AI Could Be The Next Killer App In Semiconductor Design

Introduction

It goes without saying that designing complex semiconductors is an extremely challenging engineering process. While the chips themselves are small, and the individual features on the chips are tiny, as small as 7 nanometers (seven billionths of a meter), the intricacy of the design process is astronomical.

Electronic Design Automation (EDA) tool providers and semiconductor manufacturers are turning to machine learning (ML) to help improve engineer productivity and produce superior chip designs. Emerging from ad-hoc individual tool assistance, the state of the art is now entering a new phase to automate entire physical design workflows. These workflows are being optimized using recent advancements in ML, such as reinforcement learning, to shorten development time and meet aggressive PPA (power, performance, and area) targets.

The Challenges Of Optimizing the Broader Design Space

Consider the physical design workflow: once the logic of the chip has been completed, the project's physical design team is faced with determining the optimal layout, or floorplan, which can have a significant impact on the product's performance and cost. Their workflow is iterative, with team members changing many design parameters to try to find a good layout, then changing parameters again and again to find a better alternative. And any subsequent changes in chip logic can essentially reset the physical design process, often restarting it at the beginning (Figure 1).

Chip designers must derive a layout that will produce the best balance of design targets, like power consumption, performance, and die area (PPA), for each partition of their design, typically requiring several engineers some 20 to 30 weeks to accomplish. In fact, the team faces a massive "search" problem, as floorplan exploration alone can encompass $10^{90,000}$ possibilities of placing and routing sections of the chip. Clearly, it would be impossible for a team to explore even a small fraction of these configurations to find an optimal solution, and traditional programming techniques are largely inappropriate.

AI is already changing many industries and processes. The challenges in physical design are not dissimilar to those that face chemical engineers attempting to sort through trillions of alternatives in modern drug design, for example. Research by the Universities of Florida and North Carolina has demonstrated that using AI can result in an efficiency improvement by up to six orders of magnitude in screening proteins for new drugs. Applying these techniques to searching for optimal floor plans in IC design seems like a natural fit.

FIGURE 1: TRADITIONAL PHYSICAL DESIGN SPACE EXPLORATION

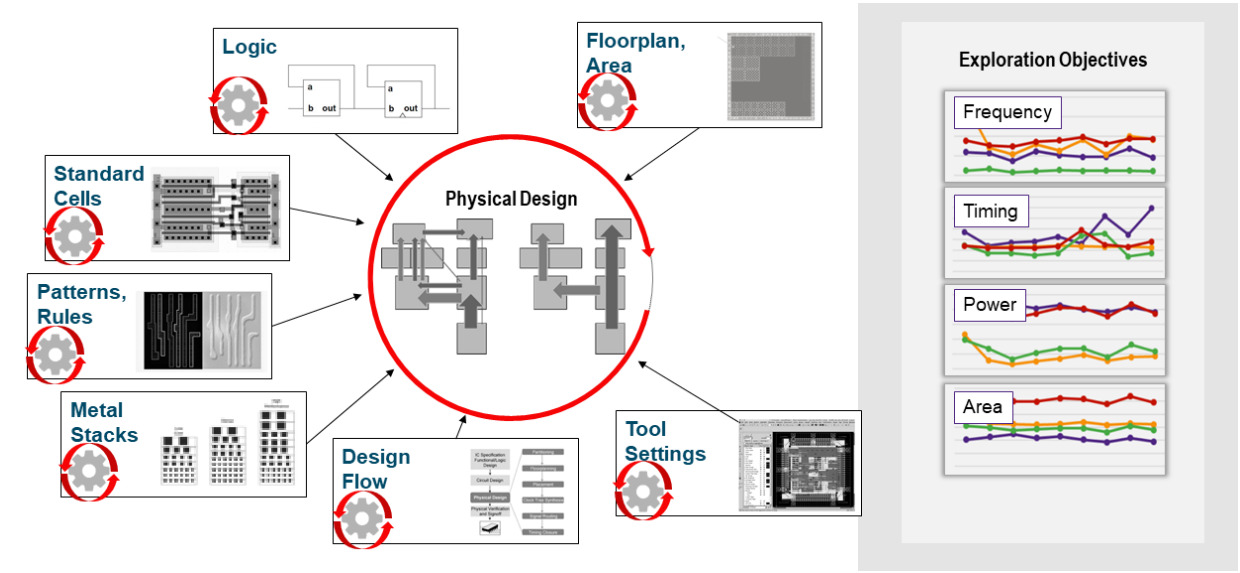

The traditional approach to physical design exploration is a tedious and time-consuming process that requires many expert engineers and can take months to complete.

Source: Synopsys[1]

One branch of ML that enables conducting these massive searches is reinforcement learning (RL), where software determines actions based on some sense of cumulative reward, such winning a game or achieving balanced PPA. Unlike supervised learning, RL does not require a massive tagged dataset to train the neural network. However, RL does require a large computational capacity to learn the "game." The application of AI in semiconductor design is not unlike that of autonomous driving, with infinite variables and countless tradeoffs that must be evaluated every millisecond to make critical

[1] All graphics in Figures 1-4 in this paper are the property of Synopsys.

decisions. And RL is only one of many techniques being explored to help improve the electronic design workflow.

The Evolution of ML in Electronic Design

Over the past 10-15 years, ML has been used as part of standard workflows in chip design, using supervised and more recently unsupervised learning to evaluate optimizations based on an offline repository of design data (scenario 2.1 in Figure 2). While very helpful, this approach still leads to a time-consuming trial-and-error process, and requires access to large volumes of data. Consequently, EDA vendors are beginning to embed ML into design tools to provide real-time refinement of design steps, like place and route. ML in individual tools provides significant benefits, modeling localized optimizations so the tools can optimize in-situ (scenario 2.2 in Figure 2). However, the engineering team still has to evaluate workflows manually, iterating over a myriad of slowly converging designs.

FIGURE 2: MACHINE LEARNING IN PHYSICAL DESIGN

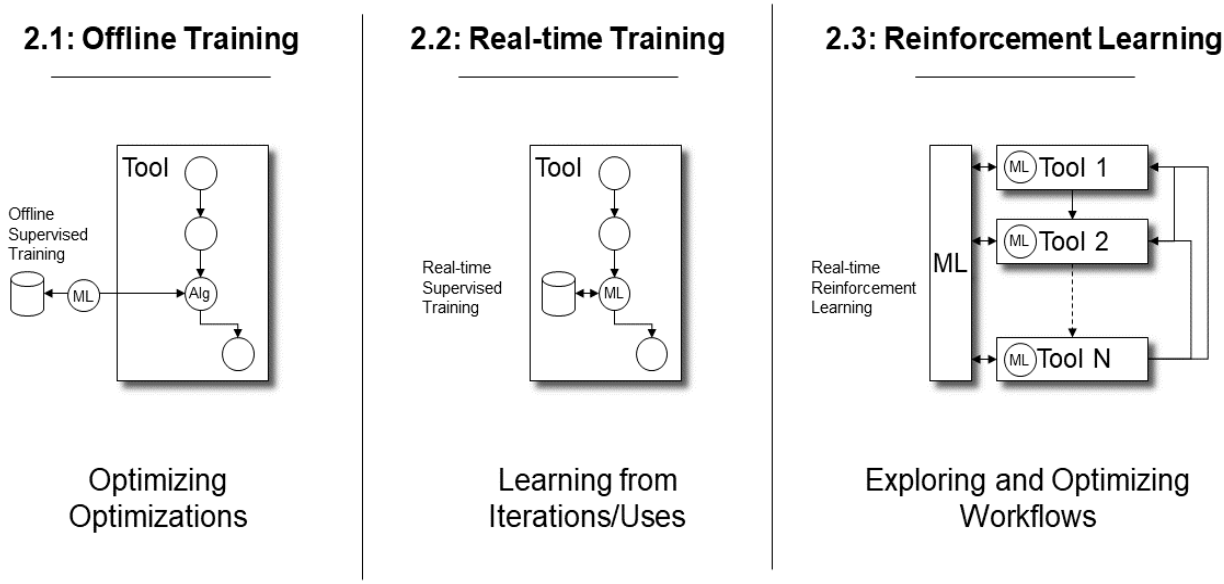

The use of ML in physical design has evolved from ancillary tools to localized optimization and most recently to providing global workflow optimization across many tools using reinforcement learning. Early results are quite promising.

A New Wave of AI in Chip Design Workflows

As mentioned earlier, new approaches to using ML in physical design are just beginning to emerge. Recently we ran into a new method that uses recent advancements like RL

to effectively search broad design spaces for a global near-optimal mix of design targets to meet a project's specific goals (scenario 2.3 in Figure 2). The Synopsys DSO.ai (Design Space Optimization AI) system is a prime example and an early innovator of this approach.

According to Synopsys, DSO.ai searches a vast space of potential physical design solutions for an optimal layout, saving time and typically producing a solution that improves performance and uses less power and less space. While this approach is very new, results are promising, producing better designs with far less work and time spent in physical design.

FIGURE 3: THE SYNOPSYS DSO.AI DESIGN SPACE OPTIMIZATION SYSTEM

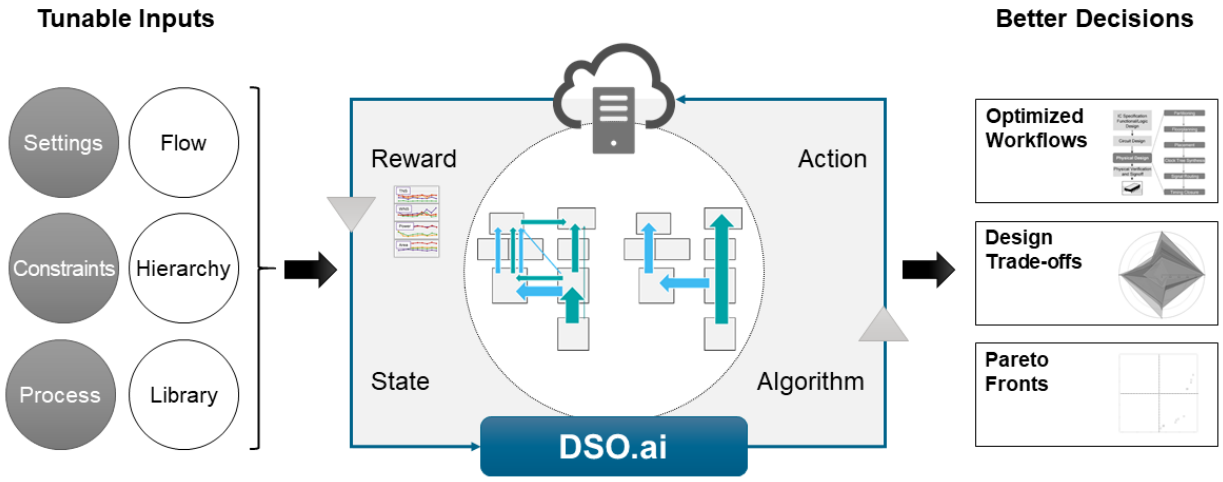

Synopsys DSO.ai shifts the burden of optimization from the traditional reliance on manual sweeps to one that relies on large computation to autonomously identify design spaces of potential solutions.

Synopsys has been developing and testing DSO.ai in close collaboration with its clients. Early results are quite impressive, spanning across a wide range of use cases including AI acceleration, networking, mobile, and automotive. In general, Synopsys claims its clients deploying DSO.ai are enjoying shortened time to results that often require fewer engineers. Clients are also producing designs that consume less power, require less area, and can out-perform chips from traditional design techniques. The primary benefit is meeting or exceeding PPA design goals more quickly and with fewer engineers (Figure 4).

FIGURE 4: SYNOPSYS DSO.AI CUSTOMER RESULTS

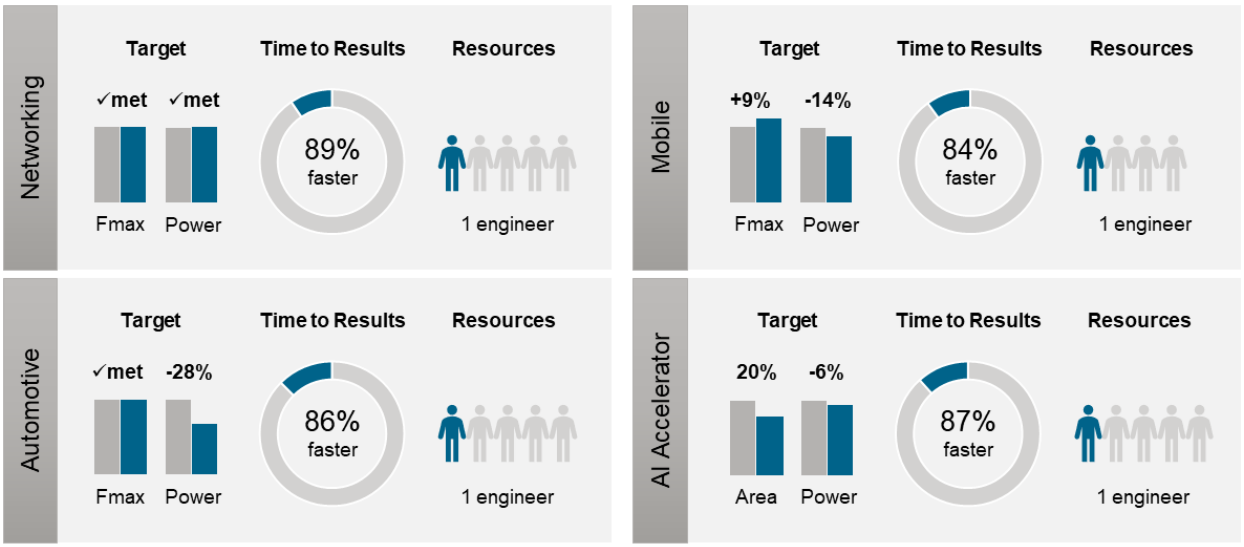

THE PATH FORWARD FOR AI IN CHIP DESIGN

Hardware design cycles can be long and expensive. The physical design phase in particular can be frustratingly lengthy and complex as the organization awaits tape-out and working silicon. AI has the potential to lower costs and produce superior products, and the journey down this path of innovation is only just beginning. The Synopsys DSO.ai effort has concentrated on using RL in physical design; however, the entire chip design workflow is ripe for ML-based optimization, which can produce dramatic savings and efficiencies and help teams get more competitive products to market faster.

RL is rapidly becoming a treasure trove of new applications for AI, but RL does require substantial computational resources. Fortunately, cloud service providers have been quick to offer GPUs. These GPU farms were recently extended to enable pools of interconnected resources for implementing large parallel models on demand, saving users time and money compared to on-prem solutions that may go largely underutilized. Large companies and startups are also bringing new domain-specific AI accelerators to market that promise dramatic performance improvements. These accelerators are lowering the cost of applying AI (including RL) and enabling more research and development in this area.

CONCLUSIONS AND RECOMMENDATIONS

The semiconductor design process is entering a new phase of innovation enabled by the application of Artificial Intelligence to streamline workflows and enable more efficient designs. In early tests, AI-driven optimization has been demonstrated to reduce time to market while meeting or exceeding PPA targets. Continued experimentation will help further explore the benefits and limitations of using RL and other ML techniques to optimize the physical design of complex chips being built on advanced process nodes. Given the rising costs of successive nodes − 7 nm design costs about three times that of 16nm chips, and 5nm design is expected to nearly double those costs yet again − the cost savings described in this paper will be welcome news to semiconductor companies. The accelerated schedules that AI could enable should help improve time to market and boost the competitive position of early adopters.

IMPORTANT INFORMATION ABOUT THIS PAPER

CONTRIBUTOR
Karl Freund, Senior Analyst at Moor Insights & Strategy

PUBLISHER
Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

INQUIRIES
Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts but must be cited in-context, displaying author's name, author's title, and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

LICENSING

DISCLOSURES
This paper was commissioned by Synopsys. Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

DISCLAIMER
The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. Moor Insights & Strategy 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 Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy 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.