Peiqing Chen

+1 857-318-8737 | pqchen99@umd.edu | College Park, MD

Professional Summary

- **Interested in system research:** Led multiple research projects in system and networking and published works in top conferences including ASPLOS, SIGMOD, IMC, NSDI, and ICDE.
- **Specialties:** System and infrastructure, Computer Networks, Databases, Parallel Computing, Robotics, Operating systems.
- **Technical Proficiency:** C++, Python, RTC, system programming, ROS.

Education

University of Maryland College Park — College Park, MD
Doctor of Philosophy in Computer Science — *Aug 2023 - Dec 2026 (Expected)*
Advisor: Zaoxing (Alan) Liu https://zaoxing.github.io/

Boston University — Boston, MA
Doctor of Philosophy in Computer Engineering — *Aug 2021 - Jul 2023*
Advisor: Zaoxing (Alan) Liu

Peking University — Beijing, China
Bachelor of Science in Data Science and Big Data Technology (GPA: 3.8, Top 7%) — *Sep 2017 - July 2021*

Work Experience

Software Engineer Intern — **Meta**
Team: Real Time Communication — *May. 2025 - Aug. 2025*

- Led a project to improve the Quality of Experience (QoE) of GenAI call by lowering AI-interrupt-human rate and AI-response latency.
- Designed a personalized semantic endpointer (C++) that can tailor the AI's response speed to the user's speech rate. The GenAI bot responds more quickly for fast speakers and waits longer for slow speakers.
- Applied to 95% of public users' GenAI calling traffic in Facebook, Instagram, and Meta AI. Improved the interrupt rate by 17% and response latency by 350ms.

Research Interests

Telemetry over Real Time Communication (RTC) calling apps
Collaborator: Woo Xie (Meta)

- Study protocol discrepancy issues in common RTC applications to support unified RTC framework.
- Improve RTC call quality via measurement and analysis of relay server infrastructure.

Packet-level network telemetry using sketches
Advisor: Zaoxing Liu (University of Maryland)

- Design probabilistic algorithms (*i.e.* sketches) with low memory overhead and rigorous error bound.
- Support telemetry (packet count, latency distribution, heavy hitters, etc.) in high-speed data stream.
- Utilize sketches for monitoring and streaming systems, improving both throughput and latency.

Professional experience & research work

RTC calling app relay infrastructure study

- The RTC call reliability, quality (latency), and scalability depend highly on the relay server infrastructure.
- Measure the relay server distribution, reachability, and fall-back strategy of 5 RTC calling apps (Zoom, Discord, Messenger, Telegram, and WhatsApp) by packet-level measurement using WireShark.
- Building foundations for the relay infra on the server end and providing insights for GenAI calling infrastructure.

RTC protocol measurement

- A high cost in both engineering and cross-application communication is aroused by the discrepancy between standard RTC protocols (TURN, STUN, RTP, RTCP, QUIC) and their real-world implementations.

- Used WireShark to conduct packet-level measurements over 5 popular RTC app calls (Zoom, Facetime, WhatsApp, Discord, Messenger).
- Building foundations for next-generation unified RTC standards and cross-RTC application communication.

Precise runtime error estimation for sketch-based network telemetry

- The performance of network telemetry is degraded due to loose error bounds in SOTA sketch algorithms.
- Designed an online error estimator for 3 popular sketch algorithms, providing near-optimal runtime error estimation with theoretical analysis support, improving algorithm performance.
- Implemented our error estimator on a sketch-based caching system and demonstrated resource efficiency by demanding less cache memory over-provisioning.

Accelerating the mapping system for Unmanned Aerial Vehicle (UAV) navigation

- The performance of UAV autonomous navigation is limited by the slow updates and queries in mapping systems.
- Developed a parallelized software cache layer (C++) to accelerate updates and queries in the mapping system.
- Implemented the cached-mapping system on a ROS-based UAV autonomous navigation simulation platform and achieved up to 45% saving in end-to-end workflow runtime, raising flight speed by 26%.

Enabling real-time, continuous network telemetry over multiple cores

- The throughput of sketch-based software network telemetry systems is bottlenecked by single-core computation.
- Designed a parallel computing framework that scales sketch algorithms to multiple cores for higher throughput.
- Applied the framework to nine sketches on three software platforms (CPU, DPDK, and EBPF XDP) and raised throughput by over 7.2x.

Publications

[1] Peiqing Chen, Peng Qiu, Guowu Xie, Zaoxing Liu
Protocol Compliance in Popular RTC Applications
ACM Internet Measurement Conference (IMC 2025)

[2] Peiqing Chen, Minghao Li, Zishen Wan, Yushun Hsiao, Minlan Yu, Vijay Janapa Reddi, Zaoxing Liu
OctoCache: Accelerating Mapping System for MAV Navigation
ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS 2025)

[3] Peiqing Chen, Yuhan Wu*, Tong Yang, Junchen Jiang, Zaoxing Liu
Precise error estimation for sketch-based flow measurement
ACM Internet Measurement Conference (IMC 2021)

[4] Peiqing Chen, Dong Chen, Lingxiao Zheng, Jizhou Li, Tong Yang
Out of Many We are One: Measuring Item Batch with Clock-Sketch
ACM Special Interest Group on Management of Data (SIGMOD 2021)

[5] Yinda Zhang, **Peiqing Chen**, Zaoxing Liu
OctoSketch: Enabling Real-Time, Continuous Network Monitoring over Multiple Cores
USENIX Symposium on Networked Systems Design and Implementation (NSDI 2024)

[6] Yuhan Wu, Shiqi Jiang, Yifei Xu, Siyuan Dong, Kaicheng Yang, **Peiqing Chen**, Tong Yang
Unbiased Real-time Traffic Sketching
IEEE Transactions on Network Science and Engineering (TNSE 2023)

[7] Yuhan Wu, Hanbo Wu, Chengjun Jia, Bo Peng, Ziyun Zhang, Tong Yang, **Peiqing Chen**, Kaicheng Yang
Scalable Overspeed Item Detection in Streams
IEEE International Conference on Data Engineering (ICDE 2024)

[8] Yuhan Wu, Aomufei Yuan, Zhuoran Shi, Yuanpeng Li, Yikai Zhao, **Peiqing Chen**, Tong Yang, Bin Cui
Online Detection of Outstanding Quantiles with QuantileFilter
IEEE International Conference on Data Engineering (ICDE 2024)

[9] Siyuan Dong, Zhuochen Fan, Tianyu Bai, Tong Yang, Hanyu Xue, **Peiqing Chen**, Yuhan Wu
M4: A Framework for Per-Flow Quantile Estimation
IEEE International Conference on Data Engineering (ICDE 2024)

[10] Jonatan Langlet, **Peiqing Chen**, Zaoxing Liu, Ran Ben Basat, Michael Mitzenmacher, Gianni Antichi
Per-Flow Quantile Estimation using M4 Framework
IEEE Transactions on Knowledge and Data Engineering (TKDE 2025)

[11] Zhuochen Fan, Yalun Cai, Siyuan Dong, Qiuheng Yin, Tianyu Bai, Hanyu Xue, **Peiqing Chen**, Yuhan Wu, Tong Yang, Bin Cui
Spatiotemporal Sketch Disaggregation: Streaming Analytics with Heterogeneous Resources
IEEE International Conference on Data Engineering (ICDE 2026)