Sana Mahmood

smahmo12@jhu.edu — +1 (781) 628-8614 — LinkedIn: smahmood12 — Github: sana-mahmoood

RESEARCH SUMMARY

My research focuses on improving the performance, scalability, and programmability of modern datacenter networks, particularly those used in AI and cloud infrastructures. My current work addresses emerging challenges in AI datacenters, including mitigating bandwidth inefficiencies in large-scale training workloads and ensuring scalable, high-throughput collective communication. I have also worked on enhancing RDMA flexibility to enable advanced performance-optimizing techniques such as adaptive load balancing and congestion mitigation, and on alleviating burstiness in datacenter traffic to improve stability and utilization. Broadly, my interests lie in building practical, architecture-aware network mechanisms that deliver efficiency and reliability at hyperscale.

EDUCATION

Johns Hopkins University, Aug 2021 — May 2027
Doctor of Philosophy (PhD) in Computer Science

National University of Computer & Emerging Sciences, Aug 2014 — May 2018
Bachelor of Science: Computer Science (CUM LAUDE)

SKILLS

- **Computer Networks:** RDMA (RoCEv2), Datacenter Networks, Network Topologies (Clos, Fat-Tree, Spine-Leaf), Load Balancing, Networking Protocols (TCP/IP, BGP, OSPF), Routing, Traffic Engineering, SDN (Software-Defined Networking)
- **Network Tools:** Network Simulation (ns-3, OMNeT++), SDN Controllers (OpenFlow, RYU), Traffic Analysis (Wireshark), Kubernetes.
- **Programming:** C, C++, Go, C#, Java, Python, OCAML, MySQL, Shell Scripting, Kernel Programming.

RESEARCH EXPERIENCE

Bandwidth Efficient Collective Communications for AI Datacenters Aug 2025 — Present
Advisor: Dr. Soudeh Ghorbani (Johns Hopkins University)

- Demonstrated the inefficiency of unicast-based collective communications in AI datacenters, highlighting its impact on bandwidth utilization and training time.
- Proposed multicast transmission as a practical and efficient alternative for collective operations, challenging conventional assumptions about its feasibility.
- Designed a scalable substrate that leverages the AI datacenter topology to enable multicast-based collective communication, leading to improved bandwidth utilization and reduced training time.

Making RDMA Datacenter networks flexible through ordering support Oct 2022 — May 2025
Advisor: Dr. Soudeh Ghorbani (Johns Hopkins University)

- Enhanced the performance of RDMA-enabled datacenter networks, removing in-order delivery constraints to support advanced techniques like fine-grained load balancing, irregular topologies, failure recovery, and incast mitigation.
- Designed an on-NIC ordering layer that dynamically allocates memory for handling packet reordering, and validated design feasibility through FPGA implementation.
- Validated the design through ns-3 simulations, demonstrating a 91% performance improvement compared to the state of the art design.

Property Decomposition for Network Verification Jun 2024 — Present
Advisor: Dr. Soudeh Ghorbani (Johns Hopkins University) and Dr. Krishan Sabnani (Johns Hopkins University)

- Designed a tool to decompose network specifications into safety and liveness properties using an automata-theoretic approach.
- Enabled verification of safety and liveness properties, unlike existing verifiers focusing on only one property type.

Handling Burstiness in Datacenter Networks Jan 2022 — Oct 2022
Advisor: Dr. Soudeh Ghorbani (Johns Hopkins University)

- Analyzed micro-bursts in datacenter networks, and identified traffic synchronization as a contributing factor.
- Validated through simulations (ns-3 and OMNET++) that breaking traffic synchronization reduces bursts and improves performance by 65%, compared to existing techniques.

Resilient and Efficient Load Balancing for Datacenter Networks Jun 2019 — Jun 2021
Advisor: Dr. Ihsan Ayub Qazi (Lahore University of Management Sciences, LUMS) and Dr. Fahad Dogar (Tufts University)

- Developed an SDN-based load balancer with adaptive rate limiting to distribute packets across micro-slices—a logical abstraction dividing a link into smaller, manageable segments.
- Implemented the design on a testbed using OpenFlow with the Ryu controller and in simulations using ns-3, achieving up to 54% improvement over existing load balancing techniques.

Internet Traffic Classification Aug 2017 — Jun 2018
Final year project of BSCS

- Designed and implemented an in-network traffic classifier using supervised machine learning, enabling real-time separation of traffic for 10 different applications.

INDUSTRY EXPERIENCE

ByteDance Ltd. San Jose, USA
Software Engineering Intern (SDN) May 2025 — Aug 2025

- Designed and deployed a measurement framework for characterizing anycast traffic across 40+ datacenters, addressing deployment challenges in Docker and Kubernetes environments.
- Devised and implemented an optimization algorithm to improve anycast routing performance based on empirical data.
- Gained experience deploying networked systems at hyperscale across globally distributed datacenters.

i2c Inc Lahore, Pakistan
Software Engineer Jun 2018 — Jun 2019

- Designed and implemented a runtime fraud detection engine for the card service industry using machine learning.
- Collaborated with a team of 15+ to drive the successful implementation of the product.

Techlogix Lahore, Pakistan
Software Engineering Intern Jun 2017 — Aug 2017

- Designed and implemented a data visualization android application using d-3 library.

PUBLICATIONS

Orderly Management of Packets in RDMA by Eunomia
Sana Mahmood, Jinqi Lu and Soudeh Ghorbani. APNet, 2025
Efficient Datacenter Load Balancing with Microslices
Ramzah Rehman*, Sana Mahmood*, M. Irteza, Fahad Dogar and Ihsan Ayub Qazi. IFIP, 2025
Burstiness in data center topologies
Sougol Gheissi, Sana Mahmood, and Soudeh Ghorbani. CoNEXT-SW, 2022

HONORS AND AWARDS

Jun Wu and Yan Zhang Fellowship
Awarded a fellowship by the Department of Computer Science at Johns Hopkins, for which I was nominated by the department. This fellowship provided partial funding for the first year of my PhD 2021

LEADERSHIP ROLES AND COMMUNITY ENGAGEMENTS

Student Council Chair (CS Department)
Johns Hopkins University 2022 — Present

- Lead a 25+ member student council for the Computer Science PhD community at Johns Hopkins University, advocating for student interests and organizing academic and social initiatives.

Member of GRACE (Organization for Women in CS and ECE)
Johns Hopkins University 2023 — 2024

- Organized academic and social events for GRACE, fostering a supportive and connected community for women in Computer Science and ECE.