Pinglei Guo

Vancouver, BC, Canada

Email: plguo002@gmail.com GitHub: https://github.com/at15 LinkedIn: https://linkedin.com/in/at1510086

Work Experience

AWS: Software Engineer - CloudWatch Agent Vancouver, Canada Sep 2020 - Present

- Working on the open source CloudWatch Agent.

Google: Software Engineer - Batch job on Kubernetes Sunnyvale, CA May 2019 - Nov. 2019

- Implemented gRPC server and client in Go and Python backed by Spanner, deployed on Borg and Kuberentes.
- Improved reliability of custom batch job controller for Kuberentes, eliminated job lost during controller service restart.
- Troubleshoot and reduced startup latency of debug session from 2min to 40s, created metrics for performance regression. Reduced latency for large batch job input (GB) by order of magnitude, i.e. 5min to 100ms.

PayPal: Software Engineer - Multi cluster container orchestration platform in Go San Jose, CA May 2018 - May 2019

- Worked on internal multi cluster container orchestration platform using Apache Mesos, Aurora and Docker.
- Introduced new deploy strategy (e.g. canary) and implemented readiness gate, reduced rollout failure rate by 20%.
- Built REST API and cli tools in Go to collect log, metrics, distributed traces and profile data.
- Built Kubernetes operator for running stateful services like database with in house underlay container network solution.

PayPal: Software Engineer Intern - Admin Server & Dashboard in Go San Jose, CA June 2017 - Sep. 2017

- Built API gateway with RBAC for internal container orchestration platform using Go.
- Enhanced dashboard using Angular 4, wrote a new dynamic table component from scratch.

Dongyue Web Studio: (Part-time) Full stack web developer & Tech lead Shanghai, China Sep. 2013 – Jan. 2016

- Led web and mobile team. Rewrote online event booking website tongqu.me, used by thousands of students.
- Utilized Redis as cache and rate limiter, increased QPS by 120%, reduced database load by 40%, filtered out most bot traffic.
- Migrated frontend from jQuery to AngularJS, reduced page loading speed by 60% using Ajax.

Project Experience

Distributed database benchmark system github.com/benchhub UCSC Nov. 2017 – March. 2018

- Designed a specification for running database benchmark for RDBMS and TSDB in distributed environment.
- Implemented a continuous integration service that stores benchmark results in databases to detect performance regression.

Distributed Time Series Database github.com/xephonhq/xephon-k UCSC Nov. 2016 – Present

- Implemented a distributed time series database on top of Cassandra in Go. Support both JSON and Protobuf via HTTP/2.
- Designed a columnar storage engine modeled after Parquet and InfluxDB with high compression and less write amplification.
- Created benchmark suite for Xephon-K, OpenTSDB, KariosDB, InfluxDB using a generic client for different TSDB.
- Surveyed popular TSDB design and implementation, made an interactive online report awesome-time-series-database.

GPU accelerated in-memory time series processing github.com/at15/ts-parallel UCSC Apr. 2017 – June 2017

- Expanded benchmark suite for different C++ GPU computing framework on CUDA and OpenCL, Thrust, Boost, ArrayFire.
- Implemented OLAP queries like top-K, group by for multi dimensional time series data on both CPU and GPU backends.

Education

MS. Computer Science	University of California Santa Cruz	GPA 3.9	Sep. 2016 – Mar. 2018
BS. Materials Science	Shanghai Jiao Tong University	GPA 3.3	Sep. 2012 – June 2016

Skills

Language	Go, C++, Rust, JavaScript, Java, Python, SQL, PHP, Shell
Database	Cassandra, MySQL, Elasticsearch, TiDB, Redis, KairosDB, OpenTSDB, InfluxDB, Prometheus, Graphite
DevOps	AWS, Kubernetes, Docker, Terraform, GCP