

# AMS-IO-Bench and AMS-IO-Agent: Benchmarking and Structured Reasoning for Analog and Mixed-Signal Integrated Circuit Input/Output Design

Zhishuai Zhang<sup>\*1†</sup>, Xintian Li<sup>\*2</sup>, Shilong Liu<sup>3</sup>, Aodong Zhang<sup>1</sup>, Lu Jie<sup>2</sup>, Nan Sun<sup>1</sup>

<sup>1</sup>Department of Electrical Engineering, Tsinghua University

<sup>2</sup>School of Integrated Circuits, Tsinghua University

<sup>3</sup>Princeton AI Lab, Princeton University

arcadia16zsz@gmail.com

## Abstract

In this paper, we propose *AMS-IO-Agent*, a domain-specialized LLM-based agent for structure-aware input/output (I/O) subsystem generation in analog and mixed-signal (AMS) integrated circuits (ICs). The central contribution of this work is a framework that connects natural language design intent with industrial-level AMS IC design deliverables. *AMS-IO-Agent* integrates two key capabilities: (1) a structured domain knowledge base that captures reusable constraints and design conventions; (2) design intent structuring, which converts ambiguous user intent into verifiable logic steps using JSON and Python as intermediate formats. We further introduce AMS-IO-Bench, a benchmark for wirebond-packaged AMS I/O ring automation. On this benchmark, AMS-IO-Agent achieves over 70% DRC+LVS pass rate and reduces design turnaround time from hours to minutes, outperforming the baseline LLM. Furthermore, an agent-generated I/O ring was fabricated and validated in a 28 nm CMOS tape-out, demonstrating the practical effectiveness of the approach in real AMS IC design flows. To our knowledge, this is the first reported human-agent collaborative AMS IC design in which an LLM-based agent completes a nontrivial subtask with outputs directly used in silicon.

**Code** — <https://github.com/Arcadia-1/AMS-IO-Agent>

**Datasets** — <https://github.com/Arcadia-1/AMS-IO-Bench>

## Introduction

Input/output (I/O) subsystems are a fundamental component of analog and mixed-signal (AMS) integrated circuits (ICs), providing signal interfacing, power delivery, and electrostatic discharge (ESD) protection. While digital I/O cells can typically be placed using scripts and routed via standard digital flows, the implementation of AMS I/O remains largely manual due to intricate, project-specific requirements. These include diverse signal types, multiple power domains, power integrity constraints, sensitive analog signal routing requirements, and layout restrictions imposed by fabrication and packaging rules. Simple scripting approaches lack the reasoning capability to handle such complexity, leaving much of the design effort to human engineers.

<sup>\*</sup>These authors contributed equally.

<sup>†</sup>Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

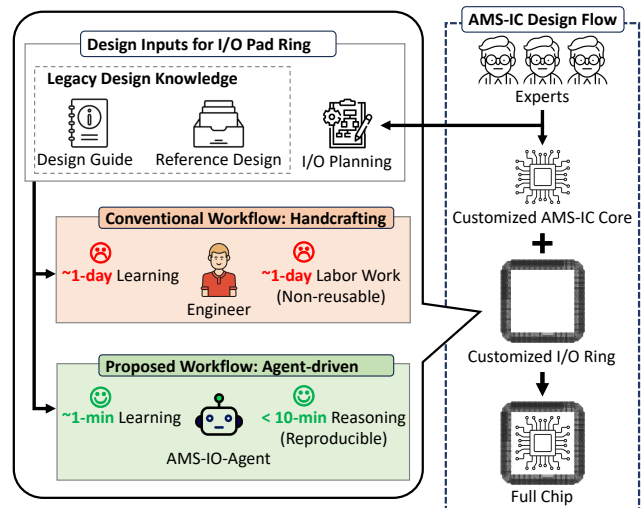


Figure 1: Comparison of conventional and agent-driven I/O ring design in a real AMS IC design workflow.

As a result, AMS I/O design is labor-intensive, with most effort rarely reusable. In wirebond-packaged chips (Fig. 1), a novice engineer may spend one or two days on studying, manually assembling, and verifying I/O placement and connections. Iterative pin changes often lead to disruptive rework risks, which intensify as tape-out approaches. These challenges highlight the need for intelligent automation, for which recent advances in large language models (LLMs) offer a promising foundation.

While LLMs have shown promising capabilities across several hardware design tasks (Chen et al. 2024; Fang et al. 2025), their application to AMS I/O design remains underexplored due to three key challenges: (1) the lack of accessible domain knowledge, which is typically confined to team-specific practices and scattered internal documents, (2) the absence of standardized task interfaces, as interaction still relies on GUI or domain-specific languages unfamiliar to pretrained LLMs, and (3) the unavailability of public benchmarks, which hinders systematic evaluation.

To bridge this gap, we propose AMS-IO-Agent, a domain-specialized LLM-based agent for structure-aware AMS I/O generation. It integrates two core capabilities: (1) a domain

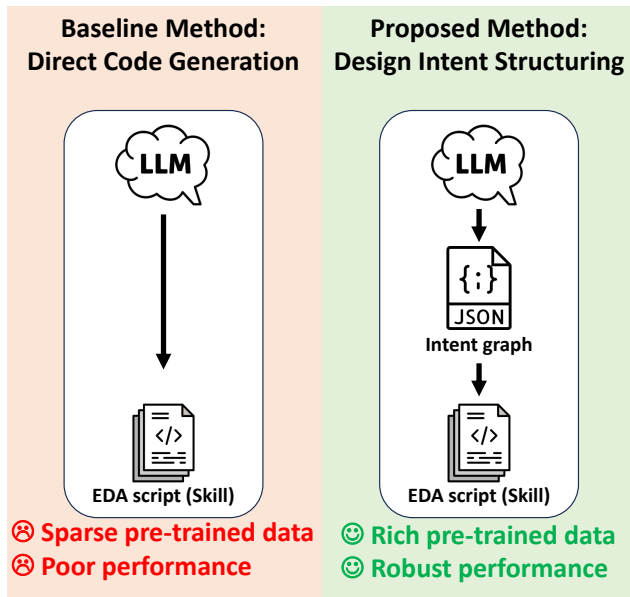


Figure 2: Comparison of baseline direct code generation and the proposed structured design intent.

knowledge base built from fragmented engineering practices, capturing reusable constraints and layout conventions, curated from real training materials developed by a professional AMS IC design team of more than 10 engineers and validated through over 50 successful tape-out cases; (2) design intent structuring, which converts ambiguous design intent into verifiable logic steps as an intermediate format, as shown in Fig. 2. Together, these components allow the agent to adapt across projects by grounding generation in prior examples and domain constraints.

To support consistent evaluation, we also introduce AMS-IO-Bench, a benchmark for wirebond-packaged AMS I/O rings. It covers key tasks such as I/O ring assembly and validation, and is used to assess the correctness, adaptability, and efficiency of generated designs. To the best of our knowledge, this is the first work to propose an LLM-based agent and comprehensive evaluation benchmark specifically targeting real-world AMS IC I/O ring design automation, and the first LLM-based agent for AMS IC design capable of direct integration into tape-out workflows as a key module collaborating with human engineers.

Experiments on AMS-IO-Bench show that AMS-IO-Agent achieves over 70% DRC+LVS pass rates, surpassing the baseline LLM by a large margin and reducing design turnaround time from days to just minutes per case. Notably, we validated our system in real industrial tape-out projects: I/O rings generated by AMS-IO-Agent were seamlessly integrated into commercial AMS IC flows and successfully fabricated on silicon. This demonstrates not only the practicality and robustness of our approach, but also its readiness for deployment in real-world chip design pipelines.

In summary, our work makes the following contributions:

- We propose an automatic pipeline for AMS IC I/O gen-

eration, formalizing the task into structured steps: intent interpretation, constraint resolution, and generation of EDA scripts, enabling practical automation.

- We introduce a novel agent architecture that combines domain-specific knowledge base and structured intent reasoning, supporting generalization across diverse AMS I/O design contexts.
- We develop a benchmark for wirebond-packaged AMS I/O ring automation, and show that the agent approach consistently delivers practical designs, reducing manual workload and design turnaround time.
- We further validate the agent in a real 28-nm CMOS tape-out, achieving the first reported demonstration of an LLM-based agent directly contributing to a nontrivial AMS IC design task.

## Related Work

### LLM-based Approaches for IC Design Automation

LLMs have been applied to a variety of EDA tasks, from knowledge-based question answering (Shi et al. 2024; Skelic et al. 2025) and register-transfer level (RTL) code generation (Chang et al. 2023; Blocklove et al. 2023; Thakur et al. 2023; Liu et al. 2023; Xu et al. 2024; Liu et al. 2024b; Tsai, Liu, and Ren 2024; Thakur et al. 2024; Fu et al. 2025) through netlist synthesis (Lai et al. 2024), circuit optimization (Yin et al. 2024; Ghose et al. 2025), to EDA script generation (Chen et al. 2025). To extend beyond one-shot code synthesis, recent studies adopt agent architectures that integrate LLMs into iterative design loops (Wu et al. 2024; Ho and Ren 2024; Liu et al. 2025, 2024a; Ghose et al. 2025). By invoking tools, observing results, and refining their outputs, these agents show promise for multi-step workflows.

Nevertheless, these advances have not alleviated the repetitive, constraint-heavy routines in AMS IC design, as they lack systematic adaptation to real design flows and tape-out validation. As a result, constraint-rich, project-specific tasks remain largely manual, and the potential of LLMs to automate them is still untapped.

### Benchmark for AMS IC agents

The scarcity of domain-specific datasets and the inherent difficulty of precisely defining AMS IC design tasks make quantitative evaluation particularly challenging. For relatively straightforward tasks involving only a few to a few dozen operation sequences, datasets can scale to over one thousand instances (Liu et al. 2025). In contrast, for complex AMS IC tasks such as netlist or testbench generation, existing datasets contain only a few dozen highly intricate examples (Table 1). This highlights the need for benchmarks that closely reflect practical AMS design requirements, consistent with how practicing chip engineers learn from a limited number of carefully curated examples.

### Automation in I/O Ring Assembly

Existing automation methods for I/O or padding construction require designers to perform extensive low-level specification through detailed configuration tables or metadata files

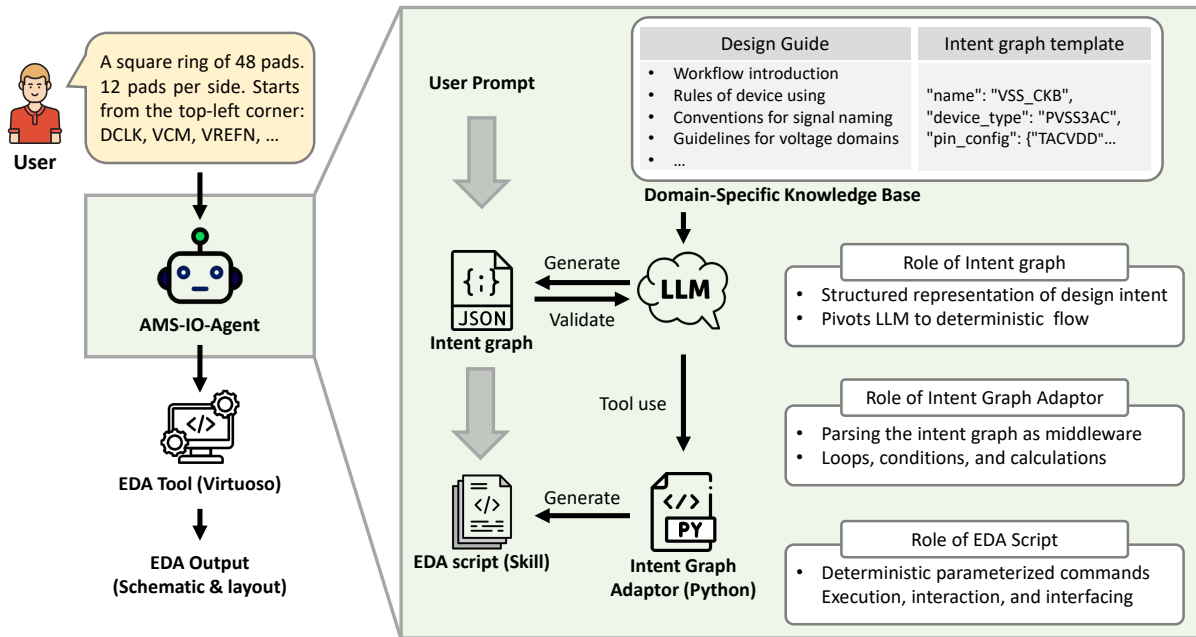


Figure 3: Overview of the proposed AMS-IO-Agent architecture, illustrating the workflow and the component roles.

Task	Data Size
Netlist Generation (Ho and Ren 2024)	17
Netlist Generation (Chen et al. 2025)	24
AMS-IO-Bench (Ours)	30

Table 1: Data sizes used in recent LLM-based agent studies for complex IC design tasks.

(Moseley, Barzic, and contributors 2025; Chen et al. 2021; Morita, Takiguti, and Noije 2015; Chen et al. 2010), making these workflows nearly as labor-intensive as manual padding assembly. These approaches are largely tailored for digital SoCs and offer limited support for AMS-specific needs such as power domain separation and custom analog pad cells; accommodating such constraints often requires direct modification of low-level scripts or tool code.

None of these methods incorporates semantic interpretation or design-intent understanding, relying instead on explicitly specified parameters. Their outputs are also limited to geometric files such as GDS that cannot be directly edited by designers. Consequently, these methods operate at a different abstraction level and do not offer a directly comparable baseline for AMS I/O ring assembly.

## Method

### Task Definition

AMS I/O generation is defined as transforming human-provided pin planning specifications into production-ready schematics and layouts in EDA tools, as a constraint-driven engineering task focused on fulfilling design intent and ensuring compliance with all required design rules rather than performance or area optimization.

The input consists of pin planning specifications, typically expressed as tables or textual descriptions, defining I/O ring dimensions, pin names and ordering, and possibly additional requirements in natural language such as power domain separation or custom device usage.

The output is a complete set of production-ready schematics and layouts that can be directly integrated into the AMS IC design flow, as illustrated in Fig. 3. The generated results must conform to the intended geometry and pin arrangement, maintain strict electrical correspondence verified by Layout Versus Schematic (LVS) checks, and satisfy all physical design rules and manufacturability requirements verified by Design Rule Check (DRC).

### Agent Overview

As illustrated in Fig. 3, AMS-IO-Agent is an LLM-based agent that bridges natural language design intent and executable EDA scripts through a structured pipeline of structured intent graph and intent graph adaptor. The agent consists of three key components:

- **Design Intent Structuring:** Converts natural language or semi-structured specifications, such as pin lists with textual constraints, into a machine-readable graph that explicitly represents device configurations, spatial relationships, and electrical connections.
- **Intent graph adaptor:** Parses the intent graph to resolve constraints, perform geometric calculations, and derive implementation parameters, transforming structured intent into executable procedures.
- **Domain-specific knowledge base:** Provides design rules, device specifications, and layout conventions for constraint checking and design consistency with AMS I/O practices.

This layered architecture separates high-level intent reasoning from low-level implementation, allowing the LLM to focus on understanding design intent while deterministic modules ensure constraint resolution and reproducible script generation.

### Design Intent Structuring

AMS I/O design often begins with informal specifications such as natural language descriptions, pin lists, or semi-structured spreadsheets, which contain essential design information but lack the structure required for automated processing. To address this gap, AMS-IO-Agent converts these inputs into a standardized *intent graph*.

The intent graph is a JSON-based representation that models the I/O ring as a sequence of interconnected nodes, each representing a pad or corner cell with attributes such as name, device type, spatial position, direction, and pin connections (Figure 4). Position encoding follows the I/O ring layout, preserving its physical organization.

The construction of the intent graph combines explicit completion and implicit inference. Explicit completion is applied to pins whose names are provided in the specification. In practice, designers often use conventional abbreviations such as DCLK (digital clock), VCM (common-mode voltage) and VREFN (N-side of reference voltage). By leveraging these patterns through the domain-specific knowledge base, the agent can infer additional attributes, including signal type, device type, default direction, and pin connections. These connections typically do not need to be specified explicitly by the user, although explicit user overrides are supported when required. Implicit inference is used for elements that are not mentioned at all, such as corner cells, which are automatically inserted based on standard design conventions without any user input.

This representation is fundamentally different from a netlist and serves as the basis for middleware processing. While a netlist only describes circuit connectivity, the intent graph also captures spatial relationships, semantic context, and domain knowledge. It can be directly understood by both humans and language models, and it can be efficiently parsed and processed by code. This dual accessibility makes it an effective interface between informal design inputs and automated implementation.

### Intent Graph Adaptor

The Intent Graph Adaptor serves as a middleware layer that bridges the LLM-based agent and commercial EDA tools. It supports the structured intent graph in completing EDA code generation and addresses the limitations of direct code generation for AMS I/O design.

Directly generating domain-specific language (DSL) scripts such as SKILL with an LLM is ineffective for producing correct and consistent outputs due to the lack of sufficient training data. Likewise, using these DSLs themselves to perform logical operations for parsing the intent graph is impractical because such languages are not designed for complex data manipulation. To overcome these issues, the adaptor employs Python-based middleware for deterministic processing, including structured data parsing, constraint

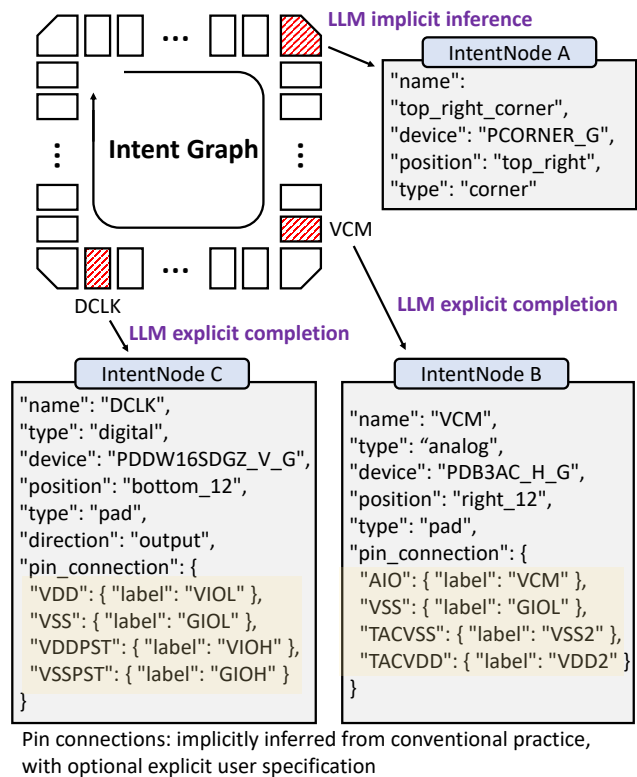


Figure 4: Illustration of the intent graph and examples of its nodes.

resolution, geometric calculations, and DSL script generation.

Although the middleware tools could, in principle, be generated by the LLM for every run, this approach suffers from low efficiency and poor robustness. Instead, implementing them as a reusable tool library allows the agent to invoke deterministic scripts for structured data parsing, constraint resolution, geometric calculations, and DSL script generation. Upon receiving the intent graph, it extracts I/O instances with their attributes and computes precise cell coordinates based on I/O design rules.

For integration with commercial EDA tools, the adaptor generates SKILL scripts to create schematics and layouts in Cadence Virtuoso and csh scripts to invoke Siemens Calibre verification tools. Together, these components form a reproducible and verifiable bridge between the LLM-generated intent graph and tape-out-ready AMS I/O design flows.

### Domain-specific Knowledge Base

The domain-specific knowledge base addresses a key challenge in AMS I/O design: design expertise is highly fragmented and often buried in informal natural-language documents such as training notes, reference guides, and internal design manuals. These materials encompass device selection practices, layout conventions, power domain rules, ESD protection requirements, naming conventions, common design techniques, and practical “know-how” accumulated

from real projects.

To consolidate this expertise, AMS-IO-Agent employs a knowledge base built from the training materials and documented practices of a professional AMS IC design team of more than 10 experienced engineers, validated through over 50 successful tape-out cases. This knowledge base reflects the same materials used to train entry-level engineers, which can typically be mastered by engineers with an undergraduate background within one to three days, ensuring both its practicality and accessibility.

Instead of converting this knowledge into rigid rule code, AMS-IO-Agent organizes it into a lightweight repository of approximately 6k tokens. This repository serves both as a reference for human designers and as an in-context knowledge source for the LLM. Because of its compact size, the LLM can directly consume it without retrieval mechanisms or model fine-tuning.

By integrating verified expertise from actual design workflows into a single accessible knowledge base, the system enables both human and machine users to leverage the same authoritative engineering knowledge, combining human-readable documentation with automated, context-driven reasoning during intent interpretation and constraint resolution.

## Benchmark

There is currently no standard evaluation framework or publicly available dataset for AMS I/O generation. To address this gap, we develop **AMS-IO-Bench**, the first benchmark suite dedicated to I/O ring generation for wirebond-packaged AMS IC chips. Although other packaging methods such as flip-chip with solder bumps exist, wirebond remains the most commonly used approach for prototype verification chips. Because wirebond requires peripheral I/O rings along the chip boundary, it provides a clear and well-defined context for evaluating I/O planning, making it the most practical choice for our benchmark.

AMS-IO-Bench is derived from the I/O planning of 10 real tape-out projects collected over the past 5 years. From these projects, we construct a benchmark of 30 cases by simplifying, augmenting, and transforming the original designs, ensuring that each case preserves the core constraints and design patterns of production flows. This approach enables both realism and reproducibility. Each benchmark instance provides a structured pad location list specifying signal assignments, power domains, and routing hints.

The benchmark is organized into three difficulty levels, as shown in table 2.

Level	Number	Realism	Features
Simple	10	Simplified	Small size, single domain
Medium	10	Default	Standard size, multiple domains
Hard	10	Complex	Large size, staggered, customized cells

Table 2: Three difficulty levels in AMS-IO-Bench.

- **Simple.** Simple cases use a single signal domain, thereby removing the need for implicit reasoning about isolation and local ESD power supply. They are derived from typical AMS IC designs through simplification.
- **Medium.** Medium-difficulty cases reflect the default complexity of typical AMS IC designs. Each case corresponds to a standard MPW chip outline of approximately 1mm × 1mm with a single-row I/O ring. The I/O ring is partitioned into multiple power domains, including digital and analog, each governed by distinct device types, design rules, filler cells, and isolation cells. Solving these cases requires the agent to reason over domain knowledge and contextual constraints, closely matching the real-world challenges encountered in tape-out projects.
- **Hard.** Hard cases represent advanced scenarios observed in complex tape-outs. They include dual-row or partially dual-row I/O rings (staggered pads), chips with enlarged outlines (1.5× to 2× the default size or larger), custom I/O cells (e.g., analog I/O cells with reduced ESD capacitance), and designs with highly specialized power domain configurations such as localized ESD power delivery. These cases stress the adaptability of agents to highly customized design requirements.

By focusing on real-world constraints and vertical domain specificity, AMS-IO-Bench establishes a practical and scalable platform for evaluating AMS I/O design agents under realistic chip development conditions.

## Experiments

### Setup

We implement and evaluate the proposed AMS-IO-Agent using the Python-based automation framework smolagents. The backbone language model is accessed via API. The agent runs on a workstation and communicates with chip-design server workstations via SSH and socket connections.

The tools invoked by the agent are implemented in Python. For integration with commercial EDA toolchains, the agent receives design intent descriptions as input, generates SKILL scripts for schematic and layout creation, and executes them in Cadence Virtuoso. Layout verification and physical rule checks are performed by invoking Siemens Calibre through csh scripts. The evaluation follows the AMS-IO-Benchmark methodology described earlier.

### Evaluation Metrics

We evaluate AMS-IO-Agent using a five-stage pipeline covering the entire process from intent interpretation to layout verification. Each metric targets a distinct stage, enabling systematic diagnosis and quantitative quality assessment.

**Metric 1: Intent Graph Pass Rate** checks whether the agent correctly converts natural language specifications into valid Intent Graphs, including proper pad naming, device type assignment, and attribute completion. Failures indicate misunderstanding of design requirements.

**Metric 2: Shape Score** quantifies layout similarity using a Vision-Language Model (VLM), which performs a binary evaluation (pass or fail) based on structural alignment with

Method	IG (%)	Shape (%)	DRC (%)	LVS (%)	DRC+LVS (%)	Time (min)	Token (k)
Human	100	100	100	100	100	≈ 480	–
LLM (GPT-4o)	0	0	0	0	0	0.2	1k
AMS-IO-Agent (GPT-4o)	100	100	76.67	66.67	63.33	4.1	160k
AMS-IO-Agent (Claude-3.7)	100	100	93.33	76.67	76.67	4.2	96k
AMS-IO-Agent (DeepSeek-V3)	100	100	93.33	76.67	76.67	5.1	105k

Table 3: AMS-IO-Agent achieves perfect intent translation and shape validity, high DRC/LVS pass rates, and reduces design time from hours to minutes, outperforming a baseline LLM while remaining far more efficient than expert manual design.

KB	IG	Adaptor	IG (%)	Shape (%)	DRC (%)	LVS (%)	DRC+LVS (%)	Time (min)	Token (k)
×	×	×	–	0	0	0	0	0.2	1
✓	×	× <sup>†</sup>	–	0	0	0	0	16.2	1098
✓	×	✓ <sup>††</sup>	–	100	20.00	0	0	27.4	2036
✓	✓	×	100	0	0	0	0	2.3 <sup>†††</sup>	18
✓	✓	✓	100	100	93.33	76.67	76.67	5.1	105

<sup>†</sup> LLM directly generates SKILL code without structured intent reasoning.

<sup>††</sup> LLM generates Python code for SKILL generation without structured intent reasoning.

<sup>†††</sup> Reduced runtime because LVS/DRC tools cannot be invoked.

Table 4: Ablation results using DeepSeek-V3 show that both the knowledge base and structured intent adaptor are essential: removing either component eliminates signoff correctness.

the reference. It captures visual and topological correctness beyond rule-based evaluation.

**Metric 3: DRC Pass Rate** reports the percentage of layouts passing all foundry-specified design rules (e.g., spacing, width, enclosure), reflecting physical manufacturability.

**Metric 4: LVS Pass Rate** checks electrical equivalence between schematic and layout, ensuring correct connectivity without opens, shorts, or mismatches.

**Metric 5: DRC+LVS Pass Rate** represents production-ready quality and serves as the overall effectiveness metric.

Together, these metrics provide a comprehensive evaluation framework spanning high-level intent interpretation and low-level physical validation.

## Main Results

We evaluate AMS-IO-Agent on AMS-IO-Bench, with results summarized in Table 3. Across different backbone models, AMS-IO-Agent achieves 100% intent graph translation and shape validity, while maintaining high DRC and LVS pass rates. Even for full DRC+LVS signoff, the agent reaches up to 76.7% (23 out of 30 cases), demonstrating substantial progress toward production-ready layout generation with minimal human effort. Compared to manual design, which typically requires around 480 minutes per task, the agent reduces turnaround time to only a few minutes and significantly outperforms baseline LLMs, with token usage remaining below 200k, making it practical for industrial deployment. These results show that AMS-IO-Agent delivers a major efficiency boost while substantially improving functional design correctness over baseline methods.

## Error Handling

Given the high cost (greater than \$10k per mm<sup>2</sup>) and long fabrication time (about three months) of tape-out, human

engineers primarily serve as reviewers. Infeasible cases detected during execution or through DRC and LVS can be resolved within minutes via prompt refinement or minor EDA edits, illustrating the advantage of generating editable intermediate deliverables rather than fixed geometric outputs.

## Ablation

We further conduct an ablation study to assess the contribution of each component of AMS-IO-Agent, as shown in Table 4. Human time was estimated from interviews with 16 PhD level designers. Without the knowledge base (KB), intent graph (IG), or adaptor, the agent completely fails to generate valid outputs, highlighting the necessity of combining all three components. Using only the knowledge base enables some basic shape assembly but results in low DRC and no LVS signoff, indicating that design rules alone are insufficient without structured reasoning. Similarly, employing the intent graph without the adaptor allows the agent to generate syntactically valid intent representations but fails to translate them into signoff-quality implementations. Only the full configuration, integrating KB, IG, and the adaptor, achieves 100% in intent graph, shape, 93.33% in DRC, and 76.67% in DRC+LVS, while reducing runtime to 5.1 minutes with modest token usage (105k). These results confirm that the three components are complementary and jointly essential for practical AMS I/O design automation.

Table 5 shows results across difficulty levels. All models succeed on simple cases, but their performance declines on medium and hard cases, validating that the benchmark’s difficulty scaling reflects real design complexity.

## Case Study

To demonstrate the real-world applicability of AMS-IO-Agent, we conducted a case study on a prototype tape-out

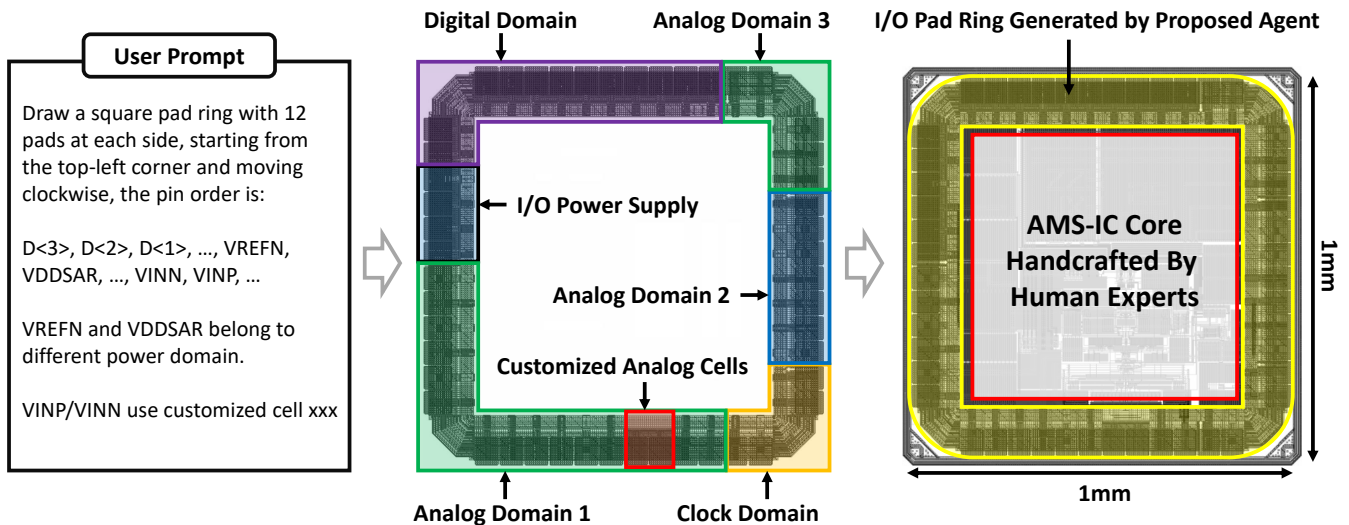


Figure 5: Application of the proposed agent system to a real AMS tape-out I/O ring design.

Model	Simple	Medium	Hard
GPT-4o	10/10	7/10	2/10
Claude-3.7	10/10	9/10	4/10
Deepseek-V3	10/10	10/10	3/10

Table 5: DRC+LVS pass rates of different models across difficulty levels in AMS-IO-Bench.

project implemented in 28-nm CMOS technology, involving a 1 mm × 1 mm wirebond-packaged mixed-signal IC with 48 I/O pads (12 per side) and multiple power and signal domains. As illustrated in Figure 5, the inner AMS core was implemented by two human designers, while the surrounding I/O ring was generated by our agent. This clear division of labor highlights the agent’s seamless integration into established design workflows.

During the design process, a major pin-order change was introduced that required sliding and reordering many pads. While such a change would normally necessitate extensive manual redrawing, the agent regenerated a fully updated and verified layout within minutes. The resulting I/O ring matched expert-designed quality and was readily adopted by the design team.

The final design, combining a manually crafted AMS core with the agent-generated I/O ring, passed LVS and DRC and was successfully fabricated. Silicon measurements confirmed correct functionality. This case study demonstrates that AMS-IO-Agent can serve as a production-ready module within real tape-out workflows, enabling true human-agent collaborative design while significantly reducing manual rework effort during late-stage iterations.

## Conclusion

In this work, we introduced *AMS-IO-Agent*, a domain-specialized LLM-based agent for structure-aware AMS I/O generation, together with *AMS-IO-Bench* for systematic

evaluation. By combining a curated domain knowledge base with structured intent representation, our approach translates informal pin-planning specifications into production-ready EDA workflows.

Experiments on AMS-IO-Bench show that AMS-IO-Agent achieves 100% intent graph correctness, 100% layout shape validity, and over 70% DRC+LVS signoff pass rate, reducing design turnaround time from days to mere minutes. Moreover, a prototype tape-out confirms the agent’s capability to seamlessly collaborate with human engineers, adapt to late-stage pin-order changes, and deliver signoff-quality I/O rings suitable for real-world chip design flows.

These results establish the feasibility of deploying LLM-based agents for AMS IC I/O automation and provide a concrete foundation for human-agent co-design. The methodology applies broadly to AMS layout tasks that exhibit regular structures and semantically meaningful signal naming, and can be adapted to other technology nodes or packaging styles by substituting the knowledge base and execution adaptor. Future work will extend this approach to more complex AMS design tasks and enable deeper integration with downstream verification and routing tools.

**Limitation and Social Impact** Our approach focuses on wirebond-packaged AMS I/O rings and relies on a domain knowledge base tailored to specific design conventions, which may limit generalization to other packaging types or foundry rules. While AMS-IO-Agent reduces manual workload, it cannot fully replace expert review, especially for highly customized or unconventional designs.

From a social perspective, our method can boost productivity and lower barriers for AMS IC design, but may also shift the demand for certain engineering skills. Responsible use and human oversight remain essential to ensure safety and reliability in real-world chip development.

## Acknowledgments

This work was supported by National Science Foundation of China (NSFC) Young Scholar Basic Research Program for doctoral students (Grant 624B2081).

## References

- Blocklove, J.; Garg, S.; Karri, R.; and Pearce, H. 2023. Chip-Chat: Challenges and Opportunities in Conversational Hardware Design. In *2023 ACM/IEEE 5th Workshop on Machine Learning for CAD (MLCAD)*, 1–6. IEEE.
- Chang, K.; Wang, Y.; Ren, H.; Wang, M.; Liang, S.; Han, Y.; Li, H.; and Li, X. 2023. ChipGPT: How far are we from natural language hardware design. arXiv:2305.14019.
- Chen, L.; Chen, Y.; Chu, Z.; Fang, W.; Ho, T.-Y.; Huang, R.; Huang, Y.; Khan, S.; Li, M.; Li, X.; Li, Y.; Liang, Y.; Liu, J.; Liu, Y.; Lin, Y.; Luo, G.; Pan, H.; Shi, Z.; Sun, G.; Tsaras, D.; Wang, R.; Wang, Z.; Wei, X.; Xie, Z.; Xu, Q.; Xue, C.; Yan, J.; Yang, J.; Yu, B.; Yuan, M.; Young, E. F. Y.; Zeng, X.; Zhang, H.; Zhang, Z.; Zhao, Y.; Zhen, H.-L.; Zheng, Z.; Zhu, B.; Zhu, K.; and Zou, S. 2024. Large circuit models: opportunities and challenges. *Science China Information Sciences*, 67(10).
- Chen, S.-H.; et al. 2021. Pad Ring Generation for Integrated Circuits. U.S. Patent 11,055,457 B1. SiFive Inc.
- Chen, W.; Liu, C.; Huang, W.; Lyu, J.; Yang, M.; Du, Y.; Du, L.; and Yang, J. 2025. AnalogTester: A Large Language Model-Based Framework for Automatic Testbench Generation in Analog Circuit Design. arXiv:2507.09965.
- Chen, Y. P.; Srujana, S. T. L. N. V.; Patel, N.; Reddy, R. R. L.; Subramanian, S.; and Vallapaneni, V. R. 2010. Automated electrostatic discharge structure placement and routing in an integrated circuit. Filed Nov. 30, 2006; published Feb. 2, 2010.
- Fang, W.; Wang, J.; Lu, Y.; Liu, S.; Wu, Y.; Ma, Y.; and Xie, Z. 2025. A Survey of Circuit Foundation Model: Foundation AI Models for VLSI Circuit Design and EDA. *arXiv preprint arXiv:2504.03711*. Available at <https://arxiv.org/abs/2504.03711>.
- Fu, Y.; Zhang, Y.; Yu, Z.; Li, S.; Ye, Z.; Li, C.; Wan, C.; and Lin, Y. C. 2025. GPT4AIGChip: Towards Next-Generation AI Accelerator Design Automation via Large Language Models. arXiv:2309.10730.
- Ghose, A.; Kahng, A. B.; Kundu, S.; and Wang, Z. 2025. ORFS-agent: Tool-Using Agents for Chip Design Optimization. arXiv:2506.08332.
- Ho, C.-T.; and Ren, H. 2024. Large Language Model (LLM) for Standard Cell Layout Design Optimization. arXiv:2406.06549.
- Lai, Y.; Lee, S.; Chen, G.; Poddar, S.; Hu, M.; Pan, D. Z.; and Luo, P. 2024. AnalogCoder: Analog Circuit Design via Training-Free Code Generation. arXiv:2405.14918.
- Liu, B.; Zhang, H.; Gao, X.; Kong, Z.; Tang, X.; Lin, Y.; Wang, R.; and Huang, R. 2025. LayoutCopilot: An LLM-Powered Multiagent Collaborative Framework for Interactive Analog Layout Design. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 44(8): 3126–3139.
- Liu, M.; Ene, T.-D.; Kirby, R.; Cheng, C.; Pinckney, N.; Liang, R.; Alben, J.; Anand, H.; Banerjee, S.; Bayraktaroglu, I.; Bhaskaran, B.; Catanzaro, B.; Chaudhuri, A.; Clay, S.; Dally, B.; Dang, L.; Deshpande, P.; Dhodhi, S.; Halepete, S.; Hill, E.; Hu, J.; Jain, S.; Jindal, A.; Khailany, B.; Kokai, G.; Kunal, K.; Li, X.; Lind, C.; Liu, H.; Oberman, S.; Omar, S.; Pasandi, G.; Pratty, S.; Raiman, J.; Sarkar, A.; Shao, Z.; Sun, H.; Suthar, P. P.; Tej, V.; Turner, W.; Xu, K.; and Ren, H. 2024a. ChipNeMo: Domain-Adapted LLMs for Chip Design. arXiv:2311.00176.
- Liu, M.; Pinckney, N.; Khailany, B.; and Ren, H. 2023. VerilogEval: Evaluating Large Language Models for Verilog Code Generation. arXiv:2309.07544.
- Liu, S.; Fang, W.; Lu, Y.; Zhang, Q.; Zhang, H.; and Xie, Z. 2024b. RTLcoder: Outperforming GPT-3.5 in Design RTL Generation with Our Open-Source Dataset and Lightweight Solution. In *2024 IEEE LLM Aided Design Workshop (LAD)*, 1–5.
- Morita, A. K.; Takiguti, R.; and Noije, W. A. M. V. 2015. Metadata based padding and pad multiplexing generation for microcontroller design. *Journal of Integrated Circuits and Systems*, 10(3): 139–146.
- Moseley, N.; Barzic, R.; and contributors, Y. 2025. padding: A padding generator for ASICs. <https://github.com/YosysHQ/padding/tree/master>. ISC License. Accessed: 2025-11-13.
- Shi, L.; Kazda, M.; Sears, B.; Shropshire, N.; and Puri, R. 2024. Ask-EDA: A Design Assistant Empowered by LLM, Hybrid RAG and Abbreviation De-hallucination. arXiv:2406.06575.
- Skelic, L.; Xu, Y.; Cox, M.; Lu, W.; Yu, T.; and Han, R. 2025. CIRCUIT: A Benchmark for Circuit Interpretation and Reasoning Capabilities of LLMs. arXiv:2502.07980.
- Thakur, S.; Ahmad, B.; Pearce, H.; Tan, B.; Dolan-Gavitt, B.; Karri, R.; and Garg, S. 2023. VeriGen: A Large Language Model for Verilog Code Generation. arXiv:2308.00708.
- Thakur, S.; Blocklove, J.; Pearce, H.; Tan, B.; Garg, S.; and Karri, R. 2024. AutoChip: Automating HDL Generation Using LLM Feedback. arXiv:2311.04887.
- Tsai, Y.-D.; Liu, M.; and Ren, H. 2024. RTLFixer: Automatically Fixing RTL Syntax Errors with Large Language Models. arXiv:2311.16543.
- Wu, H.; He, Z.; Zhang, X.; Yao, X.; Zheng, S.; Zheng, H.; and Yu, B. 2024. ChatEDA: A Large Language Model Powered Autonomous Agent for EDA. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 43(10): 3184–3197.
- Xu, K.; Qiu, R.; Zhao, Z.; Zhang, G. L.; Schlichtmann, U.; and Li, B. 2024. LLM-Aided Efficient Hardware Design Automation. arXiv:2410.18582.
- Yin, Y.; Wang, Y.; Xu, B.; and Li, P. 2024. ADO-LLM: Analog Design Bayesian Optimization with In-Context Learning of Large Language Models. In *Proceedings of the 43rd IEEE/ACM International Conference on Computer-Aided Design, ICCAD '24*, 1–9. ACM.