









Sahil Tyagi

 Email  Website  LinkedIn  GitHub  Google Scholar  ORCID

Education




- 2018 – 2024  **Ph.D., Intelligent Systems Engineering**, Indiana University Bloomington
Thesis: Towards Building Efficient Computation and Communication Models for Deep Learning Systems
Advisor: Martin Swamy
Major: Computer Engineering
CGPA: 3.77/4.0
Expected graduation date: August 2024
- 2009 – 2013  **Bachelor of Technology (B. Tech.), Electrical and Electronics Engineering**, Guru Gobind Singh Indraprastha University, New Delhi, India.
Grade: 7.6/10.0

Research Interests












-  Large-scale ML Systems
-  Deep Learning
-  Federated Learning
-  Distributed Systems
-  Edge, Cloud and High-Performance Computing (HPC)
-  Intelligent Computing

Research Publications

Journal Articles

-  1 S. Tyagi and P. Sharma, “OmniLearn: A Framework for Distributed Deep Learning over Heterogeneous Clusters,” (*under review*), 2024.
-  2 S. Chaturvedi, S. Tyagi, and Y. Simmhan, “Cost-Effective Sharing of Streaming Dataflows for IoT Applications,” *IEEE Transactions on Cloud Computing*, vol. 9, no. 4, pp. 1391–1407, 2021.  DOI: 10.1109/TCC.2019.2921371.

Conference Proceedings

-  1 S. Tyagi and M. Swamy, “Enabling Large-Batch Training via Learned Gradient Mapping,” in (*in preparation*), 2024.
-  2 S. Tyagi and M. Swamy, “On Using Large-Batches in Federated Learning,” in (*in preparation*), 2024.
-  3 S. Tyagi, “Scavenger: A Cloud Service for Optimizing Cost and Performance of DL Training,” in *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing Workshops (CCGridW)*, Los Alamitos, CA, USA: IEEE Computer Society, May 2023, pp. 349–350.  DOI: 10.1109/CCGridW59191.2023.00081.
-  4 S. Tyagi and P. Sharma, “Scavenger: A Cloud Service for Optimizing Cost and Performance of ML Training,” in *2023 IEEE/ACM 23rd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, *Accept. Rate: 21%*, 2023, pp. 403–413.  DOI: 10.1109/CCGrid57682.2023.00045.
-  5 S. Tyagi and M. Swamy, “Accelerating Distributed ML Training via Selective Synchronization,” in *IEEE International Conference on Cluster Computing (CLUSTER) 2023, Santa Fe, NM, USA, October 31 - Nov. 3, 2023*, *Accept. Rate: 25%*, IEEE, 2023, pp. 1–12.  DOI: 10.1109/CLUSTER52292.2023.00008.
-  6 S. Tyagi and M. Swamy, “Accelerating Distributed ML Training via Selective Synchronization (Poster Abstract),” in *2023 IEEE International Conference on Cluster Computing Workshops (CLUSTER Workshops)*, 2023, pp. 56–57.  DOI: 10.1109/CLUSTERWorkshops61457.2023.00023.
-  7 S. Tyagi and M. Swamy, “Flexible Communication for Optimal Distributed Learning over Unpredictable Networks,” in *2023 IEEE International Conference on Big Data (BigData), Sorrento, Italy*, *Accept. Rate: 17.5%*, Dec. 2023.

- 8 S. Tyagi and M. Swamy, "GraVAC: Adaptive Compression for Communication-Efficient Distributed DL Training" in *16th IEEE International Conference on Cloud Computing (CLOUD) 2023, Chicago, IL, USA, July 2-8, 2023*, Accept. Rate: 20%, IEEE, 2023, pp. 319–329. [DOI: 10.1109/CLOUD60044.2023.00045](https://doi.org/10.1109/CLOUD60044.2023.00045).
- 9 S. Tyagi and M. Swamy, "ScaDLES: Scalable Deep Learning over Streaming Data at the Edge," in *2022 IEEE International Conference on Big Data (BigData)*, Accept. Rate: 19.2%, Los Alamitos, CA, USA: IEEE Computer Society, Dec. 2022, pp. 2113–2122. [DOI: 10.1109/BigData55660.2022.10020597](https://doi.org/10.1109/BigData55660.2022.10020597).
- 10 S. Tyagi and P. Sharma, "Taming Resource Heterogeneity in Distributed ML Training with Dynamic Batching," in *2020 IEEE International Conference on Autonomic Computing and Self-Organizing Systems (ACSOS)*, Accept. Rate: 25%, Los Alamitos, CA, USA: IEEE Computer Society, Aug. 2020, pp. 188–194. [DOI: 10.1109/ACSOS49614.2020.00041](https://doi.org/10.1109/ACSOS49614.2020.00041).
- 11 C. Widanage, J. Li, S. Tyagi, *et al.*, "Anomaly Detection over Streaming Data: Indy500 Case Study," in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*, Accept. Rate: 20%, 2019, pp. 9–16. [DOI: 10.1109/CLOUD.2019.00015](https://doi.org/10.1109/CLOUD.2019.00015).
- 12 J. Qiu, B. Peng, R. Teja, S. Tyagi, C. Widanage, and J. Koskey, "Real-Time Anomaly Detection from Edge to HPC-Cloud," in *2018 Big Data and Exascale Computing Workshop (BDEC2)*, 2018. [URL: https://exascale.org/bdec/sites/exascale.org.bdec/files/whitepapers/Qiu_BDEC2_WP.pdf](https://exascale.org/bdec/sites/exascale.org.bdec/files/whitepapers/Qiu_BDEC2_WP.pdf).
- 13 S. Chaturvedi, S. Tyagi, and Y. Simmhan, "Collaborative Reuse of Streaming Dataflows in IoT Applications," in *2017 IEEE 13th International Conference on e-Science (e-Science)*, Accept. Rate: 36%, 2017, pp. 403–412. [DOI: 10.1109/eScience.2017.54](https://doi.org/10.1109/eScience.2017.54).

Skills

Programming	Python, C, C++, Shell scripting, MPI, OpenMP, CUDA, Java, Scala, SQL, MATLAB
Networking	Wireshark, TCPDump, Traffic Control (tc), Scapy, IPerf.
Databases	MySQL, MSSQL, SQLite, HBase, Neo4j
Frameworks	PyTorch, TensorFlow, MXNet, Keras, Hadoop, Spark, Kafka, Slurm, Cloud APIs

Teaching Experience

Associate Instructor	High-Performance Computing: Spring 2024
	Computer Networks: Fall 2023, Fall 2022
	Operating Systems: Spring 2023
	Engineering Distributed Systems: Spring 2022, Spring 2021
	Cloud Computing: Fall 2021, Fall 2020, Fall 2019

Employment History

2018 – 2024	Graduate Researcher and Associate Instructor , Luddy School of Informatics, Computing and Engineering, Indiana University Bloomington, USA.
2017 – 2018	Research Staff Member , Dept. of Computational and Data Sciences (CDS), Indian Institute of Science (IISc), Bengaluru, India.
2016 – 2016	Data Scientist , HT Media Limited, Gurugram, Haryana, India.
2015 – 2015	Data Engineer , Stayzilla, Bengaluru, Karnataka, India.
2014–2015	Software Engineer , Tatra Data Limited, New Delhi, India.

Miscellaneous

Awards and Achievements

- NSF Student Grant:** To present research at IEEE CLUSTER 2023, Santa Fe, New Mexico.

Miscellaneous (continued)

- 📌 **Luddy Dean's Graduate Student Award:** In Fall 2023 for outstanding research.
- 📌 **NSF Travel Award:** To present research at IEEE/ACM CCGrid 2023, Bengaluru, India.
- 📌 **Best early-career researcher poster award:** Awarded at IEEE/ACM CCGrid 2023.
- 📌 **Google Cloud Student Researcher (2021, 2022):** Received GCP credits for research.
- 📌 **Student Research Award:** Funded via NSF grant Data Infrastructure Building Blocks (DiBBS) 17-500, for academic year 2018-2019.

Professional Services

- 📌 **2024: USENIX OSDI (AEC), IEEE CLUSTER (TPC), USENIX ATC (AEC)**

Presentations and Talks

- 📌 **4/24:** Guest lectures, "Parallel Computing with GPUs for Distributed ML Applications", High-Performance Computing (HPC) course, Indiana University Bloomington, USA.
- 📌 **12/23:** Paper presentation, "Flexible Communication for Optimal Distributed Learning over Unpredictable Networks." 2023 IEEE International Conference on Big Data, Sorrento, Italy.
- 📌 **11/23:** Paper presentation, "Accelerating Distributed ML Training via Selective Synchronization." 2023 IEEE International Conference on Cluster Computing, Santa Fe, New Mexico, USA.
- 📌 **11/23:** Poster presentation, "Accelerating Distributed ML Training via Selective Synchronization." 2023 IEEE International Conference on Cluster Computing, Santa Fe, New Mexico, USA.
- 📌 **09/23:** Invited talk, "Towards building efficient computation and communication models for distributed deep learning systems." Mathematics and Computer Science (MCS) division, Argonne National Laboratory, Illinois, USA.
- 📌 **07/23:** Paper presentation, "GraVAC: Adaptive Compression for Communication-Efficient Distributed DL Training." 2023 IEEE International Conference on Cloud Computing, Chicago, Illinois.
- 📌 **05/23:** Paper presentation, "Scavenger: A Cloud Service for Optimizing Cost and Performance of ML Training." 2023 IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, Bengaluru, India.
- 📌 **05/23:** Poster presentation, "Scavenger: A Cloud Service for Optimizing Cost and Performance of ML Training." 2023 IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing, Bengaluru, India.
- 📌 **12/22:** Paper presentation, "ScaDLES: Scalable Deep Learning over Streaming Data at the Edge." 2022 IEEE International Conference on Big Data, Osaka, Japan.
- 📌 **07/20:** Paper presentation, "Taming Resource Heterogeneity in Distributed ML Training with Dynamic Batching." 2020 IEEE International Conference on Autonomic Computing and Self-Organizing Systems, virtual.
- 📌 **11/18:** "Real-Time Anomaly Detection from Edge to HPC-Cloud", Intel Speakerships at SC18 (Proceedings of the International Conference for High Performance Computing, Networking, Storage, and Analysis 2018), Dallas, Texas, USA.

References

Available upon request