

ARIJIT PRAMANIK

Email: apramanik3@wisc.edu | Mob No: +1 (510)-241-7934

<http://pages.cs.wisc.edu/~arijit>

EDUCATION

UNIVERSITY OF WISCONSIN-MADISON

MASTERS IN COMPUTER SCIENCE

Aug '19 - May '21

GPA : 3.95 / 4

IIT BOMBAY, INDIA

B.TECH. IN COMPUTER SCIENCE AND ENGINEERING (WITH HONORS)

Jul '15 - May '19

GPA : 9.21 / 10 (Minor in Statistics)

Semester Exchange at National Univ of Singapore (NUS) (GPA: 4.92 / 5)

KEY COURSEWORK

Advanced Computer Networks

Advanced Operating Systems

Big Data Systems

Advanced Databases

Machine Learning Optimized Systems

High-Performance Computing

Information Security & Cryptography

Machine Learning & Artificial Intelligence

Computer Vision & Image Processing

TECHNICAL SKILLS

Proficient-

•C/C++ •Python •MATLAB •Java

•Docker •P4 •SQL •Bash •Git • \LaTeX

Familiar-

•Scala •Javascript •Pytorch •SmartNIC

•Spark/Hadoop •HTML/CSS •Django

•Go •Android •OCaml •Racket •R/SAS

PATENTS & PUBLICATIONS

"Abstractive Text Summarization tailored to target characteristics"

K. Chawla, H. Singh, A. Pramanik,

M. Kumar & B. V. Srinivasan

CICLING 2019

"Method to generate a target-characteristic tuned content using a word generation model"

Filed at USPTO in Jan 2019

ACHIEVEMENTS & AWARDS

• **Institute Academic Prize** : Dept. Rank 1 in 3rd year (GPA : 9.86 / 10)

• **Cohort top 1%** in Information Retrieval, Optimization & Numerical Analysis

• **Teaching Assistant**: Computer Graphics, Programming & Computer Architecture

• **Institute Sports Citation**: Aquatics champion (4 gold, 5 silver & 11 bronze) and **Aquatics Captain** (2018-19)

KEY INTERNSHIPS

SANDBOXING OF UNTRUSTED PROCEDURES | SDE INTERN

Summer '20 | AWS Relational Database Service | Guide: Jignesh Shah

- Integrated an open-source extension PL/Container into PostgreSQL 12 and 13 to enable sandboxed execution of stored procedures inside Docker containers
- Ran performance benchmarking to identify bottlenecks in Unix sockets & grpc channels to maintain latency at par with locally running extensions
- Developed a new prototype to use a single container across all customers with a 63% memory usage reduction, while scaling to thousands of connections
- Added runtime support for Go language and created 3 separate extensions for R, Python and Go, utilizing separate containers for better isolation

CHARACTERISTICS-DRIVEN SUMMARIZATION | RESEARCH INTERN

Summer '18 | Adobe Systems | Guide: Dr. Balaji Vasani Srinivasan

- Adapted Facebook AI Research's convolutional seq2seq model for feature-driven text generation on *pointer-generator* framework with modified attention layers to focus on specific input text embeddings for topic-tuned summaries
- Altered beam search paradigm for enhancing decoder state induced word-level features with token-based learning for length constrained summarization
- Achieved a 6.4% increase in ROUGE scores with Reinforcement Learning

RESEARCH

MEMORY-EFFICIENT LEARNED INDEXES | RESEARCH ASST.

Aug '20 - May '21 | UW-Madison | Guide: Dr. Remzi & Andrea Arpaci-Dusseau

Trained a model to learn offsets from index entries (last keys) of data blocks inside SSTables to reduce lookup time from $O(\log n)$ to $O(1)$. Obtained a 53% reduction in indexing memory footprint over traditional indexes with <5% increase in point lookup latency using both *Fuzzy* and *Greedy Piecewise Linear Regression* in RocksDB

DATAPLANE-ONLY POLICY-COMPLIANT ROUTING | IND. STUDY

Aug. '19 - Dec '19 | UW-Madison | Guide: Dr. Aditya Akella

Devised a routing mechanism on programmable switches in P4 that uses search algorithms to compute the route in the data-plane without controller intervention. Handled failures through register updates utilizing failure packets in Tofino. Supports policies like middlebox-chaining, flexible WCMP with hierarchical routing

KEY PROJECTS

WORKLOAD-AWARE KEY-VALUE STORE | ADV. DATABASES

Explored Facebook's RocksDB for black-box tuning of parameters using *Bayesian Optimization* to maximize client write throughput. Obtained 2.34X improvement for industrial write bursty workloads by scheduling background memtable flushes and compactions during *idle periods* or *periods of very few writes* to the database

LEGAL CASE RETRIEVAL SYSTEM | INFORMATION RETRIEVAL

Designed a freetext search engine to rank legal case judgements (2nd/33 teams) with positional indices for proximity search and zone & field indices like court hierarchy, date. Enhanced F_1 score using query expansion: pseudo-relevance feedback with Rocchio formula, WordNet synonyms and *tf-idf* weighted co-occurrence thesaurus

PARALLELIZED IMAGE SEGMENTATION | HIGH PERF. COMPUTING

Accelerated Canny edge-detection and Fuzzy C-Means clustering leveraging GPU and hybrid OpenMP+MPI on multicores. Explored SIMD, use of templates and forced inlining with shared and unified memory. Utilized CUDA streams, dynamic parallelism and thrust library along with OMP tasks to achieve high speedup