Jin Dong

github.com/djdongjin | linkedin.com/in/jdong95 | djdongjin.github.io | djdongjin95@gmail.com | 332-262-8260

Education

McGill University — Montreal, Canada
M.S. in Computer Science, GPA: 4.00 / 4.00 — 2018/09 - 2020/10

Jilin University — Changchun, China
B.S. in Computer Science — 2013/09 - 2017/07

Skills

Languages: Go, Python, Rust, C++, Java
Technical Skills: Containerd/Docker, Kubernetes, Container Runtime/Registry, Database, Performance Tuning, Open Source

Experiences

Databricks — Seattle, WA
Software Engineer, Container Image Infrastructure (container runtime, container registry) — 2023/07 - Present

- *Image Acceleration.* Created a container image lazy download system consisting of image conversion pipeline, Kubernetes pod mutation operator, and node container runtime. Reduced the P99 pod startup latency by 80% from 6 min to 1 min without requiring migrations.
- *Container Registry Caching.* Designed a container registry caching service, which reduced registry API P99 latency by 90% and CPU usage by 80%. Implemented core features such as LRU caching and thundering herd protection. Wrote a distributed load tester, conducted intensive load testings and fixed production blockers (availability drop during restart, negative caching, etc).
- *Graviton/ARM Adoption.* Led cross-team collaboration to support Graviton CPU. Implemented multi-platform image support in all internal image pull services.
- *P2P Image Download.* Created a P2P image download agent to eliminate the bottleneck on cloud blob storage bandwidth limits. Implemented checksum validations in P2P downloads and TLS encryption for all P2P communications. Achieved 10x image download scalability in dev environments.
- *Containerd adoption.* Lead the internal adoption and maintenance of containerd. Unblocked Kubernetes upgrade by migrating from docker to containerd. Contributed 10+ performance improvements and bug fixes to containerd project.
- *Image Replication.* Built a cronjob to monitor and replicate missing container images from leader registry to follower registries. Achieved 100% eventually consistency of image replications. This helped avoid multiple production outages due to missing images.

Amazon Web Services — Jersey City, NJ
Software Engineer, Container Runtime (containerd, snapshotter, open source) — 2022/10 - 2023/06

- *Lead developer of SOCI*, a containerd remote snapshotter that speeds up container launch by lazily pulling/downloading images.
- Decoupled the snapshotter implementation and compression algorithm (gzip), enabled support to new image compression. (PR1, PR2)
- Optimized SOCI's image fetch/cache component to be thread-safe and efficient by utilizing atomic operations and mutex lock. (PR)
- Collaborated with AWS Fargate and launched SOCI on Fargate platform.
- Lead the team's contribution to nerdctl, an open source docker-compatible CLI for containerd runtime. Supported and unblocked sister teams to build Finch (a docker replacement) on top of nerdctl.

Microsoft — Redmond, WA & Vancouver, Canada
Software Engineer, Azure Kubernetes Service (Scalability) & WebEx — 2020/09 - 2022/10

- Integrated the next-gen Azure VM node pool (Flex) into AKS, reducing 50-node cluster creation from 9 minutes to 3 minutes.
- Contributed to scaling AKS to 5K nodes, by tuning kubelet/apiserver config, optimizing the scheduling of control plane components.
- Optimized AKS throttling issue by reducing and batching Azure API calls.
- Mentored an intern project enabling customers to monitor AKS control plane metrics with their Prometheus by creating a proxy server.
- Reduced the runtime of a critical MSN data pipeline from 1 hour to 20 minutes, by optimizing the data extractor and processor.

Mila (Quebec AI Institute) — Montreal, Canada
Research Assistant, Advisor: Prof. William Hamilton — 01/15/2019 - 04/01/2020

- Published two research papers about Question Answering on EMNLP conference.

Open Source

Containerd - Maintainer — 2022/11 - Present

- Maintainer and top contributor (50+ PRs) of containerd, the default container runtime of Kubernetes.
- Contributed many performance optimizations (e.g. optimize garbage collection and image unpack) and bug fixes (e.g., image pull panic due to race conditions).
- Adopted the go-native fuzz tester in containerd, enabling developers to run fuzz tests locally. (PR1, PR2, PR3)

Nerdctl - Maintainer — 2022/11 - Present

- Maintainer and top-5 contributor of nerdctl, a containerd CLI project under containerd and CNCF.
- Implemented cosign image signing/verification support for nerdctl compose.
- Created 11 docker-compose commands in nerdctl, significantly improved nerdctl's compose support.

TensorFlow - Google Summer of Code — 2020/06 - 2020/08

- Built new CLI tools that support model fine-tuning pipelines, including object detection and nearest neighbor indexing.
- Enabled users to E2E fine-tune TensorFlow Hub models via the CLI tools without writing any code.