




SYNOPSYS®

YOUR GUIDE TO FIRST-PASS SILICON SUCCESS

Mastering AI Chip Complexity

Table Of Contents

Introduction

The Rise of AI Chips & Chiplets

AI Chip Development Journey

From Big Ideas to Breakthroughs

Mapping an AI Chip Development Roadmap

Architecture Exploration

Silicon Front & Back-End Design + Cloud-Based Solutions

Advanced Packaging & Multi-Die Design

Early Software Development & System Verification


Security Considerations

Scalable End-to-End Test Solution

Silicon Lifecycle Management & Reliability, Availability, Serviceability

Conclusion

The Road Ahead with Synopsys



The Rise of AI Chips & Chiplets

AI is reshaping the world, pushing the boundaries of what's possible in computing — demanding faster, more efficient, and scalable silicon to power the next generation of breakthroughs.

To meet this growing demand, AI hardware is shifting from traditional monolithic chips to chiplet-based and multi-die design — enabling scalability, efficiency, and performance to meet the complex needs of cutting-edge AI workloads.

Market Opportunity

▲ **\$383B**


Global AI Chip Market Projected by 2032¹

▲ **38%**

CAGR Growth Rate for AI Chips Through 2032¹

Explosive AI Chip Market Growth

(2022-2032)



Source: Allied Market Research, 2023

AI Workload Complexity is Redefining Hardware Design

AI Models Are Outpacing Traditional Hardware

From AlexNet (~60M parameters) to GPT-4 (1T+ parameters)², AI model complexity has skyrocketed, outpacing the capabilities of traditional monolithic architectures. The sheer computational demands of modern AI require **multi-die integration, heterogeneous compute, and advanced interconnect strategies to sustain performance.**

Scaling Challenges in AI Hardware

As AI models scale to trillions of parameters and agentic AI becomes mainstream, the infrastructure supporting them must evolve. This means overcoming **bottlenecks in memory bandwidth, interconnects, power constraints, and thermal management** to ensure sustainable scaling without prohibitive costs or power trade-offs.




Complex technology systems and the AI silicon that powers them require a holistic approach to innovation. Synopsys is a trusted partner in helping R&D teams re-engineer how engineering is done from silicon to systems.



Sassine Ghazi
President & CEO




Exponential Growth of AI Workloads Demands Innovation in Hardware



DESIGN CHALLENGES

Overcome Silicon, Software & Packaging Complexity From Day One



Silicon Complexity

AI chip complexity is increasing to 100s of billions of transistors and performance requirements demand leading-edge process nodes. **In 2024, at least 58% of ASICs required a respin³**, driving up costs and delaying time-to-market.



Software Complexity

Growing AI model complexity demands are driving silicon and system definition for optimum performance. **Early software development and hardware-assisted verification** are now essential.



Advanced Packaging

AI chip size and economic requirements are **accelerating the adoption of multi-die designs**. Integrating multiple heterogeneous or homogeneous dies (also called chiplets) in a single package **increases design complexity** in terms of performance, thermal management, signal integrity, and security attack surfaces.

BUSINESS CONSIDERATIONS

Prepare for Rapid Cycles, Skyrocketing Costs & Talent Gaps



Shrinking Design Cycles

To meet the rapid advancement of large language model (LLM) growth, AI chip roadmaps have shifted from **18-24 months to 12 month release cycles**. This rapid acceleration **strains RTL design and verification, power analysis, and back-end layout schedules**.



Soaring Design Costs


The hardware and software cost of designing at 2nm and moving to the Angstrom era can exceed hundreds of millions of dollars per chip⁴, driven by the complexity of advanced process nodes. This also increases validation and verification demands, requiring system software to be ready when the chip returns from fab for functional proof of concept (PoC) or fast time to value (TTV) deployment.



Talent & Expert Resource Shortages

The semiconductor industry faces an acute talent gap — **AI chip expertise in design, verification, and implementation is scarce**, and hiring cycles can't keep up with demand⁵.

Maximize Power & Performance Gains Early



Source: Synopsys, 2023



Power & Performance

Optimizing power and performance can no longer be an afterthought, it's an end-to-end process. AI workloads are reshaping the energy landscape. Worldwide data centers are projected to consume 1,000 terawatt-hours (TWh) by 2026, **energy at the scale of entire nations**⁶, to meet computational demands. Semiconductors can contribute to solving this challenge through the use of **Shift-Left methodologies** — optimizing power from the start of the design cycle to make the biggest impact.



01

From Big Ideas to Breakthroughs

ASSEMBLE YOUR TEAM WITH EXPERTISE
THROUGH COLLABORATION

Build Smarter From the Start

Developing AI chips and chiplets demands more than a clear and innovative concept — it requires strategic foresight. To stay ahead, AI chip companies must anticipate challenges two generations ahead, not just solve today's problems. A well-structured approach can minimize costs, reduce risk, and accelerate time-to-market.

Before starting your journey, assemble your team with three key decisions:

ASSEMBLE YOUR TEAM: KEY DECISION #1

Foundry & Package Technology Supplier

Choosing the right fabrication and packaging partners is crucial. Collaborating with industry leaders like **Intel Foundry, Rapidus, Samsung Foundry, or TSMC** ensures access to cutting-edge process technologies and advanced packaging solutions. Key considerations include selecting the optimal **foundry and process node, leveraging shared wafer programs for cost-effective prototyping, ensuring broad IP support,** and evaluating supply chain reliability.



Synopsys' industry-leading solutions are a mission-critical on-ramp for foundries. Our close collaboration on EDA tools, IP, and support for leading edge process nodes and advanced packaging technologies tame complexity and accelerate time-to-market for our customers' innovations.



Shankar Krishnamoorthy
Chief Product Development Officer

SYNOPSYS®

Expert Collaboration With Arm


ASSEMBLE YOUR TEAM: KEY DECISION #2

Core Processor Architecture

The compute core processor architecture establishes the foundation for the overall AI system.

Synopsys partners with Arm as a part of [Arm Total Design](#), to provide robust solutions, including EDA tools, complementary IP and design services for **Arm Neoverse Compute Subsystems (CSS)** for AI workloads.

arm



Source: Arm

Custom Silicon Design

ASSEMBLE YOUR TEAM: KEY DECISION #3

Design Collaboration

Converting a concept into silicon requires deep expertise. **Synopsys System Solution Design Services** and **Synopsys Ecosystem Partners** encompass experienced design experts familiar with Synopsys' EDA tools and IP, accelerating development and ensuring an optimized final silicon, packaging and software product. You can engage with these design experts at the appropriate stage: concept, spec, IP selection, flows and methodologies, RTL, IP integration, physical implementation, package design, fab management — or have Synopsys supply your full end-to-end needs. This approach enables you to focus or augment your in-house resources on the greatest value-add and what you do best, while offloading other parts of the design to experts in those areas for fastest and most reliable time-to-market.

Synopsys provides strategic solutions for foundry partnerships, core processor architecture, and design collaboration through an integrated ecosystem to accelerate time-to-market, streamline development, and optimize performance.

Synopsys System Solution Design Services

SYNOPSYS®

Synopsys Ecosystem Partners

including



AVNET®

DNP



GUC


MegaChips

mirafra
TECHNOLOGIES

NSW

TESSOLVE
A HERO ELECTRONIX VENTURE

TOPPAN



02


Mapping an AI Chip Development Roadmap

FROM CONCEPT TO SILICON, PACKAGE & SOFTWARE

Your AI Chip Development Journey

TRANSFORM BIG IDEAS INTO BREAKTHROUGH SILICON

Start with early architecture exploration and integrate silicon, packaging, and software development in parallel to deliver comprehensive solutions.



03

Architecture Exploration

REDUCE RISK EARLY: OPTIMIZE AI CHIP
PERFORMANCE & POWER FROM DAY ONE

Architecture Exploration

Power & Performance

Security

Silicon

Silicon Front-End

Package

RTL Freeze

Package Spec


Silicon Spec

Test Spec

Software Spec

The Importance of Early Architecture Decisions

AI chip success hinges on early architecture exploration — before RTL, silicon fabrication, and costly late-stage revisions. Every decision at this initial stage carries **long-term trade-offs** in:



Performance

Compute throughput, data movement, and workload efficiency



Power

Thermal constraints, energy efficiency, and AI-specific optimizations



Cost

Die-size, manufacturing feasibility, yield, and design iteration expenses

Spreadsheets and static diagrams can't capture and analyze the complexity of modern AI workloads.

Designers must answer critical questions upfront:

- 1 What are the optimum CPU, bus bandwidth, interconnect and memory configurations?
- 2 What trade-offs exist between latency, throughput, and energy efficiency?
- 3 How does workload scaling impact multi-die design and interconnect performance?

Missteps here drive up costs, create performance bottlenecks, and delay market entry. A strategic approach ensures optimal silicon efficiency from day one.

Synopsys Architecture Analysis Tools

Platform Architect

Monolithic modeling for SoC, chiplet, and multi-die systems.


Platform Architect AI Exploration Pack (AI XP)

AI centric workload and hardware model libraries.

Platform Architect Development Kit

Web-based collaboration across architecture exploration stakeholders.

Potential Schedule Savings of up to 3-6 Months for AI Chip Designs




Mission Critical IP Selection

High-level simulation early can be used to determine the number of cores and system resources required to strike the right balance between processing cores, memory, and networking subsystems. This must be done with careful consideration of performance, power, and cost implications.

As the architecture solidifies, leverage Synopsys’ full range of high-performance, power-optimized IP for advanced AI chip implementations:

AI Accelerator Chiplet



Memory Hierarchy: Based on workload requirements, it’s key to optimize L1/L2 cache sizes and L3 memory by leveraging combinations of high-performance HBM, DDR, LPDDR or CXL.

On-Chip Interconnect: Rightsizing interconnect types, sizes, speeds, and arbitration schemes can make or break your design. In addition to the Arm Neoverse CMN S3 Coherent Mesh Network, Synopsys provides complementary interconnect IP for AMBA 4 AXI.

High-Speed Chip-to-Chip Interconnect: With large amounts of data being moved throughout your system, selecting the right type and bandwidth of high-speed interfaces is critical. Synopsys offers leading-edge PCIe 7.0 and 1.6T Ethernet IP solutions and is at the forefront of the next generation standards of Ultra Ethernet and UALink. Synopsys delivers silicon-proven high-speed interfaces up to 224G reducing your risk and improving time-to-market.

Die-to-Die Connectivity: The industry has rapidly adopted the Universal Chiplet Interconnect Express (UCIe) standard for die-to-die connectivity. Synopsys provides a complete UCIe IP solution, encompassing controller, PHY, and verification IP for advanced processes. The IP, while operating at high speeds and providing reliable die-to-die connectivity, supports standard and advanced package technologies.

Security Solution: Security architects must address protection of AI systems, data and communications from the outset to ensure users’ safety and privacy — including from new AI-related security threats such as training data poisoning and inference model theft.

SLM Strategy: Starting with the end in mind, long-term benefits of silicon lifecycle management (SLM) should be considered early to support in-field RAS deployment, yield improvements, and power monitoring and optimization.

Rain AI: Accelerating AI Chip Innovation with Synopsys

Rain AI is pioneering the most power-efficient AI compute engine, leveraging Synopsys' full suite of custom tools, silicon IP, and cloud-based solutions to streamline architecture exploration, development, accelerate verification, and optimize chip performance — all while reducing infrastructure overhead.

“People have a tendency to think they can do everything themselves. You can, but at what cost for your company? I think the success of a startup is based on getting your product out faster with high quality and focusing only on the product.”



J-D Allegrucci
VP of Engineering



04

Silicon Front-End & Back-End Design + Cloud-Based Solutions

STREAMLINED RTL-TO-SIGNOFF WORKFLOWS MEET
ON-DEMAND CLOUD INFRASTRUCTURE



SILICON FRONT-END

RTL Design Excellence: Get It Right the First Time


Rapidly advancing LLMs and domain-specific small language models combined with the increasing complexity of semiconductor design and process node scaling demands **faster design cycles, increased design productivity, and efficient verification** — all while simultaneously facing a shortage of expert talent.

**At least 58% of ASICs
require a respin³.**

Logic and functional flaws are the leading cause resulting in costly delays.

Cause of ASIC Respins*

#1 Cause = Functional/Logic Failures



Source: Synopsys Global User Survey, 2024
*more than one reply possible


SILICON FRONT-END

Synopsys Methodology: RTL Design & Verification

Synopsys integrated design and verification flow leveraging silicon-proven IP, provides the tools to optimize every step of front-end design from architecture to verification through:

- ▶ Industry-leading portfolio of Silicon-proven IP
- ▶ Software & hardware-assisted verification (HAV) solutions
- ▶ Workload-tested power & performance optimization
- ▶ A unified platform with advanced debug capabilities
- ▶ Cloud-based solutions for ease of deployment


Front-End RTL Design and Verification Methodology



SILICON FRONT-END

Quickly Build a Complete Solution with IP Integration

Designing a chip from scratch is costly and time-intensive. Synopsys provides an **industry-leading portfolio of silicon-proven IP**, allowing chip designers to **leverage pre-validated IP subsystems** while focusing on their own proprietary innovations.



Leverage Synopsys & Ecosystem Pre-Verified Subsystems while adding your own Proprietary IP

Compute: Arm Neoverse Compute Subsystems (CSS). Synopsys ARC-V

AI Accelerator: Custom. Synopsys ARC NPX NPU for Edge AI applications

Memory Interface: HBM, DDR/LPDDR, CXL

Foundation IP: Memory and Libraries: On-chip SRAM, ROM, TCAM, NVM, MRAM, RRAM, logic libraries, IOs

Bus Fabrics - NoC: Arm Neoverse Coherent Mesh Network (CMN). Arteris. Baya Networks

Complementary Bus Fabrics: Arm AMBA 4 AXI, AMBA 3 AXI, AMBA 2.0 AHB, APB

Chip-to-Chip Connectivity: PCIe, Ethernet, Ultra Ethernet, UALink, CXL

Embedded Processors: Arm System Control Processor (SCP) and Manageability Control Processor (MCP) built into Neoverse CSS. Synopsys ARC-V processors

Power Management: Custom, including dynamic voltage and frequency scaling (DVFS), clock gating

Security: Crypto IP, Trusted Execution Environments, Interface Security (PCIe, CXL, HDCP, DDR/LPDDR, MACsec, etc.), Security Protocol Accelerators

Sensors: PVT sensors for power optimization and silicon lifecycle management (SLM)

Die-to-Die Interface: UCIe, HBM, 3DIO

Peripheral Interfaces, Basic Peripherals, GPIOs: USB, HDMI, UART, I3C, SPI, GPIOs, and more


SILICON FRONT-END

Smarter, Faster RTL: Reduce Iterations & Eliminate Bugs

By combining **Synopsys' extensive IP portfolio** with **AI-driven design and verification automation**, chipmakers can:

- ✓ **Improve design productivity** for RTL development and verification.
- ✓ **Ensure seamless massively distributed simulation** for multi-die design verification between chiplets and testbench units.
- ✓ **Accelerate functional verification** and hardware/software development using [hardware-assisted verification \(HAV\)](#) for **emulation** and **prototyping**.

Synopsys Silicon to Systems Verification Continuum



SILICON BACK-END

Physical Design to Optimize Power, Performance, and Area (PPA)

The back-end of semiconductor design is beset with increasing challenges as silicon complexity grows, process nodes shrink, and design cycles accelerate. As a result, the back-end is often under immense pressure to deliver high-quality results within compressed timelines.

Critical Challenges:

- ▶ **PPA Optimization:** Balancing power, performance, and area while addressing AI processor-specific challenges.
- ▶ **High-Density & High-Bandwidth Interfaces:** Managing the complexity of HBM & UCIe interconnects.
- ▶ **Time to Market:** Meeting aggressive, increasingly compressed production schedules without compromising quality.

AI Processor Design Challenges

	Architecture Requirements	Key Design Challenges
Compute	• 100B+ transistors, 800mm² designs • Massive parallelism, replicated topologies • All-to-all connections • Dense yet flexible power grid (PG)	• Capacity and turnaround time constraints • Complex global clocking and multiple power domains • Signal routing congestion • Clock routing through dense PG structures
Memory	• Large number of local memories • High on-chip and chip-to-chip bandwidth at low latency	• Congestion and transistor-heavy layouts • Macro timing path closure challenges • EM effects, parasitics, via ladders
Power	• Memory access consumes 30–50% of budget • 100Ks MACs/cycle with high concurrent switching	• Clock vs. datapath power optimization • IR drop issues • Robust power grid requirements
I/O	• High-bandwidth interconnects (PCIe, NVLink, CCIX, UALink)	• Sparse logic • Low utilization areas

▶ **Synopsys Solutions Address these Challenges**

SILICON BACK-END


Synopsys Hyper-Convergent Design Flow

Synopsys' hyper-convergent design methodology streamlines the mammoth task of physical implementation, integrating RTL design and synthesis, place and route, and signoff into a unified flow, through the use of a common data model.

Exceeding System Requirements with Silicon-Proven Best Practices

- ▶ **Hyper-Convergent Database:** A key productivity enhancer that enables efficient design scaling.
- ▶ **High density HBM & UCIe interfaces:** Optimized for seamless, high-bandwidth interconnectivity across die-to-die and chip-to-chip implementations.
- ▶ **Arm Core PPA Optimization:** Proven power, performance, and area (PPA) optimization across leading foundries through Synopsys System Solution Design Services and Synopsys Ecosystem Partners.

Hyper-Convergent Flow



SILICON BACK-END

Optimize PPA with Synopsys' Back-End Tools

To address the complexities of physical design and signoff, Synopsys offers a comprehensive suite of tools for **RTL optimization, power reduction, and automated implementation.**

Essential Back-End EDA Tools

- ▶ **RTL Architect & Design Compiler:** AI-driven RTL synthesis and optimization for faster convergence.
- ▶ **Fusion Compiler:** Unified RTL-to-GDSII solution for power-efficient, high-performance designs.
- ▶ **IC Compiler II & Ansys RedHawk:** Robust PPA optimization through automated [IR drop mitigation](#) to enhance power integrity, performance, and reliability.

AI-Powered Optimization

- ▶ **DSO.ai (Design Space Optimization):** AI-driven exploration for optimal PPA trade-offs.
- ▶ **Automated Power Reduction:** Leveraging AI to minimize dynamic and leakage power.

SILICON BACK-END

Signoff: A Unified Solution for Unmatched Reliability

Ensuring manufacturability and reliability at advanced process nodes requires rigorous verification. As the only provider of a unified signoff system, Synopsys delivers hyper-convergent solutions that drive first-pass silicon success while optimizing power, performance, and area (PPA).

Key Signoff Considerations

Signal Integrity (SI) / Power Integrity (PI): Advanced design methodologies to ensure robust performance under high-speed and high-density conditions, addressing noise, power distribution, electromigration (EM) effects and IR drop in advanced nodes.

Timing Closure: Achieving final-stage verification to ensure reliable silicon operation.

Why Synopsys Leads in Signoff

- ▶ **Trusted Signoff Standard:** Ensuring first-pass silicon success on the world’s most advanced tape-outs.
- ▶ **10X Faster Performance:** High performance timing and power verification accelerate signoff with AI-driven optimizations.
- ▶ **Breakthrough ECO Innovations:** Faster final-stage iterations with advanced constraint management, ensuring robust design closure and optimized PPA for faster tape-out.

Unified Signoff Solutions

Synopsys integrates timing, power, and physical verification into a single, scalable system:

Timing Closure & Power Analysis

- PrimeTime®
- PrimeShield™
- PrimeClosure
- Timing Constraints Manager

Parasitic Extraction & Physical Verification (DRC, LVS)

- StarRC™
- IC Validator™

Power Integrity (PI), Analysis, Optimization, Signoff

- PrimePower
- PrimePower RTL
- ZeBu® Empower

Low-Level Equivalence and Analysis*

- PrimeLib
- NanoTime
- ESP

*Library Characterization (PVT), Transistor-Level Timing, Signal Integrity (SI), Process Variation Analysis, Equivalence Checking

CLOUD-BASED OPTIONS

Optimizing Compute Performance & Infrastructure Efficiency

For semiconductor companies managing peak compute demands or startups accelerating proof of concept without CapEx investment, cloud-based EDA and pre-silicon emulation solutions offer the **scalability, efficiency, and cost optimization** required for modern chip design.


On-Prem vs. Cloud: AWS Usage Data

As illustrated by Amazon Web Services (AWS), typical on-premises chip design compute resources spend most of their time underutilized, followed by intense periods of peak usage that exceed capacity. These spikes result in high scheduler wait times, delayed bug detection and fixes, longer design iterations, slower design closure, and ultimately, delayed innovation.

Why Cloud Makes the Difference

Cloud-based approaches enable chip engineering firms to deploy resources on demand, meet peak requirements, or even run the entire design cycle in the cloud — eliminating the need for physical deployment facilities, added personnel and up-front CapEx.

Compute Demand Across EDA Project Timeline



Source: AWS, SNUG 2025

Cloud-Based Solutions

Synopsys Cloud – Instant On-Demand EDA Tools & IP Access

Synopsys Cloud is a complete EDA license management automation platform that offers industry-leading capabilities for **BYOC (Bring Your Own Compute) deployments**, as well as a **browser-based EDA-as-a-Service experience** for chip designers who are seeking to deploy completely on the cloud. It includes both EDA tools and seamless access to Synopsys semiconductor IP. With its **unique FlexEDA pay-per-use**, customers can access unlimited EDA licenses by the minute to significantly accelerate time to market.

Accelerated Setup Time: Ready to use on day one with pre-configured compute and pre-optimized flows, Synopsys Cloud is faster and easier than traditional on-premises installations and new on cloud EDA deployments.

Faster Time-to-Results with Per-Minute Licensing Automation: FlexEDA pay-per-use licensing automates license management and provides instant, on-demand access to unlimited EDA licenses by the minute.

Increased Foundry Confidence: Foundries must validate the security and regulatory compliance of the EDA environment before providing PDKs and collateral for design tapeouts. SOC 2 Type 2 compliant, and pre-approved by 7 of the world's top foundries, Synopsys Cloud enables chip designers to easily access foundry PDKs and significantly accelerate time-to-tapeout.

ZeBu Cloud – High-performance Pre-Silicon Emulation

ZeBu Cloud is a Synopsys-hosted solution that accelerates AI chip innovation at scale with flexible, turnkey pre- and post-silicon emulation. It enables faster RTL verification and regression, software bring-up, performance validation, power analysis, and system validation for complex AI chip designs.

- ▶ **High Performance:** Industry-leading emulation performance.
- ▶ **Highly-Secure:** ISO27001 certified.
- ▶ **Reliable:** Established reputation for being a proven resource for up-time during weeks-long, high-volume generative AI and HPC software workload runs.

Use Cases

For Enterprises: Manage peak overflow without over-investing in on-prem capacity.

For Startups: Quick deployment without CapEx investment or in-house IT resources.


Scaling for the Future

Synopsys is tripling U.S. capacity to meet growing demand **and has certified a data center in Europe for ZeBu Cloud** to address concerns about data sovereignty and compliance with regional regulations.

Why Chip Design in the Cloud Changes Everything

Unlimited Resources, Streamlined Design, Faster Results.

“When we ran into an issue, we could increase our storage 2x or 4x without any disruption. It literally took 5 minutes to go from 99% usage back down to 50% and that was really beneficial.”



Jitendra Mohan
CEO, Co-Founder



“We completed our chip design in less than a year on AWS, compared to the years it would have taken using an on-premises HPC environment.”



Sanjay Gajendra
President and COO, Co-Founder



Overcoming Challenges to Unlock Multi-Die Opportunities

The semiconductor industry is rapidly adopting multi-die designs to meet the growing compute performance demands of AI workloads. **This approach enables manufacturers to push beyond reticle limits and Moore's Law**, improve performance and system power, accelerate time-to-market, and allow flexibility in product variations.

However, it also introduces many challenges — **Synopsys delivers a comprehensive and scalable multi-die solution to ease this transition**, and is actively partnering with leading ecosystem players to help customers solve their most advanced challenges and move innovation forward.

Addressing these challenges to realize the multi-die opportunities requires a complete and scalable solution

Opportunities

- ▶ Accelerating the scaling of system functionality
- ▶ Moving beyond reticle limits
- ▶ Reducing risk and time-to-market by reusing proven designs/dies
- ▶ Lowering system power while increasing throughput
- ▶ Creating new product variants for flexible portfolio management

Challenges

- ▶ Heterogeneous integration
- ▶ Power, signal integrity (SI), EM/IR, and thermal management
- ▶ System pathfinding
- ▶ Die-to-die and chip-to-chip connectivity
- ▶ Software development and modeling
- ▶ Layout-versus-schematic (LVS) checking
- ▶ Capacity and performance for system verification
- ▶ System signoff analysis
- ▶ Hierarchical test and repair
- ▶ Reliability and health
- ▶ System yield
- ▶ Memory utilization and coherency

05

Advanced Packaging & Multi-Die Design

BREAKING THE RETICLE LIMIT: FAST HETEROGENEOUS INTEGRATION
FROM EARLY ARCHITECTURE EXPLORATION TO MANUFACTURING





Many of these applications are working hard for training of large language models, huge data models, which have trillions of parameters. To really train such models, you need a lot of compute power, a lot of memory resources. One fundamental thing that Synopsys is really focused on is as customers are moving on this journey, how we can accelerate it for them, how we can make it easier for them to transition from system on chip to this new system of chiplets. To be able to do that, you have to look at the whole solution in a very holistic way.



Shekhar Kapoor

Executive Director, Product Line Management




Shekhar Kapoor at Chiplet Summit 2025

A Comprehensive and Scalable Solution for Multi-Die Design

Synopsys is at the forefront of this fundamentally new approach to chip package design, delivering solutions that **seamlessly integrate early architecture exploration, automated die-to-die routing, and silicon lifecycle management.**

Synopsys eases the shift to multi-die design with a comprehensive and scalable solution. The solution, including EDA and IP products, enables **early architecture exploration, rapid software development** and system validation, efficient die/package co-design and optimization, robust die-to-die and chip-to-chip connectivity, and improved manufacturing and reliability.



Comprehensive & scalable solution for fast heterogeneous integration

Early Architecture Exploration	Software Development & Validation	Design Implementation	Silicon IP	Manufacturing & Health
Optimize performance, power & thermal with early exploration & partitioning	Rapid software development & validation with high-performance emulation & prototyping	Efficient die/package co-design & optimization with unified exploration-to-signoff platform	Best-in-class, silicon-proven IP portfolio for 2.5D & 3D standard & advanced packaging technologies	Improve health & reliability from production to in-field with test, analysis & repair solutions

Power & Thermal Management in Advanced Packaging

As AI workloads scale, power and heat dissipation become primary constraints in multi-die system performance and reliability. Designing for **efficient power delivery and thermal regulation is critical to prevent performance throttling and long-term degradation.**

Synopsys 3DIC Compiler delivers an end-to-end solution for early-stage thermal and power-aware design, integrating pre-silicon modeling, and AI-powered optimization.

To **solve system-level electrical and thermal coupling interactions** of 2.5D/3D packages in full detail, Synopsys has teamed in [close collaboration with Ansys, part of Synopsys](#), on multiphysics analysis and design solutions based on **Ansys RedHawk SC Electrothermal**. The tightly integrated solution provides comprehensive power integrity, signal integrity, thermal integrity and mechanical stress simulation for 2.5D/3D multi-die systems.



Power Challenges

Power Distribution Efficiency

Minimizing IR drop and power noise across multiple dies requires advanced voltage regulation strategies.

Power delivery networks (PDN) must support high-current demands while minimizing parasitic losses in dense 3DIC environments.

Decap & VR Design Optimization

Decoupling capacitor (decap) placement is essential to **stabilize voltage fluctuations.**

Voltage regulation (VR) architectures must **support dynamic AI workloads without introducing excessive power overhead.**

Thermal Constraints

Heat Dissipation & Hotspot Management


High-density stacking of dies increases **localized heat concentration**, requiring enhanced thermal conductivity solutions such as:

- **Silicon bridges with integrated thermal pathways**
- **3D vapor chambers**
- **Active liquid cooling strategies**

Cooling Strategies for Multi-Die

High-power AI accelerators demand breakthrough thermal management.

Embedded microfluidic cooling, thermal vias (TSVs), and ML-driven thermal optimization work together to dissipate heat efficiently — minimizing hotspots in real time.



Industry Collaboration for a Robust Multi-Die Ecosystem

Industry Initiatives Driving Multi-Die Design

UCle Leadership

Synopsys is a contributor member of the UCle™ Consortium, helping drive industry-wide adoption of open, interoperable die-to-die connectivity.

Foundry Partnerships with Intel Foundry, Samsung Foundry, and TSMC

Supporting leading foundries' advanced package and process technologies

Arm Chiplet System Architecture (CSA)

Synopsys plays a key role in developing and supporting Arm CSA, currently at rev 1.0, an emerging chiplet system partitioning standard within the Arm ecosystem.


Synopsys & AMD: Partnering to Drive Multi-Die Innovation Forward

“Our long term relationship and close partnership with Synopsys was an essential factor for AMD’s ability to address the rapidly increased complexity of high performance chip design. Synopsys technology was integral to the creation of the MI300. The AMD MI300 leveraged Synopsys IP and EDA suite to create the world’s most advanced multi-die package. Our flow included Ansys simulation and analysis solutions.

This cutting edge EDA flow was critical to ensure first-pass silicon success and launching the MI300 on schedule to meet customer demand.”



Mark Papermaster
CTO and Executive Vice President



06

Early Software Development & System Verification

ACCELERATING SOFTWARE READINESS & DEBUGGING
EFFICIENCY FOR HIGHER-QUALITY PRODUCTS



Shift-Left: Why Early Software Development Matters

In the era of AI-driven silicon design, the demand for faster time-to-market has been impacted by:

Compressed
Design Cycles =

<12 months

Delayed Software
Readiness:

61%


High Debugging Costs:

Up to 40% of project time

Virtual Prototyping & Hardware-Assisted Verification

To meet these challenges, early software development must begin long before silicon is available. Synopsys enables this with Virtual Prototyping and Hardware-Assisted Verification (HAV) — **a Shift-Left approach that allows software teams to develop and test code in parallel with hardware design.** Traditionally, software engineers had to wait for fabricated silicon, leading to delays, costly debugging, and limited optimization. Now, with virtual prototyping, teams can start writing firmware, OS, and application-level software months in advance of silicon, closing the gap between software and silicon readiness.

Use Virtual Prototypes to Shift-Left Software Development



Synopsys Solutions: Virtualizer & Virtualizer Development Kits (VDKs)

Synopsys Virtualizer™ is a **tool suite to enable software development to start earlier by using virtual prototypes**. The output of Virtualizer is a Synopsys Virtualizer Development Kit (VDK). The VDK is an electronic digital twin of an electronic system combined with software productivity tools meant for development and testing. A VDK is executing unmodified binary production software with closest real-hardware behavior.

By enabling teams to develop, test, and optimize software before silicon is ready, Synopsys VDKs help reduce risk, improve quality, and accelerate time-to-market.

Industry-Leading Virtual Hardware Solution

- **Fast**, multi-abstraction virtual hardware simulation with multi-threading technology
- **Largest model library** for verified IP & subsystems
- Unified system-level debug and analysis for **maximum debug productivity**

Optimized for Arm-Based AI Chip Design

- Extensible **virtual prototype** based on Arm Neoverse CSS with Synopsys IP and AI subsystems, supporting full **virtual to hybrid emulation** set up
- Full **early software stack and tool development** including boot code & kernel, OS & application, driver, and SDKs
- Supports Arm hardware with **native execution** for fast simulation performance productivity

Trusted by Hundreds of Customers


- **Successfully deployed at companies of all sizes** for AI driven architectures such as HPC and automotive
- **12 months before silicon** to large parallel SW validation in CI/CD workflows

Hardware-Assisted Verification (HAV): Emulation & Prototyping

The Challenge of Compounding Complexities

Compounding complexities in the AI era — spanning software, hardware, interfaces, and workload-optimized architectures — are creating tremendous challenges for those developing today's most advanced AI accelerator chiplet, system on chip (SoC), and multi-die designs. Verification and software development processes, in particular, can require **quadrillions of test and validation cycles** across simulation, emulation, and prototyping stages. As a result, there has been a distinct need for faster HAV solutions that accelerate the development and validation of advanced AI chip designs and help ensure optimized functionality of the software-defined systems they enable.

Verification Challenge:
Quadrillions of Cycles



Addressing Verification Challenges

IP to Multi-Die Verification is a critical bottleneck — investing in HAV early prevents costly delays and ensures software and hardware function seamlessly from day one.

Why Hardware-Assisted Verification is Essential for AI & Multi-Die Systems

Software-Ready Months Earlier — Software development runs in parallel with hardware, enabling immediate execution upon silicon return.

Accelerated Debug & Verification — Virtual models provide a controlled environment for early bug detection.


Faster Time-to-Market — Reduces costly post-silicon debugging and redesign cycles.

Multi-Die Scalability — Purpose-built for AI-driven, chiplet-based architectures.

Cost-Effective — Reduce expensive late-stage fixes.

System Test Generation — Identifies difficult architecture bugs

From ISA to System: Verification Across CPUs, Clusters, SoCs & Chiplets



Synopsys HAV Solutions

Synopsys' HAV Portfolio enables early, parallel validation to reduce risk and accelerate development

Emulation: Synopsys ZeBu® emulation systems (ZeBu Server 5, ZeBu EP, ZeBu-200) deliver the performance required for verification teams and software developers working on advanced chips to succeed.

Prototyping: Synopsys HAPS® (HAPS-100, HAPS-200) is the industry's highest performance prototyping solution for software and system validation for IP blocks to complete SoCs.



Built for AI Verification at Scale:


- ▶ **Seamless Emulation & Prototyping:** Synopsys Emulation and Prototyping Ready (EP-Ready) Hardware allows teams to transition between emulation and prototyping, optimizing ROI.
- ▶ **Hybrid Technology:** Combines virtual prototypes with hardware emulation for fast OS bring-up.
- ▶ **Modular and Scalable for Multi-Die Systems:** Designed to handle IP-to-Multi-Die Verification, supporting increasing complexity in AI architectures.
- ▶ **Cost-Efficient Deployment:** A single hardware platform can be reconfigured for different verification tasks, increasing ROI.
- ▶ **Future-Proofed Verification:** Adaptable architecture ensures compatibility with evolving workloads, including next-gen AI chips.

Power & Performance Simulation & Analysis with Real-World Use Cases

With real-world workload emulation, ZeBu Empower provides **accurate power profiles and activity data that feed directly into the implementation flow**. This enables design teams to optimize PPA, sign off power and IR drop with confidence, and debug using actual system behavior. Combined with tools like Fusion Compiler™ and PrimePower, ZeBu Empower ensures designs are free from power bugs and ensure that designs stay within the expected power budget.



ZeBu Empower



Case Study: NVIDIA & Shift-Left Success

“With the increasing market requirements for handling large AI computational data sets driving the need for enormous GPU and CPU computational power, the development time for NVIDIA’s next generation AI systems have become highly compressed to a yearly release cycle, necessitating best-in-class prototyping solutions. Synopsys HAPS-200 offers the fastest prototyping speed in the industry. The 50 MHz performance we have been able to achieve with HAPS-200 has been key to boosting productivity of our software development teams. We are looking forward to scaling our HAPS-200 deployment to take full advantage for our software development teams.”



Narendra Konda

Vice President, Hardware Engineering



[Narendra Konda speaks at a Synopsys HAV event, Feb. 2025](#)

07

Security Considerations

END-TO-END SECURITY FOR AI & SILICON



AI-Specific Security Challenges


As the AI-driven era of pervasive intelligence propagates, security concerns have become more critical than ever. AI-specific hardware, multi-die design, and interconnected chiplets introduce new vulnerabilities across design, deployment, and lifecycle stages. Ensuring robust security requires a comprehensive, end-to-end approach.

Key AI Security Threats:

- ▶ **Inference Tampering:** Malicious actors can alter model outputs, leading to misclassifications.
- ▶ **Data Poisoning:** Attackers may inject corrupted data into training sets, distorting AI performance.
- ▶ **Model Theft:** Unauthorized access or reverse-engineering of AI models compromises intellectual property.

To mitigate these risks, Synopsys provides extensive security solutions to safeguard AI chip and chiplet development — protecting **data**, **hardware integrity**, and **software reliability** from initial design through deployment.

AI Security Threats In AI Applications



Synopsys AI Security Solutions

A comprehensive portfolio of AI security solutions

Customers worldwide rely on these solutions to safeguard AI workloads across cloud, IoT, automotive, and mobile applications at every level:

Synopsys Security IP Solutions

Broad Portfolio of Certified & Standards HW & SW Compliant Solutions

Broad Portfolio of Certified Security IP	tRoot HSM	Secure Boot, Debug & Authentication	Interface Security	Arm Confidential Compute Architecture (CCA) Support
Cryptographic cores, PUFs, security accelerators, and trusted execution environments that ensure certified hardware and software integrity	Provides a secure enclave for processing sensitive data, ensuring root-of-trust integrity	Prevents tampering at the hardware and firmware levels	Protects data in transit across critical interfaces, including: • PCIe & CXL IDEs • DDR/LPDDR IME • Ethernet MACsec	Multiple interface security IP blocks that support Arm CCA standards, including PCIe IDE, CXL IDE, and DDR/LPDDR IME

Multi-Die Security & Future-Proofing with Post-Quantum Cryptography

As security challenges evolve, Synopsys is advancing solutions to protect against emerging threats like multi-die vulnerabilities and quantum computing risks. Leveraging Synopsys' holistic security approach, AI chipmakers can protect their products against present and future threats, ensuring resilient and scalable architectures.

Multi-Die Security

- ▶ **Per-chiplet Authentication & Secure Boot** – Ensures trusted execution by verifying the integrity of each chiplet.
- ▶ **Data Integrity & Confidentiality** – Protects die-to-die communication from unauthorized access.
- ▶ **Supply Chain Security Tracking** – Prevents tampering by verifying silicon sourcing across vendors.

Post-Quantum Cryptography (PQC)

- ▶ Quantum computing **threatens traditional encryption** methods.
- ▶ Synopsys offers **Agile PQC solutions** to enable quantum-resistant key generation and digital signature algorithms, ensuring compliance with next-gen cryptographic standards.
- ▶ Other cryptographic cores (**symmetric, hash, TRNG, PUF**) continue to be quantum safe.

08

Scalable End-to-End Test Solution

OPTIMIZING TEST COSTS FOR MANUFACTURING, IN-FIELD AND IN-SYSTEM APPLICATIONS




Key Test Challenges with AI and HPC Based Chips and Multi-Die Chiplets

Key Test Challenges:


- ▶ Early RTL testability analysis and optimization are essential to identify and resolve DFT issues, especially given implementation challenges from numerous physical cores, abutment-based floorplans, and heterogeneous die integration.
- ▶ Limited tester access to individual die GPIOs and test IOs, drastically increasing test costs for logic, memory and IO tests.
- ▶ Escalating test data volume and pattern count driven by design complexity and advanced technology nodes.
- ▶ Test requirements for reliability and aging concerns of complex interconnect testing, including high-speed interfaces
- ▶ Need to extend structural testing beyond manufacturing to in-system and in-field use

Synopsys TestMAX™ delivers a scalable end-to-end test solution across manufacturing, in-system, and in-field phases of silicon, by leveraging advanced scan compression, high-throughput fabric architectures, and efficient parallel core testing, resulting in significantly reduced tester cycles and test data volume.

Turn-around Time




Test Cost and Quality



Synopsys TestMAX: Scalable End-to-End Test Solution

Unified Compression Architecture in SEQ (Sequential Compression)




73%
Fewer
Patterns

84%
Fewer
Cycles

88%
Less
Data

5x*
Faster
Runtime

High Throughput Streaming Fabric (SF)



5x*
Lower Test
Time

Test on High Speed Functional I/O with HSAT (High-Speed Access and Test)



10x*
Lower Test
Time

Achieve test cost reduction, higher test quality, and engineering efficiency with TSO.ai, AI-driven test space optimization

09

Silicon Lifecycle Management (SLM) & Reliability, Availability, and Serviceability (RAS)

MAXIMIZE SILICON PERFORMANCE & RELIABILITY



The Importance of In-Field Monitoring & Analytics


Why Silicon Lifecycle Management Matters

Silicon Lifecycle Management (SLM) optimizes silicon metrics across the entire system lifecycle, from initial design to in-field monitoring. As AI and high-performance computing (HPC) applications grow increasingly complex, ensuring long-term silicon performance and reliability has never been more critical.

Key Challenges:

- ▶ **Mission-Critical Reliability for AI & HPC:** Advanced workloads demand continuous monitoring to prevent silent data corruption (SDC) and maintain operational integrity including multi-die systems.
- ▶ **Power & Performance Optimization:** Vmin and voltage scaling across different silicon process variations and in-field environmental conditions.
- ▶ **Manufacturing Variability & Test Escapes:** Small inconsistencies during fabrication can lead to unpredictable failures in the field.
- ▶ **Aging & Degradation:** Chips experience wear over time, impacting performance and longevity.

In-Field Monitoring: A Lifecycle Approach To Reliability, Performance & Predictive Insights




Synopsys SLM Solutions for Optimized Performance, Power & Long-Term Reliability

Synopsys enables silicon and systems vendors with the instruments to tackle these challenges head-on. This includes embedded in-chip monitor IP, production analytics solutions, and in-field analytics software that track real-time conditions, detect anomalies, and deliver actionable insights.

Key features that improve reliability, performance, and efficiency include:

Silicon Lifecycle Management



Deploying SLM for In-Field Reliability, Availability & Serviceability (RAS) in AI-Based Systems

Synopsys and Amazon Web Services (AWS) have [collaborated closely over the last few years](#) to create and implement a unique solution that leverages an existing functional protocol-based high speed interface for testing and monitoring which has provided a consistent portable method to test silicon throughout its lifecycle.

“Synopsys SLM High-Speed Access & Test IP allowed us to conduct scan test over a functional interface. This enabled running scan test through the entire silicon lifecycle including — Wafer Sort, Final Test, System-Level Test, and In-System Test.”



Amit Pandey
Principal Engineer ASIC



Architecture For AWS ML Server





Figure shows the high-level architecture of AWS machine learning (ML) server based on AWS Trainium SoC. Synopsys SLM HSAT IP is integrated into the SoC and can be accessed via an SPI or PCIe interface.



10

The Road Ahead with Synopsys

TURNING CHALLENGES INTO SILICON SUCCESS

CONCLUSION

Key Takeaways for Chip Innovators



A Massive Market Opportunity: The AI chip market is growing fast, but success requires navigating increasing complexity.



Design Challenges Are Multiplying: Silicon, packaging, software, power, and performance are all becoming more complex — resulting in potential for silicon respins.



Your Roadmap To Success Is Parallelization: Start with architecture exploration to set the direction, then design and verify silicon design, package planning, and early software development as parallel, integrated tracks.




Address Power & Performance Early for Maximum Impact: Optimize power and performance at the beginning of the design cycle with real-world workloads for better results, rather than addressing them later.



System Verification & Validation At Scale: Validate across quadrillions of cycles using real workloads to ensure performance and reliability.



Prepare For Security Threats: Increasing security threats in both training and inference, as well as the multi-die supply chain, must be met with a holistic strategy.



Scalable End-to-End Test Solution with Multi-die Test Support: Optimizing test costs for manufacturing test, in-system and in-field test applications.



SLM Offers Immediate & Long-Term Benefits: Power optimization, in-field monitoring, and RAS capabilities are key to long-term success.

Synopsys AI Chip Development Solution Benefits



Enhanced Productivity:

- Power-optimized, silicon-proven IP
- Industry-leading AI-enabled EDA solution
- Industry's fastest HAV platforms
- Easy-access cloud-based scaling options



Multi-Die Leadership & Expertise:

- Comprehensive and scalable solutions
- Proven EDA and IP products for fast heterogeneous integration from early architecture to manufacturing



Pre-Silicon Software Validation at Scale

- Simulate billions of real software cycles to fine-tune power, performance, and area before tapeout, ensuring higher quality and first-pass success

CONCLUSION


How Synopsys Can Support Your AI Chip Development

As an AI chip innovator, you provide your unique IP (the secret sauce) — Synopsys and our ecosystem partners deliver everything around it, including the tools to get you to Proof-of-Concept or market rapidly, productively, and reliably.

Connect with our experts for a consultation on customizing your AI chip development roadmap for first-pass silicon success.

GET STARTED TODAY

Synopsys AI Chip Development Product Portfolio



 Available on the cloud.

Further Resources

The Rise of AI Chips & Chiplets

[Re-engineering Engineering™: AI-Powered Innovation](#)

[Low-Power SoC Design for HPC Applications](#)

[Energy-Efficient Chips Solutions](#)

Mapping an AI Chip Development Roadmap

[Designing Chips for the AI Age with Synopsys | AMD TechTalk Podcast](#)

Architecture Exploration

[SoC Architecture Analysis and Optimization for Performance and Power](#)

[Optimize Multi-Die Chip Designs with Arm CoreLink CMN-700 & Platform Architect](#)

[Data Center IP Solutions](#)

Silicon Front-End

[Introducing Next-Generation Verdi Platform for AI-Driven Debug and Verification Management](#)

[Accelerating Coverage Closure with AI-Based Verification Space Optimization](#)

[VC Formal Datasheet](#)

[New and Faster Standards Driving Data Center's Paradigm Shift](#)

Silicon Back-End

[Fusion Compiler Datasheet](#)

[Lower Process Nodes Drive Timing Signoff Software Evolution](#)

[Synopsys DSO.ai Demo Video](#)

Cloud-based Solutions

[How AI Chip Startups Use Cloud EDA Tools](#)

[Boosting Cloud-Based AI Chip Design Emulation](#)

Multi-Die Design

[Innovate Faster with Synopsys Multi-Die Solution](#)

[Multi-Die Design Start Guide](#)

[Ansys-Synopsys Technology Update: Advances in Multi-Die Design](#)

Early Software Development

[Virtualizer: VDK Creation & Deployment Tools](#)

System Verification & Validation

[Verifying AI Designs: Solving the Challenge of Quadrillions of Verification Cycles](#)

[Introducing ZeBu-200 & HAPS-200: Faster Hardware Verification](#)

Security

[Security IP Brochure](#)

[Secure Interfaces for Critical Semiconductor Applications](#)

[SRAM PUF - The Secure Silicon Fingerprint](#)

Test

[TestMAX: Test all the Limits](#)

[TSO.ai: Fewer Patterns, Higher Coverage, Reduced TAT](#)

[Static Verification of RTL DFT Connectivity – Getting it Right the First Time!](#)

SLM & RAS

[Silicon Lifecycle Management \(SLM\): Optimize Silicon Health & Performance](#)

[Silent Data Corruption \(SDC\) / Reliability, Availability and Serviceability \(RAS\)](#)

[Panel Discussion | Synopsys/Arm/Google/Microsoft](#)

[In-Field RAS Challenges with Chiplets and AI-based Systems](#)

Sources

Learn more about Synopsys AI chip development solutions at synopsys.com.
Access additional case studies, white papers, and technical resources.

Sources & Additional Information

¹[Allied Market Research](#)

²[List of large language models](#)

³Synopsys Global User Survey, 2024

⁴[Firm Estimates a 2nm Chip Now Costs \\$725 Million to Design](#)

⁵[KPMG's 2024 Semiconductor Industry Outlook](#)

⁶[IAE Electricity 2024 - Analysis and forecast to 2026](#)

The background is a complex, isometric digital landscape. It features numerous glowing cubes and rectangular blocks in shades of purple, blue, and teal. These blocks are interconnected by a dense network of thin, glowing lines that resemble circuit traces or data paths. Some lines are straight, while others are bent at right angles. Small, bright points of light, resembling stars or data nodes, are scattered throughout the scene, adding to the futuristic and high-tech aesthetic. The overall composition is layered, with some elements appearing closer to the viewer than others, creating a sense of depth.

synopsys®

©2025 Synopsys, Inc. All Rights Reserved.

Ver 2.0, 9/2025