




Scaling Submodular Optimization Based Techniques for Epidemic Intervention

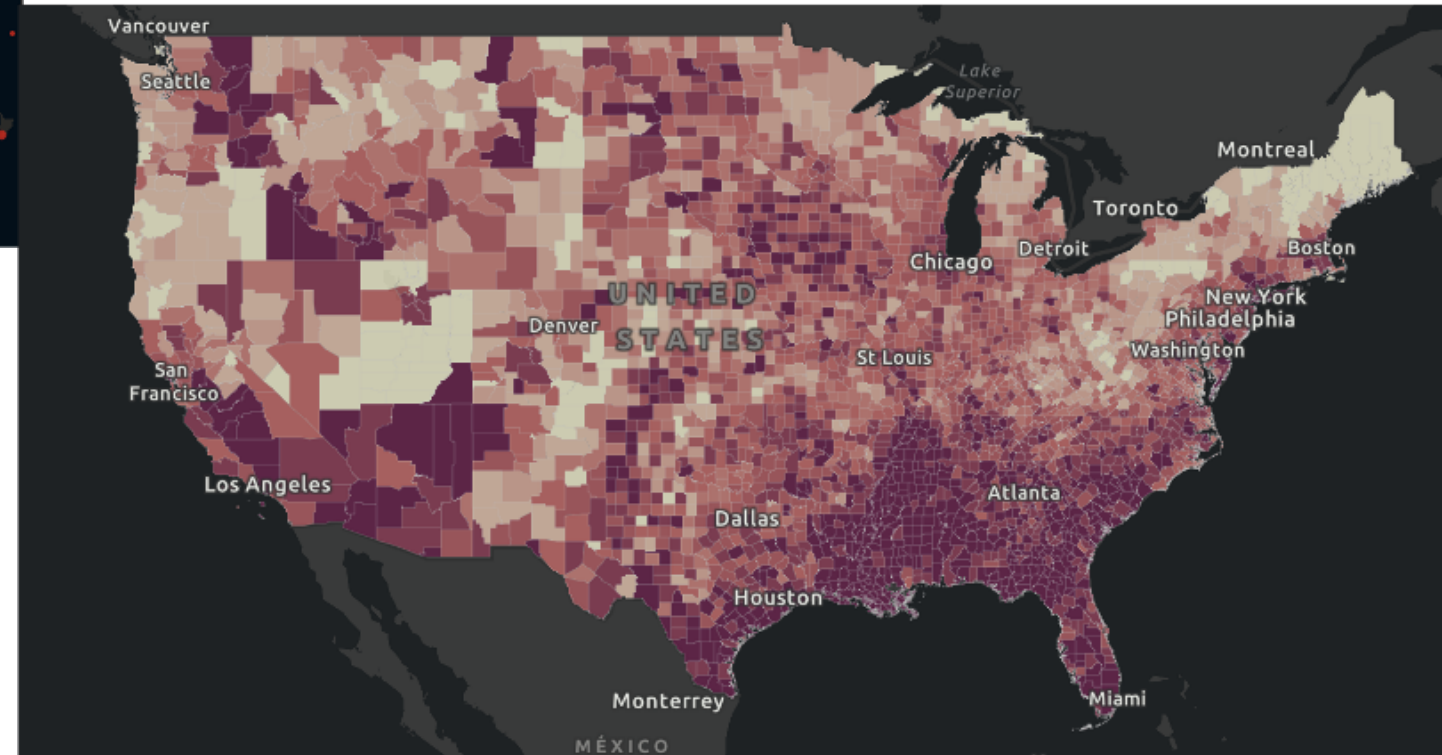
Marco Minutoli, Prathyush Sambaturu,
Mahantesh Halappanavar, Antonino Tumeo,
Ananth Kalyanaraman, Anil Vullikanti



PNNL is operated by Battelle for the U.S. Department of Energy



COVID-19: The reason we are remote today




Source: <https://coronavirus.jhu.edu/>

Epi-Control: Problem Statement

Def (Epi-Control): Given a contact network $G=(V, E, w)$, a diffusion process M , a budget k , and a set B of initially infected nodes, the Epi-Control problem is to choose a set S of size k such that the expected number of infected nodes through M is minimized.

- Assumes that B is known
 - Usually solved with randomized experiments
- The objective is to minimize infection spread
- **Susceptible-Infected-Recovered (SIR)**



Epi-Control Revisited


- The number of infections given B and S

$$\tilde{\sigma}_{G_i}(B, S) = \left| \bigcup_{b \in B} \mathcal{R}(b, S, V) \right|$$

- The number of lives saved:

$$\tilde{\lambda}_{G_i}(B, S) = \tilde{\sigma}_{G_i}(B, \emptyset) - \tilde{\sigma}_{G_i}(B, S),$$

- Def (Epi-Control):** Given a contact network $G=(V, E, w)$, a budget k , and a set B of initially infected nodes, the Epi-Control problem is to choose a set S of size k such that the expected number of lives saved is maximized.



Sub-modularity of Preempt in Rooted Trees

-

Def. (Sub-modularity): Let S be a finite set. A set function $f: 2^S \rightarrow \mathbb{R}$ is sub-modular if for any subset $X \subseteq Y \subseteq S$ and $x \in S \setminus Y$ then

$$f(X \cup \{x\}) - f(X) \geq f(Y \cup \{x\}) - f(Y).$$

In Words: f has the property of diminishing returns.

Theorem: if G_i is a rooted tree then $\tilde{\lambda}_{G_i}(B, S)$ is a sub-modular function of S .

Sketch of Proof: Sub-modularity of PREEMPT in rooted trees

Theorem: if G_i is a rooted tree then $\tilde{\lambda}_{G_i}(B, S)$ is a sub-modular function of S .

Notation:

- $\Lambda(x)$ denotes the set of vertices reachable x .
- $U_{x,X}$ is the set of vertices reachable from x that is also reachable from X .

Let's consider $X = \{z\}$


$\Lambda(x) = \{x, y, z, t, u, v\}$

$U_{x,X} = \{u, v, z\}$

We know: $\tilde{\lambda}(B, X \cup \{x\}) - \tilde{\lambda}(B, X) \leq |\Lambda(x)|$

More precisely: $\tilde{\lambda}(B, X \cup \{x\}) - \tilde{\lambda}(B, X) = |\Lambda(x)| - |U_{x,X}|$

Let's consider a bigger set $Y = \{z, y\}$



PREEMPT: Scalable Epidemic Interventions Using Submodular Optimization on Multi-GPU Systems

- EpiControl is submodular on rooted trees
 - E.g. Sexually transmitted diseases*
- Not valid in the general case but...
- **Def (PREEMPT):** Given a contact network $G=(V, E, w)$, a budget k , choose a set B of size k to vaccinate such that the expected number of infection from B is maximized

Algorithm 2: PREEMPT-HC: Selects a set of nodes S of size at most k that maximize the influence, $\sigma(\cdot)$.


Input : $(G = (V, E, \omega), k, \eta)$
Output: S

```

1  $SG \leftarrow \text{Sampling}(G, \eta)$ 
2  $S \leftarrow \emptyset$ 
3 while  $|S| \leq k$  do
4    $\forall v \in V (count[v] = 0)$ 
5   for  $v \in V \setminus S$  do in parallel
6     // Counting Phase
7     for  $G_i \in SG$  do
8        $count[v] \leftarrow \sigma_{G_i}(S \cup \{v\}) - \sigma_{G_i}(S)$ 
9     // SeedSelect Phase
10     $s_{best} \leftarrow \arg \max_{v \in V} (count[v])$ 
11     $S \leftarrow S \cup \{s_{best}\}$ 

```

- Possible strategies are:
 - Hill-Climbing Algorithm
 - Reverse Influence Sampling




Minutoli, Marco, et al. "cuRipples: influence maximization on multi-GPU systems." *Proceedings of the 34th ACM International Conference on Supercomputing*. 2020.

Minutoli, Marco, et al. "Fast and Scalable Implementations of Influence Maximization Algorithms." *2019 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 2019.

* Bearman P, Moody J, Stovel K: Chains of affection: The structure of adolescent romantic and sexual networks. *AJS* 2004, 110: 44–91.




Running Example on Zachary's Karate Club

Samples=10 K=2 p=0.2











If you can't get it right on this network, then go home!*




Select First Seed











Samples=10 K=2 p=0.2
1: S = {1} E[spread]=7.4




Select Second Seed



Samples=10 K=2 p=0.2
1: S = {1} E[spread]=7.4
2: S = {1, 34} E[spread]=12.2

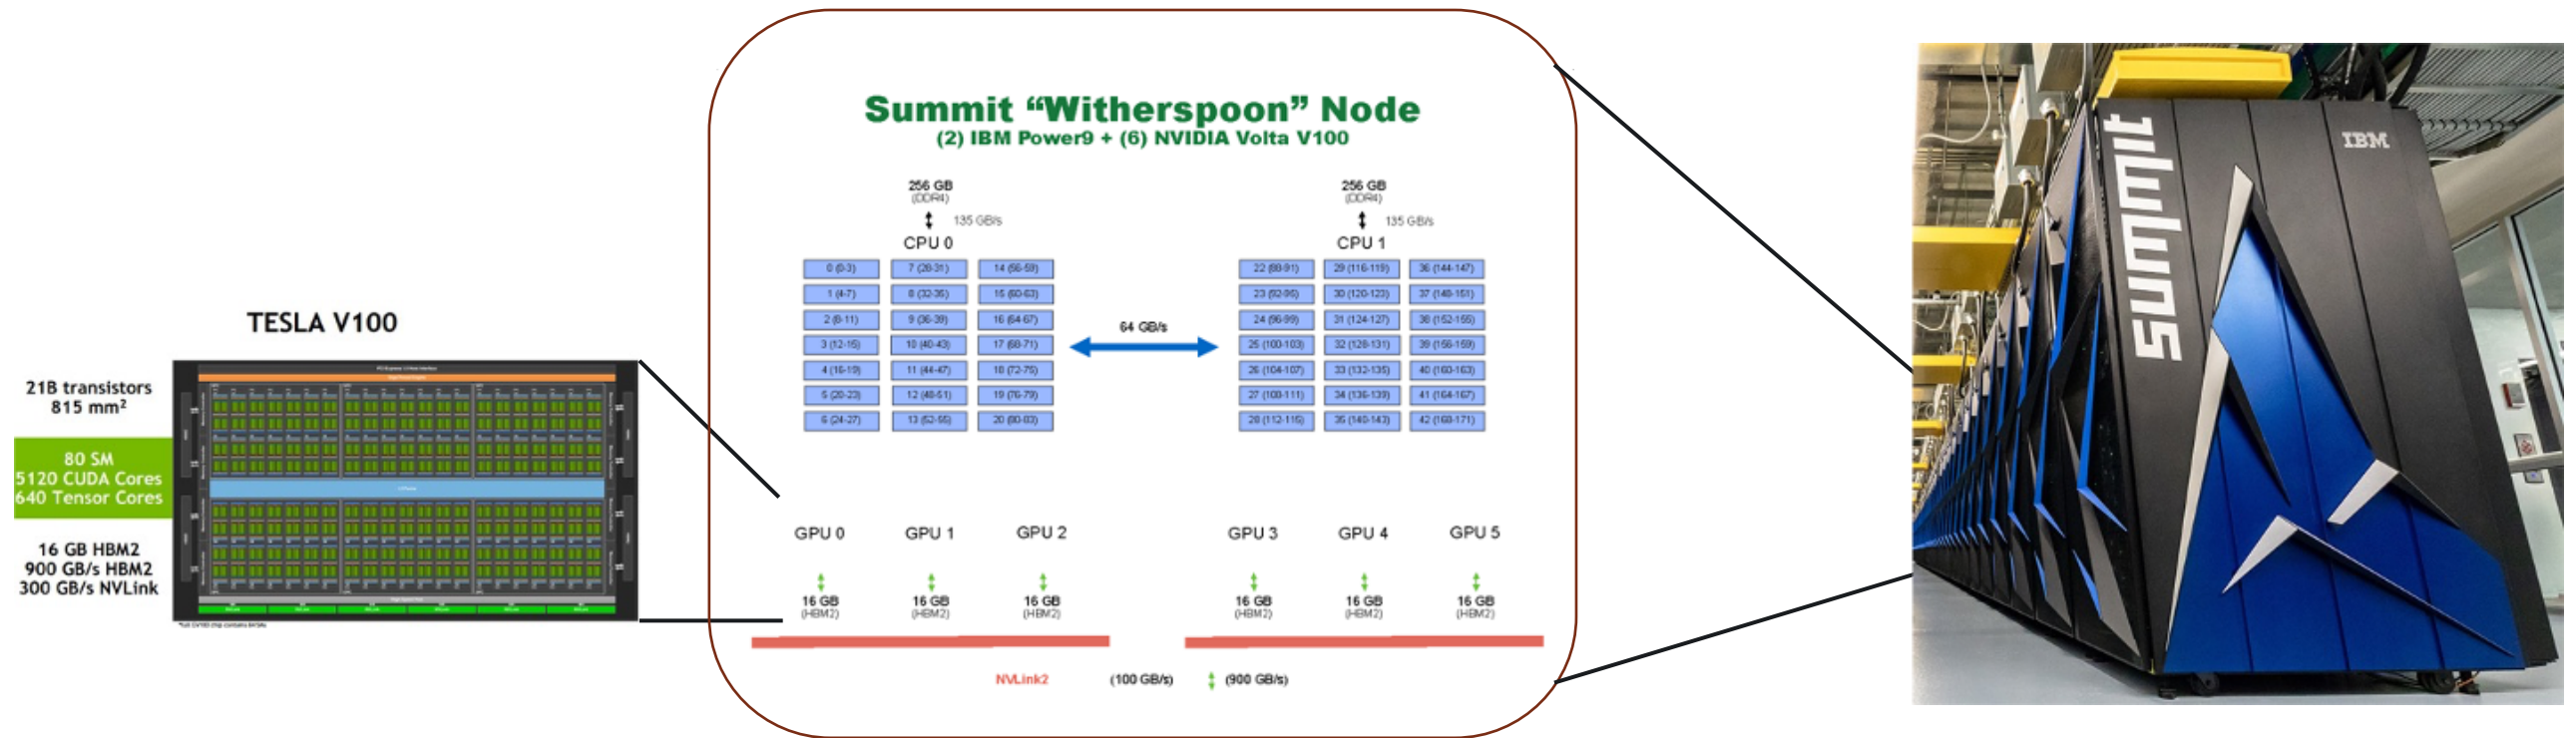


Running Example on Zachary's Karate Club



Our Intervention is $S = \{ 1, 34 \}$

Summit




Single GPU
(2048 x 80 threads)

Single Node
(6 GPUs)


Distributed Multi-GPU Cluster
(4608 nodes)

$$\frac{2048 \text{ Threads}}{\text{SM}} \times \frac{80 \text{ SMs}}{\text{GPU}} \times \frac{6 \text{ GPUs}}{\text{Node}} \times 4608 \text{ Nodes} = 4.5 \text{ Billion GPU Threads}$$


The Parallel Algorithm



Dispatching Engine



Dispatching Engine



The CPU-GPU Dispatching Engine

- The engine instantiates a thread pool
 - Usually 1 thread per core on the CPUs of system
- Each GPU has a dedicated CPU thread offloading work with the possibility to over-subscribe
 - More than 1-thread pushing work to the same device (Hyper-Q)
- The engine builds a representation of the topology of GPUs
 - To structure reductions between GPUs
 - Topology built query the CUDA runtime
- CPU and GPU workers grab from the same “task queue”


Experimental Setup

- Experiments run on Summit
 - From 2 up to 128 nodes
- Hard limit to 2 hours
- Social networks and contact networks
- Social networks will provide an idea on how the technique use will scale in other application contexts.

Graph	Nodes	Edges	Avg. Degree	Max Degree
HepPh	34,546	421,578	24.41	846
Slashdot	77,360	905,468	23.41	5,048
Epinions	75,879	508,837	13.41	3,079
DBLP	317,080	1,049,866	6.62	343
Google	875,713	5,105,039	11.66	6,353
BerkStan	685,230	7,600,595	22.18	84,290
LiveJournal	4,847,571	68,993,773	28.47	22,889
Orkut	3,072,441	117,185,083	76.28	33,313
Portland	1,501,209	21,155,681	13.78	487
Portland-141k	63,543	141,450	2.23	25
Portland-295k	131,624	295,281	2.24	25
Portland-415k	182,967	415,434	2.27	26

Demographics

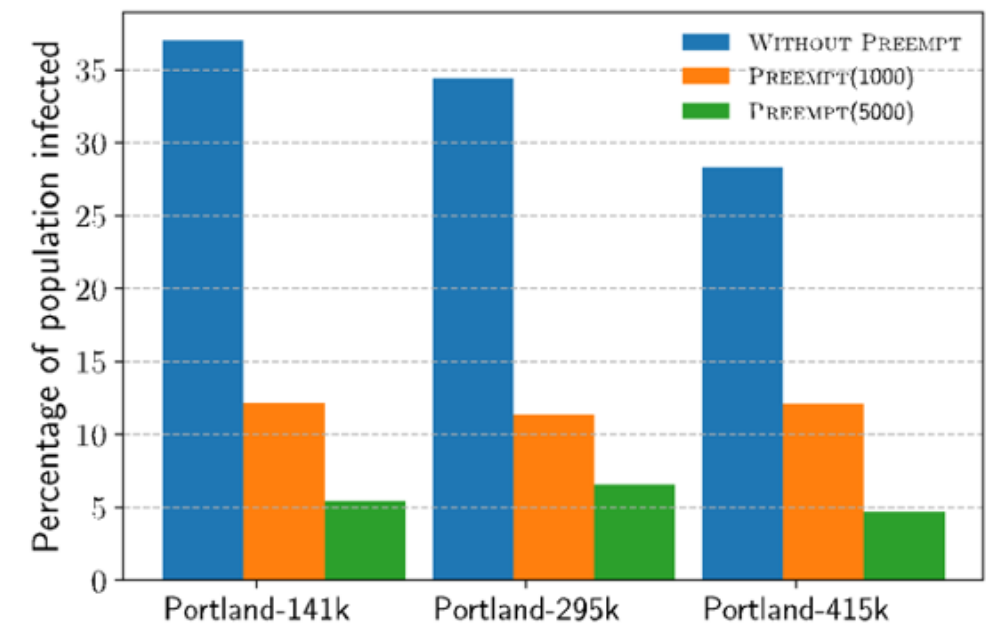

- Age and Degree of the seeds selected in the Portland-141k
- Significant spread on the degree
- Significant spread across the ages
- Suggests that focusing on specific groups like children and the elderly is sub-optimal



Seeds from Portland-141K

Quality of Results

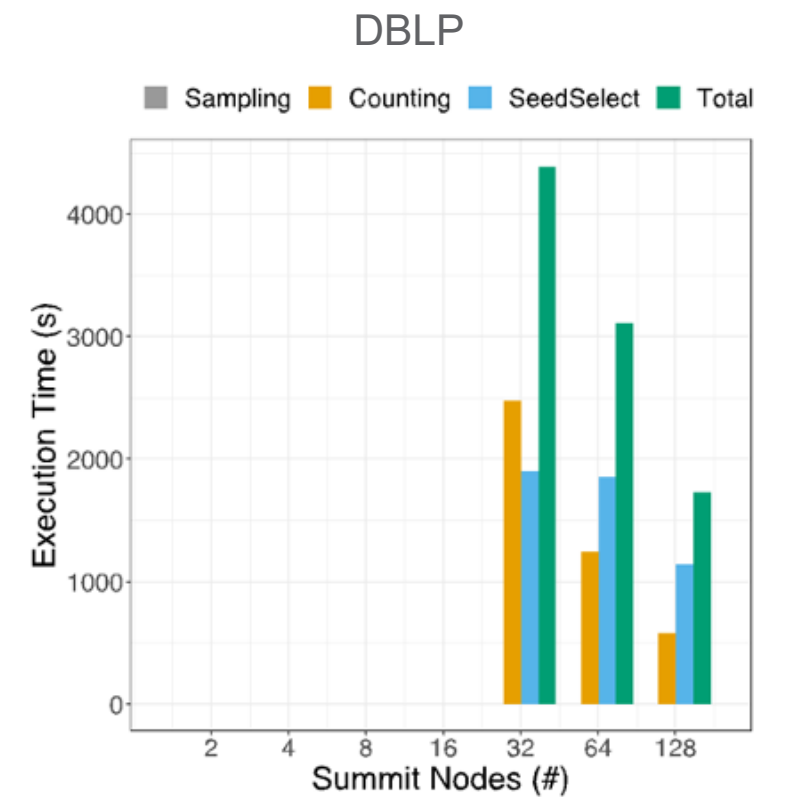



- **Control:** No action taken.
- **Baseline:** SaaRound*
- **PREEMPT:**
 - 2x—9x reduction in number of infections
 - PREEMPT-HC outperforms PREEMPT-RR
 - Up to 98.57% reduction in peak
- **Performance:**
 - SaaRound vs PREEMPT-HC comparable
 - SaaRound does not scale with problem size
 - PREEMPT-HC scales well



* Sambaturu, Prathyush, et al. "Designing Effective and Practical Interventions to Contain Epidemics." *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 2020.

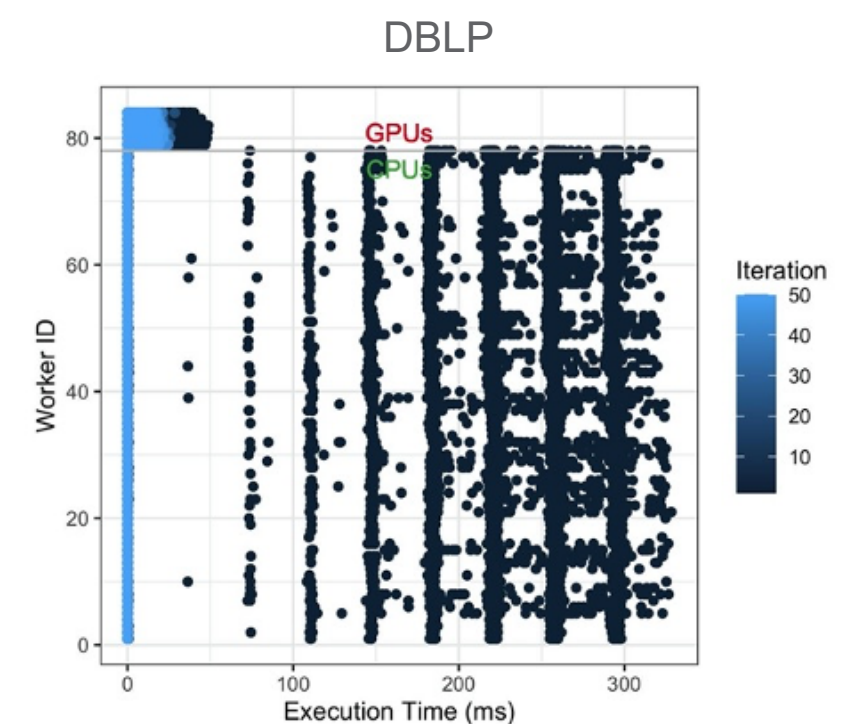

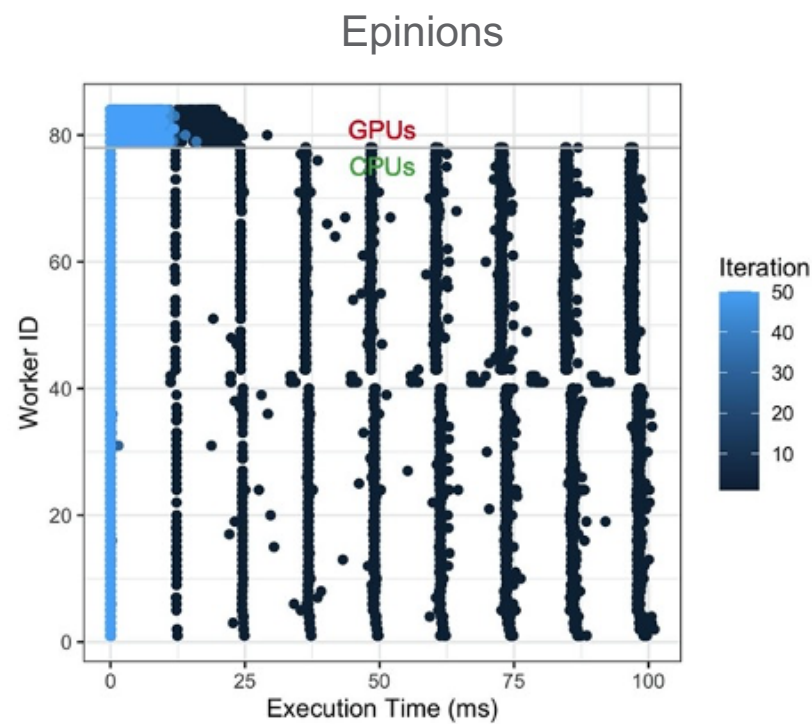

Strong Scaling Study

- **Overall:** Up to 33x compared to 2 nodes baseline
- **Sampling:** up to 155x



Dynamic Task Scheduling


- First Iterations more computation intensive
 - Up to 6.91x
- CPUs become faster later in the computation
 - Due to the offloading cost



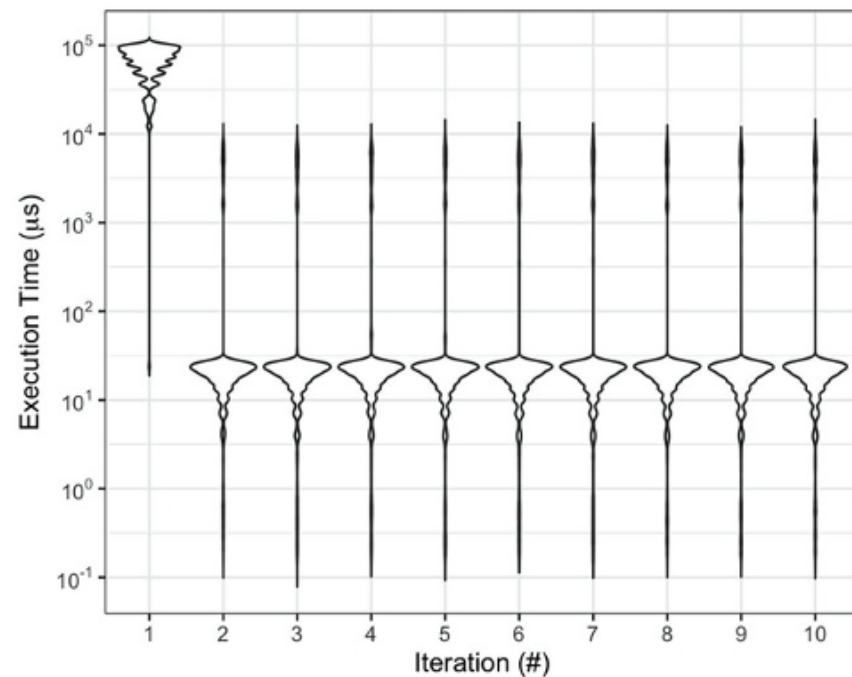
Execution Time of Tasks

- The distribution of the task duration is very skewed
- Pruning effective at bringing down the duration of a task


Portland




Epinions



Slashdot




DBLP



Future Research Direction

- **Improve Method Performance**
 - Overcome some limit of the implementation
 - Vertex Reordering to improve Locality
- **Reduce System Size**
 - HW Solution: Non-Volatile Memories
 - SW Solution: Graph Sparsification and Summarization
- **Epidemic Control**
 - Strategies for Social Distancing?
 - Fairness objectives?





Pacific Northwest
NATIONAL LABORATORY

Thanks to our Sponsors:

- The ExaGraph project at PNNL
- NSF grants CCF 1815467, OAC 1910213, and CCF 1919122 to WSU
- NIH Grant 2R01GM109718-07 and NSF grant numbers IIS-1955797, IIS-1633028, CCF-1918656, and DTRA subcontract/ARA S-D00189-15-TO-01-UVA at UVA

Thank you

