Georgia Tech School of Electrical and Computer Engineering
College of Engineering

SYNERGY

http://synergy.ece.gatech.edu

STRA SIM

Conclusion and Next Steps

Taekyung Heo
Postdoctoral Fellow, School of ECE
Georgia Institute of Technology
taekyung@gatech.edu

Acknowledgments: William Won (GT), Srinivas Sridharan (Facebook), Sudarshan Srinivasan (Intel)

Motivation of this Tutorial

- Large model distributed training is an ongoing open-research area
- Many emerging supercomputing systems being designed specifically for this problem!
  - Cerebras CS2
  - Tesla Dojo
  - NVIDIA DGX + Mellanox SHARP switches
  - Intel Habana
  - IBM Blueconnect
  - Facebook Zion
  - …
- Co-design of algorithm and system offers high opportunities for speedup and efficiency

ASTRA-sim: Status and On-going Development

✓ *Released* ➢ *In progress*

- ✓ Supports Data-Parallel, Model-Parallel, Hybrid-Parallel training loops
- ✓ Extensible to more training loops
  - ➢ Graph-based input from PyTorch

- ✓ Ring based, Tree-based, AlltoAll based, and multi-phase collectives
- ✓ Variety of scheduling policies
- ✓ Compute times fed via offline system measurements or compute simulator

- ✓ Various topologies, flow-control, link bandwidth, congestion control
- ✓ Plug-and-play options
  - ✓ Analytical (roofline)
  - ➢ Analytical with congestion
  - ✓ Garnet (credit-based)
  - ➢ NS3 (TCP, RDMA)

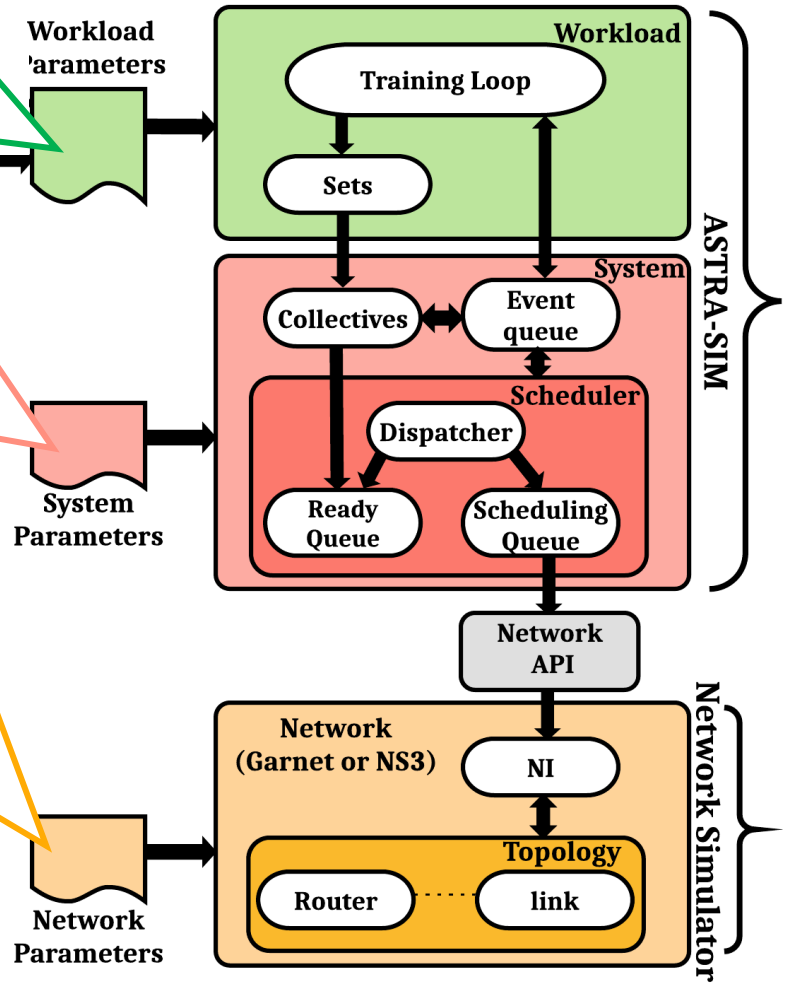

http://github.com/astra-sim/astra-sim

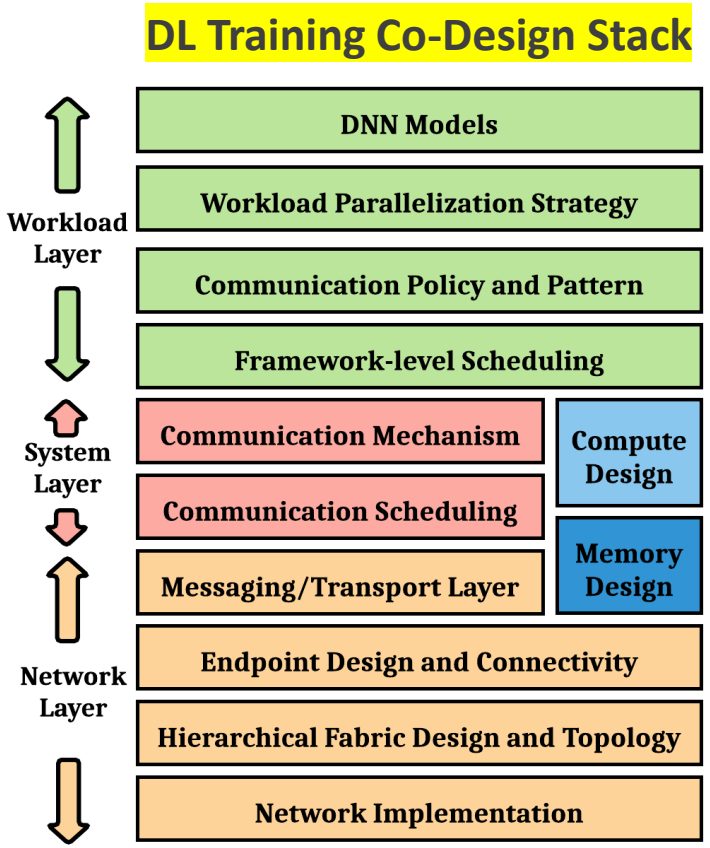

*S. Rashidi et al.,"**ASTRA-SIM: Enabling SW/HW Co-Design Exploration for Distributed DL Training Platforms**", ISPASS 2020*

*S. Rashidi, et al., "**Scalable Distributed Training of Recommendation Models: An ASTRA-SIM + NS3 case-study with TCP/IP transport**", Hot Interconnects 2020*

ASTRA-sim: Status and On-going Development

✓ *Released* ➢ *In progress*

- ✓ Supports Data-Parallel, Model-Parallel, Hybrid-Parallel training loops
- ✓ Extensible to more training loops
  - ➢ Graph-based input from PyTorch

- ✓ Ring based, Tree-based, AlltoAll based, and multi-phase collectives
- ✓ Variety of scheduling policies
- ✓ Compute times fed via offline system measurements or compute simulator

- ✓ Various topologies, flow-control, link bandwidth, congestion control
- ✓ Plug-and-play options
  - ✓ Analytical (roofline)
  - ➢ Analytical with congestion
  - ✓ Garnet (credit-based)
  - ➢ NS3 (TCP, RDMA)

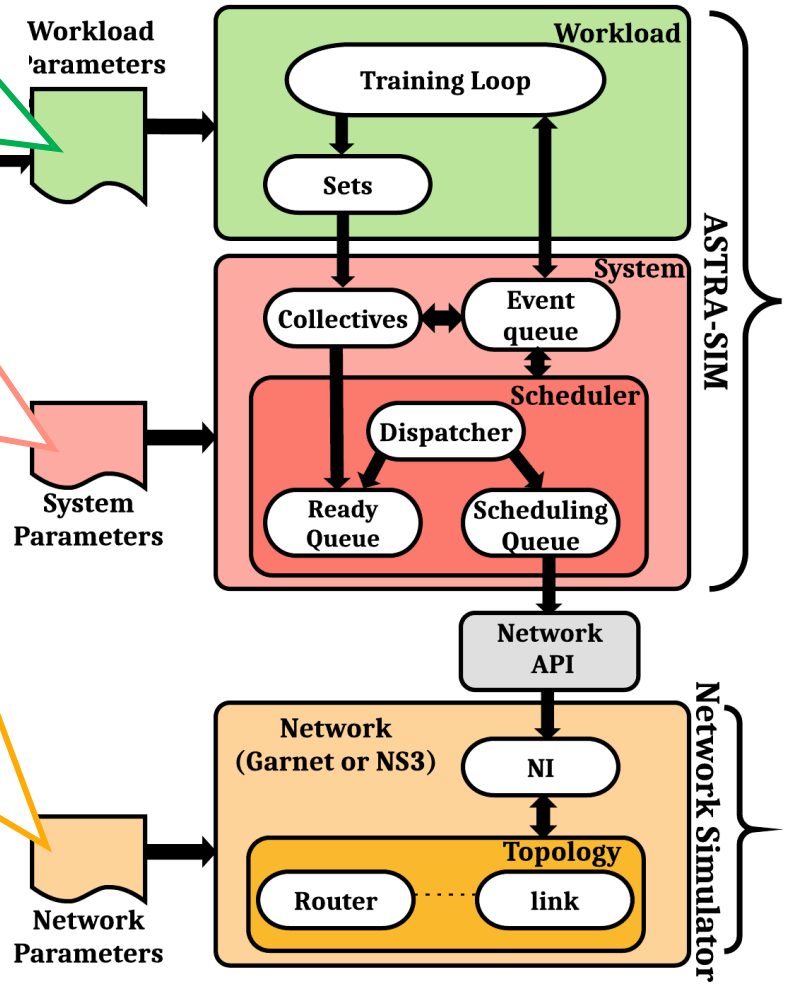

http://github.com/astra-sim/astra-sim

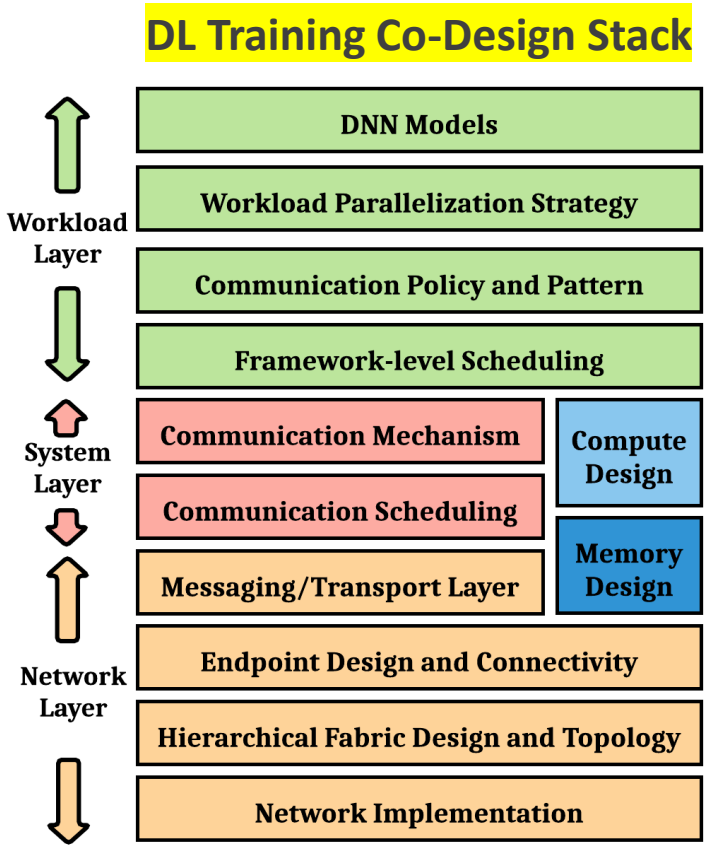

*S. Rashidi et al.,"**ASTRA-SIM: Enabling SW/HW Co-Design Exploration for Distributed DL Training Platforms**", ISPASS 2020*

*S. Rashidi, et al., "**Scalable Distributed Training of Recommendation Models: An ASTRA-SIM + NS3 case-study with TCP/IP transport**", Hot Interconnects 2020*

Modeling Real-world Execution Traces

- **Limitation**: ASTRA-sim cannot model complex training loops (pipeline)
- **Solution**: Run ASTRA-sim with execution graphs

Current ASTRA-sim

Execution-graph-based ASTRA-sim

Will be released by the next tutorial at MLSys 2022

ASTRA-sim: Status and On-going Development

✓ *Released* ➢ *In progress*

- ✓ Supports Data-Parallel, Model-Parallel, Hybrid-Parallel training loops
- ✓ Extensible to more training loops
  - ➢ Graph-based input from PyTorch

- ✓ Ring based, Tree-based, AlltoAll based, and multi-phase collectives
- ✓ Variety of scheduling policies
- ✓ Compute times fed via offline system measurements or compute simulator

- ✓ Various topologies, flow-control, link bandwidth, congestion control
- ✓ Plug-and-play options
  - ✓ Analytical (roofline)
  - ➢ Analytical with congestion
  - ✓ Garnet (credit-based)
  - ➢ NS3 (TCP, RDMA)

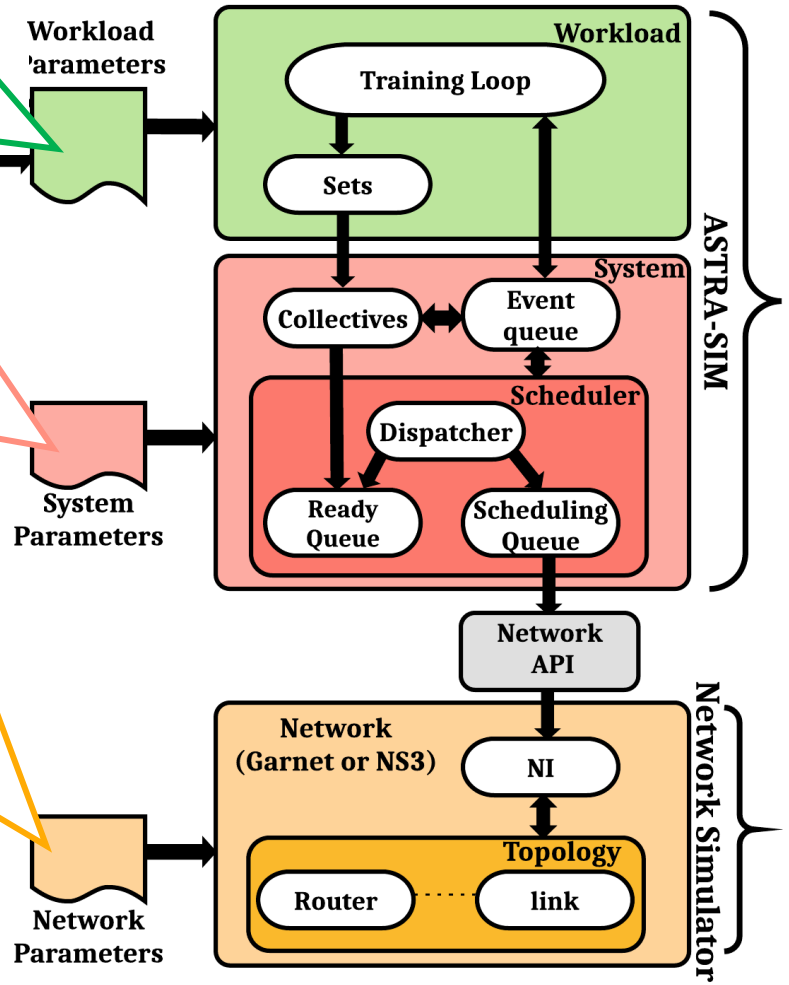

http://github.com/astra-sim/astra-sim

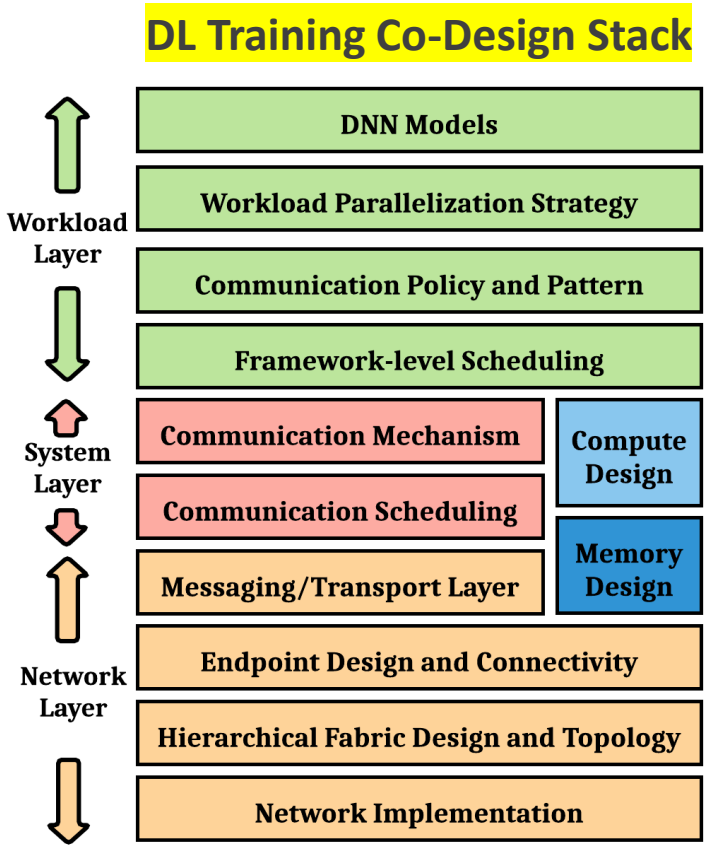

*S. Rashidi et al.,"**ASTRA-SIM: Enabling SW/HW Co-Design Exploration for Distributed DL Training Platforms**", ISPASS 2020*

*S. Rashidi, et al., "**Scalable Distributed Training of Recommendation Models: An ASTRA-SIM + NS3 case-study with TCP/IP transport**", Hot Interconnects 2020*

Network Backends

- **Analytical:**
  - Fastest backend.
  - Models rich set of hierarchical networks.
  - Accurate for congestion-less topology/comm patterns.
- **Garnet:**
  - Credit-based flow control modeling.
  - Most accurate for NOCs and chiplet-based interconnects.
- **Analytical + Congestion Modeling (under development):**
  - Same as Analytical but performs message-level congestion modeling.
  - Expected to model patterns with congestions with 10-20% error rate.
- **NS3 (under development):**
  - Models RDMA over converged ethernet (RoCE) comm protocol.
  - Supports several congestion control schemes (DCQCN, HPCC, Timely, etc.).

Contribution and Participation

- **ASTRA-sim is open-source!**
  - Feel free to raise github issues and contribute via pull-requests

- **Next Tutorial(s):**
  - MLSys 2022 (August 31st, 2022) in Santa Clara, CA

Organization Team

Contact any/all four of us if any questions

Presenters

Tushar Krishna
Associate Professor, School of ECE
Georgia Institute of Technology
tushar@ece.gatech.edu

Saeed Rashidi
PhD Student, School of ECE
Georgia Institute of Technology
saeed.rashidi@ece.gatech.edu

Will Won
Ph.D. Student, School of CS
Georgia Institute of Technology
william.won@gatech.edu

Taekyung Heo
Postdoctoral Fellow, School of ECE
Georgia Institute of Technology
taekyung@gatech.edu

Collaborators

Srinivas Sridharan
Research Scientist, Meta

Sudarshan Srinivasan
Research Scientist, Intel

Thank you!