Jiaxin Lei

CONTACT INFORMATION

Department of Computer Science and Engineering
The University of Texas at Arlington
655 W Mitchell St
Arlington, TX, 76010
https://jiaxinraylei.github.io

Phone: (607) 232-1934
Email: jiaxin.lei2@uta.edu

RESEARCH INTERESTS

Cloud Computing, Computer Networks, Reconfigurable Hardware, Systems for Machine Learning, Datacenter Infrastructure.

EDUCATION

The University of Texas at Arlington (transfer) — Aug. 2023 - Present
Ph.D. Candidate in Computer Science — Arlington, TX
Advisor: Professor Hui Lu

State University of New York at Binghamton — Jan. 2019 - Aug. 2023
Ph.D. Candidate in Computer Science — Binghamton, NY
Advisor: Professor Hui Lu

State University of New York at Binghamton — Sep. 2017 - Dec. 2018
M.S. in Computer Science — Binghamton, NY

Beijing University of Posts and Telecommunications — Sep. 2013 - May. 2017
B.E. in Telecommunications Engineering with Management — Beijing, China

PUBLICATIONS

Jiaxin Lei, Manish Munikar, Hui Lu, Jia Rao, "SmartNIC-assisted Network Packet Zero-Copying", *In Submission*, 2023.

Manish Munikar, **Jiaxin Lei**, Hui Lu, Jia Rao, "ECON: Expedited Container Overlay Network", *Under Review (OSDI '24)*, 2023.

Jiaxin Lei, Manish Munikar, Hui Lu, Jia Rao, "Accelerating Packet Processing in Container Overlay Networks via Packet-level Parallelism", In *37th IEEE International Parallel and Distributed Processing Symposium (IPDPS '23)*, St. Petersburg, FL, USA.

Manish Munikar, **Jiaxin Lei**, Hui Lu, Jia Rao, "PRISM: Streamlined Packet Processing for Containers with Flow Prioritization", In *42nd IEEE International Conference on Distributed Computing Systems (ICDCS '22)*, Bologna, Italy.

Jiaxin Lei, Manish Munikar, Kun Suo, Hui Lu, Jia Rao, "Parallelizing packet processing in container overlay networks", In *16th ACM European Conference on Computer Systems (EuroSys '21)*, Virtual.

Yu Sun, **Jiaxin Lei**, Seunghee Shin, Hui Lu, "Baoverlay: a block-accessible overlay file system for fast and efficient container storage", In *11th ACM Symposium on Cloud Computing (SoCC '20)*, Virtual.

Jiaxin Lei, Kun Suo, Hui Lu, Jia Rao, "Tackling parallelization challenges of kernel network stack for container overlay networks", In *11th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud '19)*, Renton, WA, USA.

Research Experience

SmartNIC-assisted Kernel-User Network Packet Zero-copying

Kernel-user data copying contributes one of the largest overheads in the network packet processing. I introduced a SmartNIC-based zero-copying framework utilizing Nvidia Bluefield-2. This system splits packet headers and payloads, storing them in separate memory regions. The header is processed through kernel network stack, while the payload is 'zero-copied' directly to user applications. By eliminating the copying overhead, we can significantly improve network performance.

Expedited Container Overlay Networks

Container overlay networks, compared to bare-metal host networks, experience performance degradation due to multiple asynchronous stages. Our development, ECON, streamlines this by having all packets in a flow follow a consistent path, bypassing non-critical processing stages after recognizing a flow's ultimate destination. This approach has demonstrated improvements in container throughput by up to 121%, a reduction in average latency by up to 61%.

Streamlining Packet Processing with Flow Prioritization

In highly utilized systems, short-lived, latency-sensitive network flows can suffer from extended queuing delays due to the kernel network stack's inability to differentiate packet performance needs. We proposed PRISM – an innovative in-kernel method, which identifies the packets' priority early on and establishes a distinct, streamlined processing path for high-priority flows. As a result, we reduce latency for these flows by over 50% in heavily loaded systems.

Parallelizing Packet Processing at Packet-level

Even with state-of-the-art solutions like FALCON (device-level flow pipelining), a saturated single CPU core can still bottleneck a network flow. To address this, I developed MFLOW – a novel in-kernel packet-level steering technique. MFLOW evenly distributes a high-demand flow into several lightweight micro-flows across multiple CPU cores for parallel processing. Micro-benchmark tests reveal throughput improvements by 81% for TCP and 139% for UDP.

Pipelining Packet Processing at Device-level in Overlay Networks

Container overlay networks often underperform compared to host networks due to prolonged data processing paths. In response, I designed FALCON – a fast and balanced in-kernel solution that distributes the excessive software interrupts associated with various network devices onto multiple CPU cores for processing. Falcon enhances throughput by 300% for web serving applications and cuts tail latency by 53% for data caching applications.

Block-accessible Overlay File System for Container Storage

Overlay file systems often face extended write latencies as write operations on read-only container images require copying to a separate writable layer. Our solution, BAOVERLAY, is a block-accessible overlay file system that partitions files into fine-grained blocks, optimizing Copy-on-Write operations to involve only certain blocks instead of entire files. Tested with Linux Ext4 as the backing file system, BAOVERLAY markedly improved applications' I/O performance with up to 32x faster performance for 1KB files and 64x faster for 4MB files.

Talks

"Accelerating Packet Processing in Container Overlay Networks
via Packet-level Parallelism"

— In *IPDPS '23*, St. Petersburg, FL, USA — May. 2023

"Accelerating Packet Processing in Container Overlay Networks"
— In *SUNY Binghamton Computer Science Department Seminar*
Binghamton, NY, USA — Nov. 2022

"Parallelizing Packet Processing in Container Overlay Networks"
— In *EuroSys '21*, Virtual — Apr. 2021

"Tackling parallelization challenges of kernel network stack for container overlay networks"
— In *HotCloud '19*, Renton, WA, USA — Jul. 2019

Teaching Experience

Instructor – Lab Session and Partial Lectures — Fall 2023
CSE 3320 Operating Systems (Undergraduate)
The University of Texas at Arlington — Arlington, TX

Instructor – Lab Session and Partial Lectures — Spring 2023
CS 350 Operating Systems (Undergraduate)
State University of New York at Binghamton — Binghamton, NY

Instructor – Lab Session — Fall 2022
CS 350 Operating Systems (Undergraduate)
State University of New York at Binghamton — Binghamton, NY

Instructor – Lab Session — Spring 2022
CS 350 Operating Systems (Undergraduate)
State University of New York at Binghamton — Binghamton, NY

Teaching Assistant — Fall 2021
CS 452/552 Introduction to Cloud Computing (Undergraduate/Graduate)
State University of New York at Binghamton — Binghamton, NY

Instructor – Lab Session — Spring 2021
CS 350 Operating Systems (Undergraduate)
State University of New York at Binghamton — Binghamton, NY

Teaching Assistant — Fall 2020
CS 452/580 Introduction to Cloud Computing (Undergraduate/Graduate)
State University of New York at Binghamton — Binghamton, NY

Teaching Assistant — Spring 2020
CS 480/580 Advanced Topics in Cloud Computing (Undergraduate/Graduate)
State University of New York at Binghamton — Binghamton, NY

Instructor – Lab Session — Fall 2019
CS 350 Operating Systems (Undergraduate)
State University of New York at Binghamton — Binghamton, NY

Teaching Assistant — Spring 2019

CS 550 Operating Systems (Graduate)
State University of New York at Binghamton — Binghamton, NY

Grant	Student Travel Grant for attending *IPDPS '23*	2023
	Student Travel Grant for attending *NSDI '23*	2023
	Student Travel Grant for attending *seL4Summit '19*	2019
	Student Travel Grant for attending *VEE '19*	2019
	Student Travel Grant for attending *ATC '19*	2019