

Annus Zulfiqar

Graduate Student and Research Assistant,
Computer Science and Engineering (CSE),
University of Michigan, Ann Arbor, MI

+1 (765) 746-9458
zulfiqaa@umich.edu
<https://annuszulfiqar.com>
<https://www.linkedin.com/in/annuszulfiqar/>

EDUCATION

University of Michigan

Ph.D. in Computer Science & Engineering

Dissertation: Rearchitecting the End Host Network for the Terabit Per Second Era

Advisor: Muhammad Shahbaz

Ann Arbor, MI

Sep 2021 - May 2026

National University of Sciences and Technology (NUST)

Bachelor of Electrical Engineering

Thesis: Forest Cover Detection and Change Estimation using Deep Learning

Advisors: Muhammad Shahzad, Faisal Shafait

Islamabad, PK

Sep 2015 - May 2019

HONORS AND AWARDS

- Finalist in the 3-Minute Thesis (3MT) competition at University of Michigan 2026
- Broadcom Research Award for Celeris, CAL 2025 2025
- Selected as mentor for P4 Language Consortium, Google Summer of Code (GSoC) 2025
- Distinguished Artifact Award for Homunculus, ASPLOS 2023 2023
- Conference travel grants: ASPLOS 2022, SIGCOMM 2022, NSDI 2025, SIGCOMM 2025 2022-25
- Ross Fellowship at Purdue University 2021
- National P@SHA Information and Communication Technology (ICT) Award Winner, Pakistan 2021
- Travel award for graduate EEcamp at KAIST, South Korea 2018
- One-year internship offered at DFKI, Kaiserslautern, Germany (passed) 2018
- DAAD-funded internship at Technical University of Kaiserslautern (TUK), Germany 2018
- NUST merit scholarship for top academic performance (4.00/4.00 GPA) 2015-19

PUBLICATIONS

Conference Papers

NSDI. Murayyiam Parvez*, *Annus Zulfiqar**, Sylee Beltiukov, Shir Landau-Feibish, Arpit Gupta, Paper/Code
Walter Willinger, and Muhammad Shahbaz. SPLIDT: Partitioned Decision Trees for Scalable
Stateful Inference at Line Rate. 2025. *Acceptance rate: 22.1% (*co-primary author)*

MICRO. Gerasimos Gerogiannis, Charles Block, Dimitrios Merkouriadis, *Annus Zulfiqar*, Paper
Muhammad Shahbaz, and Josep Torrellas. NETSPARSE: Hardware Acceleration for Distributed
Sparse Kernels. 2025. *Acceptance rate: 20.7%*

ASPLOS. *Annus Zulfiqar*, Ali Imran, Venkat Kunaparaju, Gianni Antichi, Ben Pfaff, and Paper/Code
Muhammad Shahbaz. GIGAFLOW: Pipeline-Aware Sub-Traversal Caching for Modern
SmartNICs. 2025. *Acceptance rate: 19%*

ASPLOS. Tushar Swamy, *Annus Zulfiqar*, Muhammad Shahbaz, Luigi Nardi, and Kunle Olukotun. Paper/Code
HOMUNCULUS: Auto-Generating Efficient Data-Plane ML Pipelines for Datacenter Networks.
2023. *Acceptance rate: 20%*

Distinguished Artifact Award

Journal Articles

IEEE CAL. Ertza Warraich, Ali Imran, *Annus Zulfiqar*, Shay Vargaftik, Sonia Fahmy, and Paper
Muhammad Shahbaz. Reimagining RDMA Through the Lens of ML. 2025

Broadcom Research Award

SIGCOMM CCR. *Annus Zulfiqar*, Gianni Antichi, Ben Pfaff, William Tu, and Muhammad Shahbaz. The Slow-Path Needs an Accelerator Too! 2023 Paper

Journal of Applied Remote Sensing (JARS). *Annus Zulfiqar*, Muhammad M. Ghaffar, Muhammad Shahzad, Christian Weis, Muhammad I. Malik, Faisal Shafait, and Norbert Wehn. AI-ForestWatch: Semantic Segmentation Based End-to-End Framework for Forest Estimation and Change Detection using Multi-Spectral Remote Sensing Imagery. 2021 Paper/Code

Preprints

ARXIV. Ertza Warraich, Ali Imran, *Annus Zulfiqar*, Shay Vargaftik, Sonia Fahmy, and Muhammad Shahbaz. OPTINIC: A Resilient and Tail-Optimal RDMA NIC For Distributed ML Workloads. 2025 Paper

Under Review

Regan McDonald, Marilyn Rego, Ertza Warraich, *Annus Zulfiqar*, Muhammad Shahbaz. Towards Network-Efficient Cross-Regional Inference via Learned Activation Compression. 2026

Annus Zulfiqar, Ben Pfaff, Gianni Antichi, and Muhammad Shahbaz. Incremental Cache Eviction for Modern vSwitches. 2026

Conference & Workshop Extended Abstracts

SIGCOMM. *Annus Zulfiqar*, Ben Pfaff, Gianni Antichi, Arpit Gupta, and Muhammad Shahbaz. KAIRO: Incremental View Maintenance for Scalable Virtual Switch Caching. 2025 Poster

NSDI. Murayyiam Parvez*, *Annus Zulfiqar**, Sylee Beltiukov, Shir Landau-Feibish, Arpit Gupta, Walter Willinger, and Muhammad Shahbaz. BRANCHPIPE: Scalable Decision Trees for Stateful Processing at Line Rate. 2025 Poster

NSDI. *Annus Zulfiqar*, Ali Imran, Venkat Kunaparaju, Gianni Antichi, Ben Pfaff, and Muhammad Shahbaz. A Smart Cache for a SmartNIC! Rethinking Caching, Locality, & Revalidation for Modern Virtual Switches. 2025 Poster

SRC TECHCON. Marilyn Rego, Murayyiam Parvez, *Annus Zulfiqar*, Sylee Beltiukov, Shir Landau-Feibish, Arpit Gupta, Walter Willinger, and Muhammad Shahbaz. SPLIDT: Partitioned Decision Trees for Scalable Stateful ML Inference at Line Rate. 2025

Hot Chips. *Annus Zulfiqar*, Ali Imran, Venkat Kunaparaju, Gianni Antichi, Ben Pfaff, and Muhammad Shahbaz. A Smart Cache for a SmartNIC! – Scaling End-host Networking to 400Gbps & Beyond. 2024 Poster

SRC TECHCON. Venkat Kunaparaju, *Annus Zulfiqar*, Ali Imran, Gianni Antichi, Ben Pfaff, and Muhammad Shahbaz. GigaFlow: A Scalable and Efficient Hardware Fast-Path for Open vSwitch. 2024

INVITED TALKS, DEMOS, AND POSTERS

Gigaflow: Pipeline-Aware Sub-Traversal Caching for Modern SmartNICs

- Intel Corporation, San Jose Oct 2025
- Google Networking Summit Oct 2025
- P4 Workshop Demo At Open Compute Project (OCP) [\[Link\]](#) Oct 2025
- ACE Center for Evolvable Computing — Demo at Annual Meeting [\[Link\]](#) Oct 2025
- P4 Developer Days Event [\[Link\]](#) Jun 2025
- NetSyn Lab, Princeton University Apr 2025
- IBM Thomas J. Watson Research Center Apr 2025
- Networked Systems Group (NSG), ETH Zurich Apr 2025
- ACM ASPLOS Conference Apr 2025
- Network Operations and Internet Security Lab, University of Chicago Mar 2025

- Systems Seminar, University of Michigan Mar 2025
- Politecnico di Milano Mar 2025
- ACE Center for Evolvable Computing — Demo at Annual Meeting [\[Link\]](#) Oct 2024
- ACE Center for Evolvable Computing — Demo at Spring Meeting [\[Link\]](#) Mar 2024

Homunculus: Auto-Generating Efficient Data-Plane ML Pipelines for Datacenter Networks

- ACE Center for Evolvable Computing [\[Link\]](#) Jul 2023
- PurNet Seminar, Purdue University Sep 2023

The Slow Path Needs an Accelerator Too!

- VMware Research Group (VRG) Aug 2022

OPEN SOURCE PROJECTS AND CONTRIBUTIONS

- **aisuite** [★ 13.7K, 📄 1.4K] <https://github.com/andrewyng/aisuite>
- **AI-ForestWatch** [★ 31, 📄 8] <https://github.com/annusgit/ForestCoverChange>
- **Taurus In-Network ML** <https://gitlab.com/dataplane-ai/tutorials/sigcomm22>
- **Gigaflow vSwitch** <https://github.com/gigaflow-vswitch>
- **Homunculus Framework** <https://gitlab.com/dataplane-ai/homunculus/artifact-asplos23>

EXPERIENCE

Next-Generation Architectures Lab, University of Michigan Ann Arbor, MI
 Graduate Student Research Assistant Jan 2025 - Present

Advisor: Muhammad Shahbaz

Rearchitecting the end host network for the terabits per second era

VMware Research Group Palo Alto, CA
 Research Intern May - Aug 2022

Mentor: Ben Pfaff

Characterized the Open vSwitch slow path performance bottlenecks and proposed an accelerator for the SDN slow path

Next-Generation Architectures Lab, Purdue University West Lafayette, IN
 Research Assistant Aug 2021 - Dec 2024

Advisor: Muhammad Shahbaz

Explored architectures for the slow path in SDN; Built a Neural Architecture Search (NAS) framework, Homunculus, for data plane ML

Pervasive Parallelism Laboratory, Stanford University Stanford, CA
 Remote Researcher Sep 2020 - Jan 2021

Mentor: Muhammad Shahbaz

Designed discrete-event network simulations for data center load balancing algorithms

Center for Advanced Research in Engineering Islamabad, PK
 Design Engineer Jun 2019 - Jul 2021

Designed Ethernet/WiFi/LTE/BLE-capable, PoE-enabled, IoT Sensor Networks for industrial machine sensing and telemetry

Technical University of Kaiserslautern Kaiserslautern, DE
 Research Intern Jun - Sep 2018

Advisors: Norbert Wehn, Christian Weis

Worked on multi-temporal forest cover change detection to analyze the largest afforestation drive in Pakistan using remote sensing and deep learning

TUKL Lab, NUST Islamabad, PK
 Research Intern Jun 2017 - May 2019

Advisors: Faisal Shafait, Muhammad Shahzad

Worked on document processing and land cover classification problems using object detection and sequence learning techniques from deep learning

PROFESSIONAL SERVICE

- MemNetAI Workshop @ **SIGCOMM** – Program Committee (PC) member 2026
- **NSDI** – Program Committee (PC) Member for Artifact Evaluation (AE) 2026
- YArch (co-located with **ISCA**) – Program Committee (PC) member 2026
- University of Michigan – Ph.D. Admissions Committee volunteer reviewer 2025
- **SIGCOMM** – Program Committee (PC) Member for Artifact Evaluation (AE) 2025
- Google Summer of Code (GSoC) – Mentor, P4 Language Consortium 2025
- University of Michigan – Ph.D. Admissions Committee volunteer reviewer 2024

MENTORING EXPERIENCE

- Advay Singh, undergrad at University of Michigan – Cloud Infrastructure 2025 - Present
- Murayyiam Parvez, Ph.D. student at Purdue University – ML for Systems 2024 - Present
- Ali Imran, Ph.D. student at University of Michigan – SmartNICs, ML Systems 2024 - Present
- Venkat Kunaparaju, undergrad at Purdue University – Cloud Infrastructure 2023 - 2024

TEACHING EXPERIENCE

SIGCOMM. Tushar Swamy, *Annus Zulfiqar*, Alex Rucker, Muhammad Shahbaz, Kunle Olukotun. [Link/Code](#)
In-Network Machine Learning using Taurus. 2022

Purdue University. CS 38100 – Introduction to the Analysis of Algorithms (GTA). Fall 2023

CERTIFICATIONS

Intel Connectivity Academy – Level 1A/B: Tofino Native Architecture (TNA) & P4 [Link](#)

REFERENCES

1. **Muhammad Shahbaz**
Assistant Professor of Computer Science and Engineering (CSE) msbaz@umich.edu
University of Michigan
2. **Gianni Antichi**
Associate Professor of Computer Science gianni.antichi@polimi.it
Politecnico di Milano
3. **Ben Pfaff**
Chief Engineer/Co-Founder blp@cs.stanford.edu
Feldera