

Sacheendra Talluri

Amsterdam | Dutch citizen | Willing to relocate or be remote sacheendra.t@gmail.com | [Website](#) | [GitHub](#) | [Google Scholar](#)

Experience

Ph.D. Candidate, Vrije Universiteit Amsterdam - Amsterdam, the Netherlands July 2019 – May 2025

- Implemented the world's fastest GPU-based JSON query engine. It achieves 2.9x speedup over other state-of-the-art engines [1].
- Evaluated scheduler APIs and architectures for distributed data processing. Expressive APIs lead to a runtime reduction of 24% [4, 5].
- Modeled failures, GPUs, data- and carbon-aware schedulers in the OpenDC simulator. Achieved 5.5-3.6% MAPE error [GitHub].
- Proposed a novel crowdsourced data collection technique to evaluate the reliability of cloud and LLM services [2, 3].
- Published work in top-tier (VLDB, TPDS) and high-quality (CCGRID, FGCS, ICPE) venues.
- Supervised 6 MSc and 6 BSc theses, resulting in 4 (2 high-quality) publications, one best thesis award, and a best thesis runner-up.

Research Visitor, IBM Research - Haifa, Israel July 2023 – September 2023

- Added dynamic dataflow analysis to the Go compiler. Support for all language constructs except maps [GitHub].
- Designed and implemented a system to associate Kubernetes errors with the configuration that caused the error.

DevOps Engineer, IPBurger - Remote October 2018 – May 2019

- Designed and implemented a highly available proxy management software.
- Decreased proxy connection time (10x) and improved stability using strategic caches and replicas.

Database Internals Intern, Databricks - Amsterdam, the Netherlands February 2018 – May 2018

- First published analysis of storage-access patterns of large data analytics clusters. Analyzed 600TB of Apache Spark logs [Paper].
- Implemented and simulated state-of-the-art (LHD, Hyperbolic, etc.) cache policies for Apache Spark.
- Published a result that simple cache policies perform as well as complex ones at large cache sizes [Paper].

Part-time Infrastructure Engineer, ReactiveSearch.io - Remote January 2014 – April 2016

- Designed and implemented a successful automated system to deploy an Elasticsearch cluster per customer.
- The deployed systems could handle loads of over 100k ops/second.
- Designed and implemented a stream processing system and libraries to access ReactiveSearch's streaming API.

Education

Vrije Universiteit Amsterdam Amsterdam, The Netherlands
Ph.D. in Computer Science (Distributed Systems) (Thesis Completed) Expected November 2025

Delft University of Technology Delft, The Netherlands
M.Sc. in Computer Science – GPA: 8.0 (10.0 Max) December 2018

Dhirubhai Ambani Institute of Information and Communication Technology Gandhinagar, India
B.Tech. in Information and Communication Technology – GPA: 8.6 (10.0 Max) May 2016

Other Experience

Secretary, SPEC (of SPEC CPU fame) Cloud research group November 2019 – present
Lead TA and Lab Coordinator, Distributed Systems at VU Amsterdam November 2019 – February 2025
Teaching Assistant, Storage Systems at VU Amsterdam September 2020 – November 2022
Treasurer, TU Delft Debating Club August 2017 – July 2018
Teaching Assistant, Big Data Processing & Distributed Systems at TU Delft November 2017 – March 2018
Teaching Assistant, System Software & Object-oriented Programming at DA-IICT August 2015 – April 2016
Reviewer, ACM/SPEC ICPE, IEEE/ACM CCGRID, IEEE TPDS

Skills

Programming Languages: Python, Go, Java, Kotlin, JavaScript, Node.js, CUDA, C, C++, Scala, Smalltalk, ARM Assembly
Platforms: Linux (Storage and Network APIs), Kubernetes, Apache Spark, AWS, GCP
Data Management: Elasticsearch, Redis, PostgreSQL, MongoDB, Apache Kafka, HDFS, Hive, MinIO, Zookeeper, etcd
Data Analysis: Numpy, Pandas, PyTorch
DevOps: Terraform, Docker, GitHub Actions, Prometheus, InfluxDB

Selected Publications

- [1] **Sacheendra Talluri**, Guido Walter Di Donato, Luca Danelutti, Koen R. Vlaswinkel, Marco Arnaboldi, Arnaud Delamare, Marco D. Santambrogio, Daniele Bonetta. “GpJSON: High-performance JSON Data Processing on GPUs.” VLDB 2025.
- [2] **Sacheendra Talluri**, Dante Niewenhuis, Xiaoyu Chu, Jakob Kyselica, Mehmet Cetin, Alexander Balgavy, and Alexandru Iosup. “Cloud Uptime Archive: Open-Access Availability Data of Web, Cloud, and Gaming Services.” Revision TPDS.
- Ritul Satish, **Sacheendra Talluri**, Sudarsan Sivakumar, Matthijs Jansen, Alexandru Iosup. “Performance Characterization of Data Store Event Trigger Mechanisms for Serverless Computing.” CCGRID 2025.
- [3] Xiaoyu Chu, **Sacheendra Talluri**, Qingxian Lu, Alexandru Iosup. “An Empirical Characterization of Outages and Incidents in Public Services for Large Language Models.” ICPE 2025.
- [4] Aratz Manterola Lasa, **Sacheendra Talluri**, Tiziano De Matteis, Alexandru Iosup. “The Cost of Simplicity: Understanding Datacenter Scheduler Programming Abstractions.” ICPE 2024.
- [5] **Sacheendra Talluri**, Nikolas Herbst, Cristina L. Abad, Tiziano De Matteis, Alexandru Iosup. “ExDe: Design Space Exploration of Scheduler Architectures and Mechanisms for Serverless Data-processing.” FGCS 2024.

Awards

AWS Cloud Research Credit 5k and 10k USD	2020, 2022
Best Artifact Award at ACSOS’21	2021
Graduate travel grant by ACM SIGSOFT and UCC	2019, 2022

References

Prof. dr. Alexandru Iosup, Full Professor, Vrije Universiteit Amsterdam	a.iosup@vu.nl
Dr. Animesh Trivedi, Senior Research Scientist, IBM Research Zurich	animesh.trivedi@ibm.com
Dr. Daniele Bonetta, Assistant Professor, Vrije Universiteit Amsterdam	d.bonetta@vu.nl