

KAHFI S. ZULKIFLI

✉ kwf3wv@virginia.edu [🌐 kahfizulkifli.github.io](https://kahfizulkifli.github.io)

EDUCATION

Bandung Institute of Technology (ITB), Indonesia
Bachelor of Science, Computer Science

April 2024

PUBLICATIONS

HEIMDALL: Accurate and Efficient I/O Admission Policy with Extensive Machine Learning Pipeline

Anonymous Authors. *Under submission, Manuscript available upon request*

EVStore: Storage and Caching Capabilities for Scaling Embedding Tables in Deep Recommendation Systems

Daniar H. Kurniawan, Ruipu Wang, [Kahfi S. Zulkifli](#), Fandi Wiranata, John Bent, Ymir Vigfusson, Haryadi S. Gunawi. *In the Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2023*

RESEARCH EXPERIENCE

International Research Collaboration on ML for Systems (HEIMDALL)

Sep 2022 – Apr 2024

- Collaborating with Prof. **Haryadi Gunawi** of The University of Chicago and Prof. Achmad Imam Kistijantoro of Bandung Institute of Technology
- Research on reducing tail latency in solid state drives (SSDs) with machine learning
- Designed **16 machine learning models** that has **accuracy between 70%-90%** with AutoML for predicting storage performance in **FEMU**, an SSD emulator
- Reduced the inference latency of machine learning models down to **10ns, 50x faster** than existing models
- Modified **Ceph**, an existing object-storage library with **machine learning models** that show **up to 40% improvement at p99 latency** and deployed **10 clusters**, each with **2 OSDs** across **20 nodes in FEMU**
- Supported as an undergraduate research assistant by the **MoECRT ACE Open Research** program

International Research Collaboration on Systems for ML (EVStore)

Aug 2021 – Aug 2022

- Collaborated with Prof. **Haryadi Gunawi** of The University of Chicago
- Research on optimizing model inference latency of deep recommendation systems (DRS)
- Designed **new caching algorithms** based on **groupability**, a novel property that measures the probability of an embedding value present in an inference request, which increased **perfect hit rate up to 30%**
- Helped **integrate EVStore system**, which **reduced average latency by 23%** and **p90 latency by 27%**, **increased the throughput by 4x** at **only 0.2% loss in accuracy**

PROJECTS

Multidisciplinary National Research Program (Indonesian Ministry of Education)

- Designed '*Bandung Smart Transportation*' for Bandung city so that it can accomplish one of the UN's Sustainable Development Goals (SDGs) on **Sustainable Cities and Communities**
- Proposed a **magnetic-levitation transport system** for Bandung City to **reduce carbon footprint** and **increase pedestrian mobility** within Bandung and between other cities surrounding Bandung

National Land Agency Business Intelligence Dashboard (Webgis Indonesia)

- This project aims to analyze huge land property dataset by utilizing a business intelligence dashboard
- Maintained the **Extract-Transform-Load (ETL) pipeline** reaching **hundreds of gigabytes** from **Apache Hive** to the **ELK Stack (Elastic Search, Logstash, and Kibana)** and analyzing the data in real time with **Superset**
- Helped integrate the whole architecture, from data warehouse to the business intelligence dashboard, which helped **2 national government agencies** to analyze huge amounts of data

National Geospatial Data Warehouse (Webgis Indonesia)

- Purpose of this project is to analyze **real-time climate and earthquake data** in Indonesia
- Maintained an **Extract-Transform-Load (ETL) pipeline** using **Open Talend Studio** and **PostgreSQL**
- Built a dashboard that analyzes climate data from three provinces using **MapLibre JavaScript**, currently being used by **multiple national government agencies**

Mosaik.id (Radya Laps Harapan Bangsa)

- Designed a Safe Internet Mobile App to promote Safe Internet usage for students in remote areas
- Developed the Backend Architecture, database, CI-CD and deployed the API on **Azure single-handedly**

Ventilator-Indonesia (Collaboration between top national universities and industries)

- Helped build low-cost ventilators for **Covid-19 patients** in **more than 50 hospitals across Indonesia**
- Spearheaded the **inventory management** of the ventilator production which consists of **more than 100 different types of electrical components**
- Collaborated with **professors, researchers** and **engineers** from **more than 10 universities, high-tech companies, local government agencies and NGOs**

HONORS AND AWARDS

Nominee of Multidisciplinary National Research Program Funding (Indonesian Ministry of Education)

- Program promotes multidisciplinary research to accomplish the UN's Sustainable Development Goals
- **191/20k proposals funded from top universities across Indonesia**

3rd Best Novice Speaker of Ganesha Open (Student English Forum of ITB)

- Ganesha Open is an annual British Parliamentary Open Debate Competition hosted by ITB
- Competed against **more than 20 teams from 10 top universities in Indonesia**

TEACHING ASSISTANT

Operating Systems, Networking

- Volunteered to become an academic tutor for midterm and end-of-term exams preparation

Introduction to Computing Laboratory

- Responsible for the grading and teaching of computational skills in Python

ORGANIZATION AND VOLUNTEERING

Student Support Association (ITB)

Personal Character Development Trainer

- This organization aims to support new students in adapting to the university's academic life
- Led a two-day **character development training session** for new students on being successful in campus

Student Cabinet (ITB)

Head of Research and Development of Qui

- Qui is a project dedicated to facilitate **multi-discipline collaboration** between students in the university to accomplish their goals in research, competitions, business, or social projects
- **Spearheaded the user research on 150 students** across different course programs regarding their ultimate goals while studying in the campus and **directed UI/UX testing** of the prototype QUI website

TECHNICAL SKILLS

Programming Languages Python, Java, C, C++, C#, Prolog, Kotlin, Javascript, Typescript, SQL

Systems Meta's Deep Learning Recommendation Model, Linux

Database MariaDB, Ceph

Data Warehouse Talend Open Studio, Logstash

DevOps Gitlab CI-CD

Cloud Computing Azure, Chameleon Cloud Testbed

Big Data Elastic Search, Kibana Dashboard