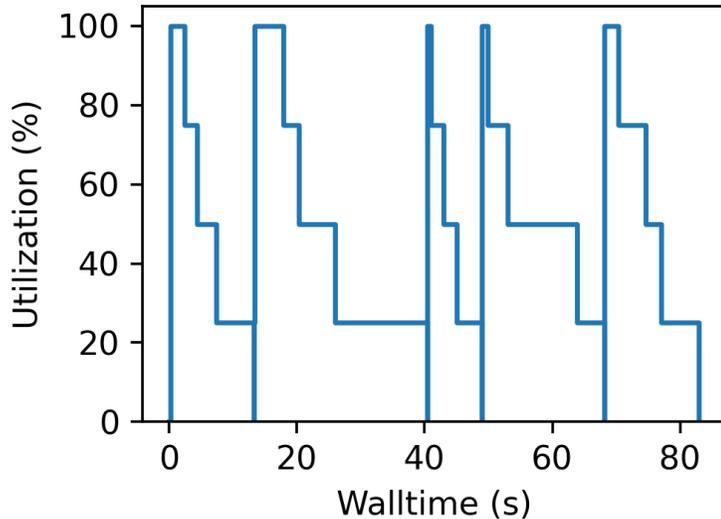
14 September 2022

Strategies for steering computational campaigns are complicated

Parallel Optimizers: A “simple” example with no optimal strategy

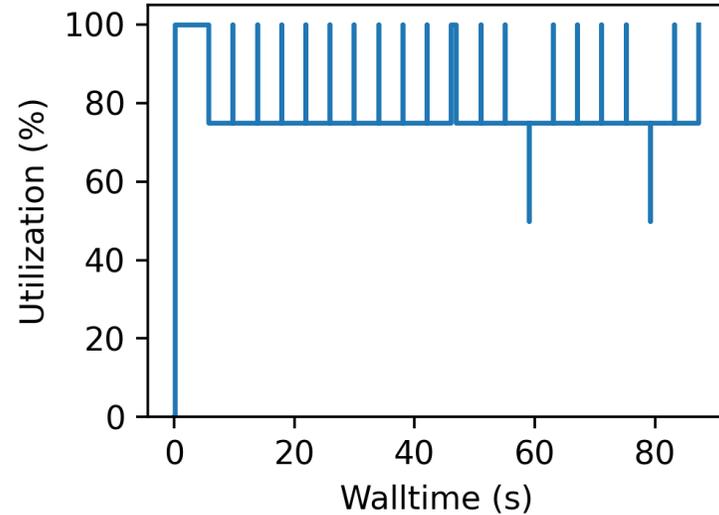


Batch Optimizer



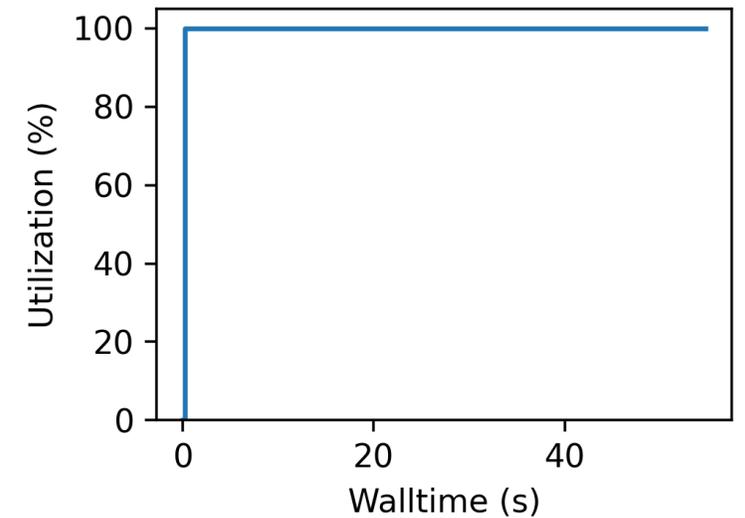
Wait for N tasks to complete, then pick next batch

Streaming Optimizer



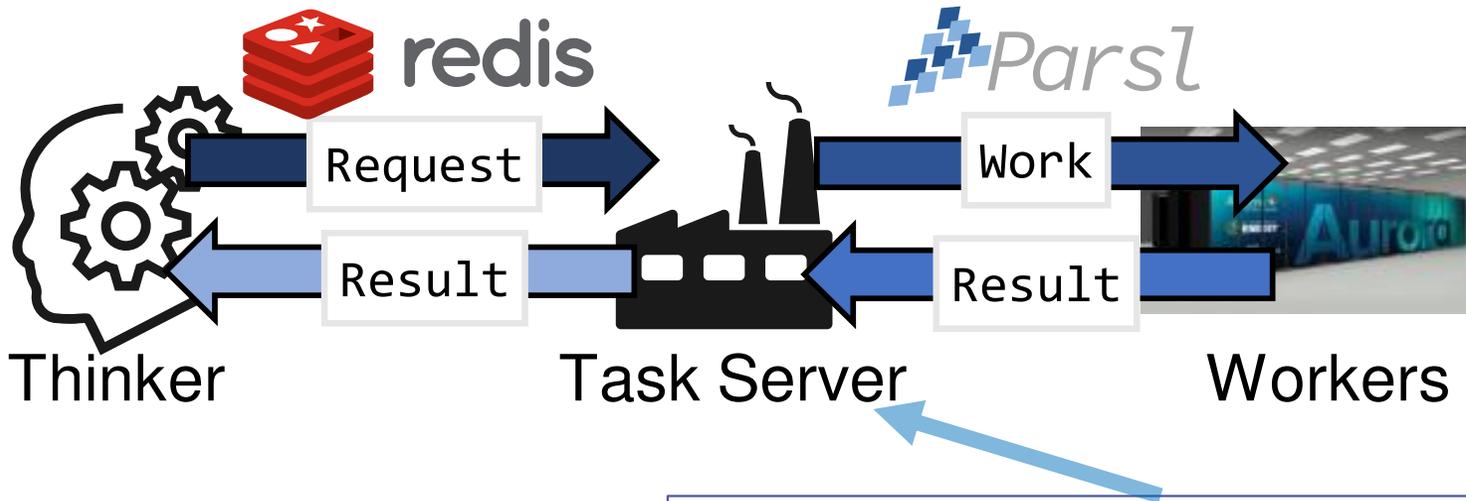
Pick new tasks as soon as one completes

Interleaved Optimizer



Maintain a task queue

Colmena is a wrapper over Exascale Workflow tools



Programming Model: Task Queues

```
# Primitive Units
queue.send_inputs(1)
result = queue.get_result()
```

Programming Model: Agents

```
class Thinker(BaseThinker):
    @agent
    def make_work(self):
        self.queue.send_inputs(1)
```

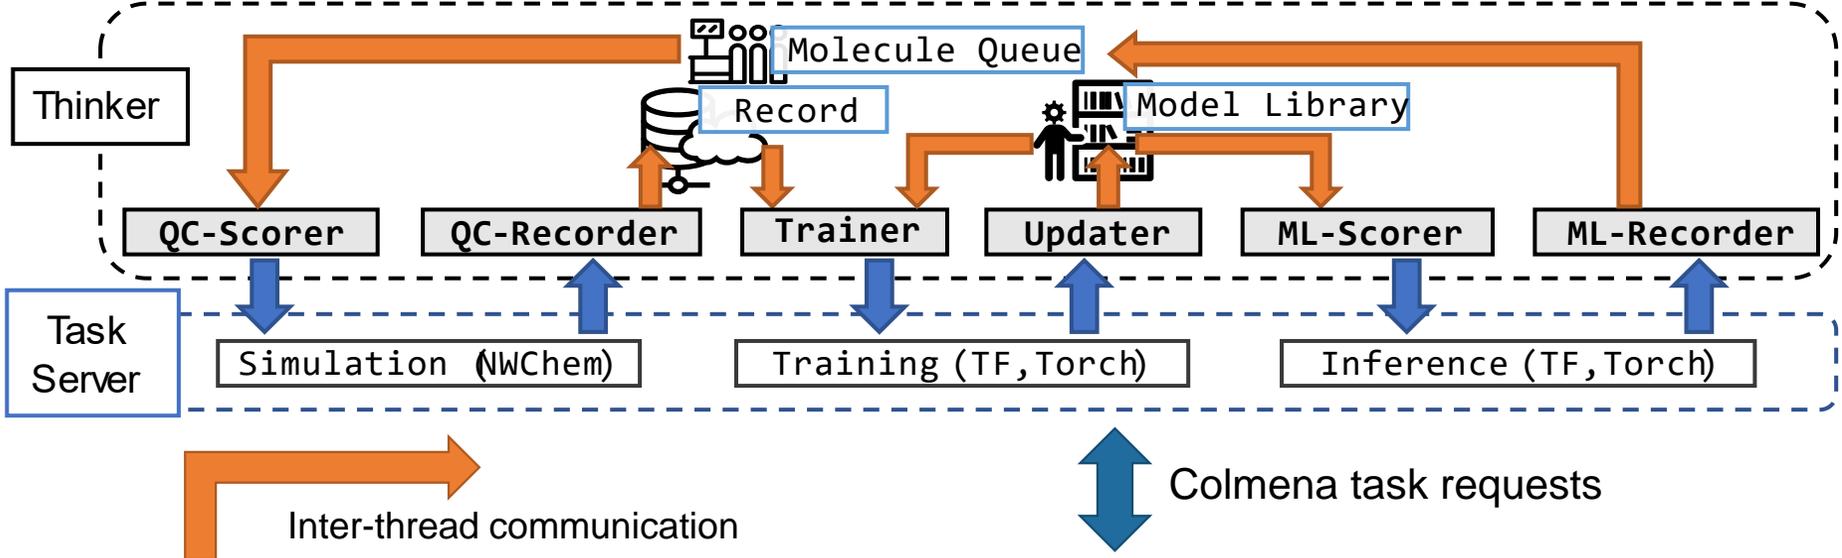
Task Server:

- Dispatches work requests to compute
- Communicates results back to thinker

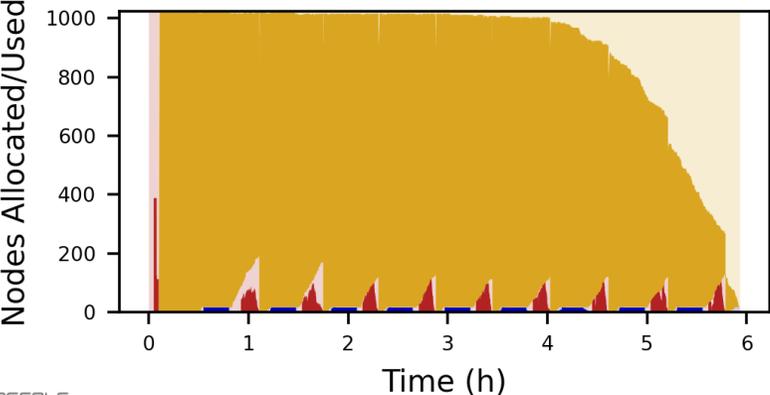
Backend:

- Supports most HPC and cloud services
- Easily configure multiple worker types, multi-site workflows
- Limited support for ensembles of MPI applications
- *Future:* Balsam, FuncX, RCT

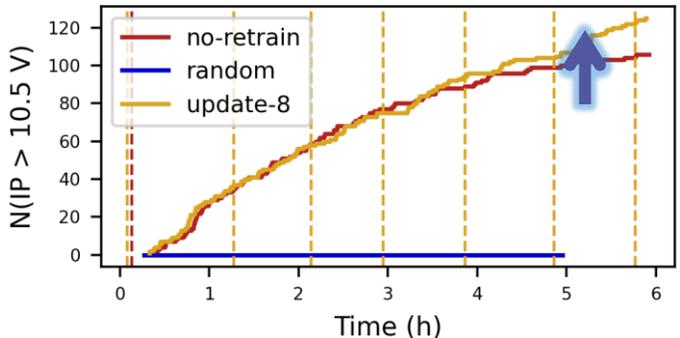
Example application: "Interleaved," AI-in-the-loop optimizer



Retasking nodes between jobs...



...yields more science per compute-hour.



So, what's new in '22?

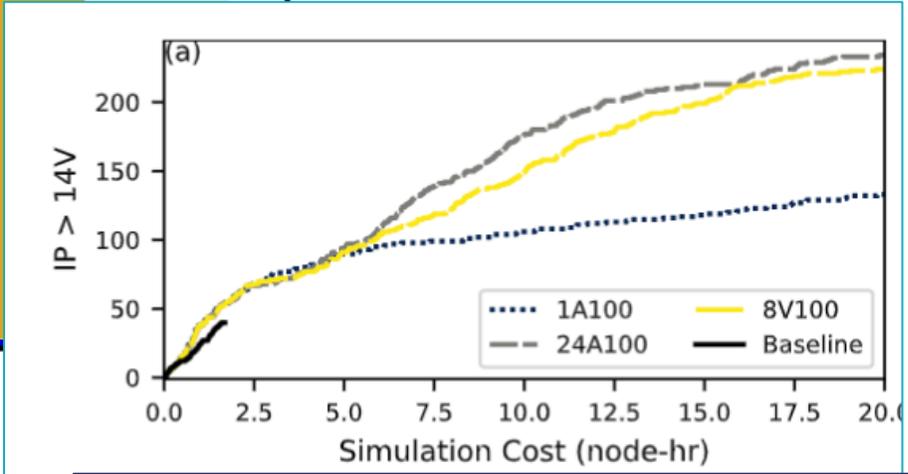
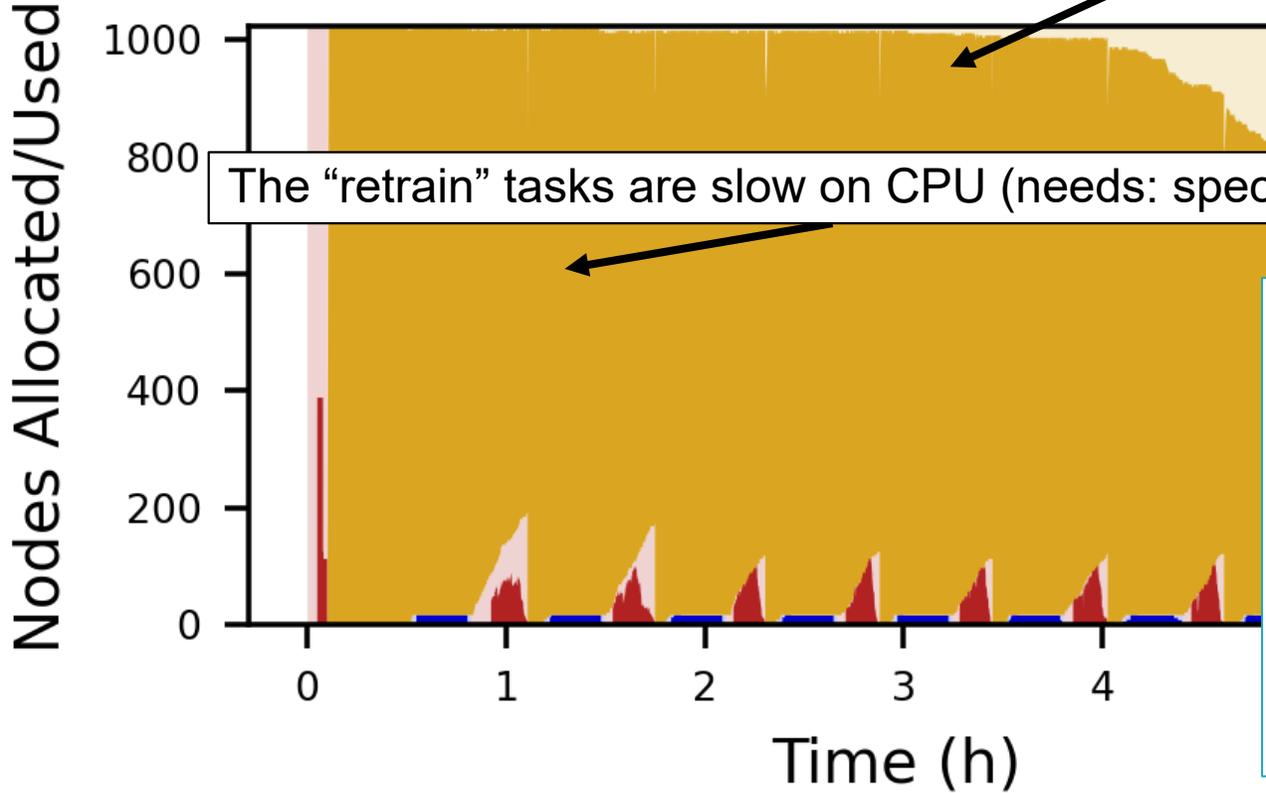
Multi-site Campaigns!



Why multi-site? *Moving compute onto best hardware*

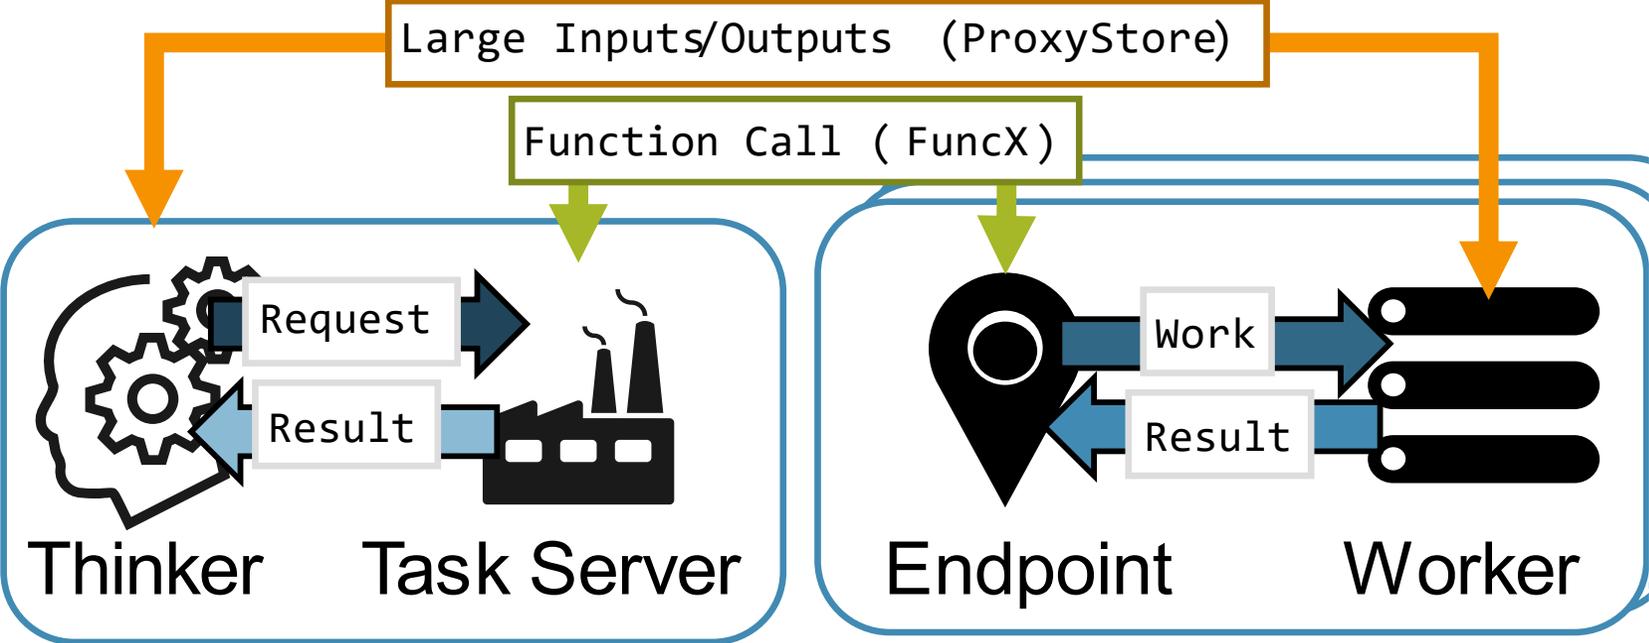
Most of our work is CPU-only (needs: many, cheap)

The "retrain" tasks are slow on CPU (needs: specialized)



Faster retraining means better steering

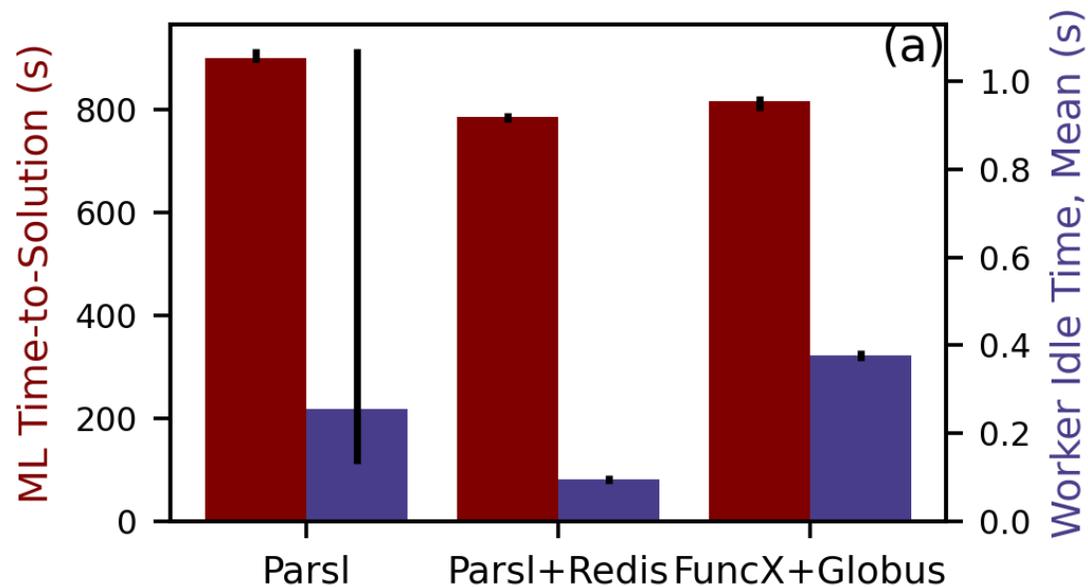
How multi-site? *FuncX*



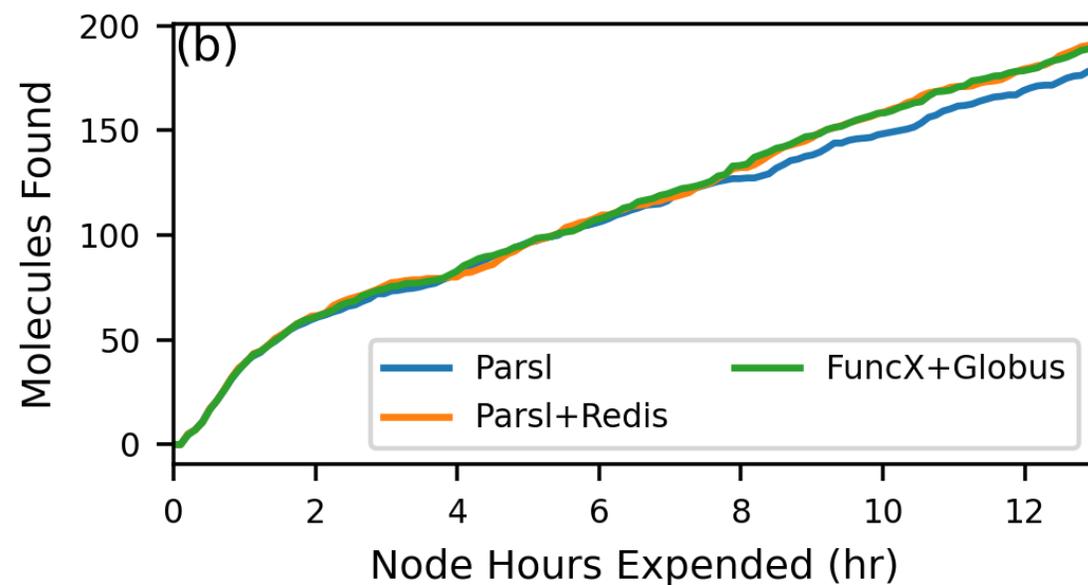
It's just FuncX. We use the "FuncXExecutor," so it acts like Parsl

How good multi-site? *Same performance, less port headache*

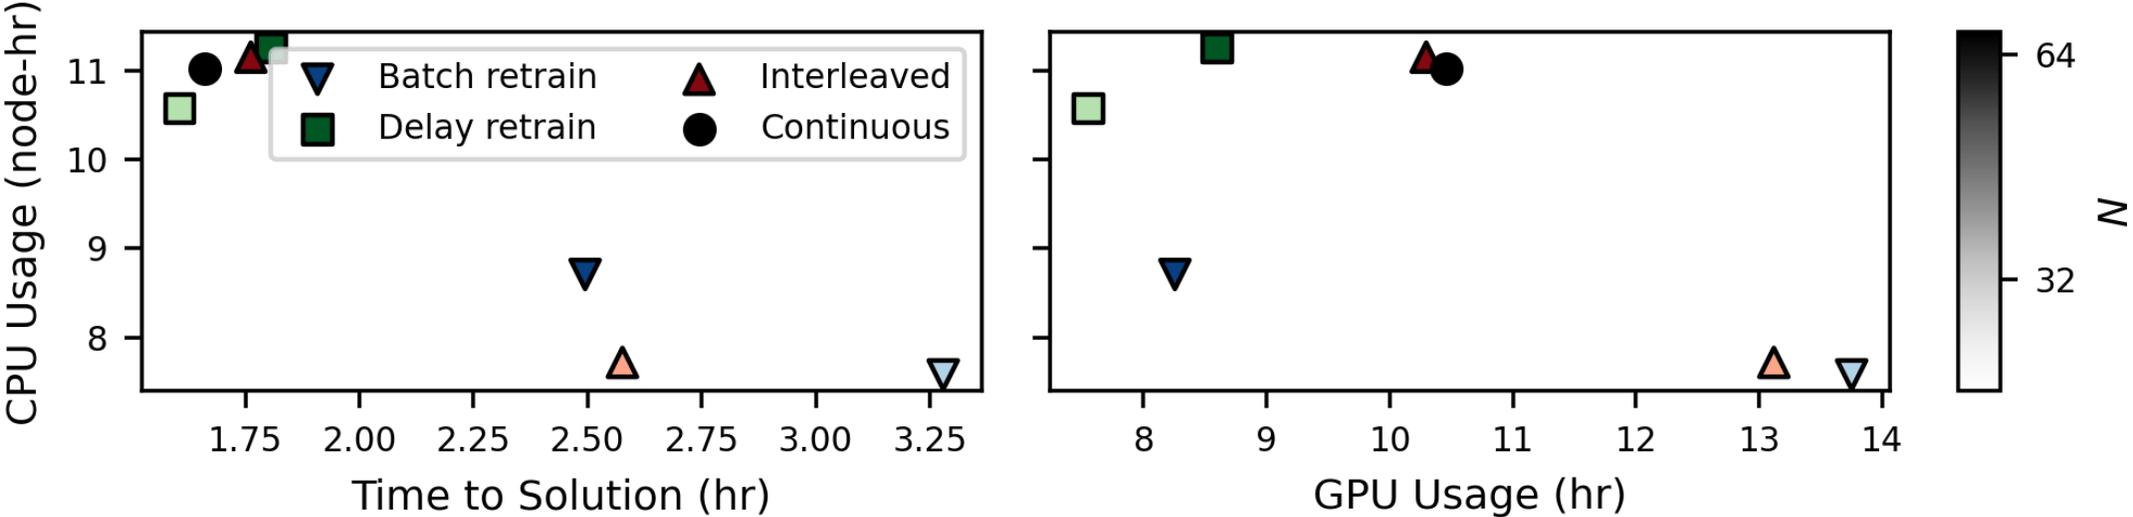
Machine learning tasks take only 3% longer than best-effort with SSH tunnels



Scientific outcomes are identical



Colmena lets you explore computational cost tradeoffs

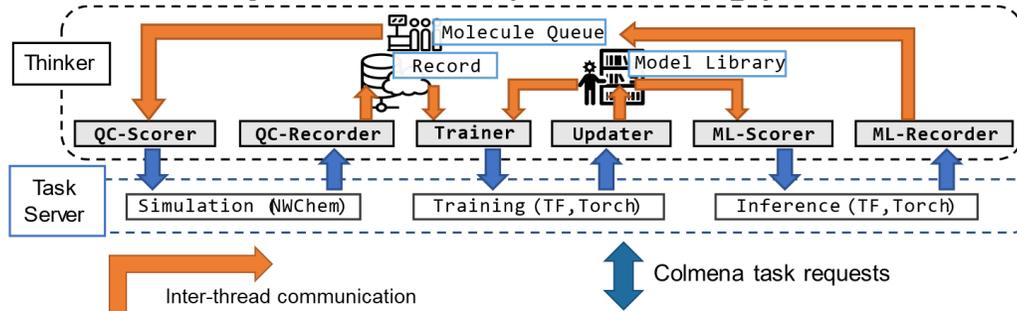


Steering policies tradeoff between time to solution, GPU time, and CPU time

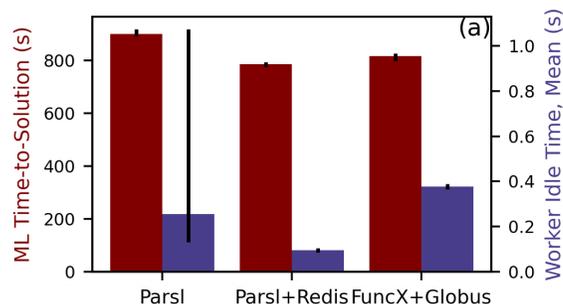
Conclusions and Future for Colmena

What did we cover today?

- Colmena lets you build complex steering policies



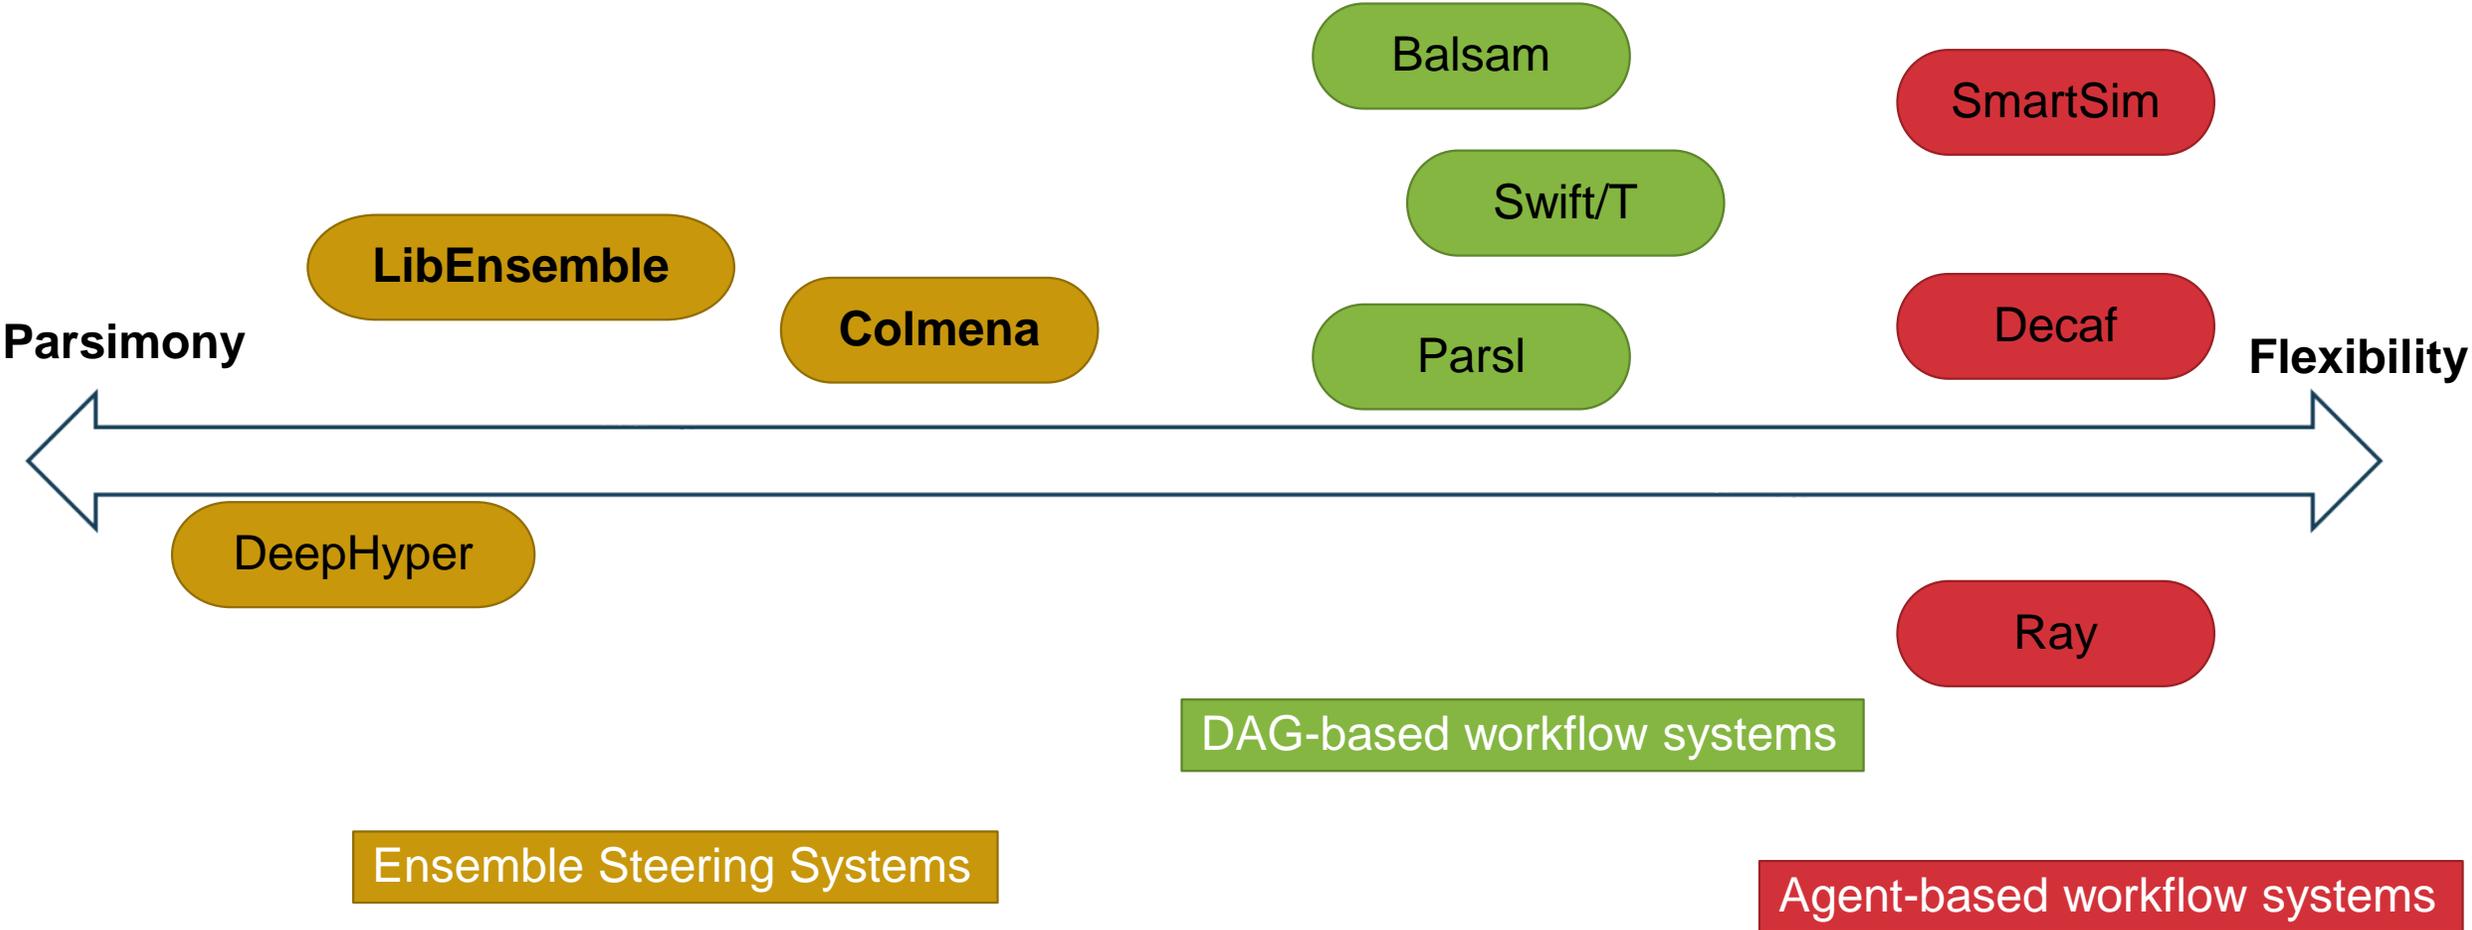
- FuncX lets policies span compute systems



What to watch for next year?

- This work published (at least on ArXiv!)
- A perspective on ensemble steering toolkits
 - “How are libE and
- More Colmena applications
 - Fitting machine-learned surrogates for simulations
 - Coordinating simulation self-driving laboratories
 - Rapid screening of HPC
- Integration with more workflow engines (e.g., RCT!)

Got opinions about what Colmena is? Join our interest group



Acknowledgements: The (growing!) team

Argonne: ExaLearn – Using AI with HPC
Yadu Babuji, Ben Blaiszik, Ryan Chard, Kyle Chard,
Ian Foster, Greg Pauloski, Ganesh Sivaraman,
Rajeev Thakur

Argonne: JCESR – Molecular modeling for batteries
Rajeev Assary, Larry Curtiss, Naveen Dandu,
Paul Redfern

MoISSI – Workflows for quantum chemistry
Lori A. Burns, Daniel Smith, Matt Welborn,
many other open-source contributors

PNNL: ExaLearn – Graph algorithms for learning
Sutanay Choudhury, Jenna Pope

BNL: ExaLearn – Optimal experimental design
Frank Alexander, Shantenu Jha, Kris Reyes, Li Tan,
Byung Jun, *and more*

Argonne ALCF – AI, Data and Simulation on HPC
Murali Emani, Alvaro Vazquez-Mayagoitia,
Venkat Vishnawath

FuncX – Seamless multisite deployment
Kevin Hunter Kesling, Kyle Chard, Ryan Chard,
Ben Clifford, *and more*

ExaWorks – Interfacing to HPC
Ayman Alsaadi, Matteo Turilli,
Shantenu Jha, Kyle Chard

Ensemble Group – Defining ensemble needs
John-Luke Navarro, Jonathon Ozik, Tom Peterka,
Stephen Hudson, Orçun Yildiz, Alex Brace,
Arvind Ramanathan, *and more*