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)
Systems for ML (improving scalability and performance of ML pipeline)
HPC I/O (improving I/O scalability and performance on HPC systems)

Education

Sep 2021 – Jan 2027 — **University of Chicago** – Chicago, IL, USA
Ph.D. in Computer Science
Master's in Computer Science (GPA: 3.63)
Advisors: Haryadi S. Gunawi

Jul 2015 – Jul 2019 — **Institut Teknologi Bandung** – Bandung, Indonesia
B.S. in Computer Science
Advisors: Achmad I. Kistijantoro, Ph.D. and Dr. Eng. Ayu Purwarianti

Employments

Sep 2021 – Present — **Research Assistant** at **University of Chicago** – Chicago, IL, USA
- Current project: **Taming I/O Optimization for Deep Learning at Scale** (with ANL and LLNL). Enables systematic, low-cost I/O optimization for large-scale deep learning pipelines delivering **2× end-to-end performance improvements** compared to baselines.
- Current project: **Analysis Tools for TB-Scale Complex Application and I/O Traces** (with ANL and LLNL). Provides scalable analysis, search, and filtering over **terabyte-scale I/O and application traces**, uncovering performance bottlenecks in complex DL pipelines.
- Past project: **Heimdall** (**EuroSys '25**). An ML-driven storage admission system achieving **15–35% lower average I/O latency** than state-of-the-art solutions, deployable across Linux kernel and KV stores (RocksDB).
- Past project: **Drift Mitigation and Storage Optimization** (**MLFS'22** with ANL). Improves storage efficiency using **ML-based optimization** that **adapts to workload drift** in production systems.
- Past project: **LIBROS** (**IEEE Cloud '22**). A cross-layer tail-latency mitigation framework spanning library, runtime, and OS layers, delivering **5–70% speedups** for multi-storage applications starting at the 90th percentile.

Summer 2025 — **Computing Graduate Scholar Intern** at **Lawrence Livermore National Laboratory** – Livermore, CA, USA
- Reproduced 7 complex AI4Science I/O workloads on the GPU-free DLIO Benchmark used by MLCommons Storage
- Built a multi-layer I/O behavior comparator to analyze deep learning workloads against benchmarked versions
- Developed a gzip indexer enabling random reads 35× faster and index storage 1000× smaller than competitor
- Added advanced dataset generation features to the DLIO Benchmark, improving usability for complex I/O workloads
- Standardized multi-layer I/O tracing for deep learning, now integrated into the DLIO Benchmark
- Created a visual tracing dashboard that exposes I/O events in deep learning workloads for intuitive analysis

Summer 2024 — **W. J. Cody Associate** at **Argonne National Laboratory** – Lemont, IL, USA
- Developed optimizations for a Transformer-based weather forecasting pipeline within the ALCF Polaris ecosystem
- Successfully **reduced** the pipeline runtime **by 30%** from the original implementation by enhancing I/O bandwidth and designing a more efficient data access pattern

Summer 2023 — **Research Aide** at **Argonne National Laboratory** – Lemont, IL, USA
- Bootstrapped project with aim to improve order robustness of continual learning
- Researched about continual learning and its usability in computer systems

Jan 2019 – Aug 2021
Remote Research Assistant at **GIK Lab** – Bandung, Indonesia
- Remote mentorship program in collaboration with the 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** 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**

Jul 2019 – Jun 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

Summer 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**)

June 2016 – July 2019
Teaching Assistant at **Institut Teknologi Bandung** – Bandung, Indonesia
- **Set up projects** for multiple database-related undergraduate courses
- **Lead** the **Database Laboratory** at Institut Teknologi Bandung

Publications

SC'25
Väinö Hatanpää, Eugene Ku, Jason Stock, Murali Emani, Sam Foreman, Chunyong Jung, Sandeep Madireddy, Tung Nguyen, Varuni Sastry, **Ray A. O. Sinurat**, Huihuo Zheng, Sam Wheeler, Troy Arcomano, Venkatram Vishwanath, Rao Kotamarthi. **AERIS: Argonne Earth Systems Model for Reliable and Skillful Predictions**. *In The SC'25: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis.* ***Finalists of SuperComputing 25 ACM Gordon Bell Prize for Climate Modeling***.

EuroSys'25
Daniar Kurniawan, Maharani A. P. Irawan, Kahfi Zulkifli, **Ray A. O. Sinurat**, Hongzhen Liang, Peiran Qin, Sandeep Madireddy, Janki Bhimani, Achmad Imam Kistijantoro, Haryadi S. Gunawi. **Heimdall: Optimizing Storage I/O Admission with Extensive Machine Learning Pipeline**. The 20th ACM European Conference on Computer Systems. 2025.

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*, 2023.

Projects

TBScaleAnalytics
Tools to analyze TB-scale App and I/O traces

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

DL I/O Opt	Optimizing Deep Learning for Science pipeline on HPC environment
CacheCL	Improving the cache eviction model to prevent it from degrading due to system and changes.
ClusterOPTIM	Detecting performance changes and optimize clusters utilization.
HEIMDALL	**[ACM EuroSys'22]** ML-based storage admission control
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	Prepared the first Indonesian dataset captions and implemented the first deep-learning based Indonesian automated image captioning using Semantic Compositional Networks in partnership with Prosa AI and Microsoft Indonesia.

Teaching Assistantship

Win 25	ADMN 30100: SYAIR-2025 (Online Research Preparation Course at University of Chicago)
Win 24	CMSC 144: Systems Programming II (University of Chicago)
Aut (21, 23-25)	CMSC 230: Operating Systems (University of Chicago)
Spr 23	CMSC 332: Topics in Operating Systems (University of Chicago)
Aut 22	CMSC 154: Introduction to Computer Systems (University of Chicago)
2018	IF 3140: Database Management (Institut Teknologi Bandung)
2017	IF 2240: Databases (Institut Teknologi Bandung)

Student Mentorship

2023-2025	**William Nixon** (CS Undergrad at Institut Teknologi Bandung, CS PhD student at University of Chicago)
2024	**Richard Tjokroutomo** (CS Undergrad at The Chinese University of Hong Kong)
2024	**Joanna Cheng** (CS Undergrad at John Hopkins University)
2023	**Jax Alemu** (Wylie High School, Texas; DSI Lab Research Assistant, now Student at UW Madison)
2022-2023	**Kangrui Wang** (Master of CS at University of Chicago)
2021-2022	**Nathanael Timothy** (B.Eng. in Electrical Engineering at Universitas Pelita Harapan Jakarta)

Awards

2024	UU Fellowship
2023	FAST '23 Travel Awards
2021	**Crerar Fellowship** (University of Chicago) *Identified as one of the strongest Ph.D. applicants.*

Open Source Maintainer

DLIO Benchmark [LINK]
DFTracer [LINK]
PyDFTracer [LINK]
DFTracer Utils [LINK]

Skills

AI	PyTorch, Keras, Tensorflow
Testbed	Emulab, Chameleon Cloud, HPC Clusters
Systems	**Hacking** Cassandra, Hadoop, Kafka, HBase **Using** ZooKeeper, HDFS, Kafka, Cassandra, MongoDB, MapReduce, Docker, Kubernetes, Lustre Parallel Filesystems
Runtime	**Hacking** JVM (Hotspot, Garbage Collection, JNI Agent)
OS	**Hacking** Linux Kernel
PL	C, C++, Python, Java, [Type/Java]script, 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