Thanos Ariyanayagam

REDACTED PHONE # | thanoshan2001@gmail.com | linkedin.com/in/thanoshan | github.com/Thanoshan

Experience

Google Sunnyvale, CA
Software Engineer *Aug. 2024 - Present*
- **Saved an estimated $53.2 SWE-years annually** by architecting a **next-generation C++ node runtime framework** to enable **ML compute fungibility** (CPU/GPU/TPU) across Google's global, heterogeneous hardware fleet.
- **Hardened Google's workload security** by driving the **isolation and sandboxing commitment** for the **TI-VM infrastructure**, leading performance optimization initiatives to reduce virtualization overhead and scale system reliability.
- Led the **architectural design and C++ implementation** of the **resource management parity initiative** for the new runtime, enabling **secure, multi-tenant execution** of untrusted ML models and other workloads at scale.

Google Sunnyvale, CA
Software Engineering Intern *May 2023 - Aug. 2023*
- Engineered a high-throughput, multithread-safe **C++ API in Borglet** for task lifecycle management, leveraging **lock-free programming techniques** to minimize contention and improve reliability at Google-scale.
- Designed and deployed a novel data-sharing solution for modularizing fleet-wide features into autonomous services, enabling secure, hierarchical dependency sharing within the **Borg cluster manager**.

Intel Corporation Toronto, ON
Software Engineer Intern *Sept. 2022 - Apr. 2023*
- **Boosted** ML model inference throughput by **11%** by implementing advanced **C++ compiler optimizations**, including graph-level operator fusion, memory layout transformations, and enabling **mixed-precision (INT8) inference**.
- Authored a patent-pending **C++ compiler extension** (pub. #20230237014) for a proprietary IR to support 3D graphs, enabling **full on-chip acceleration** for 3D deep learning models.

Google Waterloo, ON
Software Developer Intern *May 2022 - Aug. 2022*
- Deployed a high-impact ensembling library in **C++ and Go** that **boosted** a production NLP API's F1 score by **4.28%** and **cut** annotation errors by **50%**.
- Designed and automated a **scalable ML experimentation pipeline on GCP** to rigorously evaluate and benchmark model quality, enabling rapid iteration and data-driven improvements to the core NLP API.

Google Waterloo, ON
STEP Intern *May 2021 - Aug. 2021*
- Enhanced **C++ infrastructure** to capture ML anomaly scores in SpannerDB, **increasing** bad actor suspensions by **8%** and **reducing** scan quota by **10%**.

Projects

OpenStreetMaps GIS | *C++, Multi-threading*
- Achieved a **>56% runtime improvement** in a custom C++ GIS by architecting a concurrent design utilizing **thread pools** to manage asynchronous tasks and parallelize data parsing and pathfinding.
- Implemented and optimized thread-safe pathfinding algorithms, including simulated annealing and greedy heuristics, securing a **top-6 placement for route quality** against competing solutions.

NEPIADA | *Python, Reinforcement Learning*
- Designed novel **multi-agent RL algorithms (DQN, PPO)** that **outperformed** state-of-the-art methods by converging to a Nash equilibrium in adversarial, partial-information environments.
- Engineered a custom, high-fidelity multi-agent RL environment using **PettingZoo** to simulate drone swarm behavior, providing a robust platform for algorithmic research.

AI Reversi Player | *C, Algorithms*
- Developed a **top-5% ranked Reversi AI** in C (out of 400+) by implementing a highly optimized Minimax algorithm with **alpha-beta pruning** and move ordering.
- Applied advanced **game-tree search techniques**, including **transposition tables (memoization)**, to drastically reduce the search space and enable deeper strategic analysis under tight time constraints.

Technical Skills

Languages: C++, C, Python, Go, SQL, Bash (Modern C++: 17/20/23)
AI/ML: TensorFlow, PyTorch, Scikit-Learn, Pandas, NumPy, Model Pruning, RLlib
Infrastructure & Tools: GCP, AWS, Docker, Kubernetes, Git, CI/CD, Spanner, BigQuery, gRPC, Borg, CMake, Bazel
Performance & Debugging: GDB, Perf, Valgrind, Google Test (gtest), PyTest
Core Competencies: Systems Design, Compiler Optimization, High-Performance Computing, API Development, Multi-threading, Lock-Free Programming, Systems Programming

Education

University of Toronto Toronto, ON
Bachelor of Applied Science in Computer Engineering *Sept. 2019 - Apr. 2024*
- Graduated with High Honours (CGPA: 3.81/4.0); Minor in Engineering Business; Certificate in Artificial Intelligence