

# Shibo Chen

Staff Architect | Chiplet NoCs, coherent fabrics, memory systems, performance modeling, and AI-assisted hardware design

Santa Clara, CA / [chshibo@outlook.com](mailto:chshibo@outlook.com) / [shibo-chen.github.io](https://shibo-chen.github.io) / [LinkedIn](#) / [GitHub](#)

## SUMMARY

---

- Staff Architect at Tenstorrent, working across chiplet-based AI/ML accelerators and RISC-V CPU SoCs.
- Expertise in non-coherent and coherent (CHI) NoCs, D2D links, AXI/PCIe fabrics, and GDDR/LPDDR memory systems, backed by performance models correlated with RTL.
- Builds fabric generators, executable specs, and formal checks that turn architecture into verified RTL.

## CORE SKILLS

---

- **Architecture:** chiplet AI/ML accelerators and RISC-V CPU SoCs; non-coherent and coherent (CHI) NoCs, D2D, AXI/PCIe; routing, QoS, topology, RAS/debuggability
- **Memory Systems:** GDDR/LPDDR/DDR controllers, channel sizing, address interleaving, bandwidth analysis
- **Modeling:** cycle-accurate simulators and performance models correlated with RTL; workload characterization
- **Design Automation and Verification:** fabric RTL generators, executable specifications, formal verification, AI-assisted RTL

## EXPERIENCE

---

### Tenstorrent Inc. · Santa Clara, CA

May 2022 – Present

#### Staff Architect

Oct 2026 – Present

- Lead the in-house coherent fabric (CHI): transport, chip-to-chip links, and generators.
- Fabric lead for a RISC-V CPU chiplet SoC: integrate CPU clusters, the coherent fabric, IOMMU, timers, and an AI compute engine; firmware runs end to end in simulation.
- Lead AI-assisted hardware design: executable specs that generate golden models, RTL, and formal checks.
- Set performance-modeling direction across the AI NoC and the coherent CPU fabric.
- Mentor architects and interns; align modeling, RTL, and verification teams.

#### Senior Architect

May 2025 – Oct 2026

- Chiplet NoC specs: owned non-coherent NoC specs for IO and LPDDR chiplets (topology, address maps, interleaving, KPIs, LPDDR controller width).
- NoC RTL and debug: integrated trace, remap, and timeout features in an IO-chiplet fabric; fixed head-of-line blocking to meet bandwidth targets.
- GDDR chiplet: drove bring-up, performance analysis, and debug: bandwidth correlation, unaligned-access fixes, and telemetry and timeout checks.
- Third-party IP: ran design reviews and evaluations of commercial NoC and memory IP; wrapped vendor models for in-house simulation.
- Performance modeling: CPU multicore scaling and SLC policies, AI-accelerator models correlated with RTL, and a cycle-accurate NoC simulator.
- Fabric generators: built an AXI fabric generator adopted by SoC teams, and the in-house CHI fabric's generator and formal environment.
- Spec platform: built the internal service that builds, hosts, and reviews architecture specs (live editing, comments, AI review); teams use it for chip spec sign-off.
- RISC-V ACLINT: generated RTL from an executable spec and checked it formally against the golden model.
- Hiring and outreach: built interview frameworks, interviewed candidates, and supported university research collaborations.

- Studied multi-core, multi-chiplet fabric changes for performance and power; estimated on- and off-chip traffic for power planning.
- Defined cache and fabric verification testplans and built traffic-generation and simulation workflows.

- Contributed to architecture studies for a server-class RISC-V multicore CPU.
- Built a configurable fabric performance model for heterogeneous multicore systems.

## EDUCATION

---

### PhD in Computer Science and Engineering

2019 – 2025

University of Michigan, Ann Arbor, MI

- Thesis: *Streamlining High-Performance Heterogeneous Hardware Design*. Advised by Prof. Todd Austin.

### Graduate Certificate in Innovation and Entrepreneurship

2022 – 2025

University of Michigan, Ann Arbor, MI

### BSc in Computer Science with High Distinction

2016 – 2019

University of Michigan, Ann Arbor, MI

- GPA: 3.9/4.0

## SELECTED PUBLICATIONS

---

### Zipper: Latency-Tolerant Optimizations for High-Performance Buses

2025

ASP-DAC 2025

A bus protocol layer that hides latency; up to 8x speedup at 4.3% logic overhead.

### Security Verification of Low-Trust Architectures

2023

ACM CCS 2023 Distinguished Paper Award

Complete formal verification of a low-trust architecture, from the ISA down to RTL.

### Twine: A Chisel Extension for Component-Level Heterogeneous Design

2022

DATE 2022

Standard component interfaces and generated control logic; 3x smaller design specifications.

Full list of publications: <https://shibo-chen.github.io/publication/>

## SERVICE

---

- Program Committee Member, IEEE International Conference on Computer Design (ICCD 2025)
- Program Committee Member, Workshop on Computer Architecture Modeling and Simulation (CAMS 2025)

Last updated: October 2026