

Chenyuan Wu

E-mail: wucy@seas.upenn.edu | Mobile phone: +86-18800172721

EDUCATION

University of Pennsylvania

Ph.D., Computer Science

• Advisor: Prof. Boon Thau Loo

• Current GPA: 4.0/4.0

• Select Graduate courses: Operating Systems, Distributed Systems, Database and Information Systems, Deep Learning, Machine Learning, Software Foundations

PA, US

2020 - Present

Beijing Jiaotong University

B.E., Electronic Engineering

• General GPA: 3.9/4.0

Beijing, China

2016 - 2020

PAPERS

- Shaobo Wang, Hui Lyu, Jiachi Zhang, **Chenyuan Wu**, Xinyi Chen, Wenchao Zhou, Boon Thau Loo, Susan B. Davidson and Chen Chen, *Provenance for Probabilistic Logic Programs*, The 23rd International Conference on Extending Database Technology (EDBT '20). **Best Paper Award**.
- Haoxian Chen, **Chenyuan Wu**, Manqiu Zhang, Yahui Li, Zhijia Chen, Boon Thau Loo, Mayur Naik and Mukund Raghothaman, *Verifiable DSL Programs from SDN Controller Applications*, The 19th USENIX Symposium on Networked Systems Design and Implementation (NSDI '22). Status: Under Review.

RESEARCH INTERESTS

I'm currently interested in practically applying ML/DL techniques to solve systems problems, including optimizing memory hierarchy and storage, learning scheduling policies for distributed platforms, and developing other learning-based components in computer systems.

AWARDS AND HONORS

PhD Fellowship in University of Pennsylvania

Sept. 2020

President's Scholarship in Beijing Jiaotong University (10 out of 15000)

Nov. 2019

Chinese National Scholarship

Sep. 2019

Chinese National Scholarship

Sep. 2017

Second Prize in the National College Physics Competition

Dec. 2017

WORKING EXPERIENCES

Huawei, Central Research Institute

Research Intern, Network Technology Lab

Nov. 2019 – Jan. 2020

University of Pennsylvania

Research Intern, Distributed Systems Lab

July. 2019 – Sept. 2019

RESEARCH EXPERIENCE

Synthesizing Verifiable DSL Programs from SDN Controller Applications July. 2019 – Sept. 2019

Research Assistant, University of Pennsylvania

Advisor: Prof. Boon Thau Loo

Objective: AutoDSL is a toolchain that synthesizes SDN controller implementations in domain specific languages (DSL) from input-output execution traces, aiming to bridge the gap between legacy and DSL implementations. AutoDSL uses code analysis to extract internal states from legacy SDN implementations, and syntax-guided synthesis to generate logic programs that are consistent with the input-output example traces. The logic programs are compiled into simpler, more concise logic-based DSL implementations, which can then be used as a basis for correctness verification.

- Implemented two compiler modules of the toolchain in Scala, which could automatically translate the intermediate logic programs into Flowlog rules and dynamic policies in Pyretic;

- Implemented 7 legacy programs on the FloodLight platform, and implemented an execution traces collector for FloodLight;
- Compiled logic programs of MAC learning and L2/L3 firewalls into DSL implementations, validated their functionality in Mininet;
- Compared the end-to-end difference in throughput between handwritten and synthesized DSL implementations by running perf tests over a FatTree topology in Mininet, compared their difference in response ability by running Cbench tests;
- Submitted a paper to NSDI [2].

Provenance for Probabilistic Logic Programs

July. 2019 – Sept. 2019

Research Assistant, University of Pennsylvania

Advisor: Prof. Boon Thau Loo

Objective: P3 is a novel provenance model and system for analyzing probabilistic logic programs. P3 enables four kinds of queries: traditional explanation queries, queries for finding the set of most important derivation within an approximate error, most influential queries, and modification queries that change the output probability with fewest modifications to input data or program.

- Implemented two parallel computing methods for the most influential query in P3, where each work unit is in charge of a single round of Monte Carlo simulation or computing the influence value of a single Literal;
- Set up parallel computing environments on both Nvidia GPU and Intel CPU, run experiments on large provenance graphs, and analyzed results from different methods and different hardware settings, boosting the performance by more than 20-fold faster than the sequential method at most;
- Conducted a Visual Question Answering case study, further enabling the system to find inappropriate input data and show the modification strategy with minimum cost in order to make the prediction right;
- Published a paper in EDBT [1].

PROJECTS

PennCloud

A Project of the Distributed Systems Course, University of Pennsylvania;

Collaborated with team members to develop a Google-Apps-like cloud platform from scratch that provides webmail service and storage service

- Implemented a Bigtable-like (supporting PUT, GET, CPUT, DELETE operations), distributed, replicated, consistent and fault tolerant key-value store as the backend of the cloud platform;
- Implemented the remote-write protocol for the backend to handle replication and guarantee sequential consistency;
- Implemented checkpointing and logging mechanisms together with a recovery protocol for the backend to efficiently recover from failures;
- Implemented a master node of backend that keeps track of the row range partitions and alive backend servers by using heartbeats.

PennOS Projects

Projects of the Operating Systems Course, University of Pennsylvania;

- PennShell: Implemented a bash-like shell that supports redirection, multistage pipelining, foreground and background processes, and job control; The implementation used asynchronous polling;
- PennMMU: Implemented an in-memory simulation of MMU that supports dynamic memory allocation, virtual memory address translation, page replacement (supporting FIFO, CLOCK, LRU policies), and TLB;
- PennOS (on-going): Collaborated with team members to develop a user-level UNIX-like operating system, including a kernel that supports process management and priority scheduling, a FAT file system, and a basic shell and programming APIs for a user to interact with the PennOS.

TEACHING AND ACTIVITIES

Conference Reviewer: served as the secondary reviewer of the International Conference on emerging Networking Experiments and Technologies (CoNEXT' 2020).

Conference Reviewer: served as the secondary reviewer of the 2nd Workshop on Resource Disaggregation and Serverless (WORDS' 2021).