Johan Jino

johanjinoj@gmail.com +447436567710 linkedin.com/in/johan-jino-student
github.com/johanjino johanjino.github.io

PROFESSIONAL EXPERIENCE

System Software Engineer 04/2024 – 09/2024 –
NVIDIA 08/2025 – present
Reading,
United Kingdom

- Shipped SMPTE ST 2110-40 ancillary data streaming in NvDsUdp (open source, DeepStream) on the Rivermax SDK, including a new SMPTE 291 RTP payloader/depayloader used in real-time AI workflows.
- Hardened NvDsUdp's kernel-bypass transmit and receive path, fixing jitter edge cases and reconnection failures in live broadcast pipelines.
- Presented at GStreamer Conference 2025 as primary author, covering GPU-to-NIC transfer, packet-paced transmission and NMOS connection management. Contribution page
- Contributor and reviewer on open source NVIDIA NvNmos , including Rust GStreamer elements. Worked on custom Kubernetes operators in Go, and Wireshark dissectors for new ST 2110 protocols.
- Designed and shipped the H4M Installer v1 from scratch (FastAPI, React, asyncssh, Ansible). Desktop app that provisions bare-metal OpenShift or Kubernetes GPU clusters, with a docs-grounded assistant on NVIDIA NIM.
- *Intern:* Cut GPUs needed by more than 50% for 4K60 ST 2110 streams while staying broadcast compliant, using MIG, time-slicing and CUDA MPS.
- *Intern:* Built and shipped to production a metrics toolchain for Kubernetes applications (Prometheus, Grafana, Helm), showing GPU and SR-IOV network metrics in one view.

Software Engineer Intern 07/2023 – 09/2023
Qualcomm London,
United Kingdom

- Worked in C++ firmware for the RISC-V hardware root of trust in Snapdragon SoCs. Added call-stack retrieval during on-device testing, which significantly reduced debugging time.
- Rebuilt the on-device profiling and timing-analysis tool, moving it from TCL to Python and adding statistical analysis.

Research Engineer Intern 07/2022 – 10/2022
Ivy London,
United Kingdom

- Built transpiler backend features for an open-source library that converts code between ML frameworks (PyTorch, TensorFlow, JAX, NumPy). Wrote property-based tests with Pytest and Hypothesis, and reviewed community pull requests. https://github.com/unifyai/ivy

EDUCATION

MEng Electronic and Information Engineering 09/2021 – 06/2025
Imperial College London London,
United Kingdom

- First Class Honours and Dean's List. Overall 78.50%
- Final Thesis: Compiler Analysis and Profile Assisted Memory Dependency Prediction. Achieved 86%. Presented at IET REACH Conference 2025.
- Modules of interest: Machine Learning (83%), Advanced Computer Arch (82%), Digital Systems (85%), Software Systems (76%).
- Developed ExecEngine (ANSI C compiler) for RISC-V architecture, rated second best compiler in the cohort of 150 students.
- Designed an FPGA mathematical function accelerator on Intel DE-1 SoC, which ranked top 3 in execution time.

RESEARCH AND TALKS

PG-MDP: Profile-Guided Memory Dependence Prediction for Area-Constrained Cores Salt Lake City, Utah, USA

HPCA 2027 · Paper, Second Author

- Hardware-software co-design that labels loads that never conflict, so they skip the memory dependence predictor. Cuts false dependencies by 84% and predictor queries by 80%, raising IPC on a small core by 4.6% without a larger predictor. Accepted to **HPCA 2027**. Preprint .

Profiling Assisted Memory Dependency Prediction Using Aliasing and Distance Information London, UK

IET REACH 2025 · Poster, First author

Showed that store distance, the number of stores between a dependent load and store, is the most effective profiling signal for labelling loads. Built the toolchain with LLVM passes, SVF and gem5 on SPEC CPU2017. Poster

Hardware Accelerated Live Broadcasting London, UK

GStreamer Conference 2025 · Speaker

- Presented NVIDIA's open-source GStreamer plugins for GPU and NIC accelerated ST 2110, covering GPU-to-NIC transfer, packet pacing, kernel bypass and NMOS connection management. Talk

SKILLS

Technical Skills

- Languages: C++17/20, Python, C · working knowledge of Rust, Go, Bash, SQL, SystemVerilog
- GPU and systems: CUDA MPS, MIG, GPUDirect, Nsight Systems, perf, gdb, Valgrind, CPU pinning, NUMA
- Networking: Rivermax kernel bypass, RDMA, SR-IOV, PTP, ST 2110, RTP, NMOS, Wireshark
- Infrastructure: Kubernetes operators, GPU and Network Operators, Helm, Docker, Prometheus, Grafana, DCGM
- Compilers: LLVM, SVF, gem5

AWARDS

Awards

- Dean's List Award for top 10% of the cohort, Imperial College London, EEE department.
- National Academic Excellence Award - Abu Dhabi
- Prefect of the Year Award – GEMS United Indian School
- Gold Medalist for Year 12, 11 and 10 for best performance among the all students in the cohort.
- Joint Entrance Examination India (JEE) – 88.719 percentile. Qualified for Advanced (2021)

PROJECTS

ExecEngine 02/2023 – 04/2023

ANSI C Compiler

- Developed C (C90) to RISC-V assembly compiler in a pair, with full support for integer and floats, pointers, arrays, char and structs.
- Added support for most keywords including while, switch, break, continue, enum, built-in functions like sizeof as well as multiple nested and recursive functions.
- Created Bash Scripts for test automation and Yaml for GitHub automation config.
- Scored a 80% out of all the seen and unseen test cases, putting **our compiler 2nd among all 150 students**.

RISC-V CPU 11/2022 – 12/2022

SystemVerilog CPU Design

- Designed a multi-stage pipe-lined CPU compliant to RISC-V 32-bit integer instruction set architecture in SystemVerilog.
- Built both pipelined and data cached version of the CPU
- Includes an independent C++ test-bench which can be used to assert the functional correctness of any given CPU.
- Worked on RTL Design, TestBenching, Workflow automation