Chenhe Yuan

London, United Kingdom | chenhe.yuan.22@ucl.ac.uk | sdjnych@gmail.com | chenheyuan.com

Education

University College London Sept 2022 - Jun 2026
MEng (Integrated Master's) in Electrical and Electronic Engineering

- GPA: 3.79/4.0
- **Relevant Coursework**: Advanced Digital Design, DSP and Systems, Programming and Control Systems.

The University of Texas at Austin Aug 2024 - May 2025
Undergraduate Exchange, Electrical and Computer Engineering

- GPA: 3.5/4.0
- **Coursework**: Computer Architecture, Algorithms, Compilers, Digital Electronics.

Systems & Engineering Projects

Low-Latency FPGA Market Data Parser Apr 2026
Personal Project, Xilinx VU13P

- Designed a deterministic AXI-Stream UDP packet parser in SystemVerilog targeting HFT market data price extraction; achieved fmax ≈ 612 MHz (WNS +0.367 ns at 500 MHz constraint) on Xilinx UltraScale+, using 298 LUTs and 198 FFs (<0.03% device utilisation).
- Implemented strictly-ordered fail-fast header filtering (EtherType → Protocol → Dst Port) with byte-counter-driven parsing, discarding uninteresting packets at the earliest valid beat to eliminate downstream logic activity.
- Designed a host-reconfigurable runtime format table with masked discriminator lookup, decoupling payload format changes from FPGA re-synthesis and enabling live feed format updates without hardware rebuilds.
- Engineered a cross-beat price assembly datapath supporting arbitrary 4-byte alignment across AXI-Stream beat boundaries, with per-lane tkeep gating and single-cycle price-valid output upon full field capture.
- Verified with cocotb (Verilator): directed header-filter and boundary tests, bubble injection under non-contiguous tvalid, and a 100-packet randomised sweep with Python reference model across all price offsets and bubble probabilities.

LiDAR-Camera Fusion Sparse Inference Architecture for FPGA Classification Oct 2025 – Jun 2026
Team Project, advised by Prof. R. Killey, University College London

- Proposed a LiDAR-camera fusion approach that pruned distant background pixels before inference, converting dense visual input into a sparser computation problem.
- Designed a sparse tiled-convolution architecture to exploit spatial continuity in foreground regions, improving tile occupancy and DSP utilization under on-chip resource limits.
- Implemented the Jetson-based comparison pipeline and used it to evaluate trade-offs against the FPGA-oriented design.

LC-3B ISA Simulators and Assembler in C Sep 2024 – Dec 2024
ECE 460N Computer Architecture, The University of Texas at Austin

- Built LC-3B simulators in C at the instruction-level, microarchitectural, virtual-memory, and cycle-accurate levels to validate behavior across abstraction layers.
- Extended the microarchitecture with virtual memory support, including user/system address-space separation.
- Implemented an assembler translating LC-3B assembly into executable binaries with correct starting-address handling.

Research Experience

Streaming Sparse 3D CNN Accelerator on FPGA May 2025 – Dec 2025
Advisor: Prof. Lizy K. John, The University of Texas at Austin

- Designed a streaming FPGA accelerator for sparse 3D convolution on Xilinx VU9P, using four-stage pipelines and FIFO buffering to overlap adjacent layers without global synchronization.
- Built a hierarchical BRAM-based occupancy-bitmap pipeline to skip empty voxel regions online, improving performance by up to 4.4× on spatially sparse layers.
- Implemented DRAM-backed kernel tiling with URAM accumulation buffers and identified off-chip READ bandwidth as the primary bottleneck through DMA/compute throughput analysis.

Compute-in-Memory Architecture for In-Sensor CNN Evaluation on FPGA May 2024 – Oct 2024
Advisor: Prof. Chao Li, Shanghai Jiao Tong University

- Designed a digital emulation of an analog crossbar memristor array for CNN evaluation, translating compute-in-memory behavior into a configurable FPGA prototype.
- Implemented the design for a Xilinx Zynq-7000 target through synthesis, implementation, and bitstream generation across multiple configurations.

- Evaluated trade-offs between architectural configurations and workload characteristics to guide further exploration.

Embedded Acceleration Module for CNN Convolution Jun 2023 – Oct 2023

Advisor: Prof. Tinghuan Chen, The Chinese University of Hong Kong (Shenzhen)

- Contributed to a parameterized Winograd-based convolution IP generator in Chisel/SystemVerilog, reducing multiplication count in CNN inference layers; authored the templated compute logic in Chisel, enabling configurable array sizes across generated IP variants.
- Proposed a pipelined synchronization scheme for correct handshaking across pipeline stages in the generated hardware.

Skills

Hardware Description: SystemVerilog, Chisel (Scala-embedded), Vitis HLS
Simulation & Verification: Cocotb, Verilator, Icarus Verilog
FPGA Vendor Tools: Vivado, Quartus
Architecture Research Tools: gem5, ChampSim, Accel-Sim, Timeloop/Accelergy, MATLAB
Programming Languages: C, C++, Java; experienced with typed and functional programming paradigms (Scala/Chisel, Java generics, C++ templates)