Synopsys Is Raising The Bar With The Next Generation Of AI For Chip Design

COULD WE REALLY DESIGN CHIPS IN 24 WEEKS INSTEAD OF 24 MONTHS?

INTRODUCTION

The Semiconductor Industry is enjoying renewed growth despite chip shortages plaguing everything from cars to kitchen appliances. But while the chips themselves continue to get faster and smarter, the chip design process itself hasn't changed that much in 20+ years. It typically takes 2-3 years to design a chip with a large engineering team and tens or hundreds of millions of dollars to get a chip from ideas to fabrication. But now, change is coming in the form of Artificial Intelligence, which has recently demonstrated significant improvements in optimizing layouts for power, area (cost), and performance.

Using a Reinforcement Learning (RL) based approach, similar to the one that beat the world's GO champion way back in 2016, Samsung announced that they now have a chip back from their factory that was optimized by the Synopsys DSO.ai platform we discussed in May 2020. As far as we know, this is the industry's very first working chip whose layout was designed by AI.

We expect a large share of the semiconductor industry will begin using these AI platforms; the impact is too large to ignore. In fact, EDA vendor Cadence Design Systems recently followed Synopsys' lead and has introduced AI tools that can dramatically improve performance, power, cost, and design. There is a lot of money to be made helping design teams produce better chips in significantly less time.

While Synopsys appears to enjoy at least a 16-month heads start over the competition, the company is already taking the next step, laying out the company's vision in a keynote address to the annual Hot Chips conference attendees. Synopsys plans to deploy new AIs to optimize other elements of the workflow in addition to the physical design aspects already in place with DSO.ai. The results could likely change the industry, as these advances are another nugget in what will prove a rich vein of value yet to be mined. This research paper investigates the new capabilities Synopsys has announced and explores the possible future of fast AI chips designing chips faster and better than our finest engineering teams can muster with today's design tools.

THE EDA LANDSCAPE

As we explored in a recent piece on Forbes, the technology approach that Synopsys has used in "place and route" has been picked up by NVIDIA, Google, and most recently by competitor Cadence Design Systems. The research undertaken at NVIDIA and Google is quite promising. Now, we suspect most semiconductor companies will soon use AI to improve quality and time to market, while those that don't will fall behind.

The landscape we explored, however, just barely scratched the surface of what is possible here. Looking at Figure 1, one can see that the work done by NVIDIA and Google focused on floor planning (the red circle) significantly improved the power, area, and frequency of their chips compared to what designers could achieve through manual experimentation. Floor-planning after all is one of the many NP-hard problems in EDA and cannot be solved analytically due to the large search space of potential solutions.. Depending on the approach, this search can take an incredible amount of computing. While the game of GO was conquered by evaluating 10^{360} possible moves, the game of chips can have $10^{90,000}$ possible ways to place and route. But physical design presents just the tip of the iceberg when it comes to optimization domains. Synopsys wants to consider the possibilities when one raises our sights to include more of the design process. If the Synopsys vision is correct, and we believe it is, we are standing on the verge of a revolution in chip design.

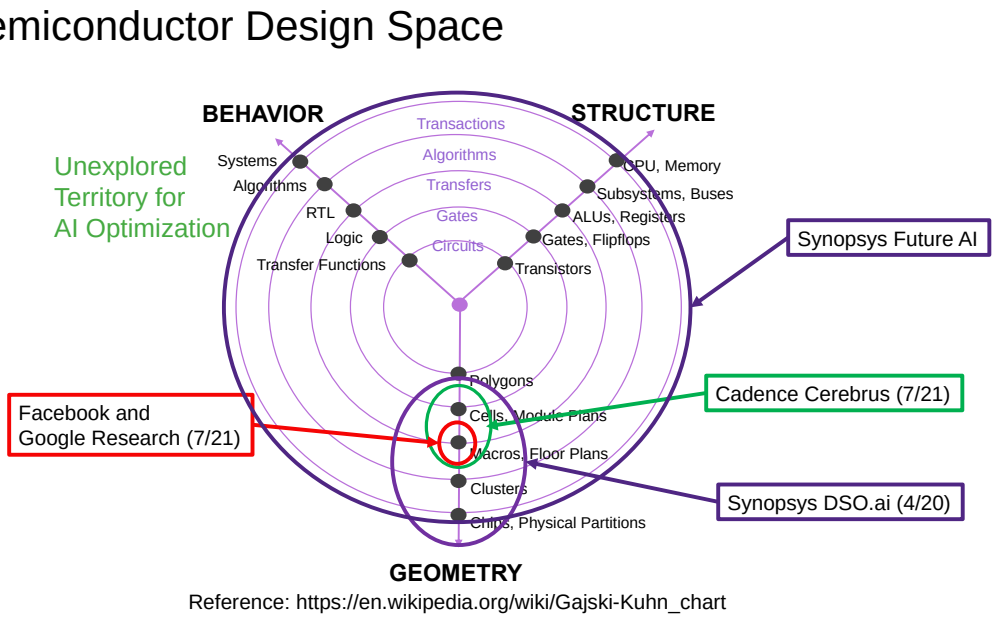

Figure 1: The chip design process aligns along three axes: Behavior (what the chip should do), Structure (how the chip will do it), and Geometry (how the chip is instantiated in masks for production).

How Far Can One Apply AI in Chip Design?

Synopsys and various researchers believe that the future will see many more aspects of design being assisted by AIs. Dr. Aart de Geus, Co-CEO of Synopsys, asserted at his keynote presentation that the company is expanding its AI-based optimization portfolio to address the needs of algorithmic (functional) and architectural (structural) exploration. In doing so, Dr. de Geus described a cyclone of AI-inspired assistance to circuit designers, expanding beyond physical design (placement) to help optimize all three axes concurrently, adding architectural and behavioral optimization as examples of what can be done. Synopsys' vision of using AI to concurrently optimize across domains that is completely new to the industry.

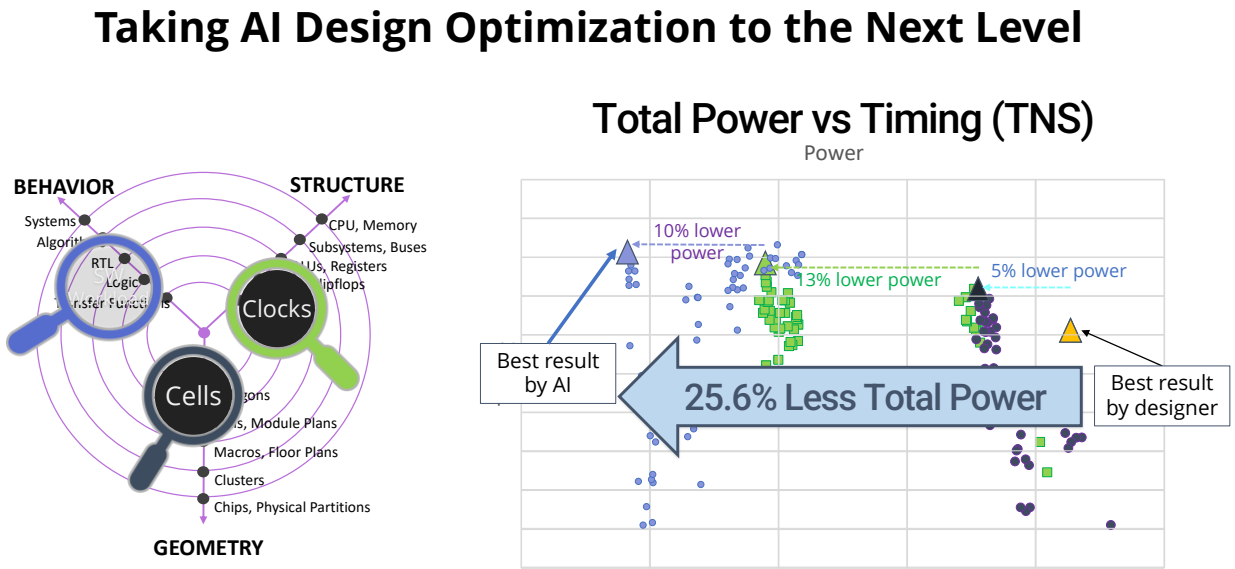

Figure 2: By applying concurrent AI optimization across all three axes of design, Synopsys was able to achieve a 25% reduction in power consumption. Source: Synopsys

In a real-world design case study shared at the Hot Chips event, Synopsys demonstrated significantly lower power using AI across the entire design space. Optimizing the first domain (geometry) with Synopsys' first generation DSO.ai product can reduce power by as much as 5% over what a design team could achieve without AI. Here is where the second-generation technology comes in: adding structural exploration to geometry, essentially concurrently exploring different architectural choice points and associated physical layouts for each, power can be reduced by an additional 13%. Finally, the AI turns to *behavior*, using the software workload itself to optimize the chip's power consumption, again concurrently with architecture and layout, by an additional

10%. Add it all together, and a design team could reduce power by an almost unbelievable 26% compared with the best result delivered by a talented chip designer.

A Concurrent Design "Cyclone"

Traditional design approaches limit interaction to successive steps in a largely linear workflow. The architectural design dictates the structural design, which can then provide a feedback loop to the architecture team. Similarly, this structural design determines the requirements of the logic design. Issues and opportunities uncovered during logic design choices can influence tradeoffs back with the structural team, and so forth. This approach, while well understood and embraced throughout the semiconductor industry, limits the options being considered across the chip, and consequently limits the extent of optimality being considered. The Synopsys approach essentially replaces this model for an end-to-end holistic design approach assisted by AI.

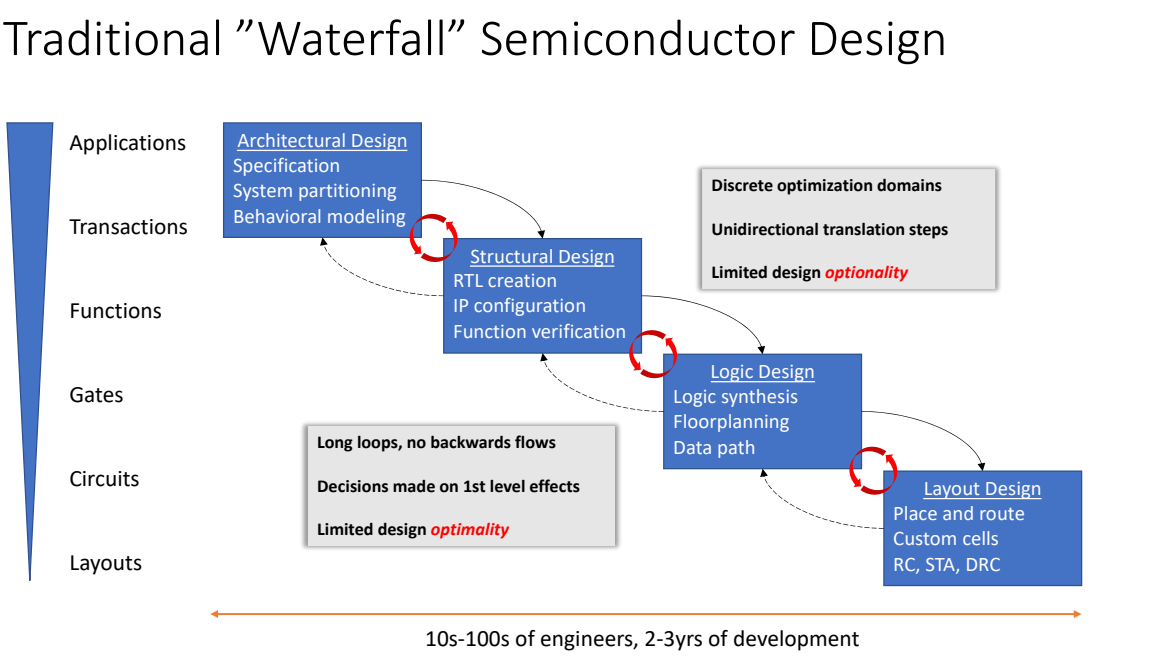

Ref: Turnbull 1992

Synopsys refers to a new "Cyclone" approach which essentially creates a cycle of concurrently evaluated design options. This "optionality" feeds a wide set of solutions which are concurrently modeled with AI to evaluate performance, power, and die area impacts. This approach has been pioneered by Synopsys, and appears to produce better, globally optimized results (lower power, higher performance, smaller die size). Where the new approach really benefits semiconductor companies, however, is in the significantly reduced engineering time. As the approach is further implemented and

refined, we can expect to see more design teams embrace the potential of AI design assistance.

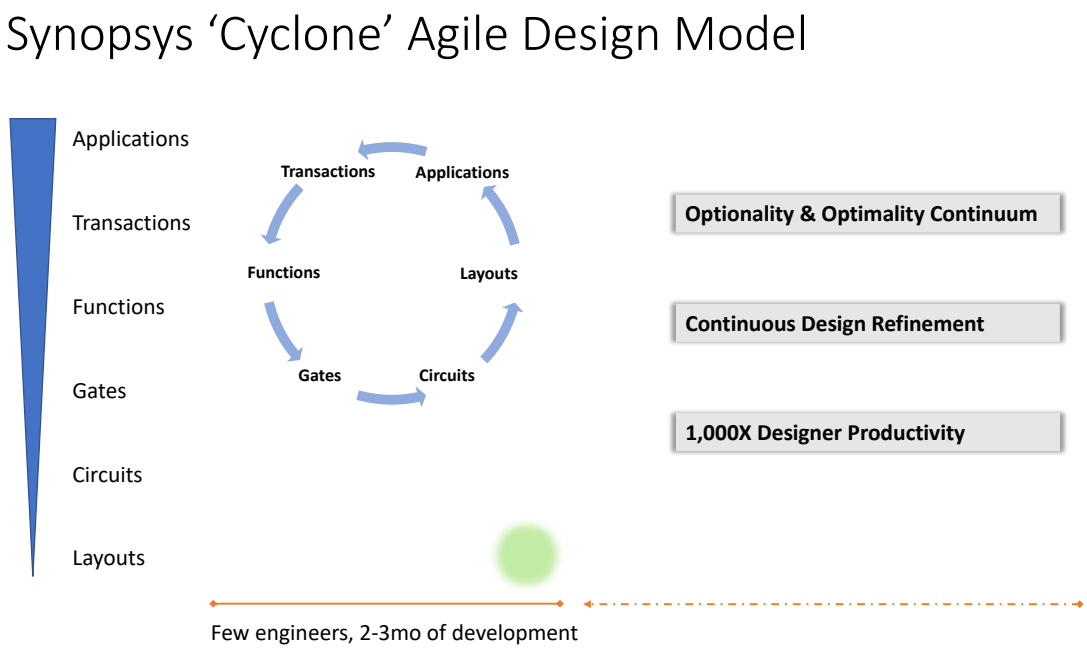

Figure 3: The Synopsys "Cyclone" began with Physical design (layout) and is expanding across other elements of the design process to expand the options being considered and determine the optimal approach and implementation.

IMPLEMENTING THE DESIGN CYCLONE

Tackling such a massive problem & solution set requires the designers to guide the potential decision space into a manageable size. A brute force method could work but would perhaps cost millions of dollars to run such a huge model. Instead, Synopsys is using reinforcement learning to iterate on a set of options for software, structural, and physical design choice points. This then presents the design team with alternatives as to what attributes they want to optimize and at what benefit and expense.

Architectural search and power optimization are the first waves of this approach. Every circuit behaves differently under various workloads. Using reinforcement learning, a design's structure can be evaluated and optimized for performance and energy consumption against a variety of software loads. Similarly, circuits can be evaluated against a variety of designs and optimized together with the placement and routing. It is this co-design and co-optimization that Synopsys is ushering as the next wave of AI-led or AI-assisted chip design.

It is important to point out that these AI platforms do not just output an answer to “do this”. The design team is presented with alternatives that can optimize performance, power, cost, or more likely a combination of all three. In one example, the team could choose to increase performance by 15%, or reduce size (and cost) by 18%. Or the team could blend the two designs and increase performance by 8% while still reducing die size.

FROM SOFTWARE-DEFINED TO SOFTWARE-*DESIGNED* HARDWARE

Software is eating the world, but it is AI that is now eating the software. As more and more software applications become data-driven, and neural networks are already crossing the trillion-neuron mark, how is the semiconductor industry going to deliver the petaflop-months of compute required for AI from the data center to the edge?

Software-defined hardware has been proposed by industry luminaries as an elegant solution. It is based on the premise that chip could become personalized to the needs of specific applications, putting software in direct control of the instruction set architecture (ISA), chip structure (microarchitecture), and implementation method (silicon technology). *Personalizing* chips could deliver 1,000X better performance and energy efficiency but here comes the problem: it currently takes 2-3 years to put a new idea into an actual socket.

AI could be the answer. AI-driven design systems like DSO.ai have delivered the productivity to accelerate months-long design tasks down to days. With more global system level optimization on its way, we could have the ability to create new, personalized chips, in just weeks. More AI in the design process could indeed enable the expansion of the software-*defined* hardware concept to software-*designed* hardware, making it both possible and economically attractive to deliver many flavors of acceleration to match the needs of the most intricate applications.

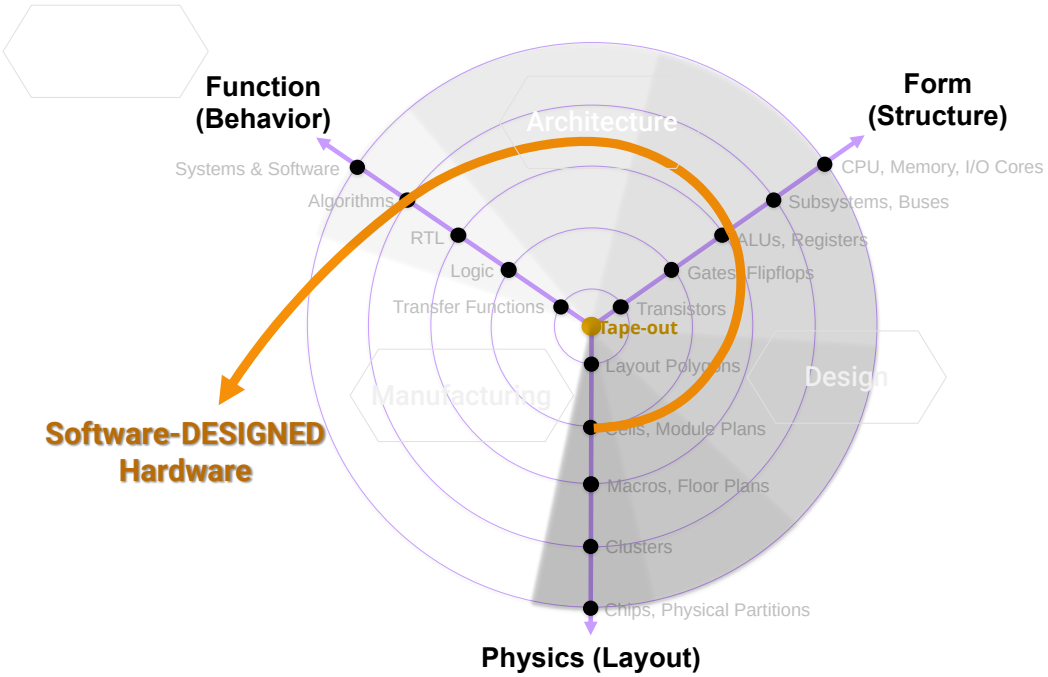

Figure 4: The path to Software-Designed Hardware: Creating personalized silicon tailored to the exact needs of the software application (workload), made possible by the advent of autonomous design

CONCLUSIONS

While the world becomes ever more dependent on semiconductors, and increasingly concerned with power consumption in data centers, the promise of AI to shorten development time while improving cost, performance, and power offers an irresistible lure. Giving design teams a platform for continual exploration to produce a better design will transform the workflow and the chips themselves. All this magic comes at the expense of a tremendous amount of computation, but those investments will return faster and more energy efficient products, and do so in less time and expense.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR: Karl Freund, Founder, and Principal Analyst at Cambrian-AI Research

INQUIRIES:

Contact us if you would like to discuss this report, and Cambrian-AI Research will respond promptly.

CITATIONS

This paper can be cited by accredited press and analysts but must be cited in-context, displaying the author's name, author's title, and "Cambrian-AI Research." Non-press and non-analysts must receive prior written permission from Cambrian-AI Research for any citations.

LICENSING

DISCLOSURES

This document was developed with Synopsys funding. Although the document may utilize publicly available material from various vendors, including Synopsys, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.

DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Cambrian-AI Research disclaims all warranties as to the accuracy, completeness, or adequacy of such information and shall have no liability for errors, omissions, or inadequacies in such information. This document consists of the opinions of Cambrian-AI Research and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Cambrian-AI Research provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.