

WORKING EXPERIENCE

- **Netflix** Los Gatos, CA
Dec 2019 - Present
Staff Software Engineer, Data Platform

I am one of the leads on the data platform team at Netflix, spearheading projects in the data processing space that are both deep and wide in scope.

 - **Mantis:** [Mantis](#) is a stream processing engine developed at Netflix, designed to address the unique challenges posed by operational data. Mantis handles 50k+ jobs/day, processing 100 PBs of data on 30k+ containers. Over the past couple of years, I have been leading its development.
 - Transitioned Mantis from Apache Mesos to Kubernetes with an innovative architecture. Guided a team of senior engineers from prototype to production, year-long migration and open-sourcing.
 - Reduced annual compute costs by millions through ML-based container optimizations.
 - Championed Mantis adoption to other companies including Stripe.
 - **Flink:** I have also contributed to the Flink ecosystem at Netflix - a stream processing engine for analytical needs.
 - Created a [system](#) enabling users to backfill Flink pipelines using Apache Iceberg without code changes, eliminating the need to maintain separate batch jobs during outages.
 - Designed the system to mimic Kafka properties when reading from Data Lakes, ensuring effortless integration.
 - The system has been adopted by hundreds of pipelines within Netflix. Helped open-source the project, now utilized by other companies using Apache Iceberg.
- **Uber** San Francisco, CA
May 2016 - Dec 2019
Senior Software Engineer II, Developer Platform

I was one of the leads on the Developer Platform team at Uber, responsible for building systems that enable 1000s of engineers to develop, test, and deploy code at scale.

 - **SubmitQueue:** 1000s of engineers committing changes concurrently to a repository leads to frequent master breakages. Explored & conceived a new system called *SubmitQueue* that guarantees an **always-green** master at scale. At Uber, *SubmitQueue* handles 1000s of commits/hr submitted by 1000s of engineers every day.
 - Led a team of 5 engineers to build the system: reading papers on state-of-the-art techniques used in similar domains such as Databases, experimented with various approaches to find a scalable solution, & architected the system to handle 1000s of changes/hr.
 - Published a research paper presenting the design & implementation of *SubmitQueue* at [Eurosys'19](#).
 - **uCI:** Because existing open-source CI systems such as Jenkins did not scale to Uber's needs, I helped build *uCI* - a distributed system to handle reliable execution of millions of stateful tasks every day on 1000s of CI machines.
 - Led a team of 6 engineers to design a state-of-the-art cluster scheduler that handles faults gracefully (*reliability*), exploits data locality to come up with optimal placements (*performance*), scales horizontally on every layer (*scalability*), and finally guarantees isolation at task/resource levels.
- **Baidu Research Silicon Valley AI Lab** Sunnyvale, CA
Jan 2016 - May 2016
Software Engineer
 - **Speech Recognition:** Designed & productionized deep-learning based Speech Recognition APIs which power Android apps such as [TalkType](#).
- **Twitter Inc** San Francisco, CA
Jun 2014 - Jan 2016
Software Engineer
 - **AddressBook Infrastructure:** Engineered a system to store and retrieve contacts from the phone books of Twitter's 300M+ Monthly Active Users (MAUs).

EDUCATION

- **Stanford University** Stanford, CA
Sep. 2012 - Jun. 2014
Master of Science in Electrical Engineering; GPA: 3.9/4.0
- **College of Engineering, Guindy, Anna University** Chennai, India
Aug. 2008 - June. 2012
Bachelor of Engineering in Information Technology; GPA: 9.32/10.0
[Mitacs Outstanding Undergraduate Research Award](#)

ANCIENT HISTORY

- **Microsoft** Redmond, WA
Software Engineering Intern, Kernel Core *Jun 2013 - Sep 2013*
- **Google Summer of Code** Chennai, India
Worked on Metalink Support for Google Chrome *Jun 2012 - Sep 2012*
- **University of Waterloo** Waterloo, Canada
Research Intern - Worked on design & application of One-Instruction Processors *Apr. 2011 - June. 2011*

SELECTED PUBLICATIONS

- [1] **Sundaram Ananthanarayanan**, Masoud Saeida Ardekani, et al. “Keeping Master Green at Scale”. *EuroSys Conference 2019, Dresden, Germany.*
- [2] Dario Amodei, **Sundaram Ananthanarayanan**, et al. “Deep Speech 2 : End-to-End Speech Recognition in English and Mandarin”. *International Conference on Machine Learning, ICML 2016.*

SELECTED TALKS

- **Backfilling Streaming Data Pipelines using Kappa Architecture**
 - *Databricks Data + AI Summit, June 2022*
 - *LinkedIn, March 2022*
 - *Flink Forward, Nov 2021*
- **Keeping Master Green at Scale**
 - *Twitter, Jan 2022*
 - *Google Journal Club, May 2019*
 - *Facebook, Jan 2019*

SKILLS

- **Languages:** Java, Python, Scala, C++
- **Interests:** Distributed Systems, Stream Processing, Machine Learning, Reinforcement Learning