

# The EDA Primer: From RTL to Silicon

## EDA 入门指南：从 RTL 到硅片

Laying the Groundwork of the Current Chip Design Paradigm

奠定当前芯片设计范式的基石

GERALD WONG, DYLAN PATEL, AND SRAVAN KUNDOJJALA

GERALD WONG, DYLAN PATEL, 以及 SRAVAN KUNDOJJALA

MAY 12, 2026 2026 年 5 月 12 日 · PAID 付费

 44



 3

Share 分享



AI demand has been driving the explosion in compute over the past few years, resulting in chip designs getting ever more complex, with silicon area and power per package seeing continued growth as designs push for even greater performance. With each successive generation, new process nodes with more design rules and restrictions further increase chip design costs.


在过去几年中，AI 需求一直推动着计算能力的爆炸式增长，导致芯片设计变得愈发复杂。随着设计不断追求更高性能，单个封装内的芯片面积和功耗也在持续增长。随着每一代技术的更迭，具有更多设计规则和限制的新工艺节点进一步推高了芯片设计成本。

At the same time, the rush to bring compute into the market as quickly as possible has put design teams under immense pressure to compress timelines and speed up validation cycles from years to months. If you're not fast, you will get lapped up and beaten by your competitors. Even a 3 month delay means billions of dollars.

与此同时，为了尽快将算力推向市场，设计团队面临着巨大的压力，必须压缩时间线，将验证周期从几年缩短至几个月。如果你不够快，就会被竞争对手超越并击败。哪怕仅仅延迟 3 个月，也意味着数十亿美元的损失。

## Semiconductor companies are facing key challenges

### Increasing engineering hours




### Shrinking talent base



Shortage of **semiconductor workers** in US by 2030

### Increasing development costs



Estimated first-time **total SoC design cost** or 3 nm node

Source: Siemens 来源: Siemens


All this is happening while the engineering talent base is shrinking. Lucrative salaries and flexible working arrangements have enticed most students into the Software and Information Systems tracks, resulting in a dwindling number of Electrical Engineer graduates that could enter the chip design workforce. Siemens presented the engineer-hours demanded of these numerous complex AI accelerator designs that far outstrip the engineering talent coming into the workforce.

这一切都发生在工程人才库不断萎缩的背景下。丰厚的薪资和灵活的工作安排吸引了大多数学生进入软件和信息系统专业，导致能够进入芯片设计行业的电气工程毕业生人数不断减少。Siemens 展示了这些众多复杂的 AI 加速器设计所需的工程师工时，这已远远超过了进入劳动力市场的工程人才数量。

One-third of the current U.S. semiconductor workforce is over 55. The pipeline of new graduates is nowhere close to filling that gap. Even Apple is actively funding education programs to encourage uptake in engineering. While their New Silicon Initiative has contributed to increasing interest and number of EE graduates, it barely moves the needle compared to the explosion in manpower requirements as transistor

count grows at a Moore's Law pace.

目前美国半导体从业人员中有三分之一已超过 55 岁。应届毕业生的储备量远不足以填补这一缺口。即便是苹果公司也在积极资助教育项目，以鼓励更多人投身工程领域。虽然其“新硅计划”(New Silicon Initiative) 有助于提高人们的兴趣并增加电子工程 (EE) 专业的毕业生人数，但与晶体管数量按摩尔定律速度增长所带来的巨大人力需求相比，这几乎是杯水车薪。



来源：苹果

With this trifecta of increasing chip complexity, compressed design timelines and a shortage of engineers, a massive bottleneck has formed at the design stage. The latest AMD MI455X packs 320 billion transistors across 12 logic dies on 2nm and 3nm processes with advanced Hybrid Bonding 3D die stacking, HBM4 memory integration and high speed 224G SerDes. Designing something at this scale is not a matter of hiring more engineers or buying more verification servers. It tests a company's tooling, methodology, and human capital organization as to whether the design succeeds or

fails.

在芯片复杂度提升、设计周期缩短以及工程师短缺这三重压力的共同作用下，设计阶段已经形成了一个巨大的瓶颈。最新的 AMD MI455X 在 2nm 和 3nm 工艺上集成了 12 个逻辑芯片，拥有 3200 亿个晶体管，并采用了先进的混合键合（Hybrid Bonding）3D 芯片堆叠技术、HBM4 内存集成以及高速 224G SerDes。设计这种规模的产品，已经不再是单纯增加工程师人数或购买更多验证服务器就能解决的问题。它考验着一个公司的工具链、方法论和人力资源组织能力，决定了设计的成败。

After spending hundreds of millions of dollars on a new SoC design, there is no guarantee the chip will work. Multiple steppings are usually required that need new mask sets, with A0 rarely going into production. When a single advanced mask set costs tens of millions of dollars, every respin is a gut punch to the balance sheet. Furthermore, it adds months to the schedule for high volume production start.

在为一个新的 SoC 设计投入数亿美元后，并不能保证芯片一定能正常工作。通常需要多次改版（Stepping）并更换新的掩模组（Mask set），A0 版本极少能直接投入量产。当一套先进的掩模组成本高达数千万美元时，每一次重新流片（Respin）都是对资产负债表的沉重打击。此外，这还会使大批量生产的启动时间推迟数月之久。

As designs get more complex, testing is becoming more important to ensure all modules within a chip are interoperable and locally sound. **Verification**, the process of proving a design does exactly what it should before committing it to silicon, now consumes up to 70% of total project effort, depending on the design. Verification engineers are the fastest-growing job category in chip development, and the industry still cannot hire them fast enough.

随着设计变得越来越复杂，测试变得愈发重要，以确保芯片内的所有模块都能互操作且局部功能完备。验证是指在将设计交付硅片制造之前，证明其功能完全符合预期的过程；根据设计的不同，验证目前占据了项目总工作量的 70%。验证工程师是芯片开发中增长最快的职位类别，而行业对这类人才的需求依然供不应求。






Figure 8. Mean peak number of IC/ASIC engineers.

Source: Siemens 来源: Siemens

While chip complexity grows at roughly 50% per year, driven by new nodes and larger SoCs, design productivity improves only about 20% each year. This design productivity gap means every new generation of silicon demands exponentially more engineering effort, more compute, and more sophisticated automation.

随着新工艺节点和更大规模 SoC 的推动，芯片复杂度每年增长约 50%，而设计生产率每年仅提高约 20%。这种设计生产率差距意味着，每一代新芯片都需要呈指数级增长的工程投入、计算资源和更先进的自动化技术。

The semiconductor industry's ability to keep building more powerful chips depends not on physics or lithography alone, but on EDA (**Electronic Design Automation**) software. These tools effectively translate human intent into manufacturable silicon. Without EDA, no chip designed after the mid-1980s would exist.

半导体行业持续制造更强大芯片的能力，不仅取决于物理学或光刻技术，还取决于 EDA（电子设计自动化）软件。这些工具能够有效地将人类的设计意图转化为可制造的芯片。如果没有 EDA，20 世纪 80 年代中期以后设计的任何芯片都不可能存在。

This primer is your guide to EDA in the semiconductor industry. In this first part, we will walk the entire journey from **RTL (Register Transfer Level)** code, the high-level hardware description language that engineers actually write, all the way to manufactured, packaged silicon. We will name the tools, explain the tradeoffs, and show why EDA is one of the most consequential and underappreciated sectors in technology.

本入门指南是您了解半导体行业 EDA 的向导。在第一部分中，我们将全程走过从 RTL（寄存器传输级）代码——即工程师实际编写的高级硬件描述语言——一直到制造完成并封装好的硅片的整个旅程。我们将列举相关工具，解释其中的权衡取舍，并展示为什么 EDA 是科技领域中最具影响力却又最被低估的部门之一。

In part 2, our EDA Market Primer dives deep into the business of EDA, profiling the major companies (Synopsys, Cadence, Siemens) and their revenue and business models. We provide comprehensive market analysis and monitoring the Chinese EDA effort, as well as IP licensing and outsourcing to design partners and the transition to Customer Owned Tooling (COT) with hyperscaler ASIC designs.

在第二部分中，我们的 EDA 市场入门指南深入探讨了 EDA 业务，剖析了主要公司（Synopsys、Cadence、Siemens）及其营收和商业模式。我们提供了全面的市场分析，并监测中国 EDA 的发展进程，同时涵盖了 IP 许可、向设计合作伙伴的外包业务，以及随着超大规模集成电路（hyperscaler）ASIC 设计而向客户自备工具（COT）模式的转型。

Part 3 then assesses how AI is disrupting the EDA industry, covering the full gamut from startups and engineer dashboards to agentic chip design flows from NVIDIA and the big three. The concept of using AI accelerators to create superhuman designs that go into future AI accelerators is the most exciting development that our industry has seen in decades. Stay tuned as we cover the incoming revolution in chip design.

第三部分随后评估了人工智能如何颠覆 EDA 行业，涵盖了从初创公司和工程师仪表板，到来自 NVIDIA 和三大巨头的代理式芯片设计流程的方方面面。利用人工智能加速器来创造超人类的设计，并将其应用于未来的人工智能加速器，这一概念是数十年来我们行业所见过的最令人兴奋的发展。请继续关注我们对即将到来的芯片设计革命的报道。

# A Brief History: From X-Acto Knives to the Big Three

简史：从美工刀到“三巨头”时代



Source: Intel 



In the 1960s and 1970s, designing an integrated circuit meant drawing it by hand. Engineers sketched layouts on graph paper, and technicians transferred those sketches onto sheets of **Rubylith** — a red cellophane film laminated onto clear Mylar. Using X-Acto knives and light tables, they cut away sections of the film to define each layer of the chip. The finished masters were then photo-reduced up to 100 times to create production photomasks. A single slip of the blade could ruin weeks of work. This was the standard design process up to and including the Intel 8080 with its Rubylith pictured above.

在 20 世纪 60 和 70 年代，设计集成电路意味着手工绘图。工程师在坐标纸上绘制版图草图，技术人员再将这些草图转移到 Rubylith 片基上——这是一种层压在透明聚酯薄膜上的红色玻璃纸。他们利用美工刀和透光台，切掉部分薄膜以定义芯片的每一层。完成的母版随后会被缩小拍摄多达 100 倍，以制作生产用的光罩。刀片的一次打滑就可能毁掉数周的工作成果。直到 Intel 8080（其 Rubylith 版图如上图所示）及其之前的芯片，这都是标准的电路设计流程。

The first step toward automation came in 1971, when **Calma** shipped its **Graphic Design System (GDS)** to Intel, allowing engineers to digitize and edit layouts on minicomputers. In 1978, Calma released **GDS II**, whose stream file format became the de facto standard for exchanging mask data. Remarkably, **GDS II remains the dominant interchange format today, nearly five decades later**, alongside its modern successor OASIS.

迈向自动化的第一步始于 1971 年，当时 Calma 公司向 Intel 交付了其图形设计系统（GDS），使工程师能够在小型机上对版图进行数字化处理和编辑。1978 年，Calma 发布了 GDS II，其流文件格式成为了交换掩模数据的行业事实标准。值得注意的是，在近五十年后的今天，GDS II 及其现代继任者 OASIS 仍然是主流的交流格式。

The EDA industry as we know it was born in **1981**, when three companies launched within months of each other: **Daisy Systems, Mentor Graphics, and Valid Logic Systems**. Known collectively as “**DMV**,” they introduced computer-aided engineering to the front end of the design flow, schematic capture, simulation, and logic verification, running on dedicated workstations. By the late 1980s, all three had migrated to standard Unix workstations from Apollo and Sun Microsystems,

establishing the software-centric business model that defines EDA today.

我们所熟知的 EDA 行业诞生于 1981 年，当时有三家公司在几个月内相继成立：Daisy Systems、Mentor Graphics 和 Valid Logic Systems。它们被统称为“DMV”，将计算机辅助工程引入了设计流程的前端，包括原理图捕获、仿真和逻辑验证，并运行在专用工作站上。到 20 世纪 80 年代末，这三家公司都已迁移到来自 Apollo 和 Sun Microsystems 的标准 Unix 工作站，确立了定义当今 EDA 行业的以软件为中心的商业模式。

## The Big Three Emerge 三大巨头崛起

The modern EDA landscape is dominated by three companies. Synopsys, founded in 1986 by Aart de Geus and colleagues from General Electric's research group, introduced **Design Compiler** in 1987, the first commercial **logic synthesis** tool. Logic synthesis automated the translation of high-level hardware descriptions into optimized gate-level netlists, a breakthrough that enabled the leap from thousands of hand-placed transistors to the billions we design today. **Cadence Design Systems** formed in 1988 through the merger of SDA Systems and ECAD, quickly becoming the leading provider of IC layout and place-and-route tools. And **Mentor Graphics**, one of the original DMV trio, was acquired by **Siemens** in 2017 for **\$4.5 billion**, rebranding as **Siemens EDA** in 2021 and bringing deep verification and physical design expertise into the Siemens Digital Industries portfolio.

现代 EDA 领域由三家公司主导。Synopsys 由 Aart de Geus 及其通用电气研究小组的同事于 1986 年创立，并于 1987 年推出了首款商业逻辑综合工具 Design Compiler。逻辑综合实现了将高级硬件描述自动转换为优化的门级网表，这一突破使得设计规模从数千个手动放置的晶体管飞跃到了如今的数十亿个。Cadence Design Systems 于 1988 年由 SDA Systems 和 ECAD 合并而成，迅速成为 IC 布局和布线工具的领先供应商。而作为最初“DMV”三巨头之一的 Mentor Graphics，于 2017 年被西门子以 45 亿美元收购，并于 2021 年更名为 Siemens EDA，将深厚的验证和物理设计专业知识引入西门子数字化工业软件产品组合中。

Compared to the early Rubylith days, logic synthesis not only speed up design, it fundamentally changed what was possible. By abstracting away manual gate placement, it unlocked a multi million-fold increase in design complexity to form

today's multi-billion-transistor SoCs.

与早期的 Rubylith 时代相比，逻辑综合不仅加快了设计速度，还从根本上改变了设计的可能性。通过抽象掉手动的门级放置，它实现了设计复杂度数百万倍的增长，从而造就了当今拥有数十亿晶体管的 SoC。


# The Chip Design Waterfall

## 芯片设计瀑布流

Building a chip is a multi-year relay race with thirteen legs. Miss a handoff and the whole schedule slips, by months, or even by quarters. The diagram below lays out the full flow from a blank whiteboard to volume production. This article will go through the stages where EDA tools are used in the design flow.

制造芯片是一场耗时数年的十三棒接力赛。一旦错过任何一次交接，整个进度就会延误数月甚至数个季度。下图展示了从空白白板到量产的全流程。本文将重点介绍在设计流程中使用 EDA 工具的各个阶段。

# Chip Design Flow



Source: SemiAnalysis SemiAnalysis

- 1. Planning:** Define the product requirements, target market, and PPA (power, performance, area) goals that will constrain every decision downstream.



规划：定义产品需求、目标市场以及 PPA（功耗、性能、面积）目标，这些目标将约束后续的每一项决策。

2. **Architecture:** Design the microarchitecture: instruction set choices, cache hierarchies, bus widths, and the block diagrams that partition the chip into manageable units.

架构：设计微架构：包括指令集选择、缓存层级、总线宽度，以及将芯片划分为可管理单元的框图。

3. **RTL Design:** Write the actual hardware description code, almost always in **SystemVerilog**, that specifies every register, mux, and state machine in the design.

RTL 设计：编写实际的硬件描述代码（几乎总是使用 SystemVerilog），明确设计中的每一个寄存器、多路复用器和状态机。

4. **RTL Verification:** Exhaustively test or prove that the RTL behaves correctly across billions of scenarios. Implemented with Testbenches or formal proofs.

RTL 验证：通过测试平台（Testbench）或形式验证，详尽地测试或证明 RTL 在数十亿种场景下的行为正确性。

5. **RTL Freeze:** The design is locked. No more functional changes allowed, only bug fixes that pass a strict change control review.

RTL 冻结（RTL Freeze）：设计已锁定。不再允许进行功能性变更，仅允许通过严格变更控制审查的错误修复。

6. **FW/SW Development (Parallel):** Firmware and software teams begin bring-up on emulators and FPGA prototypes, often running in parallel with physical design to save months of schedule.

固件/软件开发（并行）：固件和软件团队开始在仿真器（Emulator）和 FPGA 原型上进行调试（Bring-up），通常与物理设计并行进行，以节省数月的进度。

7. **Physical Design: Logic Synthesis** to convert the RTL into a gate-level netlist, **Placement** (gates onto the die), **routing** (wiring them together) and **floorplanning**

(assigning areas of the die for each functional block).

物理设计：逻辑综合（将 RTL 转换为门级网表）、布局（将逻辑门放置在芯片上）、布线（将它们连接在一起）以及规划（为每个功能模块分配芯片区域）。

8. **Signoff:** Run final checks that the design meets **timing closure** (every signal arrives on time), power budgets, and **DRC/LVS** (manufacturing rule) requirements.

签核 (Signoff)：进行最终检查，确保设计满足时序收敛（所有信号准时到达）、功耗预算以及 DRC/LVS（制造规则）要求。

9. **Foundry Handoff:** The finished layout is exported as a **GDSII** file, the multi-gigabyte blueprint the foundry uses to create photolithography masks. Known as the “tapeout” milestone.

代工厂交付 (Foundry Handoff)：完成后的版图被导出为 GDSII 文件，这是代工厂用于制作光刻掩模的数 GB 级蓝图。这一阶段被称为“流片 (tapeout)”里程碑。

10. **Fabrication:** Wafers are manufactured in the fab over 3-4 months, passing through thousands of processing steps across dozens of tools.

制造：晶圆在晶圆厂中经过 3-4 个月的制造，期间需经过数十台设备的数千道工艺步骤。

11. **Post-Silicon Validation:** Real chips come back from the fab. Post-silicon bring up engineers test them on custom boards and probe cards, debug errata, and decide on binning strategies (productizing parts with varying yield and performance into different SKUs). Multiple steppings may be done in this phase. Reliability tests are done with burn-in and Final Test.

硅后验证：真实的芯片从晶圆厂返回。硅后启动工程师在定制电路板和探针卡上对其进行测试，调试勘误表 (errata)，并决定分选策略（将具有不同良率和性能的部件产品化为不同的 SKU）。在此阶段可能会进行多次步进 (stepping)。可靠性测试通过老化测试 (burn-in) 和终测 (Final Test) 完成。

12. **System Integration:** Validated chips are integrated into boards, packages and connected to devices, with drivers, BIOS, and OS support qualified with System

## Level Testing.

系统集成：经过验证的芯片被集成到电路板和封装中并连接到设备，其驱动程序、BIOS 和操作系统支持均通过系统级测试（System Level Testing）进行鉴定。

13. **Production:** Volume manufacturing ramps to meet demand, with ongoing yield optimization and supply chain coordination.

生产：批量制造规模扩大以满足需求，并伴随持续的良率优化和供应链协调。

This is a simplified “waterfall” view. In practice, many of these stages overlap heavily and iterate. Architecture bugs found during verification force RTL changes; timing failures in physical design send engineers back to re-optimize critical paths. A modern SoC program manages dozens of these feedback loops simultaneously, which is exactly why EDA tooling exists, no human team could track it all by hand.

这是一个简化的“瀑布式”视图。在实践中，许多阶段会大量重叠并不断迭代。验证期间发现的架构缺陷会迫使 RTL 进行修改；物理设计中的时序失败会促使工程师重新优化关键路径。一个现代 SoC 项目需要同时管理数十个此类反馈循环，而这正是 EDA 工具存在的意义——没有任何人工团队能够手动追踪这一切。

## 1. Planning 1. 规划

The first stage to any chip is to decide on what role the chip serves. Each design department usually specializes in a given family of chips, be it CPUs and accelerators to the more mundane system controllers and embedded sensors. The product requirements and high level specifications are defined with respect to the current generation of products in the market, along with competitive analysis of others in the target market.

芯片开发的第一阶段是确定该芯片的功能定位。每个设计部门通常专注于特定的芯片系列，从 CPU 和加速器到更常规的系统控制器和嵌入式传感器。产品需求和高级规格的定义是参考市场上当前一代产品，并结合对目标市场中其他竞争对手的分析来确定的。

Strawman concepts are proposed that evolve rapidly as Program Managers work within the insertion schedules of various IP blocks from the design teams that may be ready for integration. Learnings from Post-Mortems of previous projects are factored in, forming a knowledge base to work from on what works and what is too ambitious for a given timeframe.

草案概念被提出并迅速演变，与此同时，项目经理根据设计团队中可能已准备好进行集成的各种 IP 模块的切入时间表开展工作。前期项目的复盘经验被纳入其中，形成了一个知识库，用以评估哪些方案可行，以及哪些方案对于给定时间节点而言过于激进。

The key high-level metrics here are **PPACt**: Performance and Power consumption, usually given as a percentage improvement over the prior generation and where it might sit in the competitive landscape, the area that such a design takes up in silicon on a given process node, which translates to Cost. Time to Market is the final metric that determines whether the product is viable both from a design time and product competitiveness standpoint. In a fast growing market where performance doubles every few years, being 1 year late could spell the end of a project's success.

这里的关键高级指标是 PPACt：性能（Performance）和功耗（Power consumption），通常以相对于上一代产品的改进百分比以及在竞争格局中的地位来衡量；该设计在特定工艺节点上占用的硅片面积（Area），这直接转化为成本（Cost）；最后是上市时间（Time to Market），它是从设计周期和产品竞争力角度决定产品是否可行的最终指标。在一个性能每几年翻一番的快速增长市场中，晚一年上市可能意味着一个项目的成功化为泡影。

These feasibility studies will then need to be greenlit by management before project kickoff begins in earnest. Each company has with work within their R&D budget with finite engineering resources. Scheduling resource allocation with ongoing projects in the roadmap requires strict completion deadlines so engineers can be released to start working on the next project. Communicating early with suppliers to project the wafer, memory and packaging demands for each design is now increasingly important to

secure capacity.

这些可行性研究随后需要获得管理层的批准，然后项目才能正式启动。每家公司都必须在研发预算范围内运作，且工程资源有限。在路线图中为进行中的项目安排资源分配需要严格的完成期限，以便工程师能够抽身开始下一个项目。尽早与供应商沟通以预测每个设计的晶圆、内存和封装需求，对于确保产能正变得日益重要。


## 2. Architecture Layout 2. 架构布局

Closely tied to planning, the architectural layout is done alongside design space exploration. A high-level floorplan diagram sets the initial area bounding boxes for each logic and I/O block design team to work within. Each functional block is broken down into smaller elements that are easier to design and can be repeated multiple times across the design. These area budgets may increase over the design cycle based on features that may be added later that take more area. For example, a feature update in an Instruction Set Architecture (ISA) with additional computing elements to support new instructions. On the AI accelerator side, this equates to adding dataflow accelerators and doubling Matrix Multiplication engine widths.

架构布局与规划紧密相关，是在设计空间探索过程中同步完成的。高级别的平面规划图（floorplan）为每个逻辑和 I/O 模块设计团队设定了初始的面积约束框。每个功能模块被进一步分解为更小的单元，这些单元更易于设计，并能在整个设计中多次复用。随着设计周期的推进，这些面积预算可能会因为后续增加的、占用更多面积的功能而增加。例如，指令集架构（ISA）的功能更新可能需要额外的计算单元来支持新指令。在 AI 加速器方面，这等同于增加数据流加速器并将矩阵乘法引擎的宽度翻倍。



**TTU:** Tensor Unit  
**TVP:** Vector engine  
**TDMA:** Tile Data Movement Engine  
**TCP:** Tile Control Processor



Source: Microsoft Microsoft

Block diagrams are drawn up with relations and Network on Chip (NoC) bandwidth requirements decided for each functional block, with memory bus widths and SRAM area budgeted based on cache hierarchy and early simulations of performance vs memory pressure. These simulations, known as **Design Space Exploration**, have traditionally been done with targeted Design of Experiments that simulate the performance impacts and interactions between each functional block, varying unit sizes, widths and bandwidths to find the lowest hanging fruit to maximize

performance gains.

绘制方框图时，需确定每个功能块之间的关系和片上网络（NoC）的带宽需求，并根据缓存层级以及性能与内存压力关系的早期仿真结果，对内存总线宽度和 SRAM 面积进行预算。这些被称为“设计空间探索”的仿真，传统上通过有针对性的实验设计（DoE）来完成，旨在模拟各功能块之间的性能影响和相互作用，通过改变单元尺寸、宽度和带宽，寻找最容易实现的优化点，从而实现性能增益的最大化。

Going forward, this step has increasingly been accelerated with AI, as the task is easily verifiable with assignable reward functions for PPA in a multi-dimensional input space. First party AI-driven exploration tools such as Synopsys' [DSO.ai](#) have followed the many internal efforts by the fabless design houses to leverage AI to accelerate pathfinding and planning decisions. An in-depth analysis on this will be featured in Part 3 of this EDA series.

展望未来，随着这一任务在多维输入空间中可以通过可分配的 PPA（功耗、性能、面积）奖励函数轻松进行验证，AI 正在日益加速这一步骤。在无晶圆厂设计公司内部利用 AI 加速路径探索和规划决策的众多尝试之后，Synopsys 的 DSO.ai 等原厂 AI 驱动探索工具也应运而生。关于这一点的深入分析将在本 EDA 系列的第 3 部分中呈现。


## 3. RTL Design 3. RTL 设计

With the architecture specified, engineers must then describe exactly what the chip does. This is done at the level of registers, data paths and combinational logic, which will later be translated into transistor implementations. This description is called RTL (Register Transfer Level) code, and it is where the design's behavior is defined in a language that both humans and synthesis tools can read. Most of the engineering hours in the chip design flow is spent writing and verifying the RTL code. Below we look at the aspects to RTL design.

在确定了架构之后，工程师必须精确描述芯片的具体功能。这一过程是在寄存器、数据通路和组合逻辑层面完成的，随后这些描述将被转换为晶体管实现。这种描述被称为 RTL（寄存器传输级）代码，它使用人类和综合工具都能理解的语言来定义设计的行为。在芯片设计流程中，大部分工程时间都花在编写和验证 RTL 代码上。下面我们来看看 RTL 设计的各个方面。

In the real world, transistors don't switch instantaneously. There is a propagation delay where it takes some time for an input change to produce a stable output. This delay has two components: the gate delay (how fast the transistors themselves switch) and the wire delay (how long the electrical signal takes to travel along the metal interconnect to the next gate). At advanced process nodes, wire delay ends up dominating gate delay as transistors switch faster while datapaths lengthen with complex designs.

# Read Waveform



SRAM Cell Read Waveform. Source: MediaTek

SRAM 单元读取波形。来源：联发科技 (MediaTek)

Digital chips use a clock signal to synchronize all operations. Two timing constraints govern correctness. **Setup time** requires that input data be stable for a minimum period before the clock edge arrives. **Hold time** requires that data remain stable for a minimum period after the clock edge. The clock period (the inverse of frequency) must be long enough to accommodate the slowest signal path in the entire design. This

worst-case path is called the **critical path**. If your critical path takes 0.2 nanoseconds and you want a 5 GHz clock (0.2 ns period), you are right at the edge, with no margin for process variability. This is why timing optimization consumes enormous effort in chip design, with many trade-offs in performance and complexity.

数字芯片使用时钟信号来同步所有操作。两个时序约束决定了正确性。建立时间（Setup time）要求输入数据在时钟沿到来之前保持稳定的最小周期。保持时间（Hold time）要求数据在时钟沿之后保持稳定的最小周期。时钟周期（频率的倒数）必须足够长，以容纳整个设计中最慢的信号路径。这条最差情况路径被称为关键路径。如果你的关键路径耗时 0.2 纳秒，而你想要 5 GHz 的时钟（0.2 纳秒周期），那么你正处于边缘，没有工艺偏差的余量。这就是为什么时序优化在芯片设计中耗费巨大精力，并在性能和复杂性之间进行多次权衡的原因。

## State Elements 状态元件

Combinational logic computes outputs from inputs, but it needs to be combined with memory to build useful functions such as a counter, a processor pipeline stage, or a protocol engine. These memory registers are implemented as **flip-flops**. A flip-flop captures and holds one bit of data on each clock edge, acting as a tiny one-bit memory. Multiple flip-flops are chained together with combinational logic to form a **Finite State Machine** (FSM). This circuit steps through a defined sequence of states, one clock cycle at a time. This is sequential logic, which forms the base for chips to compute. Thus, RTL is an abstraction that describes how data moves between registers and combinational logic on each clock cycle.

组合逻辑根据输入计算输出，但它需要与存储器结合才能构建出有用的功能，例如计数器、处理器流水线阶段或协议引擎。这些存储寄存器由触发器（flip-flops）实现。触发器在每个时钟沿捕获并保持一位数据，充当微型的一位存储器。多个触发器与组合逻辑链连接在一起，形成有限状态机（FSM）。这种电路一次一个时钟周期地按定义的序列步进。这就是时序逻辑，它构成了芯片计算的基础。因此，RTL 是一种抽象，描述了数据在每个时钟周期如何在寄存器和组合逻辑之间流动。

## Writing the RTL 编写 RTL

RTL is written in a hardware description language (HDL). The dominant choice today is **SystemVerilog**, an extension of the original Verilog language that adds features for



both design and verification. VHDL, the older alternative, still appears in aerospace and legacy applications. A designer writing RTL specifies what happens on every clock edge, where data moves between registers, arithmetic operations execute, and state machines transition. Synthesis tools (covered in the next section) then convert this description into actual gates and transistors.

RTL 使用硬件描述语言（HDL）编写。当今的主流选择是 SystemVerilog，它是原始 Verilog 语言的扩展，增加了用于设计和验证的功能。VHDL 作为较早的替代方案，仍出现在航空航天和遗留应用中。编写 RTL 的设计人员会规定每个时钟沿发生的操作，包括数据在寄存器之间的移动、算术运算的执行以及状态机的转换。随后，综合工具（将在下一节介绍）将这些描述转换为实际的门电路和晶体管。

Once written, RTL passes through linting, a static analysis that catches coding mistakes, race conditions, and syntax errors. This is done as a quick code review without requiring simulation. **VC SpyGlass** from Synopsys is the industry-standard linting tool, flagging seemingly subtle issues that could cause intermittent silicon failures. This is essentially the chip design equivalent of a compiler's warning flags, just with far costlier consequences.

编写完成后，RTL 会经过 Lint 检查，这是一种静态分析，用于捕捉编码错误、竞争条件和语法错误。这是一种无需仿真即可完成的快速代码审查。Synopsys 的 VC SpyGlass 是行业标准的 Lint 工具，它能标记出那些看似细微但可能导致芯片间歇性故障的问题。这本质上相当于芯片设计领域的编译器警告标志，只是后果要严重得多。

## IP Integration **IP 集成**

In most modern SoC (System on Chip) designs, only about 20-30% of the RTL is truly custom logic designed in-house. It is easier to reuse previous designs for non-critical components, with the rest comprised of licensed IP blocks. These are pre-designed, pre-verified modules purchased from third-party vendors. ARM provides processor cores, GPU and other IP. Synopsys **DesignWare** supplies USB, PCIe, DDR memory controllers, and hundreds of other interface blocks. Broadcom's excellent high speed IO can be used if they are handling the rest of your chip design. Meanwhile, smaller IP

vendors sell everything from GPIO interfaces to cryptographic accelerators.

在大多数现代 SoC（片上系统）设计中，只有约 20-30% 的 RTL 是真正内部设计的定制逻辑。对于非关键组件，重用以前的设计会更容易，而其余部分则由授权的 IP 模块组成。这些是向第三方供应商购买的、经过预先设计和验证的模块。ARM 提供处理器内核、GPU 和其他 IP。Synopsys DesignWare 提供 USB、PCIe、DDR 内存控制器以及数百个其他接口模块。如果 Broadcom 负责处理你芯片设计的其余部分，则可以使用其出色的高速 IO。与此同时，较小的 IP 供应商销售从 GPIO 接口到加密加速器的各种产品。

IP licensing is the result of economics. Designing a custom PCIe Gen 6 controller from scratch would require spinning up a dedicated team of I/O design and verification engineers working to prove compliance with PCI-SIG's specification. Licensing one costs a fraction of that and comes pre-verified against the spec. However, the IP integration itself can be challenging, something we will cover for our subscribers below.

IP 授权是经济效益驱动的结果。从零开始设计一个定制的 PCIe Gen 6 控制器需要组建一支专门的 I/O 设计和验证工程师团队，致力于证明其符合 PCI-SIG 规范。而授权一个控制器的成本仅为前者的一小部分，且已经过针对该规范的预验证。然而，IP 集成本身可能具有挑战性，我们将在下文中为订阅者详细介绍。

## 4. RTL Verification 4. RTL 验证

The RTL code then goes through the verification process, crucial to iron out any bugs or design errors within. This is done through simulation, which runs the design in software, applying stimulus, and checking the outputs. Three commercial simulators dominate the market, in order of ubiquity:

随后，RTL 代码将进入验证阶段，这对于消除内部的任何漏洞或设计错误至关重要。验证通过仿真完成，即在软件中运行设计，施加激励并检查输出。三款商业仿真器占据了市场主导地位，按普及程度排序如下：

- VCS (Synopsys): The market leader, known for raw simulation speed and deep integration with the rest of the Synopsys flow.

VCS (Synopsys): 市场领导者, 以极高的仿真速度以及与 Synopsys 其他流程的深度集成而闻名。

- **Xcelium** (Cadence): Cadence's simulator, competitive on multi-core performance and mixed-signal simulation.

Xcelium (Cadence): Cadence 推出的仿真器, 在多核性能和混合信号仿真方面具有竞争力。

- **Questa** (Siemens EDA): Strong in advanced debug and coverage analysis, with deep UVM support.

Questa (Siemens EDA): 在高级调试和覆盖率分析方面表现强劲, 并提供深度的 UVM 支持。

Most large chip companies license at least two of these. Running a full regression suite with tens of thousands of test cases on a complex SoC can consume thousands of CPU core-hours per run. Dedicated on-prem verification servers are usually insufficient these days, with cloud-based simulation on AWS and Azure shoring up short-term demand as teams try to burst capacity during crunch periods before tapeout. The amount of data this generates is also staggering, with multiple Petabytes of disk space required to house just a single chip's entire definition and test items.

大多数大型芯片公司会至少获得其中两种工具的授权。在复杂的 SoC 上运行包含数万个测试用例的全套回归测试, 每次运行可能消耗数千个 CPU 核心小时。如今, 专用的本地验证服务器通常已不足以满足需求, 因此团队在流片前的关键冲刺阶段, 通常会通过 AWS 和 Azure 上的云端仿真来补充短期算力缺口。由此产生的数据量也令人惊叹, 仅存储单个芯片的完整定义和测试项目就需要数 PB 的磁盘空间。


As mentioned above, you will usually find more Verification engineers than any other single role in a chip design house. With chips getting more complex, even more things need to be verified with one another, placing huge demands on the verification staff.

We will dive into what this means for chip design in reality for our subscribers below.

如上所述，在芯片设计公司中，验证工程师的人数通常比任何其他单一职位的都要多。随着芯片变得越来越复杂，需要相互验证的内容也越来越多，这对验证人员提出了巨大的需求。我们将在下文中为订阅者深入探讨这对现实中的芯片设计意味着什么。

The Verification flow takes two paths: Standard DV testing on one end, and Formal Verification with proofs on the other.

验证流程分为两条路径：一端是标准的 DV 测试，另一端是带有证明的正式验证。



Source: SemiAnalysis

## UVM Testbench UVM 测试平台

RTL simulation is structured in UVM (Universal Verification Methodology). This is an industry standard SystemVerilog library and method for building reusable testbenches. Before UVM was standardized by Accellera in 2011, every team rolled their own testbench architecture. UVM brought the industry together by defining a common set of components:

RTL 仿真采用 UVM（通用验证方法学）架构。这是一个工业标准的 SystemVerilog 库和方法，用于构建可重用的测试平台。在 2011 年 Accellera 将 UVM 标准化之前，每个团队都会开发自己的测试平台架构。UVM 通过定义一套通用的组件将业界统一起来：

- **Sequencer:** Generates sequences of transactions and feeds them to the driver. This is where test scenarios are defined.

Sequencer (序列发生器): 生成事务序列并将其发送给驱动器。这是定义测试场景的地方。

- **Driver:** Converts abstract transactions (e.g. “send a 32-byte read request”) into wiggling signals on the design’s input pins.

Driver (驱动器): 将抽象事务 (例如“发送一个 32 字节的读取请求”) 转换为设计输入引脚上的波形信号。

- **Monitor:** Passively observes signals on the design’s interfaces and reconstructs the transactions that occurred.

监视器 (Monitor): 被动地观察设计接口上的信号, 并重建已发生的事务。

- **Scoreboard:** Compares expected outputs from a reference model against actual outputs from the design. Any mismatch is flagged as a bug.

计分板 (Scoreboard): 将参考模型的预期输出与设计的实际输出进行比较。任何不匹配都会被标记为错误 (bug)。

The testbench is used for **constrained random verification**. Instead of hand-writing every test case for directed testing, the engineer defines constraints such as legal address ranges, valid packet formats and protocol rules. The tool then randomly generates millions of input combinations within those bounds. This approach finds corner-case bugs that might not be caught with directed testing. These constrained random regression tests are very resource intensive due to the large sample range, but is usually more effective and fault detection compared to writing directed tests.

测试平台 (Testbench) 用于受约束的随机验证。工程师不再为定向测试手动编写每个测试用例, 而是定义约束条件, 例如合法地址范围、有效数据包格式和协议规则。随后, 工具会在这些范围内随机生成数百万种输入组合。这种方法可以发现定向测试可能无法捕获的极端情况 (corner-case) 错误。由于样本范围巨大, 这些受约束的随机回归测试非常耗费资源, 但与编写定向测试相比, 它们在故障检测方面通常更有效。

## Formal Verification 形式验证

Formal verification takes a fundamentally different approach from simulation. Instead of applying specific inputs and checking outputs, formal tools use mathematical proof engines such as SAT solvers and model checkers to exhaustively prove that a design property holds for all possible inputs and all possible sequences of states. If the property can be violated, the tool produces a concrete counterexample showing exactly how. FV is done with properties, usually **SystemVerilog Assertions (SVA)** that define the expected behavior.

形式验证采用了与仿真截然不同的方法。形式化工具不通过施加特定的输入并检查输出来验证，而是利用诸如 SAT 求解器和模型检查器等数学证明引擎，穷尽式地证明设计属性在所有可能的输入和所有可能的状态序列下均成立。如果该属性可能被违反，工具会生成一个具体的反例，精确展示违规是如何发生的。形式验证（FV）通过属性来完成，通常使用定义预期行为的 SystemVerilog 断言（SVA）。

The leading tools are **JasperGold** (Cadence) and **VC Formal** (Synopsys). Formal verification shines for protocol compliance (e.g. the handshake signal is never asserted for more than 3 cycles), control logic correctness, and security properties (e.g. this register is reserved for software with elevated privileges). However, FV's limitation is scalability. Formal engines hit capacity limits on datapath-heavy designs with wide buses. In practice, formal and simulation are complementary. FV proves critical properties exhaustively on targeted blocks, while simulation covers the full chip at statistical confidence.

主流工具包括 JasperGold (Cadence) 和 VC Formal (Synopsys)。形式验证在协议合规性（例如：握手信号断言绝不超过 3 个周期）、控制逻辑正确性以及安全性属性（例如：该寄存器仅限高权限软件访问）方面表现卓越。然而，形式验证（FV）的局限性在于可扩展性。形式化引擎在处理具有宽总线的数据流密集型设计时会遇到容量瓶颈。在实践中，形式验证与仿真互为补充。形式验证对目标模块的关键属性进行穷尽式证明，而仿真则以统计置信度覆盖全芯片。

## 5. RTL Freeze 5. RTL 冻结

To know when verification is complete, engineers look at several coverage metrics, a quantitative measure of what each testbench has exercised. There are two categories:



为了确定验证何时完成，工程师会查看多项覆盖率指标，这是对每个测试平台（testbench）执行情况的定量衡量。覆盖率分为两类：

**Code coverage** measures structural completeness:

代码覆盖率（Code coverage）衡量结构完整性：

- Line coverage: Has every line of RTL been executed?

行覆盖率（Line coverage）：每一行 RTL 代码是否都已执行？

- Branch coverage: Has every possible branch been taken?

分支覆盖率：是否每个可能的分支都已执行？

- Toggle coverage: Has every signal toggled between 0 and 1?

翻转覆盖率（Toggle coverage）：是否每个信号都经历了 0 到 1 之间的翻转？

- FSM coverage: Has every state and transition in every FSM been visited?

FSM 覆盖率：是否已访问过每个有限状态机（FSM）中的每个状态和转换？

**Functional coverage** measures intent:

功能覆盖率衡量的是设计意图：

- Did we actually test the scenarios we care about?

我们是否真正测试了我们所关心的场景？

- Are there known corner cases that need to be focused on? (for example: concurrent writes to the same address, FIFO buffer is full while interrupt pending)

是否存在需要重点关注的已知极端情况（Corner Cases）？（例如：同时向同一地址写入数据、FIFO 缓冲区已满且中断待处理）

- What are the specific variables to sample to test these scenarios? **Covergroups** are defined in SystemVerilog that contain explicit descriptions of these test cases and,

and tracks if a regression test hits specific variables.

为了测试这些场景，需要采样哪些特定变量？SystemVerilog 中定义了覆盖组 (Covergroups)，其中包含对这些测试用例的显式描述，并跟踪回归测试是否命中了特定变量。

Coverage closure is the final step in the verification process. While 90% of test cases complete quickly, ironing out the remaining 10% with functional coverage takes serious effort, sometimes requiring weeks to write targeted tests while adding or modifying constraints and exclusions in other tests. The more specific and complex a test case is, the more esoteric the knowledge is in whether a design might be susceptible. Design houses tap on their vast history of leanings from previous designs to help inform and prioritize the most important tests.

覆盖率收敛是验证过程的最后一步。虽然 90% 的测试用例可以快速完成，但要通过功能覆盖率解决剩余的 10% 需要付出巨大的努力，有时需要花费数周时间编写针对性测试，同时在其他测试中添加或修改约束和排除项。测试用例越具体、越复杂，关于设计是否可能存在缺陷的知识就越深奥。设计公司会利用其从以往设计中积累的丰富经验，来协助指导并确定最重要测试的优先级。

When all coverage goals are met and no open bugs remain at the target severity level, the project's RTL is frozen. This formal milestone, known as **RTL Freeze**, signals that no more functional changes to the RTL are permitted. From this point forward, any modification must go through a formal process known as **Engineering Change Order** (ECO), requiring re-verification and equivalence checking. ECOs may be required in the design process at a late stage to fix a bug or tweak timings that were not caught earlier. RTL Freeze ensures that the next step, Physical Design, has a concrete base to work from, separating front-end design from back-end physical implementation.

当所有覆盖率目标均已达成，且在目标严重级别上不再存在未解决的漏洞时，项目的 RTL 将被冻结。这一正式的里程碑被称为 RTL 冻结 (RTL Freeze)，标志着不再允许对 RTL 进行任何功能性修改。从此时起，任何修改都必须经过被称为工程变更指令 (ECO) 的正式流程，并需要进行重新验证和等效性检查。在设计过程的后期阶段，可能需要通过 ECO 来修复漏洞或调整早期未发现的时序问题。RTL 冻结确保了下一步的物理设计 (Physical Design) 有一个坚实的基础，从而将前端设计与后端物理实现隔离开来。

While verification is often overlooked as the unglamorous side of chip design, it is critical to the development of new architectures. Designing a chip is easy. Knowing your design works with all possible scenarios is hard.

虽然验证常被视为芯片设计中乏味且易被忽视的一面，但它对于新架构的开发至关重要。设计芯片并不难，难的是确保你的设计在所有可能的场景下都能正常工作。

## 6. Firmware / Software Development

### 6. 固件 / 软件开发

With the chip development process already taking years, software teams cannot afford to wait for silicon to arrive before beginning to write software for it. An operating system, firmware stack, and driver suite need to be substantially ready **before** the first chip comes back from the fab. To write software concurrently with hardware development, engineers rely on **Pre-Silicon Hardware Emulation**. The chip's RTL design is mapped onto large arrays of FPGAs that execute the chips functions at 50MHz speeds. Programmable logic elements in the FPGAs are routed to roughly match the logical configuration of each design, enabling these emulators to run 1000x faster than pure software RTL simulation on a CPU.

由于芯片开发周期长达数年，软件团队无法等到芯片交付后再开始编写软件。在第一颗芯片从晶圆厂返回之前，操作系统、固件栈和驱动程序套件必须基本准备就绪。为了在硬件开发的同时进行软件编写，工程师们依赖于“硅前硬件仿真”(Pre-Silicon Hardware Emulation)。芯片的 RTL 设计被映射到大型 FPGA 阵列上，以 50MHz 的速度执行芯片功能。FPGA 中的可编程逻辑元件经过布线，以大致匹配每个设计的逻辑配置，使这些仿真器的运行速度比 CPU 上的纯软件 RTL 仿真快 1000 倍。

# Software-Defined Hardware-Assisted Verification Portfolio



SYNOPSYS®

Source: Synopsys

The two dominant platforms are **ZeBu** from Synopsys and **Palladium** by Cadence. Synopsys's latest ZeBu-200 clusters can emulate up to 23 billion gates and delivers up to 2x the runtime performance of its predecessor. Cadence's Palladium Z3 can scale to design with up to 48 billion gates and is 1.5x faster than the Z2 generation. These systems allow firmware teams to boot Linux, test firmware, and conduct software validation months before silicon arrives.

两大主流平台是 Synopsys 的 ZeBu 和 Cadence 的 Palladium。Synopsys 最新的 ZeBu-200 集群可仿真高达 230 亿个门电路，其运行性能比前代产品提升了 2 倍。Cadence 的 Palladium Z3 可扩展至支持高达 480 亿个门电路的设计，速度比 Z2 代快 1.5 倍。这些系统使固件团队能够在芯片交付前数月就启动 Linux、测试固件并进行软件验证。

## 7. Physical Design 7. 物理设计

Up to this point, the chip exists only in high-level RTL descriptions. Before Physical Design can take place, a crucial translation step must be done.

到目前为止，芯片仅以高级 RTL 描述的形式存在。在进行物理设计（Physical Design）之前，必须完成一个至关重要的转换步骤。

# Logic Synthesis 逻辑综合 (Logic Synthesis)


Logic synthesis transforms RTL code into a gate-level netlist, a connectivity map of logic gates drawn from a foundry's standard cell library. These synthesis tools parse the RTL code and determines the right combination of logic gates, connected in a certain order, to carry out the functions described by RTL. We will explain what these logic gates are below.

逻辑综合将 RTL 代码转换为门级网表 (gate-level netlist)，即由代工厂标准单元库中的逻辑门组成的连接图。这些综合工具解析 RTL 代码，并确定逻辑门的正确组合及其连接顺序，以实现 RTL 所描述的功能。我们将在下文解释这些逻辑门具体是什么。

On top of this, the synthesizer optimizes the netlist and works within the limits set by the design. It balances timing (can the gates in this circuit complete the operation within 4 clock cycles?), area (how many gates can I squeeze into the area set out by the architectural description?), and power (how many watts worth of power loss from the dynamic and static leakage of these logic gates?). These conflicting goals are somewhat mitigated by techniques such as minimizing repeated logic, sharing logic gates across multiple functions and retiming functions to reduce the load on critical paths. The tool explores thousands of alternative implementations to find the best tradeoff between these demands.

在此基础上，综合器会对网表进行优化，并在设计设定的限制范围内工作。它会平衡时序（该电路中的门电路能否在 4 个时钟周期内完成操作？）、面积（在架构描述所设定的区域内，我能挤进多少个门电路？）以及功耗（这些逻辑门的动态和静态泄漏会产生多少瓦的功率损耗？）。通过减少重复逻辑、跨多个功能共享逻辑门以及重新调整功能时序以减轻关键路径负载等技术，这些相互冲突的目标在一定程度上得到了缓解。该工具会探索数千种替代实现方案，以在这些需求之间找到最佳的折中方案。





Source: Synopsys [Synopsys](#)

The original and still dominant tool is **Design Compiler** from Synopsys, which established the entire category with multiple versions such as NXT and Ultra with greater integration and features **Design Vision**, a graphical interface for engineers to evaluate the synthesis flow. Cadence offers **Genus** as their synthesizer. Synopsys now pushes **Fusion Compiler**, which unifies synthesis with place and route in a unified flow to allow cross-probing between RTL, timing and layout. We cover these Unified EDA flows in more detail below.

最初且目前仍占据主导地位的工具是来自 Synopsys 的 Design Compiler，它开创了整个工具类别，并推出了 NXT 和 Ultra 等多个版本，具备更强的集成度，并提供了 Design Vision 这一图形界面，方便工程师评估综合流程。Cadence 则提供 Genus 作为其综合器。Synopsys 现在正力推 Fusion Compiler，它将综合与布局布线整合在统一的流程中，从而实现 RTL、时序和布局之间的交叉探测。我们将在下文中更详细地介绍这些统一的 EDA 流程。

## Equivalence Checking 等价性检查

Once RTL has been synthesized into a gate-level netlist, you have to check that the synthesis tool did not introduce any bugs. To do this, the design is proven mathematically with **Equivalence Checking**, a formal technique that verifies two representations of a design (RTL and gate netlist) are functionally identical, input for input, output for output. **Formality** (Synopsys) and **Conformal LEC** (Cadence) are the




standard tools.

一旦 RTL 被综合成门级网表，你就必须检查综合工具是否引入了任何错误。为此，需要通过等价性检查（Equivalence Checking）对设计进行数学证明。这是一种形式化验证技术，用于验证设计的两种表示形式（RTL 和门级网表）在功能上是否完全一致，即输入对输入、输出对输出均相等。Formality (Synopsys) 和 Conformal LEC (Cadence) 是该领域的标准工具。

Equivalence checking is run at every major translation step, not just after synthesis. Later gate-to-gate transformations such as clock tree insertion, scan chain stitching, routing optimization, and after every ECO. With each transformation comes a potential vector to introduce errors. Equivalence checking is thus the safety net that catches errors introduced by the tools themselves.

等价性检查不仅在综合后运行，在每一个主要的转换步骤中都会执行。随后的门级到门级转换，如时钟树插入、扫描链缝合、布线优化以及每一次 ECO 之后，都需要进行检查。每一次转换都可能成为引入错误的途径。因此，等价性检查是捕获工具自身所引入错误的安全网。

## Logic Gates 逻辑门




The synthesizer selects from a range of logic gates in a standard cell library, with each logic gate carrying out a Boolean function, translating a given set of binary inputs into an output. The permutations between inputs and outputs are listed in Truth Tables as shown above. The seven basic logic gates are INV and NAND as shown above, AND, OR, NOR, XOR and XNOR. Transistors laid out in standard cells then carry out these operations in the real world, with the output signal voltage pulled up to  $V_{dd}$  for “1” or pulled down to  $V_{ss}$  for “0”.

综合器从标准单元库的一系列逻辑门中进行选择，每个逻辑门执行一个布尔函数，将给定的一组二进制输入转换为输出。输入与输出之间的排列组合如上文真值表所示。七种基本逻辑门包括上文所示的 INV（反相器）和 NAND（与非门），以及 AND（与门）、OR（或门）、NOR（或非门）、XOR（异或门）和 XNOR（同或门）。随后，布局在标准单元中的晶体管在现实世界中执行这些操作，输出信号电压被拉高至  $V_{dd}$  表示“1”，或拉低至  $V_{ss}$  表示“0”。

## Standard Cell Libraries 标准单元库

### CMOS Density Improvement

Logic standard cell area  $\sim$  CPP x Cell height



Source: TSMC TSMC

The above gates are not designed from scratch. Instead, layouts are pulled from a standard cell library, a catalog of pre-designed, pre-characterized logic building blocks provided by the foundry or a third-party IP vendor that obeys all design rules. Each cell has a fixed height and a variable width, tiling neatly in horizontal rows across the

die like bricks in a wall. This standardized, ordered layout with two sites for transistors (usually PMOS above NMOS) and standard pitches for power delivery is crucial to enable automated place and route.

上述门电路并非从零开始设计。相反，它们的版图是从标准单元库中提取的。标准单元库是由代工厂或第三方 IP 供应商提供的预先设计、预先表征的逻辑构建模块目录，且符合所有设计规则。每个单元具有固定的高度和可变的宽度，像墙上的砖块一样整齐地平铺在芯片的水平行中。这种标准化的、有序的版图——通常包含两个晶体管位点（通常 PMOS 在 NMOS 之上）以及用于电源传输的标准间距——对于实现自动布局布线至关重要。

Apart from the basic logic gates above, more complex standard cells are offered with advanced functionality. These combine logical operations into a higher level function that will be used repeatedly throughout a design. Those who work closely with the foundry may also design their own custom standard cells for these complex operations. Examples include (in increasing complexity), AOI, OAI, Registers, Multiplexers, D-Flip Flops and Full Adders. Special purpose cells are also featured throughout the design that deal with functions beyond logical computation. Examples include IO pad cells, Tap cells, Filler and Decap cells, Tie-high and Tie-low cells, power and clock gate cells and level shifters to carry signals over multiple voltage domains. SRAM Macro compilers are also included in each node that generates Static RAM blocks which are the densest memory cells in a logic process.

除了上述基础逻辑门外，库中还提供具有高级功能的更复杂的标准单元。这些单元将逻辑操作组合成在整个设计中反复使用的高级功能。与代工厂密切合作的设计者也可以为这些复杂操作设计自定义标准单元。示例包括（复杂度递增）：AOI、OAI、寄存器、多路复用器、D 触发器和全加器。设计中还包含处理逻辑计算之外功能的特殊用途单元。示例包括 IO pad 单元、Tap 单元、Filler 和 Decap 单元、Tie-high 和 Tie-low 单元、电源和时钟门控单元，以及用于跨多个电压域传输信号的电平转换器。每个工艺节点还包含 SRAM Macro 编译器，用于生成静态随机存取存储器模块，这是逻辑工艺中密度最高的存储单元。

A modern standard cell library for an advanced node like TSMC N2 contains tens of thousands of individual cells. Each logic gate will have multiple layout options to choose from depending on wire routing and signal pin accessibility. Each gate will also come in several drive strength options, where higher drive strength cells are used when its output needs to drive multiple input signals. These are used selectively as

higher drive current transistors have greater leakage power. Modern TSMC process nodes have upwards of 6 Threshold Voltage options to enable fine tuning of signal reliability and power. One the problems with Intel's initial 18A was that there were only 3 VT options vs 6 at TSMC meaning even if Intel performs better at a given point, the design has a harder time being on the pareto optimal curve across all VTs. Intel's 18AP fixes this.

像 TSMC N2 这样先进节点的现代标准单元库包含数万个独立的单元。根据布线情况和信号引脚的可访问性，每个逻辑门都有多种版图选项可供选择。每个门还提供多种驱动强度选项，当输出需要驱动多个输入信号时，会使用高驱动强度的单元。由于高驱动电流晶体管的漏电功耗更大，因此需要有选择性地使用这些单元。现代 TSMC 工艺节点拥有多达 6 种阈值电压（VT）选项，以实现信号可靠性和功耗的精细调整。Intel 最初的 18A 工艺存在的问题之一是仅提供 3 种 VT 选项，而 TSMC 为 6 种，这意味着即使 Intel 在某个特定点表现更好，设计也很难在所有 VT 范围内达到帕累托最优曲线。Intel 的 18AP 修复了这一问题。

Synthesis tools pick the right size for each cell based on timing and power constraints. With millions of permutations to choose from, EDA tools in both Synthesis and Physical Design are mandatory to navigate this optimization problem.

综合工具根据时序和功耗约束为每个单元选择合适的尺寸。面对数百万种可能的排列组合，在综合和物理设计阶段都必须使用 EDA 工具来解决这一优化问题。

With Moore's law slowing, newer nodes with novel area scaling techniques have exacerbated the layout complexity further. TSMC's N3 FinFlex and N2 NanoFlex mix standard cells of different cell heights together, adding another layer of design constraints. Multiple library options are also offered for each node, such as HD high density and HP High Performance with different standard cell dimensions and power characteristics. Designers can mix and match different libraries for each region of a chip to extract the best PPA from the new process. Apple for example uses TSMC's 3-2 FinFlex for its high performance CPU cores, with the rest of the die adopting the



much denser and lower power 2-1 FinFlex library.

随着摩尔定律的放缓，采用新型面积缩放技术的新工艺节点进一步加剧了布局的复杂性。台积电（TSMC）的 N3 FinFlex 和 N2 NanoFlex 将不同高度的标准单元混合在一起，增加了另一层设计约束。每个节点还提供多种库选项，例如具有不同标准单元尺寸和功耗特性的 HD 高密度库和 HP 高性能库。设计人员可以为芯片的每个区域混合搭配不同的库，从而从新工艺中榨取最佳的 PPA。例如，苹果公司在其高性能 CPU 核心上使用了台积电的 3-2 FinFlex，而芯片的其余部分则采用了密度更高、功耗更低的 2-1 FinFlex 库。

The standard cell library is the foundry's main commercial interface with chip designers. It encodes not just logic functions but the foundry's manufacturing capabilities, design rules, and process characteristics into a form that EDA tools can consume. When a fabless company “ports” a design to a new foundry, migrating the standard cell library is the first and most impactful step, and the one that triggers the most re-work across the entire tool flow.


标准单元库是代工厂与芯片设计师之间主要的商业接口。它不仅包含了逻辑功能，还将代工厂的制造能力、设计规则和工艺特性编码成 EDA 工具可以处理的形式。当一家无晶圆厂（fabless）公司将设计“移植”到新的代工厂时，迁移标准单元库是第一步，也是影响最深远的一步，它会触发整个工具流程中最大规模的返工。

## Process Corners **工艺角**

Physical Design tools must also account for the effects of the real world, namely on variability in manufacturing. A cell's speed and power consumption may shift significantly depending on variations in Process, Voltage and Temperature (PVT). Libraries are therefore characterized across a range of PVT corners as shown below.

物理设计工具还必须考虑现实世界的影响，即制造过程中的变异性。单元的速度和功耗可能会随着工艺、电压和温度（PVT）的变化而发生显著偏移。因此，库会在一系列 PVT 角点（Corners）下进行特征化，如下所示。





Source: Keysight 

TT (typical-typical) is the nominal case. FF (fast-fast) means both NMOS and PMOS transistors came out faster than nominal. SS (slow-slow) means both came out slower. FS and SF capture skew between NMOS and PMOS. These cases are particularly tricky because they affect circuit balance in ways that symmetric corners do not.

TT (typical-typical) 是标称情况。FF (fast-fast) 意味着 NMOS 和 PMOS 晶体管都比标称情况快。SS (slow-slow) 意味着两者都比标称情况慢。FS 和 SF 则捕捉了 NMOS 和 PMOS 之间的失配。这些情况尤其棘手，因为它们会以对称工艺角所不具备的方式影响电路平衡。

Voltage is varied around the nominal supply. For example, 0.75 V nominal with corners at 0.675 V (-10%) and 0.825 V (+10%). Lower voltage means slower switching and at lower power. Temperature ranges depend on the application. Consumer chips might cover 0°C to 105°C, while automotive-grade libraries must be characterized from -40°C to 125°C. Designs must meet timing at the slow corner and stay within its




power budget at the leaky corner.

电压在标称电源电压附近波动。例如，标称电压为 0.75 V，工艺角设在 0.675 V (-10%) 和 0.825 V (+10%)。较低的电压意味着开关速度更慢且功耗更低。温度范围取决于具体应用。消费级芯片可能涵盖 0°C 到 105°C，而汽车级库必须在 -40°C 到 125°C 范围内进行特征化。设计必须在慢速工艺角（slow corner）满足时序要求，并在漏电工艺角（leaky corner）保持在功耗预算之内。

## Metal Interconnect 金属互连

Standard cells define the active transistor layer at the bottom of the metal interconnect. The lowest metals, M0 and M1, are used within cells to connect transistors to pins. These are thin, high-resistance wires, fine for short distances inside a cell but not for routing signals across the die. Each alternating metal layer has lines that are perpendicular in orientation, with odd layers for north-south routing and even layers for east-west routing.

标准单元在金属互连层的底部定义了有源晶体管层。最低层的金属 M0 和 M1 在单元内部用于将晶体管连接到引脚。这些是细且高电阻的导线，适用于单元内部的短距离连接，但不适用于跨芯片的信号布线。每个交替的金属层都具有相互垂直的走线方向，其中奇数层用于南北向布线，偶数层用于东西向布线。



Source: Intel Intel

Semi-global metals (M3 through M5) handle block-level routing, connecting standard cells to each other within a functional unit. These are progressively thicker and wider,

reducing resistance for longer runs. The thickest top metals above are reserved for the power distribution network and global clock trees. The total number of metal layers depends on the complexity of the chip, from 10 layers in the low-cost mobile SoC to 19 layers in a high performance AI processor. Backside Power Delivery schemes such as Intel's 18A above unlock another dimension for routing tools to utilize by segregating power and signal lines, reducing parasitic capacitance and signal interference.

半全局金属层（M3 至 M5）负责处理块级布线，将功能单元内的标准单元相互连接。这些金属层逐渐变厚变宽，以降低长距离布线的电阻。上方最厚的顶层金属则预留给电源分配网络和全局时钟树。金属层的总数取决于芯片的复杂程度，从低成本移动 SoC 的 10 层到高性能 AI 处理器的 19 层不等。如上所述，英特尔 18A 等背面供电方案通过将电源线和信号线分离，为布线工具开辟了另一个可利用的维度，从而降低了寄生电容和信号干扰。

## Process Design Kit 工艺设计套件

All of the parameters above are detailed and provided by the foundry in a Process Design Kit (PDK). This dense package of files, models, and rules translates a foundry's manufacturing process into terms that EDA tools can understand.

上述所有参数均由代工厂在工艺设计包（PDK）中详细列出并提供。这套包含文件、模型和规则的密集数据包，将代工厂的制造工艺转化为 EDA 工具可以理解的术语。

Each PDK comes with the following information, each having their own industry standard file formats:

每个 PDK 都包含以下信息，且每种信息都有其各自的行业标准文件格式：

- The LEF file (Library Exchange Format) provides the physical description of the standard cells, including pin locations, metal blockages (areas to avoid) and the cell boundary. Place and route tools use LEF as it contains sufficient geometry to position cells and connect wires.

LEF 文件（库交换格式）提供了标准单元的物理描述，包括引脚位置、金属阻挡层（避让区域）和单元边界。布局布线工具使用 LEF，因为它包含了足以进行单元定位和布线的几何信息。

- LIB files (Liberty Format) contain timing arcs, defining how long a signal takes to propagate and decay, power consumption tables, and noise immunity data. Separate .lib files are made for each PVT corner combination.

LIB 文件（Liberty 格式）包含时序弧，定义了信号传播和衰减所需的时间，以及功耗表和抗噪声能力数据。针对每种 PVT Corner 组合都会制作独立的 .lib 文件。

- SPICE models (Simulation Program with Integrated Circuit Emphasis) capture detailed analog behavior of the transistors such as threshold voltages, leakage currents, capacitances, and how they shift across process corners. Other simulation models of passive devices such as resistors, inductors and capacitors may be included.

SPICE 模型（集成电路重点仿真程序）捕捉了晶体管的详细模拟行为，例如阈值电压、漏电流、电容，以及它们在不同工艺角（process corners）下的偏移情况。此外，还可能包含电阻、电感和电容等无源器件的其他仿真模型。

- Parameterized cells (P-cells) are used in the analog I/O regions, generating transistor dimensions based on certain variables that the designer can adjust. This is vastly more efficient than drawing out every cell for each variable dimension.

参数化单元（P-cells）被用于模拟 I/O 区域，它可以根据设计者可调的特定变量生成晶体管尺寸。这比为每个变量尺寸单独绘制每个单元要高效得多。

- Parasitic extraction decks tell extraction tools how to calculate the resistance and capacitance of every wire and via in the layout. These come in multiple corners, RC-max (worst-case slow interconnect) and RC-min (best-case fast interconnect) to account for manufacturing variation.

寄生参数提取文件（Extraction decks）指导提取工具如何计算版图中每条连线和通孔的电阻与电容。这些文件包含多个工艺角，如 RC-max（最差情况下的慢速互连）和 RC-min（最佳情况下的快速互连），以考虑制造工艺的波动。


- The Design Rule Manual (DRM) contains thousands of geometric constraints such as minimum spacing, minimum width, enclosure rules, blockage and density requirements. Physical Verification parameters are provided, detailed in the next

section.

设计规则手册（DRM）包含数千条几何约束，例如最小间距、最小宽度、包围规则、阻挡层以及密度要求。手册中还提供了物理验证参数，详见下一节。

GDS (GDSII Stream) layout parameters may also be provided, giving an example of the full, exact physical layout for every transistor and metal layer, including diffusion regions, polysilicon gates, contacts, and vias. This file is what ultimately goes to the foundry for tapeout.

可能还会提供 GDS (GDSII Stream) 版图参数，给出每个晶体管和金属层的完整、精确的物理版图示例，包括扩散区、多晶硅栅极、接触孔和通孔。该文件最终将发送至代工厂进行流片。



Source: Cadence Cadence

A standard cell as seen in a GDS viewer. Each color represents a different fabrication layer. Polysilicon gate (red), metal interconnect (blue), and vertical contacts/vias



between layers (blue squares).

在 GDS 查看器中看到的标准单元。每种颜色代表不同的制造层。多晶硅栅极（红色）、金属互连线（蓝色）以及层间的垂直接触孔/通孔（蓝色方块）。

## PDK Versions **PDK 版本**

As the process is being developed, customers begin to design test chips and learn about the process characteristics. As such, PDKs go through many revisions on its way to version 1.0 for production tapeouts. The major PDK milestone releases are:

随着工艺的开发，客户开始设计测试芯片并了解工艺特性。因此，PDK 在发布用于生产流片的 1.0 版本之前会经历多次修订。PDK 的主要里程碑版本包括：

- PDK 0.1 / PDK 0.3: Built entirely from TCAD simulation (Technology Computer-Aided Design), not real silicon data. Used for early architecture exploration and pathfinding with wide uncertainty bands. Only available to anchor customers under a Joint Development Agreement (JDA), up to 2 years before volume production. Standard cell dimensions are defined but libraries at this stage are minimal.

PDK 0.1 / PDK 0.3: 完全基于 TCAD 仿真（工艺计算机辅助设计）构建，而非真实的硅片数据。用于早期架构探索和路径寻找，具有较大的不确定性范围。仅在量产前至多 2 年，通过联合开发协议（JDA）提供给核心客户。标准单元尺寸已确定，但此阶段的库非常精简。

- PDK 0.5: Incorporates early silicon test data from short-loop wafers and initial test chips. This is where real chip design typically begins, starting with test chips. Models are substantially refined, but PVT corner coverage is still incomplete. Data for process corners will be available, but not full silicon models. The standard cell library is growing but not finalized. Regular foundry customers gain access at this stage.

PDK 0.5: 包含来自短流程晶圆和初始测试芯片的早期硅片测试数据。这通常是真实芯片设计的起点，从测试芯片开始。模型经过了大幅改进，但 PVT Corner（工艺、电压、温度偏差）覆盖仍不完整。将提供工艺偏差数据，但尚无完整的硅片模型。标准单元库正在扩充但尚未定型。普通代工厂客户在此阶段获得访问权限。

- PDK 0.9: Provides full silicon characterization across all PVT corners. Near-final Physical Verification decks are prepared, validated against millions of test structures. All major EDA tool vendors will have broad support and have IP available.

PDK 0.9: 提供跨所有 PVT Corner 的完整硅片特性分析。准备好近乎最终版的物理验证（Physical Verification）规则文件，并经过数百万个测试结构的验证。所有主要的 EDA 工具供应商都将提供广泛支持并备有可用 IP。

- PDK 1.0: The final production-ready package. Complete standard cell libraries with all Vt variants and circuit models are available. This is what production tapeouts use. PDK 1.0 is a foundry milestone that signals manufacturing readiness.

PDK 1.0: 最终的可量产版本。提供包含所有 Vt（阈值电压）变体和电路模型的完整标准单元库。这是生产级 Tapeout（流片）所使用的版本。PDK 1.0 是代工厂的一个里程碑，标志着制造准备就绪。

Between each major version, the foundry issues point releases that tighten design rules, narrows model uncertainty and adds DFM (Design for Manufacturability) hotspot rules discovered from real fabrication data.

在每个主要版本之间，代工厂会发布点版本更新，以收紧设计规则、缩小模型不确定性，并增加根据实际制造数据发现的 DFM（可制造性设计）热点规则。

## PDK Licensing and Access **PDK 许可与获取**

Foundry PDKs are universally NDA-protected, and access is controlled through a rigid tiered system that reflects both business relationships and IP security.

代工厂 PDK 普遍受保密协议（NDA）保护，其获取权限通过严格的分层体系进行控制，这反映了业务关系和知识产权（IP）安全。

- Tier 1: Anchor / JDA customers (eg Apple for Mobile, AMD / Nvidia for HPC). These companies are heavily involved with process technology definition and co-development, receiving access 3+ years before production. They push the design rules and may “bankroll” the foundry to develop specific process features. In



return, they get first-mover advantage and the deepest collaboration.

第一层级：锚点/联合开发（JDA）客户（例如移动领域的 Apple，高性能计算领域的 AMD / Nvidia）。这些公司深度参与工艺技术的定义和共同开发，在量产前 3 年以上即可获得访问权限。他们推动设计规则的制定，并可能为代工厂开发特定工艺特性提供资金支持。作为回报，他们获得了先发优势和最深度的协作。

- Tier 2: EDA and IP partners. Synopsys, Cadence, Siemens EDA, and IP vendors like Arm receive early PDK access for tool validation and IP development. Their standard cell libraries, memory compilers, and interface IP must be qualified and available for fabless design houses to use, requiring several months of engineering lead time to implement.

第二层级：EDA 和 IP 合作伙伴。Synopsys、Cadence、Siemens EDA 以及 Arm 等 IP 供应商会提前获得 PDK 访问权限，用于工具验证和 IP 开发。他们的标准单元库、存储器编译器和接口 IP 必须经过验证并提供给无晶圆厂（fabless）设计公司使用，这通常需要数月的工程提前期来实现。

- Tier 3: Standard foundry customers. These companies design chips on an established process but do not influence its development. They accept the design rules as-is.

第三层级：标准代工厂客户。这些公司在成熟的工艺上设计芯片，但不会影响工艺的开发。他们直接接受现有的设计规则。

- Tier 4: Academic and research. Heavily restricted. Universities may receive limited PDK subsets for research, and usually many years after the process has matured.

第 4 层：学术与研究。受到严格限制。大学可能会获得用于研究的有限 PDK 子集，且通常是在工艺成熟多年之后。

Even then, Tier 1 customers never see the foundry's physical recipe. Chip designers see electrical characteristics, with no need to know the exact doping profiles, process temperatures, or etch chemistries. The PDK is an abstraction layer that makes the

fabless model possible, protecting IP between foundry and customer.

即便如此，一级客户也从未见过代工厂的物理配方。芯片设计人员看到的是电气特性，无需了解精确的掺杂分布、工艺温度或蚀刻化学成分。PDK 是一个抽象层，它使无晶圆厂（fabless）模式成为可能，保护了代工厂与客户之间的知识产权。

## PDK Timeline: Intel 18A

### PDK 时间线：Intel 18A

- Sep 2022: PDK 0.3 now in hands of early design customers, test chips are under design for end of year tapeout.

2022 年 9 月：PDK 0.3 现已交付给早期设计客户，测试芯片正在设计中，计划于年底流片。

- Mar 2023: Shared the engineering release of PDK 0.5 (process design kit) with our lead customers and expect to have the final production release in the next few weeks.

2023 年 3 月：向领先客户分享了 PDK 0.5（工艺设计套件）的工程版本，并预计在未来几周内发布最终生产版本。

- Sep 2023: PDK 0.9 is “imminent” and many 18A test chips and shuttles for both internal and foundry customers are running in the fab.

2023 年 9 月：PDK 0.9 “发布在即”，许多针对内部和代工客户的 18A 测试芯片及穿梭班次（shuttles）正在晶圆厂内运行。

- Jul 2024: Intel released PDK 1.0 to foundry customers

2024 年 7 月：Intel 向代工客户发布了 PDK 1.0

- Jan 2026: Launch of Intel internal Panther Lake CPUs on 18A

2026 年 1 月：基于 18A 工艺的 Intel 内部 Panther Lake CPU 发布

## Open Source PDK 开源 PDK

As foundry PDKs are closely guarded secrets, few will ever have the access to learn and use them. To date, only one production PDK has been released as fully open source. In 2020, Google partnered with SkyWater Technology to release the SKY130 PDK, a complete, manufacturable design kit for SkyWater's 130nm process. SKY130, originally developed by Cypress Semiconductor, includes SPICE models, DRC/LVS decks, standard cell libraries, and IO cells. Everything needed to go from RTL to fabricated silicon, published on GitHub under an Apache 2.0 license.


由于代工厂的 PDK 是严密保护的商业机密，很少有人能获得学习和使用它们的机会。到目前为止，只有一款量产级 PDK 实现了完全开源。2020 年，Google 与 SkyWater Technology 合作发布了 SKY130 PDK，这是一套针对 SkyWater 130nm 工艺的完整且可制造的设计套件。SKY130 最初由 Cypress Semiconductor 开发，包含 SPICE 模型、DRC/LVS 验证脚本、标准单元库和 IO 单元。从 RTL 到芯片制造所需的一切资源，均已在 GitHub 上以 Apache 2.0 协议发布。

While the technology is over 20 years old, it has been crucial for education and open source research as well as training open source models in Physical Design. The open-source EDA ecosystem includes OpenROAD for place-and-route and OpenLane for automated flow orchestration. Google has funded several multi-project wafer (MPW) shuttle runs, manufacturing community designed chips for free, enabling students to complete the entire RTL-to-GDSII flow without a single commercial license. Unfortunately Google pulled funding for these areas. Check out <https://wafer.space>.

虽然这项技术已有 20 多年的历史，但它对于教育、开源研究以及物理设计领域开源模型的训练都至关重要。开源 EDA 生态系统包括用于布局布线的 OpenROAD 和用于自动化流程编排的 OpenLane。Google 曾资助过多次多项目晶圆（MPW）班次，免费制造社区设计的芯片，使学生能够在没有任何商业授权的情况下完成从 RTL 到 GDSII 的整个流程。遗憾的是，Google 已经撤回了对这些领域的资助。请访问 <https://wafer.space> 了解更多信息。

## Physical Design Tools and Functions

### 物理设计工具与功能



Source: Synopsys Synopsys

The two dominant physical design platforms are **IC Compiler II** (Synopsys) and **Innovus** (Cadence). Together, they handle virtually every advanced-node chip taped out today. The choice between them often comes down to which tool the team has the most expertise with. Decades of institutional knowledge built over dozens of projects make switching difficult and expensive. These tools take all of the parameters and considerations above along with the synthesized netlist and foundry PDK to generate a complete GDSII layout for tapeout. Engineers use these EDA tools for:

目前最主流的两大物理设计平台是 IC Compiler II (Synopsys) 和 Innovus (Cadence)。这两款工具共同处理了当今几乎所有先进工艺节点的芯片流片。在两者之间做出选择，通常取决于团队对哪种工具更具经验。数十个项目积累下来的数十年机构知识，使得切换工具变得既困难又昂贵。这些工具结合上述所有参数和考量因素，连同综合后的网表和代工厂 PDK，生成用于流片的完整 GDSII 版图。工程师使用这些 EDA 工具进行：

- **Floorplanning:** Defines the chip's physical outline and internal organization, akin to deciding the exact position of all walls and rooms in a house. Based off the high-level architectural planning above. Large macro blocks are placed (SRAM,

PLLs, analog IP) and IO pin locations are demarcated along the chip perimeter.

底层规划 (Floorplanning): 定义芯片的物理轮廓和内部结构, 类似于确定房屋中所有墙壁和房间的具体位置。该步骤基于上述的高级架构规划。在此阶段, 需要放置大型宏单元 (如 SRAM、PLL、模拟 IP), 并在芯片周边划定 IO 引脚的位置。

- **Power Planning:** The **power distribution network (PDN)** delivers supply voltage to every transistor on the chip. Horizontal and vertical power stripes with via stacks for a power grid mesh to feed power rails for transistors to access. The PDN must handle multiple voltage domains for power efficiency regions, along with the ability to selectively turn off power to regions when idling, a technique known as **power gating**. The goal is to minimize IR drop from wire resistance. Voltage supply must be uniform across the entire chip area, especially in dense power grid hotspots such as Floating Point units and Tensor Cores.

电源规划: 电源分配网络 (PDN) 为芯片上的每个晶体管提供工作电压。水平和垂直的电源条通过过孔堆叠形成电源网格, 为晶体管接入的电源轨供电。PDN 必须能够处理多个电压域以实现能效分区, 并具备在闲置时选择性关闭特定区域电源的能力, 这种技术被称为电源门控。其目标是最大限度地减少由导线电阻引起的 IR 压降。电压供应必须在整个芯片区域内保持均匀, 特别是在浮点运算单元和张量核心 (Tensor Cores) 等高密度电源网格热点地区。

- **Placement:** Assigns physical locations to millions of standard cells within the floorplan. Global placement optimizes total wirelength and minimizes congestion hotspots. Detailed placement then legalizes these positions, locking the standard cells into defined grid rows and resolving overlaps. Modern placers are timing-driven and congestion-aware, trading off wirelength for routability. The image above shows the cell density distribution in a design. Infrastructure cells are also added such as fillers, tap cells and spare cells reserved for ECO debugging.

布局: 在规划好的平面布局内为数百万个标准单元分配物理位置。全局布局旨在优化总连线长度并减少拥塞热点。随后, 详细布局将这些位置合法化, 将标准单元锁定在定义的栅格行中并解决重叠问题。现代布局工具采用时序驱动和拥塞感知技术, 在连线长度与可布线性之间进行权衡。上图显示了设计中的单元密度分布。此外, 还会添加基础设施单元, 如填充单元 (fillers)、偏置单元 (tap cells) 以及预留用于 ECO 调试的备用单元 (spare cells)。



- **Routing:** Connects all placed cells to the metal interconnect stack above for both signal and power. Local, semi-global and global routing is handled by the PnR tool to meet both power delivery and clock timing budgets. The tool will spread out standard cells and lower area utilization in locations with high routing congestion where cell pin density is too high for the metal stack to handle.

布线：将所有已放置的单元连接到上方的金属互连层，以实现信号和电源连接。PnR 工具负责处理局部、半全局和全局布线，以满足功耗输送和时钟时序预算。在布线拥塞程度较高、且单元引脚密度过大导致金属层无法处理的区域，该工具会分散标准单元并降低面积利用率。

- **Clock Tree Synthesis:** Builds the distribution network that delivers the clock signal from a central source to every transistor in the clock domain, ensuring synchronous operation. The goal here is minimizing skew, the difference in clock arrival times between any pair of functional units. Timing budget would be eaten up to account for variations in skew. Different blocks in a chip may operate in different clock domains for power efficiency regions, requiring clock domain crossings for signals to traverse between them.

时钟树综合（Clock Tree Synthesis）：构建将时钟信号从中心源传输到时钟域内每个晶体管的分配网络，以确保同步运行。其目标是最小化时钟偏斜（skew），即任何一对功能单元之间时钟到达时间的差异。时钟偏斜的波动会消耗时序预算。为了提高能效，芯片中的不同模块可能运行在不同的时钟域中，这需要通过跨时钟域（clock domain crossings）处理来使信号在它们之间传输。

- **Design for Test (DFT):** Inserts testability structures into the physical implementation. Scan chains and DFT access pads are added with long shift registers so test equipment can load test patterns and retrieve results. MBIST (Memory Built-In Self Test) adds on-chip test logic for memory chips. ATPG (Automatic Test Pattern Generation) software computes the optimal set of input vectors for maximum test coverage.

可测试性设计（DFT）：在物理实现中插入可测试性结构。通过添加扫描链（scan chains）和 DFT 访问焊盘以及长移位寄存器，使测试设备能够加载测试向量并检索结果。MBIST（存储器内建自测试）为存储芯片添加了片上测试逻辑。ATPG（自动测试向量生成）软件则负责计算最优的输入向量集，以实现最大的测试覆盖率。



Physical design is done over multiple optimization loops. After each major step, the tools run timing optimization, inserting buffers on long or high-fanout nets, resizing gates (swapping a standard cell for a larger or smaller variant), applying intentional skew (shifting clock arrival times to help critical paths), remapping logic (restructuring gate-level topology), and hold fixing (adding delay cells to prevent data from arriving too early). The tools iterate through these techniques to arrive at the best PPA.

物理设计是通过多个优化循环完成的。在每个主要步骤之后，工具都会进行时序优化，包括在长网络或高扇出网络上插入缓冲器、调整门电路大小（将标准单元更换为更大或更小的变体）、应用有意偏斜（移动时钟到达时间以辅助关键路径）、重映射逻辑（重构门级拓扑结构）以及修复保持时间（添加延迟单元以防止数据过早到达）。工具通过迭代这些技术来达到最佳的 PPA。

## 8. Signoff 8. 签核 (Signoff)

Much like RTL design and verification is locked in with RTL Freeze, Physical designs undergo **Physical Verification (PV)** before Signoff locks the GDSII in place. Signoff proves design will actually work in both function and manufacturability, across every operating condition and power scenario. Synopsys offers **IC Validator**, Cadence offers **Pegasus**, and Siemens has **Calibre**. All three have deep foundry entrenchment that are indispensable for the fabless design houses in properly utilizing a given foundry process. The main components to clear before Signoff are:

就像 RTL 设计和验证通过 RTL Freeze（RTL 冻结）来锁定一样，物理设计在签核（Signoff）锁定 GDSII 之前，必须经过物理验证（PV）。签核证明了设计在所有工作条件和功耗场景下，在功能和可制造性方面都能切实可行。Synopsys 提供 IC Validator，Cadence 提供 Pegasus，而 Siemens 则拥有 Calibre。这三款工具都与晶圆厂有着深厚的合作基础，对于无晶圆厂（Fabless）设计公司正确利用特定晶圆厂工艺而言是不可或缺的。在签核前需要通过的主要环节包括：

- **Design Rule Check (DRC):** Verifies geometric constraints from the foundry. Minimum wire widths, spacings between shapes, imposed blockages and pattern density limits are obeyed. Advanced nodes impose thousands of such rules, with

design rule manuals exceeding 1000 pages.

设计规则检查 (DRC): 验证是否符合晶圆厂的几何约束。确保遵守最小线宽、图形间距、强制避让区 (Blockages) 以及图形密度限制。先进工艺节点涉及数千条此类规则, 其设计规则手册 (DRM) 往往超过 1000 页。

- Layout vs. Schematic (LVS): Proves that the physical layout implements exactly the intended circuit connectivity, with no unintended shorts or opens. A netlist is extracted from the physical layout and compared to the original gate-level netlist to ensure logical correctness.

版图与原理图一致性检查 (LVS): 证明物理版图精确实现了预期的电路连接, 不存在任何非预期的短路或断路。从物理版图中提取出网表, 并将其与原始门级网表进行对比, 以确保逻辑正确性。

- Electrical Rule Check (ERC): Catches floating nodes, shorted supplies and other electrical violations in the power grid. ERC checks for current density limits and ESD (electrostatic discharge) reliability.

电气规则检查 (ERC): 捕捉电源网络中的悬空节点、电源短路以及其他电气违规行为。ERC 还会检查电流密度限制和 ESD (静电释放) 可靠性。

- Static Timing Analysis (STA): Verifies that every timing path in the design meets its setup and hold constraints. Tools like **PrimeTime** by Synopsys and **Tempus** from Cadence have more advanced functionality that look at timing across PVT corners and DVFS (Dynamic Voltage Frequency Scaling) curves, known as MCMM (Multi-Corner Multi-Mode) analysis.

静态时序分析 (STA): 验证设计中的每条时序路径是否满足其建立时间和保持时间约束。Synopsys 的 PrimeTime 和 Cadence 的 Tempus 等工具具有更高级的功能, 可以查看跨 PVT 工艺角和 DVFS (动态电压频率调节) 曲线的时序, 这被称为 MCMM (多工艺角多模式) 分析。

- Power Signoff: IR drop analysis verifies that the power distribution network delivers adequate voltage to every transistor under both static (average) and dynamic (switching) conditions. Wire currents are checked for electromigration limits, where electrons gradually move copper in the metal lines causing shorts and opens over time. Synopsys **RedHawk-SC** and Cadence **Voltus** tools handle

power integrity verification.

功耗签核：IR 压降分析用于验证供电网络在静态（平均）和动态（翻转）条件下，是否都能为每个晶体管提供充足的电压。同时还会检查导线电流是否符合电迁移限制，电迁移是指电子逐渐移动金属线中的铜，随着时间的推移会导致短路和断路。Synopsys RedHawk-SC 和 Cadence Voltus 等工具负责处理电源完整性验证。

Occasionally, ECO flows may have to be implemented to clear Signoff. A functional ECO fixes logic bugs by repurposing pre-placed spare cells, only requiring new masks in the lowest metal layer. A timing ECO, guided by PrimeTime, resizes or remaps cells to close timing, and may require a full mask set change. Each type trades flexibility for cost and turnaround time, the mechanics of which will be explained in the Steppings section below.

有时，必须实施 ECO 流程以通过签核（Signoff）。功能性 ECO 通过重新利用预先放置的备用单元（spare cells）来修复逻辑漏洞，仅需在最低金属层更换新掩模。时序 ECO 则在 PrimeTime 的指导下，通过调整单元尺寸或重新映射单元来闭合时序，这可能需要更换全套掩模。每种类型都在灵活性、成本和周转时间之间进行权衡，其具体机制将在下文的“步进（Steppings）”章节中详细说明。

## Unified Flows **统一流程**

Traditionally, each stage of the chip design flow operated as a silo. The RTL team wrote code and handed it to the synthesis team. Synthesis produced a netlist and handed it to place and route. They generated a layout and handed it to signoff. Each tool ran independently, creating a coordination nightmare when issues are found downstream that needed a fix from layout or even at RTL. These inevitable bugs would throw programs off by many weeks as these verticals are brought together to implement fixes. To address this, EDA vendors have unified the physical design

process so every step is factored in concurrently rather than sequentially.

传统上，芯片设计流程的每个阶段都是孤立运行的。RTL 团队编写代码并将其移交给综合团队。综合团队生成网表并将其移交给布局布线团队。布局布线团队生成版图并将其移交给签核团队。每个工具都独立运行，当在下游发现需要从版图甚至 RTL 层面进行修复的问题时，就会产生协调方面的噩梦。由于需要整合这些垂直环节来实施修复，这些不可避免的缺陷会使项目进度推迟数周。为了解决这一问题，EDA 供应商已经统一了物理设计流程，使每个步骤都能并发考虑，而非按顺序进行。

Synopsys **Fusion Compiler** was the first major answer to this problem. Built on a single data model, it unifies synthesis, place and route, and timing analysis into one engine. Cadence responded with **iSpatial**, which embeds the Innovus placement and optimization engines directly into the Genus synthesis tool. Unified flows bring the “**Shift Left**” philosophy. Have signoff-quality analysis earlier in the design process to avoid late surprises.

Synopsys Fusion Compiler 是针对这一问题的首个主要解决方案。它基于单一数据模型构建，将综合、布局布线和时序分析统一到同一个引擎中。Cadence 则推出了 iSpatial 作为回应，该技术将 Innovus 布局和优化引擎直接嵌入到 Genus 综合工具中。统一的流程带来了“左移 (Shift Left)”理念，即在设计过程的早期进行签核级质量的分析，以避免后期出现意外。

## 9. Tapeout 9. 流片 (Tapeout)

Once every signoff check passes, the design is exported as a GDSII or OASIS file, the industry standard formats describing the geometric shapes on every layer. This file is sent to the foundry, marking the milestone known as the tapeout, a term coined from the old days where GDSII files with sent on reels of magnetic tape. With this, the first mask set is designed. OPC (Optical Proximity Correction) algorithms modify the mask patterns with SRAFs (Sub-Resolution Assist Features) to compensate for optical distortion during lithography. Then mask making begins, followed by wafer

fabrication.

一旦通过所有签核检查，设计就会被导出为 GDSII 或 OASIS 文件，这是描述每一层几何形状的行业标准格式。该文件被发送至代工厂，标志着被称为“流片”的里程碑。这个术语源于早期将 GDSII 文件存储在磁带卷上发送的时代。至此，第一套掩模组设计完成。OPC（光学邻近效应修正）算法会利用 SRAF（亚分辨率辅助图形）对掩模图形进行修改，以补偿光刻过程中的光学畸变。随后开始掩模制造，紧接着是晶圆制造。

## 10. Fabrication & Packaging

### 10. 制造与封装

The fabrication process, from tapeout to first silicon arriving, typically takes 8-12 weeks. However, these first wafers can be expedited by purchasing Hot Lots, with higher priority wafer shuttles cutting down cycle times by several weeks. This allows post-silicon bring up engineers earlier access to begin debugging the new chip.

从交付设计（tapeout）到首批芯片（first silicon）送达，制造过程通常需要 8 到 12 周。然而，通过购买“热片”（Hot Lots）可以加快这些首批晶圆的进度，高优先级的晶圆流片服务能将周期缩短数周。这使得硅后调试工程师能够更早地开始对新芯片进行调试。

Before debugging can happen, silicon needs to be packaged to protect the fragile die and break out the tiny I/O bumps on the silicon surface to a reliable socketable package. On modern advanced processors, Heterogeneous Integration (chiplets) and [Advanced Packaging](#) come into play here. Multiple dies per package, 3D die stacking and 2.5D interposers such as TSMC's CoWoS (Chip-on-Wafer-on-Substrate) technology are used to scale performance beyond the 26 x 33 mm reticle limit, the maximum die size that each lithography machine can expose.

在进行调试之前，硅片需要进行封装，以保护脆弱的裸片，并将硅片表面微小的 I/O 凸点引出到可靠的可插拔封装中。在现代先进处理器上，异构集成（小芯片）和先进封装在此发挥了作用。每个封装包含多个裸片、3D 裸片堆叠以及 2.5D 中介层（如 TSMC 的 CoWoS，即晶圆级封装技术）被用于将性能扩展到 26 x 33 mm 的光罩极限之外，这是每台光刻机所能曝光的最大裸片尺寸。




# 11. Post-Silicon Validation & Bringup

## 11. 硅后验证与上电调试 (Post-Silicon Validation & Bringup)

Automatic Test Equipment (ATE) from the likes of Teradyne and Advantest assesses every single chip coming off the production line, applying thousands of test vectors generated earlier by ATPG tools. JTAG debug interfaces within the chip provide direct access for post-silicon debugging of errata, when behavior deviates from expectations. Initial bringup features multiple rounds of testing, as more bugs are found and workarounds are implemented by updating the original firmware that was developed on FPGA emulators. In production runs, this step is known as **Final Test (FT)**, where functionality is verified with tests that return a simple pass/fail result.


来自 Teradyne 和 Advantest 等公司的自动测试设备 (ATE) 会对生产线上的每一颗芯片进行评估，应用由 ATPG 工具预先生成的数千个测试向量。当芯片行为偏离预期时，芯片内部的 JTAG 调试接口可为硅后勘误调试提供直接访问。初始上电调试包含多轮测试，随着发现更多漏洞，工程师会通过更新最初在 FPGA 仿真器上开发的固件来实现变通方案 (Workarounds)。在量产阶段，这一步骤被称为终测 (Final Test, FT)，通过返回简单的通过/失败 (Pass/Fail) 结果来验证功能。



Source: MPI Corp MPI Corp

To access every single wire coming out of the package, large breakout boards known as **Probe Cards** are used. Bringup engineers use them to isolate each signal pin from the package, hooking up oscilloscopes to the terminals to measure signal integrity and strength that meets design requirements.

为了访问从封装引出的每一根导线，需要使用被称为探针卡（Probe Cards）的大型分线板。调试工程师利用它们将封装中的每个信号引脚隔离开来，并将示波器连接到终端，以测量信号完整性和强度是否符合设计要求。




Source: Synopsys Synopsys

EDA vendors also have a hand in this space, with Synopsys TestMAX family of post-silicon testing tools providing comprehensive coverage from ATPG during the design phase, yield diagnosis, DFT structures that test patterns run on, and software to run and manage chip testing and results logging.

EDA 厂商在这一领域也占有一席之地，例如 Synopsys 的 TestMAX 系列后期测试工具，提供了从设计阶段的 ATPG（自动测试向量生成）、良率诊断、运行测试向量的 DFT（可测试性设计）结构，到运行和管理芯片测试及结果记录的软件等全方位覆盖。

## Burn-in Test 老化测试



Source: NVIDIA 

The next stage in the qualification process is for thermal reliability with **High Temperature Operating Life (HTOL)** testing. This is an accelerated burn-in where chips are tested beyond the typical operating temperature, weeding out any defective chips that may not survive the rated operating environment and thermal cycles over the expected lifetime of the chip. Removing these “infant mortality” class of chips greatly reduces the chance of customers receiving units that may fail early on in the Bathtub Curve of reliability.

资格认证过程的下一个阶段是热可靠性测试，即高温工作寿命（HTOL）测试。这是一种加速老化测试，芯片在超出典型工作温度的环境下接受测试，旨在剔除那些可能无法在额定工作环境和预期寿命内的热循环中存活的缺陷芯片。排除这些“早期夭折”类芯片，可以大大降低客户收到在可靠性“浴盆曲线”早期阶段发生故障的产品的概率。


Test durations depend on the reliability requirements of each design, though on average range from 72 to 168 hours. Reliability critical applications such as aerospace and automotive may undergo up to 1000 hours of stress testing. Lower end, cost sensitive consumer devices may only have some chips from each production batch randomly selected for extended testing. Testing requirements are standardized by



JEDEC, with JESD47 covering HTOL and JESD22 for package reliability under high humidity and temperature cycling.

测试时长取决于每个设计的可靠性要求，平均范围通常在 72 到 168 小时之间。航空航天和汽车等对可靠性要求极高的应用，可能会经历长达 1000 小时的压力测试。而低端、对成本敏感的消费类设备，可能仅从每个生产批次中随机抽取部分芯片进行延长测试。测试要求由 JEDEC 标准化，其中 JESD47 涵盖了 HTOL，JESD22 则涵盖了高湿度和温度循环下的封装可靠性。

## Steppings 步进 (Steppings)



Source: Intel 

The first chip that comes back from the fab is known as A0 silicon. The ideal case is when A0 meets performance, power and reliability targets with no tweaks required for production silicon. However, sometimes the first chip doesn't boot at all, and the fundamental flaw must be discovered. As such, most design teams budget additional development time to do multiple steppings, meaning the chip design is updated with new GDSII files sent to the fab for mask making and fabrication.

从晶圆厂返回的第一颗芯片被称为 A0 步进 (A0 silicon)。理想情况是 A0 芯片无需任何调整即可满足性能、功耗和可靠性目标，并直接用于量产。然而，有时第一颗芯片甚至根本无法启动，此时必须找出根本缺陷。因此，大多数设计团队都会预留额外的开发时间来进行多次步进 (steppings)，这意味着芯片设计将通过发送给晶圆厂的全新 GDSII 文件进行更新，以重新制作掩模并进行制造。

From Intel's example above, multiple rounds of **Engineering Sample** (ES) versions are made, with a final Qualification Sample that is validated before production wafers are run. Steppings come with many codenames such as B1 or even [E5](#) as was the case for the long delayed Intel Sapphire Rapids. Here's a guide to what the stepping number means:

从上述英特尔的例子可以看出，在正式量产晶圆之前，会制造多轮工程样品（ES）版本，并进行最终的资格认证样品（QS）验证。步进（Stepping）有许多代号，例如 B1，甚至是像长期延期的英特尔 Sapphire Rapids 那样的 E5。以下是关于步进编号含义的指南：

- Major Stepping (A0 to B0): Requires full flow from DE to DV to PV, verification engineers have to create a new coverage closure, full mask set update usually required.

大步进（Major Stepping，如 A0 到 B0）：需要从 DE（设计工程）到 DV（设计验证）再到 PV（物理验证）的完整流程；验证工程师必须创建新的覆盖率收敛；通常需要更新全套掩模（Mask Set）。


- Minor Stepping (A0 to A1): Typically a small mask change to the metal stack for bugfixes, implementing changes verified through circuit editing on the prior stepping. Minor Design and Physical Verification required.

小步进（Minor Stepping，如 A0 到 A1）：通常是对金属层堆栈进行微小的掩模修改以修复漏洞，实施通过前一步进的电路编辑（Circuit Editing）验证过的更改。仅需要少量的设计和物理验证。





## FIB Backside Chip Circuit Edit Process



Source: GlobalFoundries | GlobalFoundries

This circuit editing is done with advanced Focused Ion Beam (FIB) tools to add jumper wires or alter transistor interconnect patterns while keeping the chip fully functional so the fix can be tested. Designs can accommodate this process by having buffer or spare cells and metal lines that leave more wiggle room for fixing routing faults. As this process can be time consuming, it is mostly used to physically validate the layout changes that are planned to go into the next stepping.

这种电路编辑是通过先进的聚焦离子束（FIB）工具完成的，用于添加跳线或更改晶体管互连模式，同时保持芯片功能完整，以便对修复方案进行测试。设计可以通过包含缓冲器或备用单元以及金属线来适应这一过程，从而为修复布线故障留出更多余地。由于该过程可能非常耗时，它主要用于对计划进入下一个步进（stepping）的版图更改进行物理验证。

## Binning 分箱 (Binning)

The last step is speed binning, where each chip is tested at progressively higher frequencies, and those that pass at the highest clocks are sold as premium parts. This performance spread is the result of natural variability in manufacturing. While all the chips at this stage are fully functional, some may require slightly higher voltage than others to achieve the target operating frequency. Binning also encompasses **yield harvesting**, where some cores or sub-components in the chip may be defective. These are then fused off and sold at a discount with lower performance.

最后一步是速度分级（speed binning），即在逐渐升高的频率下对每个芯片进行测试，通过最高时钟频率测试的芯片将作为高端型号销售。这种性能差异是制造过程中自然波动的结果。虽然这一阶段的所有芯片功能都完全正常，但有些芯片可能需要比其他芯片略高的电压才能达到目标工作频率。分级还包括产量回收（yield harvesting），即芯片中的某些核心或子组件可能存在缺陷。这些部分会被熔断禁用，并以较低的性能和折扣价格出售。

Binning has allowed Intel to create its product mix, the famous Core i5, i7 and i9 branding for CPUs, while NVIDIA GPUs almost never have all SMs enabled due to yield harvesting.

分级制度使 Intel 能够构建其产品组合，即著名的 Core i5、i7 和 i9 CPU 品牌；而 NVIDIA 的 GPU 由于产量回收，几乎从不启用所有的 SM 单元。

## 12. System Integration 12. 系统集成

Validated chips are then placed in reference boards and connected to devices such as storage and networking, with drivers, BIOS, and OS support qualified with **System Level Testing (SLT)**. The chips are run through various stress testing benchmarks, and typical software applications are run to ensure end-user functionality.

经过验证的芯片随后被放置在参考板上，并与存储和网络等设备连接，同时通过系统级测试 (SLT) 对驱动程序、BIOS 和操作系统支持进行鉴定。芯片会运行各种压力测试基准程序，并运行典型的软件应用程序，以确保最终用户的功能体验。

Additionally, these boards are bundled with ES silicon and are sampled with partners and developer to receive early customer feedback. These Reference Validation Platforms (RVP) loop the customer into the development cycle. Major applications



developers can start optimizing their software to take advantage of the new chips features and ensure day 1 software support upon the chip's availability.

此外，这些电路板捆绑了工程样品（ES）芯片，并向合作伙伴和开发人员提供样片，以获取早期客户反馈。这些参考验证平台（RVP）将客户纳入开发周期。主要应用程序开发人员可以开始优化其软件，以利用新芯片的特性，并确保在芯片上市首日即可提供软件支持。

## 13. Production 13. 生产

Once the production ready stepping is qualified with satisfactory yield, the design can finally begin ramping up production and deployed with customers. However, the work does not stop. Failure analysis of defective silicon returned from customers can help iron out the last kinks in the design with minor revisions to the manufacturing flow. Design houses continue to work with TSMC as part of the Continuous Improvement Process (CIP) flows to help improve chip yields. FA engineers use tools such as **Synopsys Avalon** to map a given defect to the corresponding circuit schematic to identify the affected gates and wires in the design.

一旦具有满意良率的生产就绪步进（stepping）通过验证，设计就可以最终开始扩大生产并交付给客户。然而，工作并未就此停止。对客户退回的有缺陷芯片进行失效分析（Failure analysis），可以通过对制造流程进行微调，帮助消除设计中最后的瑕疵。设计公司继续与 TSMC 合作，作为持续改进工艺（CIP）流程的一部分，以帮助提高芯片良率。失效分析工程师使用 Synopsys Avalon 等工具将给定的缺陷映射到相应的电路原理图，从而识别设计中受影响的门电路和连线。



Source: Synopsys Synopsys

## EDA in Foundry 代工厂中的 EDA

Apart from the tools used by the fabless design houses, EDA also features extensively in the Foundry as they simulate and design the next generation of process nodes. On top of that, tight collaboration between foundry and fabless companies in the form of



DTCO and STCO are used to squeeze even more PPA from the foundry process.

除了无晶圆厂（fabless）设计公司使用的工具外，EDA 在代工厂（Foundry）中也得到了广泛应用，因为他们需要利用 EDA 来模拟和设计下一代工艺节点。除此之外，代工厂与无晶圆厂公司之间通过 DTCO（设计工艺协同优化）和 STCO（系统技术协同优化）形式进行的紧密合作，也被用于从代工工艺中榨取更高的 PPA（性能、功耗、面积）表现。

**Technology CAD (TCAD)** is the physics simulation layer that lets engineers design new transistor architectures entirely in software before spending tens of millions of dollars on silicon experiments. The dominant tool suite is **Synopsys Sentaurus**, which has two core engines working in sequence. Sentaurus Process simulates every physical fabrication step. Ion implantation, oxidation, thin-film deposition, plasma etching, lithography patterning and predicts the resulting 3D structure of the device down to nanometer-scale dopant profiles. Feed in a process recipe and it simulates the output device geometry. Sentaurus Device then takes that 3D structure and simulates its electrical behavior, with I-V curves, capacitance, leakage current and breakdown voltage. Together, they let process engineers iterate on a complex transistor designs, crucial in the Gate-All Around generation and beyond to CFET and novel materials such as Ruthenium, featured in [our IEDM 2025 coverage](#).

技术计算机辅助设计（TCAD）是物理仿真层，它让工程师在投入数千万美元进行硅片实验之前，能够完全在软件中设计新的晶体管架构。主导该领域的工具套件是 Synopsys Sentaurus，它拥有两个按顺序工作的核心引擎。Sentaurus Process 模拟每一个物理制造步骤，包括离子注入、氧化、薄膜沉积、等离子体刻蚀、光刻图形化，并预测最终器件的 3D 结构，精确到纳米级的掺杂分布。输入工艺配方，它就能模拟出输出器件的几何形状。随后，Sentaurus Device 获取该 3D 结构并模拟其电气行为，包括 I-V 曲线、电容、漏电流和击穿电压。两者结合，使工艺工程师能够对复杂的晶体管设计进行迭代，这在全环绕栅极（GAA）时代以及未来的互补场效应晶体管（CFET）和新型材料（如我们在 IEDM 2025 报道中提到的钌）中至关重要。

**Synopsys Mystic** then takes the simulated transistor and extracts compact model parameters (industry-standard SPICE models like **BSIM-CMG**) from the output. The device simulations are used to form the earliest PDK 0.1 for circuit designers to get



started months before any real silicon exists.

Synopsys Mystic 随后获取仿真后的晶体管，并从输出中提取紧凑模型参数（如 BSIM-CMG 等行业标准 SPICE 模型）。这些器件仿真被用于构建最早的 PDK 0.1 版本，以便电路设计师在任何真实芯片问世前数月即可开展工作。

Zooming into the atomic frontier, Synopsys QuantumATK is used in materials research, simulating at the scale of individual atoms. It uses **density functional theory (DFT)** and **non-equilibrium Green's function (NEGF)** methods to model quantum transport and electron tunnelling to for material interface engineering. This is particularly useful in Work Function Metallization schemes that precisely control the transistor threshold voltage.

深入到原子尺度的前沿领域，Synopsys QuantumATK 被用于材料研究，在单个原子尺度上进行模拟。它利用密度泛函理论 (DFT) 和非平衡格林函数 (NEGF) 方法来模拟量子输运和电子隧穿，从而进行材料界面工程。这在功函数金属化方案中尤为有用，能够精确控制晶体管的阈值电压。

## Design Feedback Loop 设计反馈环

Silicon measurement data such as transistor characteristics, IR drop and yield maps of chips made on the current process node feed directly back into the foundry's next-generation PDK development and the design team's planning for future node updates. **Best Known Methods (BKMs)** are refined over the lifetime of the node to help customers improve yields and lower costs for the foundry. This forms a productive feedback loop where customers help define the next generation of foundry processes.

来自当前工艺节点的硅片测量数据（如晶体管特性、IR 压降和芯片良率图）会直接反馈到代工厂下一代 PDK 的开发以及设计团队对未来节点更新的规划中。最佳实践方法（BKM）在节点的整个生命周期内不断完善，以帮助客户提高良率并降低代工厂的成本。这形成了一个高效的反馈闭环，客户借此参与定义下一代代工厂工艺。

## DTCO: Blending Process and Design





**DTCO: 工艺与设计的融合**

For decades, chip manufacturing followed a strict handoff: process engineers developed the transistor technology, characterized it, then “threw it over the wall” to chip designers. The designers worked with whatever they got. **Design-Technology Co-Optimization (DTCO)** demolishes this wall. It evaluates process options using chip-level **PPA (Power, Performance, Area)** metrics from the very first day of process development.

几十年来，芯片制造遵循着严格的交接流程：工艺工程师开发晶体管技术并对其进行表征，然后将其“扔过围墙”交给芯片设计师。设计师只能基于拿到的现有条件进行工作。而设计-工艺协同优化（DTCO）打破了这道围墙。它从工艺开发的第一天起，就利用芯片级的 PPA（功耗、性能、面积）指标来评估工艺方案。

The DTCO flow chains the entire tool stack into a feedback loop: Sentaurus TCAD (device physics) → Mystic (PDK) → SiliconSmart/HSPICE (cell characterization) → IC Compiler II/StarRC/PrimeTime (chip-level PPA evaluation), where the results feed back to the process engineers. Process knobs are tweaked to push design rules beyond the conservative cookie cutter standard cells. not in isolation, but by its impact on a real chip design. The biggest companies such as Apple, NVIDIA and AMD have dedicated foundry departments that work closely with TSMC to engineer their own custom cell libraries to push PPA by as much as 15% over standard libraries. Routability improvements translate directly to lower delay, higher performance, lower power and even reduced area from higher area utilization in congested routing zones.

DTCO 流程将整个工具链整合进一个反馈循环：Sentaurus TCAD（器件物理）→ Mystic（PDK）→ SiliconSmart/HSPICE（单元库特征化）→ IC Compiler II/StarRC/PrimeTime（芯片级 PPA 评估），其结果会反馈给工艺工程师。通过调整工艺参数，设计规则得以突破保守的通用标准单元限制。这种优化并非孤立进行，而是基于其对真实芯片设计的影响。像 Apple、NVIDIA 和 AMD 这样的大型公司都设有专门的代工厂部门，与台积电（TSMC）紧密合作，开发定制单元库，使其 PPA 性能比标准库提升高达 15%。可布线性（Routability）的改善直接转化为更低的延迟、更高的性能、更低的功耗，甚至通过提高拥塞布线区域的面积利用率来缩小芯片面积。

| M0 Optimization                                                                   | Double Height Cells                                                               | Stage Ratio Optimization                                                           | Skew Cell Optimization                                                              |
|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Cell M0 optimization for lower cap and higher speed                               | For better form factor for high-driving cells                                     | Optimized stage ratio and speed for multi-stage cells                              | Skew cells for better path timing with optimized resistance and pin cap             |
|  |  |  |  |

Source: TSMC TSMC

New transistor schemes such as TSMC's FinFLEX and NanoFlex make DTCO almost mandatory to effectively utilize alternating rows of transistors with significantly different performance and leakage characteristics. Backside power designs with Intel 18A and TSMC A16 expose another dimension of signal and power routing, potentially unlocking new standard cell routing scheme that route around both sides of the transistor layer.

新的晶体管方案（如 TSMC 的 FinFLEX 和 NanoFlex）使得 DTCO 几乎成为有效利用具有显著不同性能和漏电特性的交替晶体管行的必经之路。Intel 18A 和 TSMC A16 的背面供电设计开启了信号和电源布线的另一个维度，有望解锁新的标准单元布线方案，从而在晶体管层的两侧进行布线。


## STCO: Co-Optimizing the Entire System

### STCO: 全系统协同优化

System-Technology Co-Optimization (STCO) extends the DTCO concept up one more level: from chip and process co-design to system and package co-design. It addresses chiplet partitioning decisions, packaging technology selection, cross-chiplet bandwidth and latency trade-offs, multi-die thermal management and power integrity across the entire package. As single-die scaling hits economic and physical limits,

STCO is how the industry keeps delivering generational performance gains.

系统技术协同优化 (STCO) 将 DTCO 的概念又向上扩展了一个层级：从芯片与工艺的协同设计扩展到系统与封装的协同设计。它解决了芯粒 (chiplet) 划分决策、封装技术选择、跨芯粒带宽与延迟的权衡、多芯片散热管理以及整个封装范围内的电源完整性问题。随着单芯片缩放达到经济和物理极限，STCO 已成为行业持续提供代际性能提升的关键手段。



Source: Intel 

Intel's **Ponte Vecchio** GPU is the poster child for a chip that required extensive STCO. It integrates 47 active dies manufactured across five different process nodes, stitched together using both **EMIB** (Embedded Multi-die Interconnect Bridge, a 2.5D silicon bridge) and **Foveros** (Intel's 3D face-to-face die stacking). However, the chip faced several design challenges and was delayed for years, with final performance coming nowhere near original targets. With proper STCO and sound design practices, complex designs can now come out on schedule and on target, which is what AMD

plans to do with their MI455X GPU releasing later in 2026.

Intel 的 Ponte Vecchio GPU 是需要进行广泛系统技术协同优化（STCO）的典型芯片案例。它集成了 47 个活跃裸片，这些裸片采用五种不同的工艺节点制造，并利用 EMIB（嵌入式多芯片互连桥接，一种 2.5D 硅桥）和 Foveros（Intel 的 3D 面对面裸片堆叠技术）缝合在一起。然而，该芯片面临着多项设计挑战并推迟了数年，最终性能也远未达到最初的目标。通过适当的 STCO 和完善的设计实践，复杂的设计现在可以按时并按目标完成，这正是 AMD 计划在其将于 2026 年早些时候发布的 MI455X GPU 上所实现的。

Below, we discuss the reality of chip design in the real world, and the challenges engineers face from increased design complexity while pushing for compressed design cycles. We also take a look at the hardware that EDA tools run on in simulation and emulation.

下文我们将探讨现实世界中芯片设计的现状，以及工程师在缩短设计周期的压力下，面对日益增加的设计复杂性所迎接的挑战。我们还将了解 EDA 工具在仿真和硬件仿真实践中所运行的硬件平台。

## The Ugly Reality of Chip Design in the Real World

### 现实世界中芯片设计的残酷现状

Software evolves at a near daily pace, with tons of optimisations to squeeze out every last bit of performance from the hardware. New models and computing schemes such as Attention Feed Forward Network Disaggregation can dramatically alter the balance of computing to memory and data demands. When it takes 2 years to bring a chip from idea to mass production, designing hardware for what software might look like in 2 years is one of the biggest unknowns. The challenge of hardware “prediction” results



in some intentional design quirks.

软件的演进几乎以天为单位，通过大量的优化来榨取硬件的每一丝性能。诸如注意力前馈网络解耦（Attention Feed Forward Network Disaggregation）等新模型和计算方案，可以极大地改变计算与内存及数据需求之间的平衡。当一颗芯片从构思到量产需要两年时间时，针对两年后软件可能呈现的形态来设计硬件，便成了最大的未知数之一。硬件“预测”带来的挑战导致了一些刻意的设计特性。

Jeff Dean of Google Deepmind revealed that their TPU designs pay a small area penalty to integrate structures that may not be useful now, but are predicted to become useful as software evolves over the 5 year expected lifetime of the chip. Google deems this worth the effort and initial PPA loss if it allows older chips to remain relevant as software evolves. While they did say this functionality unlocking has happened before, they would not disclose how often their predictions have paid off, and how many dormant circuits remain in their TPUs.

Google Deepmind 的 Jeff Dean 透露，他们的 TPU 设计付出了一定的面积代价，用于集成一些目前可能无用、但预计在芯片 5 年预期寿命内随软件演进会变得有用的结构。Google 认为，如果这能让旧芯片在软件发展过程中保持竞争力，那么为此付出的努力和初始 PPA 损失是值得的。尽管他们表示这种功能解锁此前确实发生过，但并未透露其预测奏效的频率，以及他们的 TPU 中还保留着多少休眠电路。

To reduce this gap between Hardware and Software, design teams are being pushed to the limit, trying to compress the design cycle time as much as possible, from parallelizing tasks to overlapping design stages.


为了缩小硬件与软件之间的这一差距，设计团队正被推向极限，试图通过从任务并行化到设计阶段重叠等各种手段，尽可能地压缩设计周期。

## The Overlapping Waterfall: Compressed Timelines, “Shift Lefts”

**重叠瀑布流：压缩的时间线与“左移”**

In the ideal case, the concept of Shift Left is to have design things right the first time with in-line verification. Each block does not have to go through many rounds of iteration if designers can ensure it will pass the verification tests on the first try. EDA vendors offer various functionality in their tools to conduct early pre-verification checks to ensure smooth transitions to the next design stage, minimizing the probability of design errors and timing violations coming back that require respins. Virtual ATE with Verilog testbenches allow an early simulation of device test results months before tapeout, enabling designers to catch potential faults early in the design phase.

在理想情况下，“左移”(Shift Left)的概念是指通过在线验证，在第一时间就把设计做对。如果设计人员能确保每个模块在初次尝试时就能通过验证测试，那么这些模块就不必经历多轮迭代。EDA 供应商在其工具中提供了各种功能来进行早期的预验证检查，以确保向下一个设计阶段的平稳过渡，从而最大限度地降低因设计错误和时序违规而导致重新设计的概率。带有 Verilog 测试平台的虚拟 ATE 允许在流片前数月对器件测试结果进行早期仿真，使设计人员能够在设计阶段早期发现潜在故障。



Source: SemiAnalysis

In theory, the RTL Freeze is a sacred milestone where no more design changes and bugfixes can be made after coverage closure, handing the netlist off to the Physical Design (PD) team. In real world, however, the chip design waterfall overlaps. The RTL

is frozen before the verification test cases have even finished passing, with potentially broken RTL passed on to the physical design team. While verification is ongoing, Design Engineers (DE) continue to tweak and add features right up until the last minute. In reality, the RTL Freeze is just to tell DE to stop adding new features.

从理论上讲，RTL 冻结（RTL Freeze）是一个神圣的里程碑，意味着在覆盖率收敛后不再进行任何设计更改和错误修复，并将网表移交给物理设计（PD）团队。然而在现实世界中，芯片设计的瀑布流模型是相互重叠的。RTL 往往在验证测试用例尚未全部通过之前就被冻结，导致可能存在缺陷的 RTL 被传递给物理设计团队。在验证工作持续进行的同时，设计工程师（DE）会一直进行微调并添加功能，直到最后一刻。实际上，RTL 冻结只是为了告知设计工程师停止添加新功能。

While PD is underway, Design Verification (DV) is still hard at work writing test cases to catch bugs and minimise the risk of failed tape outs. DV engineers move forward with integration under the assumption that a fix implemented upstream would not cause problems elsewhere. If a bug is caught after RTL freeze, an Engineering Change Order (ECO) has to be done.

在物理设计（PD）进行期间，设计验证（DV）团队仍在努力编写测试用例以捕捉漏洞，并最大限度地降低流片失败的风险。DV 工程师在进行集成时，通常假设上游实施的修复不会在其他地方引发问题。如果在 RTL 冻结后发现了漏洞，则必须执行工程变更指令（ECO）。

To accommodate this tight timeline, DE intentionally makes parts of their design “re-programmable”. Features can be readily bypassed with software registers, thus avoiding problematic chip features that may cause catastrophic errors that require a new tape out. While this obviously hurts PPA with hidden disabled features that could not be fixed in time, this design flexibility allows hard deadlines to be met, with potential performance upside should all features be verified in time.

为了适应如此紧凑的时间表，设计工程师（DE）特意将其设计中的部分内容设为“可重编程”。通过软件寄存器可以轻松绕过某些功能，从而避免那些可能导致灾难性错误并需要重新流片的存疑芯片特性。虽然由于隐藏了无法及时修复的禁用功能，这显然会损害 PPA（功耗、性能和面积），但这种设计灵活性确保了能够按时完成硬性截止任务，并且如果所有功能都能及时通过验证，还存在潜在的性能提升空间。

# Manpower Shortages 人力短缺

As mentioned above, verification is taking up a higher proportion of engineer hours. Ideally, the ratio of DE to DV engineers would be 1:4 but no company has this. Some companies even have twice as many DE to DV engineers. In reality, DE ends up doing some of the verification themselves as they know the design best to write targeted testbenches.

如前所述，验证工作在工程师工时中所占的比例越来越高。理想情况下，设计工程师（DE）与验证工程师（DV）的比例应为 1:4，但没有哪家公司能达到这一水平。有些公司的 DE 工程师数量甚至是 DV 的两倍。在现实中，DE 最终往往会亲自承担一部分验证工作，因为他们最了解设计，能够编写出更具针对性的测试平台（testbench）。

## The Collaboration and licensed IP Black Box

### 协作与授权 IP 黑盒

Verification becomes even more challenging when using protected IP from external parties. Collaborative efforts such as the NVIDIA x MediaTek GB10 chip and the licensed ARM CPU Core IP within increase the barriers for engineers to understand the whole flow. Each company withholds their base RTL code to prevent IP theft, while only exposing the interfaces via standard of custom data protocols such as AMBA CHI.

当使用来自外部第三方的受保护 IP 时，验证工作变得更具挑战性。诸如 NVIDIA 与联发科（MediaTek）合作的 GB10 芯片以及其中集成的授权 ARM CPU 核心 IP 等协作项目，增加了工程师理解整个流程的难度。各家公司都会保留其底层 RTL 代码以防止 IP 窃取，仅通过标准或自定义数据协议（如 AMBA CHI）公开接口。

## GB10: A Successful Collaboration Between NVIDIA and Mediatek

### Merging of two architectures

- GB10 combines MTK's power efficient CPU & Memory Subsystem implementation with NVIDIA's high-performance GPU
- Possible through implementation of industry defined standard interfaces and protocols
- Extensive performance modeling of GPU memory traffic into MTK's memory subsystem

### NVIDIA IP Integration

- C2C and other NVIDIA IP delivered for integration by MTK through a well-defined IP model

### Design and Verification flow alignment

- Hierarchical functional verification approach with heavy reliance on BFM's as well as co-simulation for complex end-to-end features
- Heavy reliance on emulation to validate MTK's + NVIDIA's HW, FW and SW working together



Source: NVIDIA 

The verification engineer simply sees a black box, with documentation stating the signals and timing that is expected from and to the black box. Should there be an upstream timing conflict that is not resolved from the provider, the engineer has limited visibility to determine whether there a design fault is there or if there was an error within the testbench itself. Compressed design cycles exacerbate the problem as companies want the bleeding edge in IP that might not be fully signed off yet.

验证工程师看到的只是一个黑盒，文档中规定了该黑盒预期的输入输出信号及定时。如果出现来自供应商且未解决的上游定时冲突，工程师的可见性将非常有限，难以判断是设计本身存在缺陷，还是测试平台（testbench）内部出现了错误。由于公司追求尚未完全完成签核（sign-off）的最前沿 IP，压缩的设计周期进一步加剧了这一问题。

## FPGA and CPU Demand in Chip Design

### 芯片设计中的 FPGA 和 CPU 需求

As mentioned above, chip emulation for firmware development runs on FPGAs, specifically Xilinx Versal VP1902 processors as sold in the ZeBu servers from Synopsys



and Protium X3 prototyping systems from Cadence. Emulation was done by IBM-based CPU designs for the Palladium offerings by Cadence, with modern Palladium Z3 systems using a customized Cadence Emulation Processor.

如前所述，用于固件开发的芯片仿真运行在 FPGA 上，具体而言是 Synopsys ZeBu 服务器和 Cadence Protium X3 原型系统中销售的 Xilinx Versal VP1902 处理器。Cadence 的 Palladium 系列曾采用基于 IBM 的 CPU 设计进行仿真，而现代的 Palladium Z3 系统则使用了定制的 Cadence 仿真处理器。

Fabless design houses also turn to AWS EC2 F2 cloud FPGAaaS offerings to augment their own FPGA fleet. AWS themselves uses the hosted AMD Virtex VU47P FPGAs with HBM for their own chip design efforts.

无晶圆厂设计公司（Fabless）也会求助于 AWS EC2 F2 云端 FPGAaaS 服务，以扩充其自有的 FPGA 集群。AWS 自身在进行芯片设计工作时，也使用了托管的带有 HBM 的 AMD Virtex VU47P FPGA。

While FPGAs are used for emulation and prototyping, CPUs dominate the rest of the simulation and verification stack. RTL, netlist and physical verification are very sequential tasks, with branchy conditionals with many dependencies, waiting on results, outputs and timings from a previous circuit. While verification can be parallelized to specific modules and logic blocks within each design, higher levels need to wait for lower levels to be verified.

虽然 FPGA 被用于仿真和原型设计，但 CPU 在仿真和验证堆栈的其他部分占据主导地位。RTL、网表和物理验证是非常顺序化的任务，包含带有许多依赖关系的各种分支条件，需要等待前一个电路的结果、输出和时序。虽然验证可以针对每个设计中的特定模块和逻辑块进行并行化，但高层级仍需等待低层级验证完成。

As such, large numbers of high performance datacentre CPUs are used for this task, with engineers submitting testbenches to run overnight. Full chip verification will span multiple days, so high per-core performance is desired to reduce turnaround times to help teams identify bugs quicker. Some run their own AMD Turin or Intel Granite Rapids CPU servers, others take up cloud services like Synopsys verification cloud SaaS. CPU producers use their own CPUs internally to design future products. For example, AWS deploys 1 million Graviton cores internally to run EDA tools to

design future Graviton, Trainium and Nitro chips.

因此，大量高性能数据中心 CPU 被用于这项任务，工程师们会提交测试平台进行隔夜运行。全芯片验证将跨越数天，因此需要极高的单核性能来缩短周转时间，从而帮助团队更快地识别漏洞。一些公司运行自己的 AMD Turin 或 Intel Granite Rapids CPU 服务器，另一些则采用 Synopsys 验证云 SaaS 等云服务。CPU 生产商在内部使用自家的 CPU 来设计未来的产品。例如，AWS 在内部部署了 100 万个 Graviton 核心来运行 EDA 工具，以设计未来的 Graviton、Trainium 和 Nitro 芯片。



Recommend SemiAnalysis to your readers

向您的读者推荐 SemiAnalysis

Bridging the gap between the world's most important industry, semiconductors, and business.

架起全球最重要的行业——半导体与商业之间的桥梁。

Recommend 推荐



44 Likes 44 次赞 · 3 Restacks 3 次转推

← Previous 上一篇

Discussion about this post

关于此帖的讨论

Comments 评论 Restacks 转贴



Write a comment...

