

Ray A. O. Sinurat (Ray Andrew)

rayandrew@uchicago.edu · rayandrew.me

John Crerar Library 283, 5730 S Ellis Ave, Chicago IL, USA 60637

Research Interests

Distributed & Storage Systems (improving reliability, scalability and performance)
Machine Learning for Systems (applying ML to improve scalability in the clusters)

Education

- 2021 – 2026 **University of Chicago** – Chicago, IL, USA
Ph.D. in Computer Science
Advisors: Professor Haryadi S. Gunawi and Sandeep S. Maddireddy
- 2015 – 2019 **Institut Teknologi Bandung** – Bandung, Indonesia
B.S. in Computer Science
Advisors: Achmad I. Kistijantoro, Ph.D. and Dr. Eng. Ayu Purwarianti

Employments

- 2021 – Present **Research Assistant at University of Chicago** – Chicago, IL, USA
- Current project: **Drift Mitigation and Storage Optimization** in collaboration with Argonne National Laboratory. Improving the usage optimization of storage systems using an ML-based approach and adapting to drift in production systems.
 - Past project: **LIBROS** (published in **IEEE Cloud '22**), an ecosystem of tail-latency mitigation with supports from library, runtime, and operating system layers. **LIBROS** is able to **improve multi-storage applications speed by 5-70%**, starting at 90th percentile
- Summer 2023 **Research Aide at Argonne National Laboratory** – Lemont, IL, USA
- Bootstrapped project with an aim to improve order robustness of continual learning in several datasets
 - Researched about continual learning and its usability in computer systems
- 2019 – 2021 **Remote Research Assistant at GIK Lab** – Bandung, Indonesia
- Remote mentorship program in collaboration with the Computer Systems group at **University of Chicago**
 - Studied system-related bugs, such as **scalability, distributed concurrency and cascading failure**, focusing specifically on **scalability bugs**
 - Researched how **Java Virtual Machines (JVMs) can share memory to reduce memory usage**, especially in a virtualized environment, using **Linux system calls** such as mmap and madvise
 - Implemented **predictive model for Garbage Collection (GC) Time using live and dead objects** from OpenJDK8 ParallelGC algorithm to **reduce tail latencies**
- 2019 – 2021 **CS Researcher at Emmerich Research Center** – Jakarta, Indonesia
- Implemented **Fungi Processing Automation Systems for Leather Production**, such as: **Automated Tending Machine** and **Contamination Detection**
 - Researched **Black Soldier Fly's lifecycle**, a popular biomass for alternative protein, using **Deep Learning** approach

- 2018 **Software Engineer Intern** at **Dekoruma** – Jakarta, Indonesia
- Developed **Mobile Web Marketplace**, such as **Product Details and After Order**, using **React JS and React Native Web**
 - Implemented company's **new React infrastructure** by developing **Server Side Rendering with Code Splitting Strategy** (accessible through NodeJS library **Centarius**)
 - Developed **novel modal implementation for React Native** (accessible through NodeJS library **Modal React Native Web**)

Publications

- Manuscript Ready Anonymous Author(s). **Concept Drift Detection with Large Language Models**. *In Submission*, 2023.
- Manuscript Ready Yuyang Huang*, **Ray A. O. Sinurat***, Nanqinqin Li, Mark Powers, Michael Sherman, Kate Keahey, Haryadi S. Gunawi. **STORREP: Storage Research Experiment Patterns on Chameleon Cloud and Trovi**. 2023.
- ML for Systems'22 **Ray A. O. Sinurat**, Anurag Daram, Haryadi S. Gunawi, Robert B. Ross, Sandeep Madireddy. **Towards Continually Learning Application Performance Models**. *Workshop on ML for Systems at NeurIPS*, 2022.
- IEEE CLOUD'22 Meng Wang, Cesar A. Stuardo, Daniar H. Kurniawan, **Ray A. O. Sinurat**, and Haryadi S. Gunawi. **Layered and Uniform Contention Mitigation Capabilities for Cloud Storage**. *In the Proceedings of the 15th IEEE International Conference On Cloud Computing*, 2022.
- UChicago TR'20 Daniar H. Kurniawan, Cesar A. Stuardo, **Ray A. O. Sinurat**, and Haryadi S. Gunawi. **Notification and Prediction of Heap Management Pauses in Managed Languages for Latency Stable Systems**. *In The University of Chicago Technical Report*, 2020.

Posters

- ML for Systems'22 **Ray A. O. Sinurat**, Anurag Daram, Haryadi S. Gunawi, Robert B. Ross, Sandeep Madireddy. **Towards Continually Learning Application Performance Models**. *Workshop on ML for Systems at NeurIPS*, 2022.

Projects

- CLUSTEROPTIM Detecting performance changes and optimize clusters utilization.
- LIBROS **[IEEE CLOUD'22]** Implementing Java GC predictor to give delay prediction that is then used as cancellation mechanisms for reducing tail-latencies.
- Bug Study Studying and analyzing scalability bugs in numerous distributed systems, such as Hadoop, HBase, Cassandra, ZooKeeper, Spark, HDFS, Flume, and Storm.
- Indonesian Image Captioning Preparing the first Indonesian dataset captions and implementing the first deep-learning based Indonesian automated image captioning using Semantic Compositional Networks in partnership with Prosa AI and Microsoft Indonesia.

Teaching Assistantship

- Fall 2023 CMSC 230: Operating Systems (University of Chicago)
- Spring 2023 CMSC 230: Operating Systems (University of Chicago)
- Fall 2022 CMSC 154: Introduction to Computer Systems (University of Chicago)
- Fall 2021 CMSC 230: Operating Systems (University of Chicago)
- 2018 IF 3140: Database Management (Institut Teknologi Bandung)
- 2017 IF 2240: Databases (Institut Teknologi Bandung)

*The authors contribute an equal amount of work and are sorted alphabetically based on their last names.

Student Mentorship

2023- Summer 2023	William Nixon (CS Undergrad at Institut Teknologi Bandung)
2022-2023	Jax Alemu (Wylie High School, Texas; DSI Summer Lab Research Assistant)
2021-2022	Kangrui Wang (Master of CS at University of Chicago)
	Nathanael Timothy (B.Eng. in Electrical Engineering at Universitas Pelita Harapan Jakarta)

Awards

2023	FAST '23 Travel Awards
2021	Crerar Fellowship (University of Chicago) <i>Identified as one of the strongest Ph.D. applicants.</i>

Skills

AI	PyTorch, Keras, Tensorflow
Testbed	Emulab, Chameleon Cloud
Systems	Hacking Cassandra, Hadoop, Kafka, HBase Using ZooKeeper, HDFS, Kafka, Cassandra, MongoDB, MapReduce
Runtime	Hacking JVM (Hotspot, Garbage Collection, JNI Agent)
OS	Hacking LINUX KERNEL
PL	C, C++, Python, Java, [Type/Java]script, C#, PHP, Bash
Cloud	Google Cloud, AWS, Microsoft Azure, Heroku, DigitalOcean
Database	MongoDB, MySQL, PostgreSQL, Google Firebase, RethinkDB, SQLite, Redis
Search Engine	Algolia, Meilisearch, RediSearch
IOT	Arduino, Raspberry Pi
IaC & PaaS	Hashicorp Terraform, Docker, Docker Compose

References

Haryadi S. Gunawi (Ph.D. Advisor) Associate Professor, Computer Science Department, University of Chicago
Sandeep Madireddy Assistant Computer Scientist, Mathematics and Computer Science Division, Argonne National Laboratory
Cesar A. Stuardo Network System Engineer, ByteDance