Arjun Balasubramanian

Graduate Student in Computer Sciences at UW-Madison
Advisors: Prof. Aditya Akella and Prof. Shivaram Venkataraman
Interested in designing and building large-scale systems for emerging workloads and hardware
balarjun@cs.wisc.edu | 608-471-0054 | https://pages.cs.wisc.edu/~balarjun

EDUCATION

UW-MADISON
PH.D. IN COMPUTER SCIENCE
Expected May 2023 | Madison, WI
CGPA: 4.0

NIT TIRUCHIRAPPALLI
B.TECH IN COMPUTER SCIENCE AND ENGINEERING
July 2012 - May 2016 | Trichy, India
Department Silver Medalist
CGPA: 9.82 / 10

LINKS

LinkedIn: **balarjun**

COURSEWORK

GRADUATE
Distributed Systems
Big Data Systems
Advanced Operating Systems
Topics in Database Management
Data Science: Algorithms and Principles
Introduction to Artificial Intelligence
(Teaching Assistant)
Introduction to Computer Networks

UNDERGRADUATE
(Relevant Work)
Operating Systems
Computer Networks
Database Management Systems
Distributed Systems
Real-Time Systems
Artificial Intelligence
Data Mining
Data Structures
Algorithms

SKILLS

PROGRAMMING
Over 5000 lines:
C • C++ • Java • Android
Over 1000 lines:
Javascipt • Bash • Go • HTML • CSS
Python • MySQL • LaTeX
Technologies:
Apache Hadoop • Apache Spark
Scikit-Learn • TensorFlow • PyTorch
Apache Zookeeper • Redis

WORK EXPERIENCE

AMAZON WEB SERVICES | SOFTWARE ENGINEERING INTERN, AWS AURORA
May 2019 – Aug 2019 | Palo Alto, California
Designed a ground-up architecture to stream database updates to stream processing engines such as Apache Flink. This use-case is termed as Change Data Capture (CDC). The main challenges were to provide strict exactly-once semantics, ordering guarantees, and achieve high throughput and sub-second latency at scale, while not degrading the performance of the database platform.

AMAZON | SOFTWARE DEVELOPMENT ENGINEER, EREADER PRODUCTS
July 2016 – July 2018 | Chennai, India
Worked on Amazon's first co-branded eReader device named "Kindle Migu X" for which the OS was built from scratch. Developed the BT and Audio Framework to support playback of Audible books on Kindle eReaders. Promoted to SDE-2 in April'18. Explored software mechanisms to improve the I/O performance of accesses to the disk partition which held user content like books. Prototyped two approaches - each gave a performance gain of 5%-10%.

RESEARCH WORK

ACCELERATING DEEP LEARNING INFERENCE VIA FREEZING
VISION PAPER PRESENTED AT HOTCLOUD'19
Oct 2018 – Present | University of Wisconsin-Madison | Link to paper
Freeze Inference is a technique that performs approximate caching at each intermediate layer of a Deep Neural Network (DNN). By employing such a cache, one can avoid the need to run computations for all layers of the DNN for a sizeable portion of inference requests.

FASTPS: A FAST PUBLISH-SUBSCRIBE SERVICE USING RDMA
Sept 2019 – Present | University of Wisconsin-Madison
FastPS is a fast publish-subscribe service that utilizes RDMA both for fast transport and efficient CPU offload. We first analyze requirements from new classes of applications that have very different requirements from the use-case for which Queue Management Systems were designed. This motivates a ground-up redesign of such systems with RDMA at the heart of it.

ARCHIPELAGO: A SCALABLE LOW-LATENCY SERVERLESS PLATFORM
Sept 2018 – Nov 2019 | University of Wisconsin-Madison | Link to arxiv preprint
Worked on a project that involved building an entire serverless platform from scratch with the goals of having low scheduling overheads and minimizing the effects of cold starts incurred during the setup of sandboxed environments.

THEMIS: FAIR AND EFFICIENT GPU CLUSTER SCHEDULING
Feb 2019 – Sept 2019 | University of Wisconsin-Madison | Link to arxiv preprint
Themis is a new scheduling framework that provides both fairness and efficient utilization for ML training jobs that often involve hyper-parameter exploration. Built a configuration-based and event-driven simulator that simulated the assignment of training jobs to a large GPU cluster over time. Used the simulator to compare the proposed scheme with earlier schemes using trace-driven workloads.

HONORS AND SCHOLARSHIPS

2018	Awardee	Special CS Scholarship worth $6,000 at UW-Madison
2017	Winner	Kindle eReader Hackathon at Amazon India
2015	Awardee	Best Outstanding Student in B.Tech CSE at NIT Trichy

ACTIVITIES

PRAGYAN, TECHNICAL FESTIVAL OF NIT TRICHY | CORE MEMBER
May 2015 – May 2016 | National Institute of Technology, Tiruchirappalli
- Prominent speakers included Dr. Peter C Schultz and Dr. Gianni Di Caro.
- Co-Founder of Pragyan Youth Business Summit 2016 - Centered around the "Make In India" initiative and included talks by Ms. Kumud Srinivasan and Mr. Amit Jain.