

# From Pragmas to Partners: A Symbiotic Evolution of Agentic High-Level Synthesis

Niansong Zhang  
Cornell University  
USA

Sunwoo Kim  
Cornell University  
USA

Shreesha Srinath  
Cerebras Systems  
USA

Zhiru Zhang  
Cornell University  
USA

## Abstract

The rise of large language models has sparked interest in AI-driven hardware design, raising the question: does high-level synthesis (HLS) still matter in the agentic era? We argue that HLS remains essential. While we expect mature agentic hardware systems to leverage both HLS and RTL, this paper focuses on HLS and its role in enabling agentic optimization. HLS offers faster iteration cycles, portability, and design permutability that make it a natural layer for agentic optimization. This position paper makes three contributions. First, we explain why HLS serves as a practical abstraction layer and a golden reference for agentic hardware design. Second, we identify key limitations of current HLS tools, namely inadequate performance feedback, rigid interfaces, and limited debuggability that agents are uniquely positioned to address. Third, we propose a taxonomy for the symbiotic evolution of agentic HLS, clarifying how responsibility shifts from human designers to AI agents as systems advance from copilots to autonomous design partners.

## 1 Introduction

High-level synthesis (HLS) lifts accelerator design from RTL to high-level software specification, but achieving high performance still requires substantial manual, iterative, expertise-driven optimization: restructuring code, rewriting kernels, and adding pragmas guided by toolchain feedback to shape compute, memory, data layout, and communication.

The rise of large language models and agentic AI has sparked renewed interest in automating this process. A natural question arises: does HLS still matter, or will AI simply generate RTL directly? We argue that HLS remains not only relevant but essential in the agentic era. We do not claim that HLS alone or RTL alone solves agentic hardware design. We expect a mixed workflow where agents move between high-level intent, HLS source, and RTL, depending on the task. This paper focuses on the role of HLS in that workflow.

This position paper makes three contributions: (1) it argues that HLS is a strong abstraction layer for agentic hardware design, enabling iteration speed, portability, and design permutability beyond pure RTL generation (§2); (2) it highlights limitations in current HLS workflows, especially performance feedback, interfaces, verification, and debugging, that agents can help address (§3); and (3) it introduces an autonomy-inspired taxonomy for the symbiotic evolution of agentic HLS, clarifying capabilities and shifting human-agent responsibilities (§4) [3].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

LATTE '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

## 2 Why HLS Remains Relevant

Despite advances in AI-driven RTL generation, HLS offers fundamental advantages that make it the natural abstraction layer for agentic hardware design.

*Rapid Design Space Exploration.* Agentic workflows require many iterations to test correctness and reach design goals/constraints. HLS enables rapid design space exploration (DSE) through two mechanisms. First, it provides an executable high-level software specification that serves as a golden reference, allowing for fast software simulation before hardware generation. Second, HLS representations are highly permutable: one design point can be easily transformed into another while preserving functional semantics. In contrast, direct RTL generation lacks a clear golden reference and equivalent RTL transformations require invasive rewrites that are harder to automate and more error-prone. Operating primarily at RTL also raises the cost of iteration. RTL simulation is slower, and code and traces are more verbose, which increases the token cost for agentic loops that depend on repeated inspections and refinements. Higher abstraction is therefore even more important when optimization is driven by many agentic iterations.

*Portability.* The same high-level software specification can serve as a single source to target different FPGA families, vendors, or accelerator platforms with minimal modification. Beyond platform portability, HLS enables constraint portability: the same source can be synthesized for high throughput, low latency, or low power by adjusting directives and synthesis settings. For agentic HLS systems, learned optimizations and design patterns can be transferred across platforms and constraints, enabling reuse across projects.

*Verification and Debugging.* HLS improves verification by centering the workflow around a readable and executable specification. The high-level source serves as a golden reference for functional validation and supports fast regression testing before synthesis. This reduces the burden of debugging in RTL while still allowing RTL checks when needed.

## 3 HLS Limitations That Agents Can Address

Despite rapid DSE, portability, and ease of verification being central to the value proposition of HLS, current tools fall short in all three. We identify key gaps that agents are uniquely positioned to address.

*Performance Feedback Limits DSE.* Fast and meaningful performance feedback is crucial for effective DSE, yet current HLS tools fall short. Synthesis reports provide latency estimates, resource utilization, and scheduling information, but these are often difficult to interpret and relate back to source code. Worse, many designs exhibit data-dependent control flow, causing tools to report

**Table 1: Autonomy-inspired levels for agentic HLS.**

Level	Name	Human Effort	AI Effort	Key Technology
L0	Manual DSE	High	Low	Standard HLS compilation and reports; manual code and pragma iteration; feedback and debugging largely human-driven.
L1	HLS Copilot	High	Low	Explains HLS reports with source attribution, suggests local code and pragma edits, helps triage failures, and improves RTL readability for follow-up debugging.
L2	Autotuning Agent	Medium	Medium	Runs closed-loop DSE over pre-defined search space using tool feedback, integrates HLS IPs using standard interfaces, generates lightweight replayable traces for performance analysis, and performs basic differential testing.
L3	Guardrailed Agents	Medium	High	Performs multi-step tool use with strong guardrails, leverages verification to provide strong correctness guarantee, including equivalence checking, assertion synthesis, and counterexample explanations; constructs mixed-fidelity performance model, performs localized failure and bottleneck analysis for DSE.
L4	Domain Architect	Low	High	Handles system-level optimization and integration with customized interfaces, use complete workloads to drive performance and correctness tests.
L5	Silicon Partner	Low	High	End-to-end autonomy across architectures and tools; continual learning across projects; improves the HLS stack itself (compiler passes, scheduling heuristics, codegen, tracing and verification infrastructure).

unknown or variable cycle counts that provide little actionable guidance. Agents can address this through mixed-fidelity performance modeling: extracting operation schedule, loop-level information, and control-flow regions from HLS artifacts to construct performance estimates and replayable traces. This enables actionable feedback even for designs where traditional HLS reports provide only question-mark estimates, presenting insights at the same abstraction level where designers make optimization decisions.

*Rigid Interface Limits Portability.* A major pain point is composing accelerators with the surrounding system. Tools such as Vitis HLS [1] support standard interfaces such as AXI, but offer limited ability to customize or extend these protocols. Integrating HLS-generated IP with systems built using different frameworks, such as the LiteX SoC builder [4], requires manual writing interface adapters and bridge logic. Tools such as Catapult HLS [8] support more flexible interfaces, but defining protocols remains a complex task. This *composability gap* undermines the portability promise of HLS. AI agents can bridge this gap by understanding both the HLS design intent and the target system’s interface requirements, then generating the necessary adapter logic and protocol bridges.

*Lack of End-to-End Verification and Debugging.* HLS verification today is still split between high-level software simulation, RTL simulation, and ad hoc debugging, with limited help connecting failures back to source intent. Agents can use HLS as a golden reference for both humans and RTL, enabling differential testing, equivalence checking, assertion synthesis, and targeted counterexample explanation across abstraction levels. For many workloads, we also want testbenches written at an even higher level, such as PyTorch models, to drive software simulation and RTL co-simulation with realistic inputs and oracle outputs. This integration is not available in current flows, but an agent could mediate it by generating adapters, harnesses, and simulation drivers. Finally, agents can improve the readability of the generated RTL by refactoring the structure, naming signals, and adding documentation, which reduces the cost of debugging when designers must resort to RTL.

## 4 Levels of Automation

Recent work explores LLM assistance for HLS, including code generation and guidance [5, 10], automated repair [12], retrieval-augmented optimization [11], software-to-hardware bridging [2], and agentic DSE and multi-agent design [6, 7, 9]. These systems span a wide range of autonomy and human involvement. To characterize this spectrum of autonomy and human involvement, we propose the taxonomy in Table 1. The taxonomy is useful not just for labeling systems, but for highlighting what must improve in agentic HLS to enable greater autonomy, especially actionable performance feedback, flexible interface, improved verification, and debuggability.

At **L0 (Manual DSE)**, designers iterate manually from HLS reports. **L1 (HLS Copilot)** assists with report interpretation and localized edit suggestions, but humans still implement and validate. **L2 (Autotuning Agent)** runs closed-loop optimization over a predefined space and automates routine evaluation and integration. **L3 (Guardrailed Agents)** performs multi-step tool use with strong guardrails, with verification and mixed-fidelity feedback for targeted debugging and DSE. **L4 (Domain Architect)** shifts humans to intent definition and review, while agents drive system-level optimization and workload-level validation within a bounded domain. Finally, **L5 (Silicon Partner)** extends autonomy across architectures and tools, including improving the HLS stack itself.

Current systems operate primarily at L1–L2, with research pushing toward L3. Advancing further depends directly on addressing the composability and feedback limitations discussed in §3.

## 5 Conclusion

We argue that HLS will remain a valuable layer in the agentic era, supporting faster iteration, portability, and design exploration. We highlight performance feedback, rigid interfaces, and debuggability as key gaps in today’s HLS workflows where agents can help, and propose a taxonomy of automation levels to structure the space. We hope that this paper encourages joint progress between the HLS and AI communities.

## References

- [1] AMD. 2024. *Vitis High-Level Synthesis User Guide*. AMD. <https://docs.amd.com/en-US/ug1399-vitis-hls>
- [2] Luca Collini, Siddharth Garg, and Ramesh Karri. 2025. C2HLSC: Leveraging Large Language Models to Bridge the Software-to-Hardware Design Gap. *ACM Trans. Design Autom. Electr. Syst.* 30, 6 (2025), 96:1–96:24. <https://doi.org/10.1145/3734524>
- [3] SAE International. 2021. Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. (April 2021). [https://www.sae.org/standards/j3016\\_202104-taxonomy-definitions-terms-related-driving-automation-systems-road-motor-vehicles](https://www.sae.org/standards/j3016_202104-taxonomy-definitions-terms-related-driving-automation-systems-road-motor-vehicles) Defines Levels 0–5 of driving automation and related terminology.
- [4] Florent Kermarrec, Sébastien Bourdeauducq, Jean-Christophe Le Lann, and Hannah Music. 2020. LiteX: an open-source SoC builder and library based on Migen Python DSL. (2020). <https://arxiv.org/abs/2005.02506>
- [5] Yuchao Liao, Tosiron Adegbija, and Roman Lysecky. 2024. Are LLMs Any Good for High-Level Synthesis?. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design (ICCAD 2024)*. Association for Computing Machinery, New York, NY, USA, 29:1–29:8. <https://doi.org/10.1145/3676536.3699507>
- [6] Ali Emre Oztas and Mahdi Jelodari. 2024. Agentic-HLS: An agentic reasoning based high-level synthesis system using large language models. AI for EDA Workshop 2024 (co-located with NeurIPS 2024), arXiv preprint. (2024). <https://doi.org/10.48550/arXiv.2412.01604>
- [7] Seyed Arash Sheikholeslam and Andre Ivanov. 2024. SynthAI: A Multi Agent Generative AI Framework for Automated Modular HLS Design Generation. arXiv preprint. (2024). <https://doi.org/10.48550/arXiv.2405.16072>
- [8] Siemens EDA. 2024. Catapult HLS. (2024). <https://eda.sw.siemens.com/en-US/ic/catapult-high-level-synthesis/>
- [9] Hanyu Wang, Xinrui Wu, Zijian Ding, Su Zheng, Chengyue Wang, Neha Prakriya, Tony Nowatzki, Yizhou Sun, and Jason Cong. 2025. LLM-DSE: Searching Accelerator Parameters with LLM Agents. arXiv preprint. (2025). <https://doi.org/10.48550/arXiv.2505.12188>
- [10] Chenwei Xiong, Cheng Liu, Huawei Li, and Xiaowei Li. 2024. HLLSPilot: LLM-based High-Level Synthesis. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design (ICCAD 2024)*. Association for Computing Machinery, New York, NY, USA, 226:1–226:9. <https://doi.org/10.1145/3676536.3676781>
- [11] Haocheng Xu, Haotian Hu, and Sitao Huang. 2024. Optimizing High-Level Synthesis Designs with Retrieval-Augmented Large Language Models. In *2024 IEEE LLM Aided Design Workshop (LAD)*. IEEE, Piscataway, NJ, USA, 1–5. <https://doi.org/10.1109/LAD62341.2024.10691855> Preprint: arXiv:2410.07356.
- [12] Kangwei Xu, Grace Li Zhang, Xunzhao Yin, Cheng Zhuo, Ulf Schlichtmann, and Bing Li. 2024. Automated C/C++ Program Repair for High-Level Synthesis via Large Language Models. In *Proceedings of the 2024 ACM/IEEE International Symposium on Machine Learning for CAD (MLCAD 2024)*. Association for Computing Machinery, New York, NY, USA, 15:1–15:9. <https://doi.org/10.1145/3670474.3685953> Preprint: arXiv:2407.03889.