STEAM: System Testbed, Evaluation, and Architecture Metrics

ENLITENED Annual Meeting

July 19, 2022

Steve Crago
Associate Director, USC/ISI
Research Associate Professor, USC, ECE
crago@isi.edu, (703) 812-3729

John Paul Walters
Research Team Leader, USC/ISI
jwalters@isi.edu

David Kang
Computer Scientist, USC/ISI
dkang@isi.edu

Andy Rittenbach
Computer Scientist, USC/ISI
arittenbach@isi.edu

Peng Xie
Computer Scientist, USC/ISI
pengxie@isi.edu

Jana-Lynn Louis
Project Manager, USC/ISI
jlouis@isi.edu

STEAM: System Testbed, Evaluation, and Architecture Metrics Phase 2

Performer

USC Viterbi
School of Engineering
Information Sciences Institute

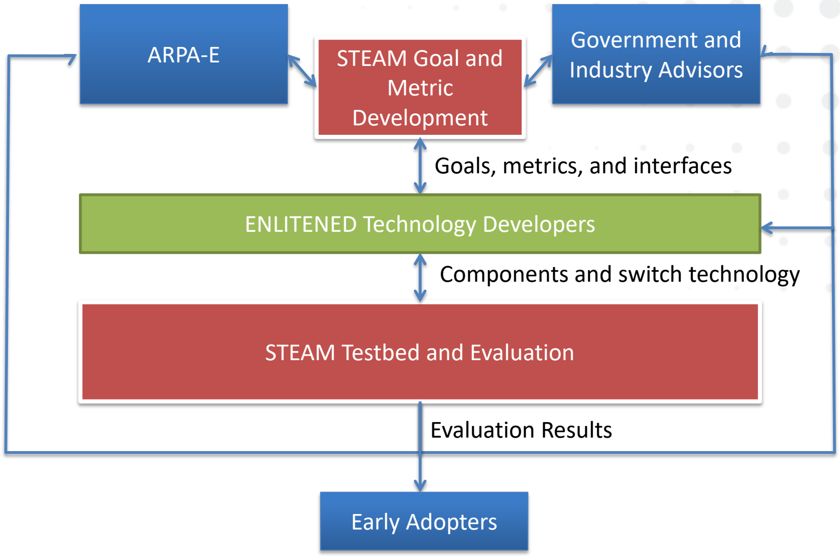

Technology Summary

STEAM will further develop our successful Phase 1 modeling and simulation capability to detailed datacenter and HPC workloads, ENLITENED components, and architectures to enable transition

Technical Approach

- Extend multi-resolution, multi-level modular testbed to characterize and demonstrate ENLITENED technology in datacenters and HPC
- Workloads based on existing and projected software stacks, focusing on hyperscale and HPC
- Amortize system-level characterization and modeling and simulation costs for ENLITENED projects

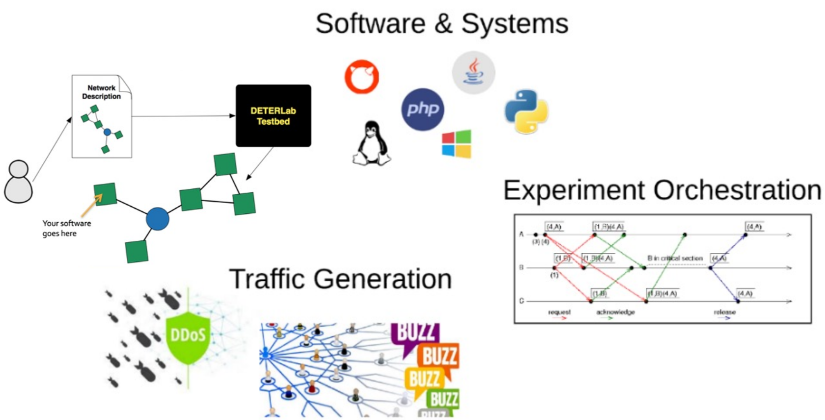

Industry Advisory Committee Membership

Name	Organization	Commercial	HPC	System	Component	User	Vendor	Amazon	Facebook	Apple	Google	Microsoft	Alibaba	DOE
Rob Stone	Facebook	X		X		X			X					
Robert Blum	Intel	X	X		X		X							
Constantinos Evanelinos	IBM		X	X		X	X							
Hitesh Billani	Microsoft	X		X		X						X		
Michael Woodacre	HPE	X	X	X			X							
Anthony Yu	GlobalFoundries	X			X		X							
Mahidhar Tatineni	SDSC		X	X		X								
Jeff Vetter	ORNL		X	X		X								X
Richard Carlson	DOE OS		X	X		X								X
Scott Hemmert	NNSA/Sandia		X	X		X								X
Vivek Raghuraman	Broadcom	X			X		X							
Cedric Lam	Google	X			X	X					X			

3 new members representing DOE/NNSA, Broadcom (vendor), Google (hyperscale datacenter operator). Quarterly meetings to resume soon.

Working Group Membership

Name	Team	Organization
Laurent Schares	MOTION	IBM
Pavlos Maniotis	MOTION	IBM
Madeleine Glick	PINE	Columbia
Max Mellette	LEED	InFocus
John Shalf	PINE	LBL
George Michelogiannakis	PINE	LBL
Manya Ghobadi	PINE	MIT
Larry Dennison	PINE	Nvidia
Adel Saleh	INTREPID	UCSB

IAC/WG Input Summary

- Systems
  - HPC systems/networks designed for the largest jobs (full systems)
  - Up to ~150K nodes today
  - Simulations of thousands of nodes needed
- Workloads
  - Large models for machine learning
  - Consider data analytics for HPC, synthetic benchmarks for irregular accesses
  - Most teams interested in single-tenant/single application workloads
  - MOTION interested in multi-tenant cloud workloads
  - Need mixed workload traces, will need to be developed by combining job traces with network traces
- Consider reliability and resilience, bursty workloads

General Workflow of Simulating Traces

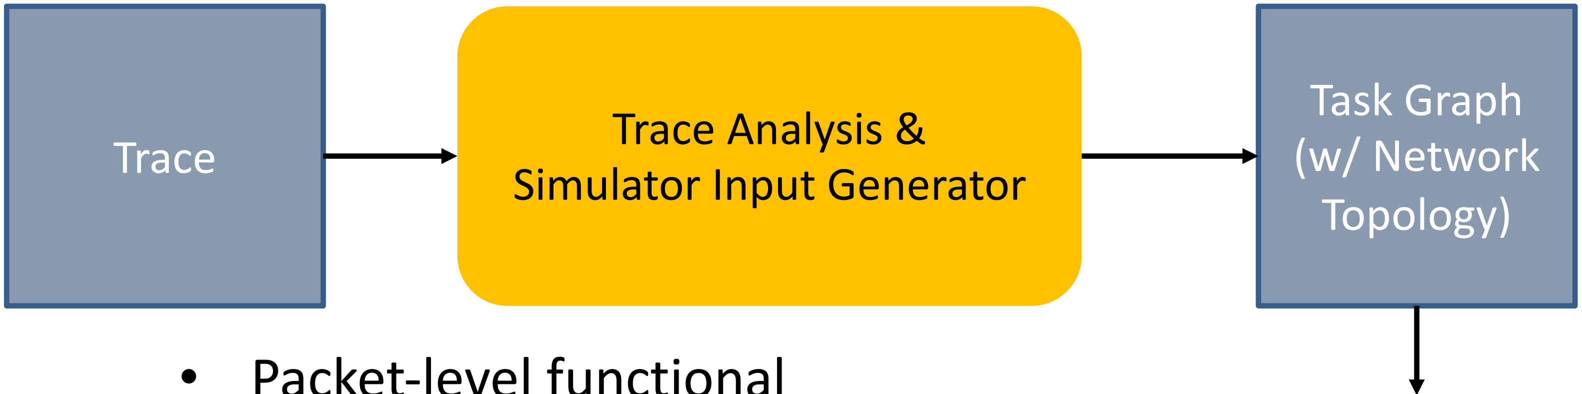

- Packet-level functional simulator
- Light-weight, fast, and highly scalable
- Used by both LEED and PINE teams
- STEAM enhancement to support dependencies and simulate INTREPID

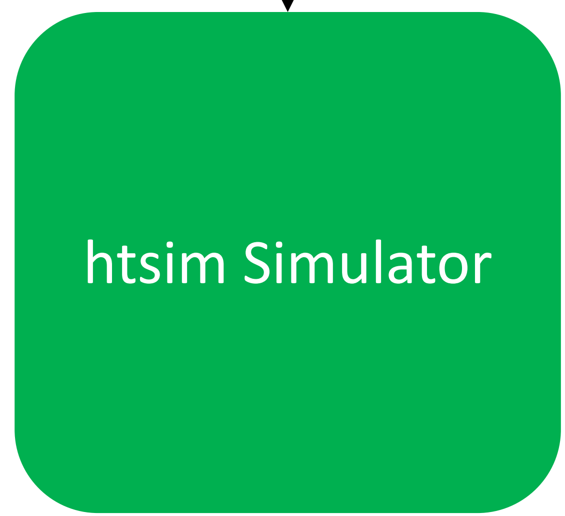

Proposed Benchmarks/Traces

After discussion with each team, we are currently planning to use the following benchmarks to evaluate each ENLITENED team.

Benchmark	Machine Learning	Storage	Big Data HPC	HPC Simulation	Multi-tenant cloud	LEED	PINE	INTREPID	MOTION
Alibaba Cluster Trace Program/cluster-trace-gpu-v2020	X				X				X
Facebook datacenter trace					X				X
DLRM	X	X	X			X	X	X	
CANDLE	X			X		X	X	X	
ML-Perf (Inference) NLP	X							X	

Red X indicates higher priority

DLRM Benchmark

- Deep learning recommendation model initially developed by Facebook for predicting whether users will interact with a particular advertisement
- Model can have 10s of trillions of parameters and consists of dense neural networks and very large (>20 GB) sparse embedding tables
- Despite having several orders of magnitude more parameters, only as computationally intensive as training popular transformer models such BERT
- Depending on how distributed training is configured, network can become training bottleneck

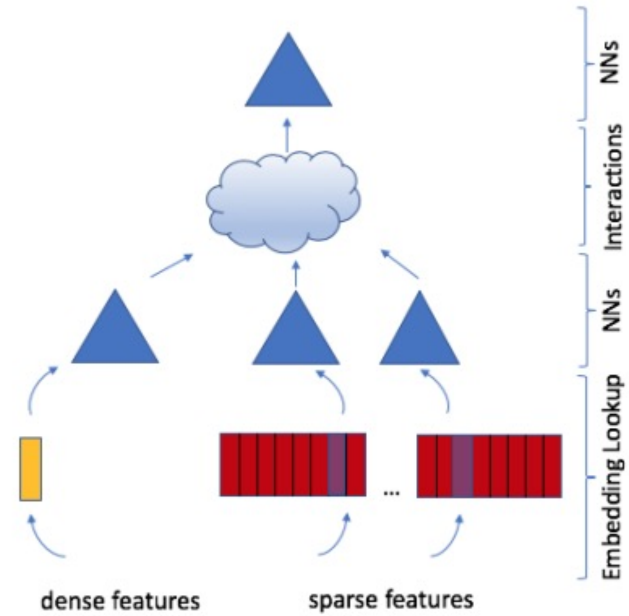

DLRM architecture consisting of dense neural networks and sparse embedding tables

Naumov, Maxim, et al. "Deep learning recommendation model for personalization and recommendation systems." *arXiv preprint arXiv:1906.00091* (2019).

CANDLE Benchmark

- Cancer Distributed Learning Environment benchmarks for training a model to predict whether tumor will interact with a particular drug
  - Initially developed as collaboration between DoE, NCI, and NIH
- Developed as part of Exascale Computing project for running on largescale HPC systems

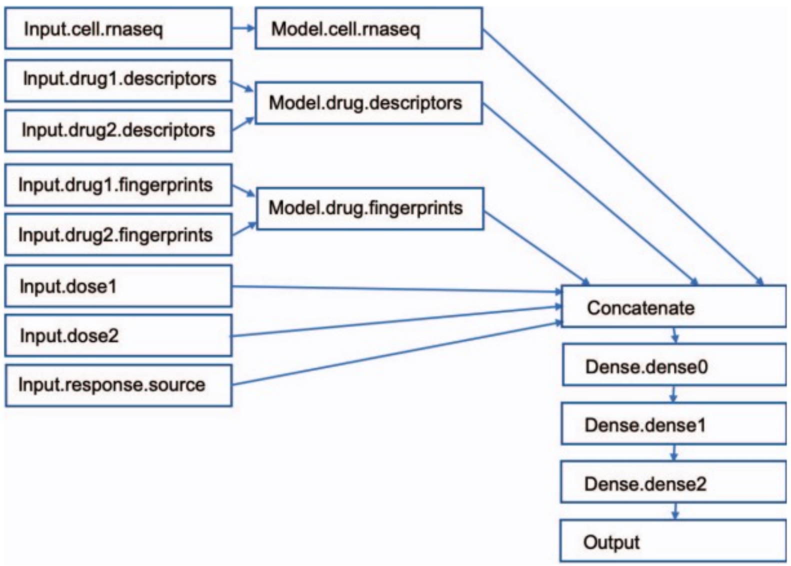

CANDLE UNO benchmark consisting of separate neural networks for RNA sequence, drug modeling, and final prediction of drug response

Wozniak, Justin M., et al. "High-bypass learning: Automated detection of tumor cells that significantly impact drug response." *2020 IEEE/ACM Workshop on Machine Learning in High Performance Computing Environments (MLHPC) and Workshop on Artificial Intelligence and Machine Learning for Scientific Applications (AI4S)*. IEEE, 2020.

Initial Benchmark Model Parameters

CANDLE

	8 Node Simulation	16 Node Simulation	32 Node Simulation
Batch size (per GPU)	**512**	**256**	**256**
# of Dense feature layers	**8**	**16**	**16**
# of weights per dense feature layer	**4096**	**16384**	**16384**
# of Dense layers	**4**	**8**	**8**
# of weights per dense layer	**4096**	**16384**	**16384**

DLRM

	8 node Simulation	16 Node Simulation	32 Node Simulation
Batch size (per GPU)	**512**	**128**	**128**
# of layers in bottom MLP	**2**	**8**	**8**
# of weights per layer in bottom MLP	**1024**	**4096**	**4096**
# of layers in top MLP	**4**	**8**	**16**
# of weights per layer in top MLP	**2048**	**8192**	**8192**
# of embedding tables	**8**	**16**	**32**
Embedding table dimensions	**256x10000000**	**128x10000000**	**128x10000000**

Preliminary PINE Simulations

- Number of nodes: N = {16, 32}
- 4 GPUs/node
- Three configurations
  - State-of-the-art commercially available hardware
    - Link speed: B = 400 Gps, # of links/node d = 4
    - Total bandwidth per node: 1.6 Tbps
  - PINE sprout
    - Link speed: B = 50 Gps, # of links/node d = 16
    - Total bandwidth per node: 0.8 Tbps
  - PINE sapling
    - Link speed: B = 50 Gps, # of links/node d = 32
    - Total bandwidth per node: 1.6 Tbps
- These are static configurations. We will be simulating network reconfigurations between benchmarks in the future.

State-of-the-art commercially available hardware

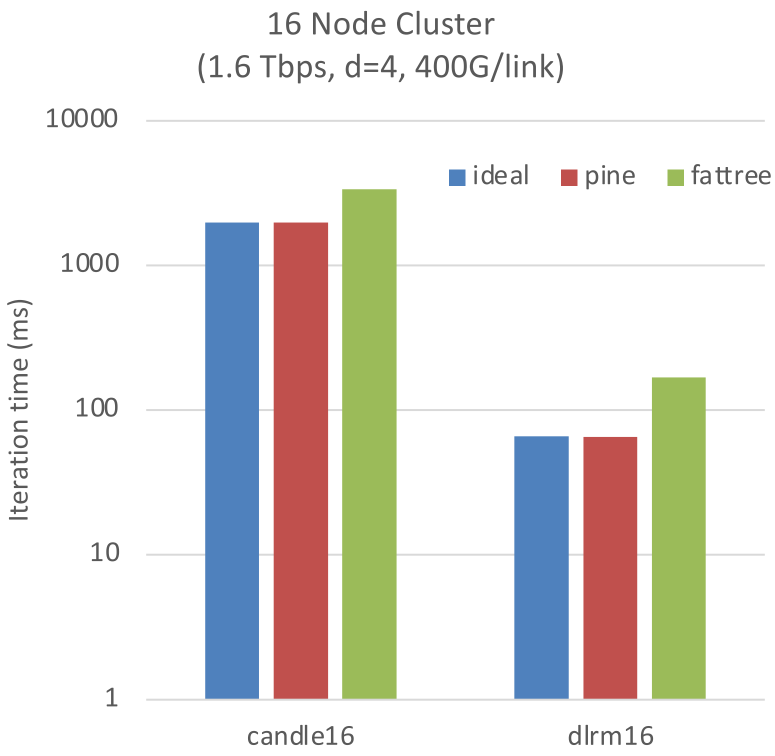

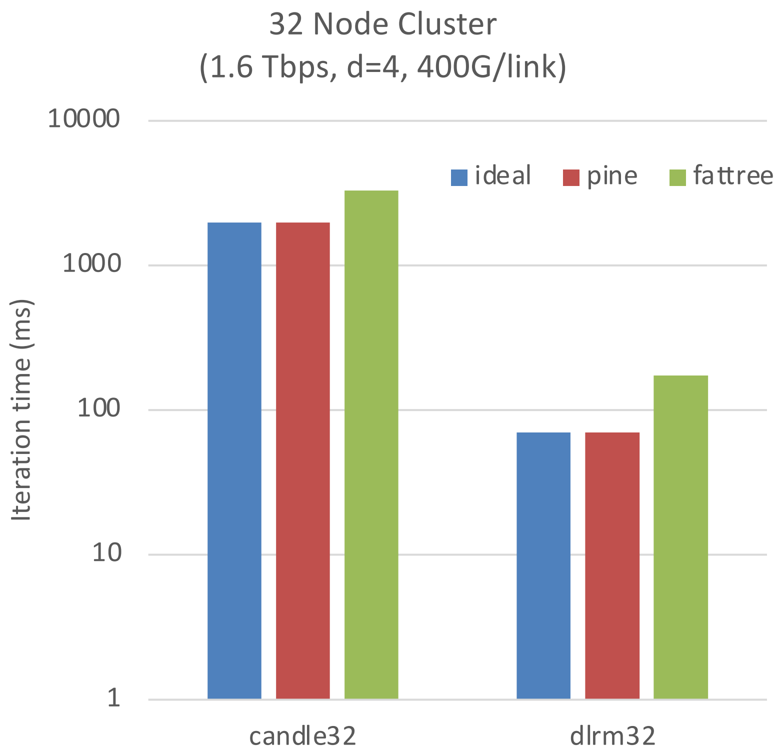

PINE Sprout

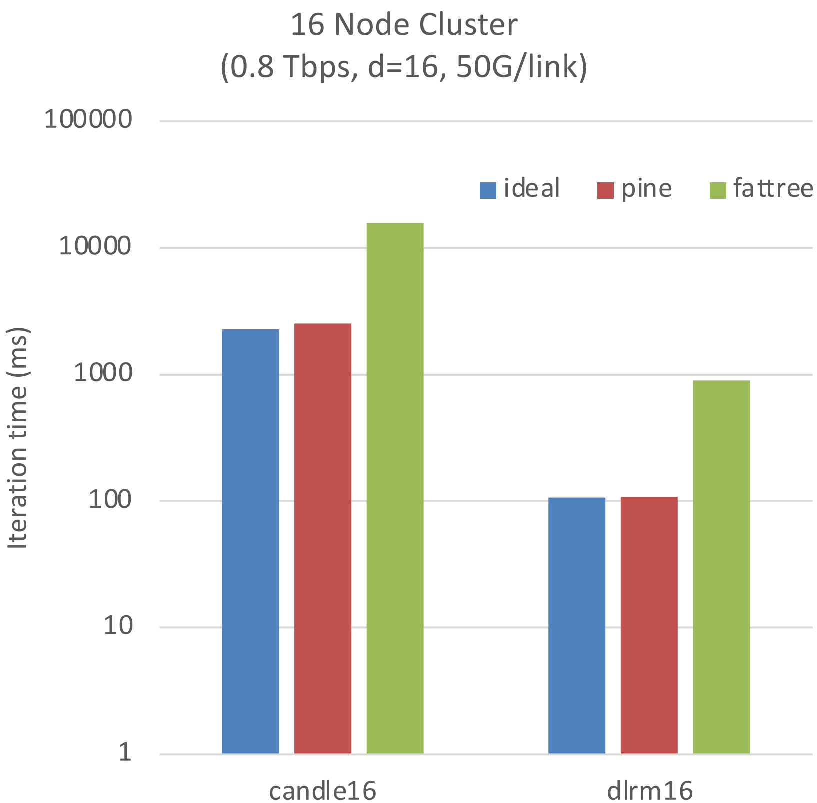
16 Node Cluster
(0.8 Tbps, d=16, 50G/link)
100000
10000
1000
100
10
1
Iteration time (ms)
ideal
pine
fattree
candle16
dlrm16

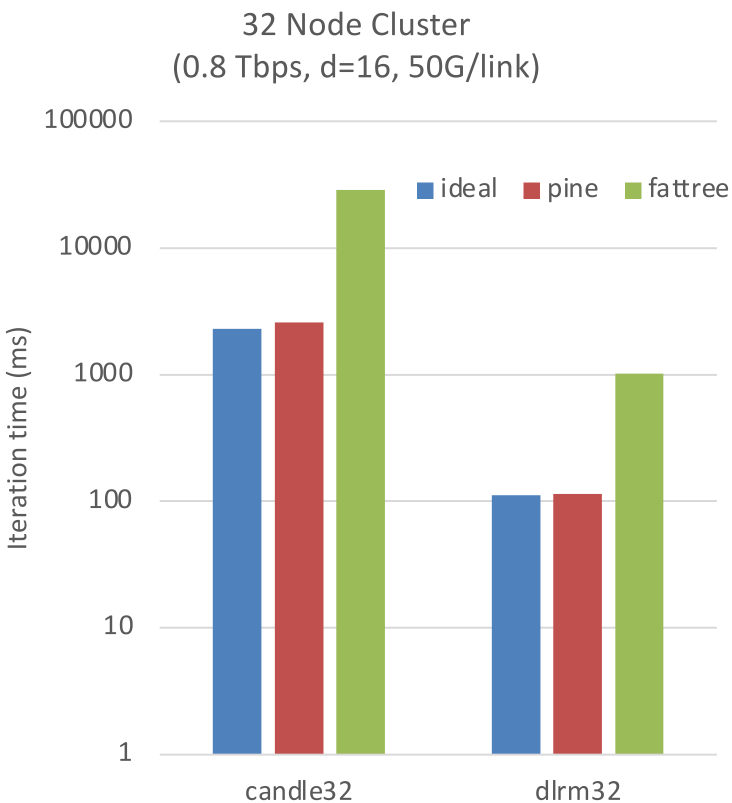
32 Node Cluster
(0.8 Tbps, d=16, 50G/link)
100000
10000
1000
100
10
1
Iteration time (ms)
ideal
pine
fattree
candle32
dlrm32

PINE Sapling

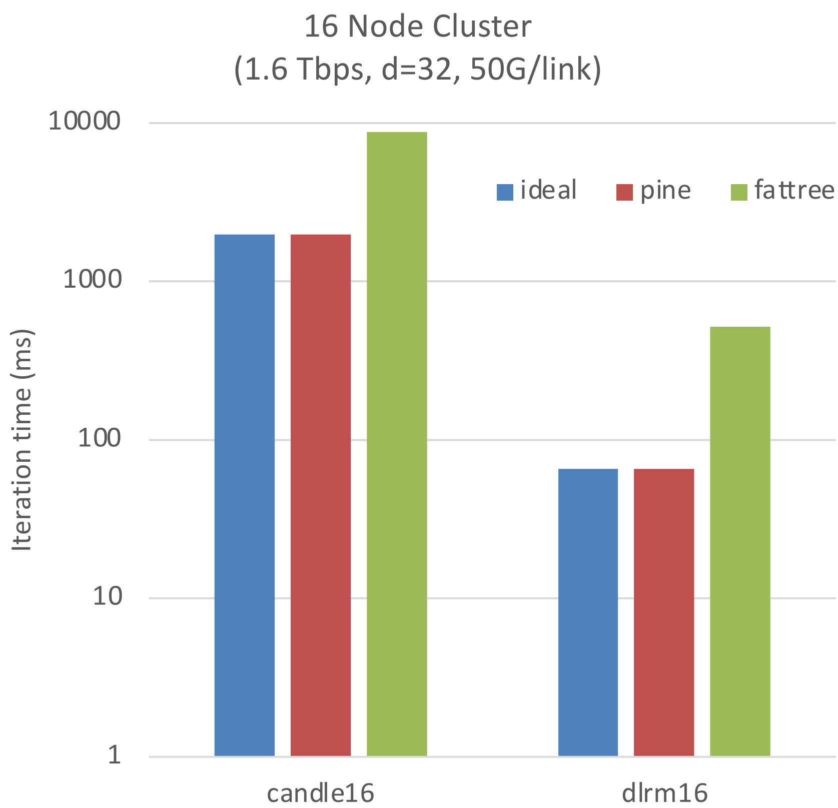

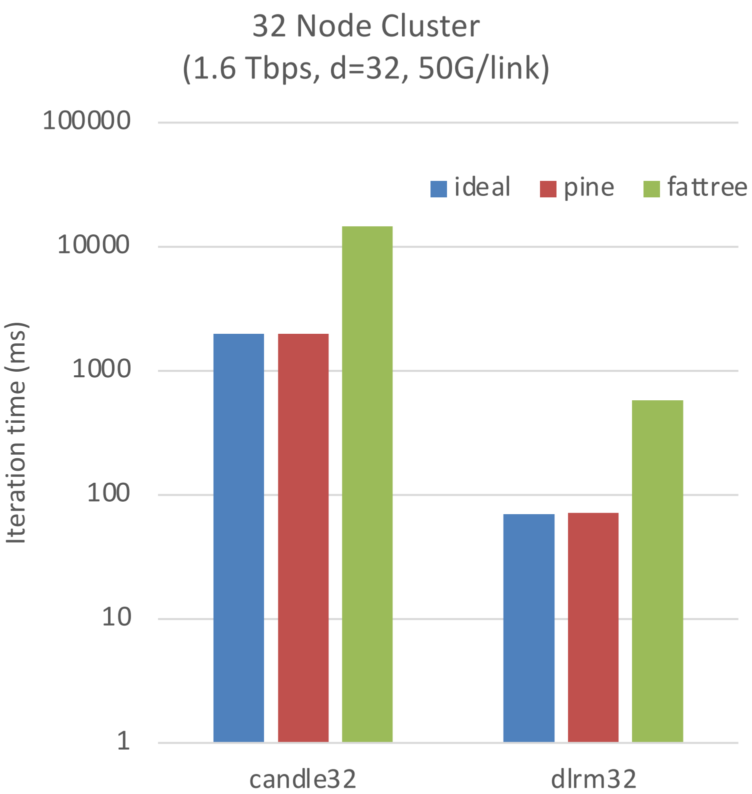

Preliminary INTREPID Simulations

- 3 simple topologies for proof of concept (8 node network)
  - Shortest path fixed routing assumed
- No attempt to optimize placement of tasks to nodes

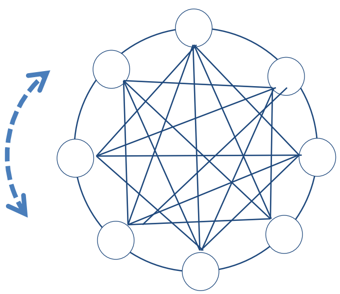

Ring 1
All-to-all
1 wavelength/connection

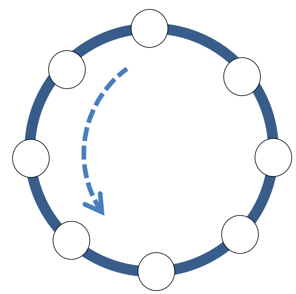

Ring 2
8 wavelengths/connection
Unidirection (counterclockwise)

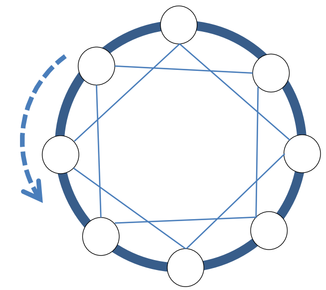

Ring 3
7 wavelength ring
1 wavelength ring
Unidirection (counterclockwise)

Preliminary Performance of 8 Node Topologies

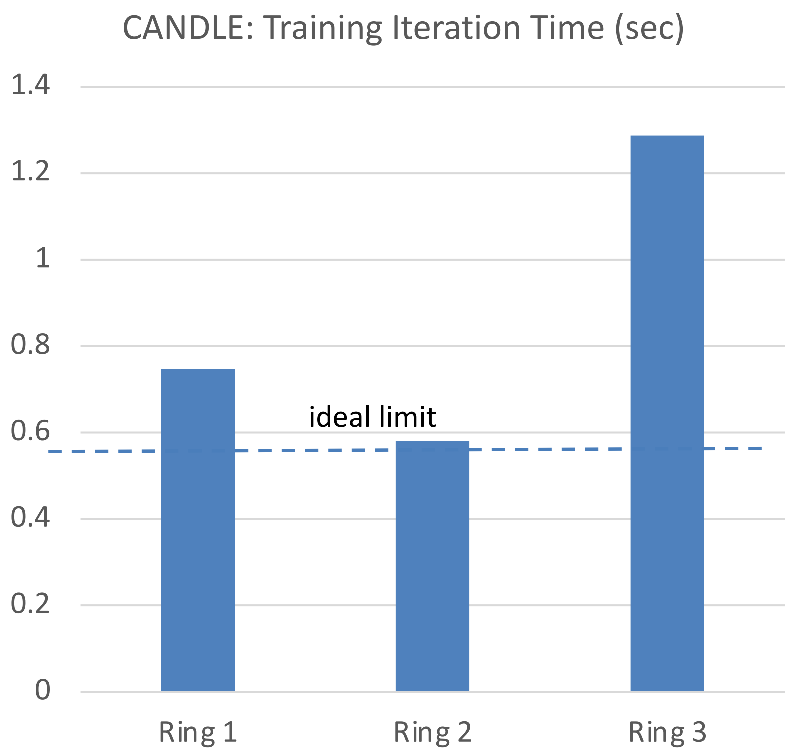

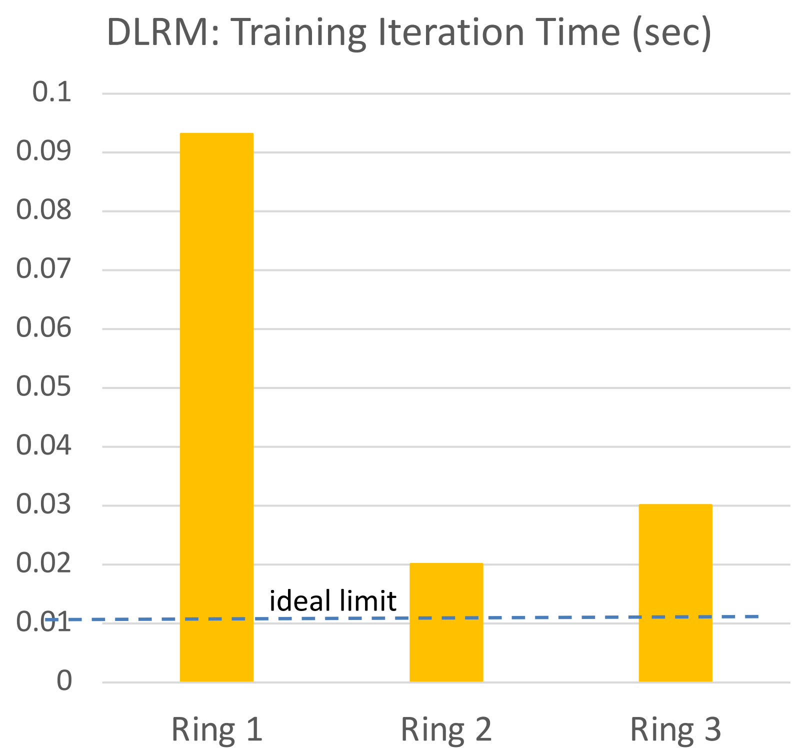

- Topology and placement critical to energy efficiency
- Need mapping tools as network size scales up

Summary

- IAC has been enhanced and quarterly meetings to resume ASAP
- Phase 2 simulation testbed operational and generating results for PINE and INTREPID
  - LEED results coming soon
  - Results are being discussed with the technology development teams
- LightCounting Study Update next
- Panel discussions after that