Evan Ram

evanram.com | evan@evanram.com | github.com/ioqr | linkedin.com/in/evanram

Research Interests

My research focuses on systems for real-time communication at scale. I build systems that make multi-party video conferencing scalable, measurable, and trustworthy, from benchmarking frameworks that drive large emulated conferences against production media servers (SFUs) to media servers that continuously screen every participant in a live call for deepfakes. My work emphasizes building real systems end to end, releasing them as open-source artifacts, and evaluating them under realistic, large-scale workloads.

Education

University of Colorado Boulder, Boulder, CO
M.S. in Computer Science (GPA: 3.739) — May 2026
Research Advisor: Professor Sangtae Ha
B.S. in Computer Science (GPA: 3.346) — May 2021

Peer-Reviewed Publications

[C1] **Evan Ram**[*], Tuan Tran[*], Youngwook Son, and Sangtae Ha. RTCBench: A Scalable Benchmarking Framework for Multi-Party Video Conferencing. In *Proceedings of the 34th ACM International Conference on Multimedia (MM)*, November 2026. ([*]Equal contribution.) Full paper; acceptance rate: 23.9% (1,703/7,103). Code: github.com/rtcbench/rtcbench

[C2] Tuan Tran, Shirin Ebadi, S. M. H. Hosseini, Woongsub Shin, **Evan Ram**, Youngwook Son, Seyeon Kim, Nam Bui, Kyunghan Lee, Eric Keller, and Sangtae Ha. DeepSFU: Scalable Deepfake Detection for Video Conferencing. In *Proceedings of the Annual Conference of the ACM Special Interest Group on Data Communication (SIGCOMM)*, August 2026. Full paper; acceptance rate: 21.2% (109/513).

Research Experience

Internet Systems Lab, University of Colorado Boulder, Boulder, CO
Researcher — March 2024–Present

- **RTCBench** [C1]: Designed and built an open-source, distributed framework that benchmarks multi-party video conferencing systems (Jitsi Meet, Janus) at scale using lightweight Go clients; reverse-engineered Jitsi's signaling protocol from Chrome HAR captures.
- **DeepSFU** [C2]: Contributed to a system for scalable, real-time deepfake detection in video conferencing that cuts per-frame processing latency by 143.7× and raises conferencing capacity by 26.2× (Java, C++, Go).
- **eXpressSFU** (USENIX NSDI '26): Contributed to a SmartNIC-accelerated media server (SFU) for large-scale video conferencing; acknowledged in the paper.
- **Other projects**: Simulated reinforcement-learning schedulers for GPU workloads on 1,000+ node Kubernetes clusters (Go, Python); built distributed EEG feature-extraction and model-training pipelines (Python).

Open-Source Software

- **RTCBench** (github.com/rtcbench/rtcbench): Benchmarking framework for multi-party video conferencing [C1]. Lead developer and maintainer.
- **Labpane** (labpane.com): React UI libraries, `@labpane/ui` (page layouts, forms, dialogs, menus, tables, and a terminal) and `@labpane/os` (tabs and resizable split panes); used in RTCBench's web interface. MIT license.

Industry Experience

Lumen Technologies, Boulder, CO
Software Engineer (Contract) June 2022–June 2023

- Led development of an edge monitoring product that tracks build status across services (Go, Nginx, Redis).
- Built Groovy/Java plugins and automated SQL reporting for a bare-metal server orchestration platform.

Amazon, Seattle, WA
Software Development Engineer, Amazon Advertising July 2021–February 2022
Software Development Engineer Intern May 2020–August 2020

- Built and operated Kindle ad-ingestion systems (Java, Go).
- Built ETL pipelines and ad-reporting workflows feeding a petabyte-scale Amazon Redshift data warehouse (Go, Python).

Startup internships, Chicago, IL
Software Engineering Intern at three startups, including Speakeasy and Detective.io 2016–2018

Teaching

University of Colorado Boulder, Boulder, CO
Course Manager, CSCI 4273/5273: Network Systems Fall 2024

Skills

Systems: WebRTC, SFUs (Jitsi Meet, Janus), SmartNICs, GPUs, Kubernetes, Docker, AWS
Languages: Go, C++, Java, Python, SQL, Bash