

# Zhewen Lyu

346-542-0735 | [zhewenlyu@yahoo.com](mailto:zhewenlyu@yahoo.com) | [linkedin.com/in/zhewen-lyu](https://www.linkedin.com/in/zhewen-lyu) | [github.com/ZavierLyu](https://github.com/ZavierLyu)

## EDUCATION

---

### Rice University

*M.S. in Computer Science — GPA:3.9/4.0*

Houston, US

*Expected Dec. 2023*

### The Chinese University of Hong Kong, Shenzhen

*B.E. in Computer Science and Engineering — GPA:3.59/4.00 — Dean's List*

Shenzhen, CN

*Jul. 2021*

## EXPERIENCE

---

### ByteDance Inc.

*Software Engineering Internship*

May 2023 – Present

*Bellevue, WA*

- Designed and orchestrated a **FaaS@Edge** (Function as a Service at the Edge) architecture on ByteDance's **CDN** (Content Delivery Network) service, allowing customers to customize the content delivered by the CDN and execute functions on edge servers closer to the viewer, reducing the response latency and origin server load
- Developed an **Nginx** module with a subrequest handler to trigger executions and an output filter to normalize responses, coordinating HTTP requests between the viewer, JavaScript engine, cache, and origin server
- Built a set of control plane including an internal reverse proxy for validating FaaS-related requests, and a central administration microservice in **Golang** for synchronizing function codes and configurations stored on **etcd** with the JavaScript engine on the edge

### Pinduoduo Inc.

*Full-time Software Development Engineer*

Jul. 2021 – Jun. 2022

*Shanghai, China*

- Engineered and maintained a high-throughput data transfer system for Pinduoduo's search engine, incorporating a comprehensive control plane website powered by **Java Spring Boot** and **MySQL**, encompassing monitoring, logging, dynamic configuration, and throttling functionalities
- Implemented automation and load balancing of **Kafka** consumers and producers using a consistent hash algorithm, leading to a significant **25%** reduction in distribution skewness
- Developed a monitoring system utilizing CAT (Central Application Tracking) and InfluxDB to effectively track and monitor memory usage and CPU usage of Kafka client nodes, leveraging statistics obtained from Kubernetes
- Achieved exceptional data pipelining performances by leveraging industry-leading big data frameworks such as **Kafka**, **Hive**, **HDFS**, **HBase**, and **Flink**, resulting in a remarkable throughput of over **100 MB/s** streaming data and **100 TB/day** offline data

## PROJECTS

---

### Operating System Implementation | *Kernel, Virtual Memory, File System*

Mar. 2023 – May 2023

- Developed a **Unix** style operating system kernel on simulated computer system hardware, encompassing context switching, virtual memory management, multi-programming, process control and teletypewriter functionalities
- Implemented kernel calls including fork, exec, wait, delay and terminal I/O based on trap handling mechanism
- Constructed a Unix style **file system**, enabling multiple users to perform read, write, mkdir, chdir, rmdir and hard link operations within a hierarchical directory through either absolute or relative paths

### Database Implementation | *B+ Tree, Buffer Manager, SQL Execution*

Mar. 2023 – May 2023

- Built a relational database supporting GB-scale storage and SQL queries including select, aggregation and join
- Developed an LRU buffer manager to enhance memory and file I/O operations, a **B+ tree** table organizer for efficient range-based queries, a syntax checker, and an SQL execution engine with an **exhaustive query optimization** algorithm, utilizing memoization and costs estimation, to accelerate complex queries processing
- Practiced standard **C++** OOP design patterns including inheritance, template, iterator, and abstract factory

### Social Media Web Application | *MongoDB, Express, React, Node.js*

Sep. 2022 – Nov. 2022

- Developed a LinkedIn-like social media web application with features including feed, comment, registration, subscription, customizable user profile, and Google OAuth 2.0
- Created an engaging front-end application with a user-friendly UI and seamless interactions by **React** and **Redux**
- Established a backend service with **Node.js** and Express framework, implementing **RESTful** APIs for data persistence in **MongoDB** and image storage in Cloudinary, and successfully run on Heroku for one year

## TECHNICAL SKILLS

---

**Programming Languages:** C/C++, Java, Python, Golang, JavaScript, Verilog, R, SQL, HTML/CSS

**Frameworks:** Nginx, Linux, Spring Boot, Hadoop, React, Node.js, Flask, Gin, Express

**Tools:** Kafka, MySQL, MongoDB, InfluxDB, Flink, Dubbo, Docker, Scikit-learn, pandas