Martin Chang

Software Developer/Team Lead

Education

2016 - 2021

Computer science and engineering (B. Sc.)

National Sun-Yet-Sen University
Special Talent Program

Graduation project: „HTM based Agent in standard RL Environment and Beefing up Framework Performance“.

Skills

C++ 13+ Yrs

Linux 11+ Yrs

Concurrent/HPC programming 7+ yrs.

OpenCL 4+ yrs.

(CERN ROOT) Big Data analysis 2+ yrs.

Language skills

Chinese Native

English (TOEIC) 815

Esperanto (Est. KER) C1

Additional skills

- Lock free programming
- Lua, Python, Common Lisp
- Qt, PostgreSQL, GNUnet
- RK3588, GGML, Tenstorrent

github.com/marty1885

clehaxze.tw

Compiled 13th January 2025 with LaTeX

Profile

I started coding at a very young age. With 15+ years of C++ experience. Now following the latest C++ standard and TS. I'm interested in HPC and systems programming. I am also an regular open source contributor. And an active maintainer of Drogon, one of the fastest web application framework. Currently working on porting AI accelerators to GGML like the RK3588 NPU and Tenstorrent processors. I'm confident in the diversity of my skill set and my ability to understand then fix freestanding bugs that I find.

Work experience

Senior Engineer **07/2024 - CURRENT**

SatLayer

Help deliver our core SDK, manage partner engineering relationship and manages a a team of 4. Does research into existing protocol implementations and discover protocol limitations and system guarantees.

Senior Software Engineer **11/2023 - 07/2024**

NVIDIA

Works on internal tools and package maintaince in the Omniverse project. Does a lot of cross-discipline communication and reading new, incoming/broken code.

Native Application Lead **09/2021 - 11/2023**

Lumina Industries, Inc

Manages and leads our software team while also contributing to the codebase. As well as handling deployment to customer environments. Does pilot implementations to reduce workload on our engineering team.

Open source

Some especially notable projects that I am involved and did major contributions.

llama.cpp fork - marty1885/llama.cpp **AUTHOR**

LLaMA fork with NPU support

Experimental fork of the popular inference framework that enables use of non traditional accelerators. Supporting the RK3588 NPU and Tenstorrent Wormhole/Grayskull processors. This work started efforts to reverse engineer the RK3588 NPU. The Tenstorrent backend is close and in talk of upstreaming.

Drogon - drogonframework/drogon **MAINTAINER**

C++14/17/20 web application framework

Drogon is one of the fastest web application framework. I maintain and contributed core subsystems for the framework. Such as C++20 coroutine, HaikuOS support, Gemini protocol integration and the DrogonTest async test framework. I aslo maintain trantor - our transport library. Overhauled the TLS support and CSPRNG subsystem. With an working HTTP/2 client implementation from scratch.

Etaler - etaler/Etaler **AUTHOR**

High performance HTM/tensor library

Etaler is a high performance implementation of Hierarchical Temporal Memory, a biologically inspired AI model by Numenta. At time of release, Etaler is more than 20x faster than the community developed HTM.core on CPU and 40x faster on a GPU. I'm also a trusted member of the HTM community.