

Conclusion and Next Steps



Taekyung Heo

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

Acknowledgments: William Won (GT), Srinivas Sridharan (Meta), 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

Contribution and Participation

- **The new features will be released soon!**
 - Please reach out to us for early access
- **ASTRA-sim is open-source!**
 - Feel free to raise GitHub issues and contribute via pull-requests

Organization Team

Contact any/all five of us
if any questions



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

Presenters



Taekyung Heo

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



Srinivas Sridharan

Research Scientist
~~Sudarshan Srinivasan~~
Research Scientist, Intel
ssrinivas@fb.com

Collaborators and Contributors

Sudarshan Srinivasan

Research Scientist, Intel

Tarannum Khan

Graduate Student, CS
UT Austin

Aditya Akella

Professor, CS
UT Austin

Thank you!