Hodjat Asghari Esfeden, Ph.D.

Senior Silicon & AI Engineer · Agentic Chip Design, TPU Silicon & Evaluation Benchmarks

Bay Area, CA · (951) 801-0376
hodjat91@gmail.com · hodjat.me
LinkedIn · Google Scholar

Senior Silicon Engineer at **Google** with **6+ years** contributing to multiple generations of **Google Tensor Processing Units (TPUs)**, cloud-scale EDA simulation infrastructure, and production LLM agent systems for hardware engineering. Builds autonomous **spec-to-silicon multi-agent flows**, simulation-in-the-loop debugging agents, and **execution-grounded evaluation benchmarks** over commercial EDA toolchains. Ph.D. in Computer Science with research in GPU microarchitecture, ML acceleration, and hardware security (**3x ISCA, 2x MICRO, ASPLOS, DAC**).

INDUSTRIAL EXPERIENCE

Google · Mountain View & Sunnyvale, CA — **Apr 2021 – Present**

Senior Silicon Engineer: Agentic Silicon & Benchmarks · Google Research (in partnership with AI & TPU Silicon) — 2026 – Present

- **Spec-to-Silicon Agents:** Building multi-agent flows spanning microarchitecture, RTL/CSR generation, UVM/formal verification, and synthesis/PPA against production silicon signoff gates.
- **Execution-Grounded Benchmarks:** Mining real production RTL and DV bugs and ground-truth fixes into multi-turn benchmarks scored via live EDA build, simulation, and verification checks (for model eval and RL).
- **Simulation-in-the-Loop Debug Agent:** Built an autonomous RTL/UVM debugging agent integrating live simulator CLI checkpointing, incremental re-elaboration, and waveform triage.

Senior Silicon Engineer: AI-Native Chip Design & DV · Google Cloud (AI & TPU Silicon) — 2024 – 2026

- **Autonomous Hardware Code Reviewer:** Built and deployed a multi-turn repo-aware RTL/UVM/Bazel review agent on live TPU repos using 7x Best-of-N sampling (80–92% cache hit rate), consensus filtering, and 100% AST syntax checks.
- **LLM Coverage Closure & Analytics:** Built a C++ Parquet/Polars coverage engine (**4x faster**, **sub-0.2s** queries across **11M+ bins** at zero EDA license cost) driving multi-agent directed stimulus generation.
- **Adaptive Seed Pruning:** Deployed a marginal-coverage and set-cover analyzer across cloud regressions to prune redundant simulation seeds, cutting compute and EDA license spend.

Silicon Engineer: TPU Verification & Cloud EDA Infrastructure · Google Cloud (AI & TPU Silicon) — 2021 – 2024

- **TPU Design Verification:** Owned functional and performance verification (SystemVerilog/UVM) for On-Chip Interconnects (OCI) and management fabrics across multiple **Google TPU** generations.
- **Cloud EDA Infrastructure:** Built runtime-aware regression schedulers, automated cloud RAM/CPU and EDA license optimization, RTL generator caching, and 3-step incremental simulation flows.
- **Execution & Leadership:** Owned **RTL Freeze** across multiple TPU programs; one of the main Google contributors to open-source `riscv-dv` (CHIPS Alliance) and lead of the Silicon US Internship Program (2022–Present).

Student Researcher & Hardware Intern: ML for Verification · Google (4 Terms) — 2019 – 2021

- Co-developed **CDG4CDG** with Google Research (**DVCon '22**), extracting coverage dependency graphs to adaptively steer constrained-random TPU verification stimulus.

TECHNICAL EXPERTISE

AI & Agentic Flows:	Multi-Agent Systems, Execution-Grounded Benchmarks, Tool-in-the-Loop Eval, Best-of-N, ML for EDA
Silicon Design & DV:	SystemVerilog, Verilog, UVM, SVA / Formal, SystemRDL, Coverage Closure, Interconnects, RISC-V, TPU/GPU
EDA & Infra:	Synopsys VCS / Verdi / UCAPI, Cadence Xcelium / Jasper, Synthesis & PPA, Verible AST, Bazel (HDL), Gerrit
Languages & Tools:	Python, C++, SystemVerilog, C, Polars / Parquet / SQL, DPI-C, Bash, Git, GPGPU-Sim, gem5

EDUCATION

University of California, Riverside · Ph.D. & M.Sc. in Computer Science (GPA: 3.91 / 4.00) — 2015 – 2021
Dissertation: *Enhanced Register Data-Flow Techniques for High-Performance, Energy-Efficient GPUs* (Advisor: Prof. Nael Abu-Ghazaleh)

Sharif University of Technology · B.Sc. in Computer Engineering · Tehran, Iran — 2010 – 2015

RESEARCH EXPERIENCE

Graduate Student Researcher: Computer Architecture & ML Systems · UC Riverside 2015 – 2021

- Researched GPU microarchitecture, compiler co-design, ML acceleration, and hardware security under Prof. Nael Abu-Ghazaleh (**8 peer-reviewed papers in ISCA, MICRO, ASPLOS, DAC, CAL; 2 US patent filings**).
- Co-designed task-based GPU execution and register-file architectures (**BlockMaestro, BOW, CORF, VPPS, RegMutex**), achieving up to 2.9x–6x workload speedups and 55% register-file dynamic energy savings.

Instructor: Computer Architecture (CS161) · UC Riverside 2018 – 2020

- Taught undergraduate processor design, pipelining, and cache/memory hierarchy; received the campus-wide **Distinguished Teaching Award (Honorable Mention)** and **Outstanding TA Award**.

Research Assistant: Computer Architecture · IPM & Sharif University of Technology 2013 – 2015

- Researched processor microarchitecture and network-on-chip interconnects at HPCAN Lab and IPM.

SELECTED PUBLICATIONS

ISCA '25 — **SpecASan: Mitigating Transient Execution Attacks Using Speculative Address Sanitization**
S. Ganjisaffar, E. M. Koruyeh, J. Zellmer, **H. Asghari Esfeden**, C. Song, N. Abu-Ghazaleh · *ACM/IEEE ISCA, 2025*

DVCon '22 — **Adaptive Test Generation for Fast Functional Coverage Closure**
A. Nazi, Q. Huang, H. Shojaei, **H. Asghari Esfeden**, A. Mirhosseini, R. Ho · *DVCon US, 2022*

ISCA '21 — **BlockMaestro: Enabling Programmer-Transparent Task-based Execution in GPU Systems**
A. Abdolrashidi, **H. Asghari Esfeden**, A. Jahanshahi, K. Singh, N. Abu-Ghazaleh, D. Wong · *ACM/IEEE ISCA, 2021*

MICRO '20 — **BOW: Breathing Operand Windows to Exploit Bypassing in GPUs**
H. Asghari Esfeden, A. Abdolrashidi, S. Rahman, D. Wong, N. Abu-Ghazaleh · *IEEE/ACM MICRO, 2020*

IEEE CAL '19 — **Locality-aware GPU Register File**
H. Jeon, **H. Asghari Esfeden**, N. Abu-Ghazaleh, D. Wong, S. Elango · *IEEE Computer Architecture Letters, 2019*

ASPLOS '19 — **CORF: Coalescing Operand Register File for GPUs**
H. Asghari Esfeden, F. Khorasani, H. Jeon, D. Wong, N. Abu-Ghazaleh · *ACM ASPLOS, 2019*

MICRO '18 — **In-Register Parameter Caching for Dynamic Neural Nets with Virtual Persistent Processor Specialization**
F. Khorasani, **H. Asghari Esfeden**, N. Abu-Ghazaleh, V. Sarkar · *IEEE/ACM MICRO, 2018*

ISCA '18 — **RegMutex: Inter-Warp GPU Register Time-Sharing**
F. Khorasani, **H. Asghari Esfeden**, A. Farmahini-Farahani, N. Jayasena, V. Sarkar · *ACM/IEEE ISCA, 2018*

DAC '17 — **RIC: Relaxed Inclusion Caches for Mitigating LLC Side-Channel Attacks**
M. Kayaalp, K. N. Khasawneh, **H. Asghari Esfeden**, J. Elwell, N. Abu-Ghazaleh, D. Ponomarev, A. Jaleel · *ACM/IEEE DAC, 2017*

PATENTS

Issued '25 — **Coalescing Operand Register File for Graphical Processing Units**
N. Abu-Ghazaleh, **H. Asghari Esfeden**, D. Wong · *Granted U.S. Patent 12,417,083, 2025*

Pending '23 — **Breathing Operand Windows to Exploit Bypassing in Graphics Processing Units**
H. Asghari Esfeden, N. Abu-Ghazaleh · *Published U.S. Patent Application 2023/0393850 A1 (App. 18/032,157)*

HONORS & SERVICE

- **Distinguished Teaching Award (Honorable Mention) & Outstanding TA Award**, UC Riverside (2020): Top 2 campus-wide.
- **Dean's Distinguished Fellowship**, UC Riverside (2015); **Ranked 17th / 100,000+** in Iran's National University Entrance Exam (2010).
- **External Review Committee / Reviewer:** Design Automation Conference (DAC), IEEE Trans. on Computers, ACM TECS, HASP.