

Preface

"Multi-die systems define the future of semiconductors" is an MIT Technology Review

When it comes to the business use cases for multi-die systems, Patrick Moorhead, founder, CEO, and chief analyst at global technology consulting firm Moor Insights & Strategy, notes that these custom designs may soon be a key differentiator for companies looking to stand out among competitors. “As more people are looking at more custom silicon as a way to differentiate what they bring to the table, that’s what businesspeople should be looking at,” he says. “Chiplets enable smaller companies with smaller pocketbooks to use semiconductors for unique competitive advantage.”

Gerry Talbot, a corporate fellow at semiconductor company AMD, boils the business value of chiplets down to the wide range of use cases for the technology. “I don’t think [business leaders] will be so excited about the technology itself,” he says, “as much as the application and the enablement of a unique user experience that can help sell their product.”

Reframing Moore's law

A critical concept in the history of semiconductor design is Moore's law: In the 1960s, businessman and engineer Gordon Moore observed that the number of transistors on an integrated circuit was doubling every 18 to 24 months. This pace of progress continued for decades. For businesses, it meant a consistent rate of chip technology advancement at a predictable price point.

In recent years, however, Moore's law has been slowing. The march to more advanced process nodes now requires more development time, and a more complex fabrication process has led to poorer chip yields (the number of functional chips produced). Compute advances can no longer be taken for granted.

According to Sassine Ghazi, president and chief operating officer at Synopsys, a semiconductor design

Semiconductor glossary

The following vocabulary is helpful for understanding the evolution of semiconductor design.

- An **integrated circuit**, or chip, is an assembly of electronic circuits on a flat piece of semiconductor material (usually silicon). This base material is usually called a wafer.
- A **transistor** is a foundational building block of all modern electronics, including integrated circuits. Transistors’ main job is to amplify or switch electrical signals.
- **Process nodes** are successive generations of the semiconductor manufacturing process. The term once correlated to the physical size of the transistor’s gate, typically measured in nanometers (e.g., “the 3 nanometer process”). Today, however, there’s little fixed relationship between a process node’s name and its physical features.

solutions company, the slowdown of Moore's law became particularly apparent about seven years ago, when the industry shifted from a process node of seven to three nanometers. "The scalability was still there, but the benefit you were getting from the process upgrade versus the predictability of silicon results and the overall cost was becoming questionable," he says.

With Moore's law being reframed to meet modern realities, business leaders across the board may be in for a rude awakening. Uri Frank, vice president of engineering at Google, says this is particularly true

Nevertheless, for many years, semiconductor companies reliably secured their market position by pushing their engineers to the next process node. Because that means of advancement is no longer reliably available, these companies now need to innovate across all aspects of chip design—leading to innovations like multi-die design, as well as a re-envisioning of the semiconductor ecosystem.

Redesigning the semiconductor ecosystem

The semiconductor ecosystem consists of many types of companies. A [Stanford analysis](#) divides them into seven core buckets: chip intellectual property (IP) core providers, electronic design automation (EDA) solution creators, specialized materials providers, wafer fab equipment makers, “fabless” chip companies (which focus on design and then contract out manufacturing of chips), integrated device manufacturers, and chip foundries.

Technology changes continually open new market opportunities, potential partnerships, and challenges across this ecosystem. Chip foundries like TSMC help foster chip innovation—although they don’t design product. “We just focus on technology and manufacturing,” explains Zhang. “We create a technology platform and allow our customers to innovate on it. We bring in different product players—including fabless design companies, system companies, for example—to develop a product that can be seamlessly integrated with our underlying technology and create innovative products to deliver significant benefits to the end users.” To successfully manufacture chiplet-based systems for its customers, TSMC

works closely with its ecosystem partners to facilitate collaboration between players in chip design, materials, testing, and packaging.

At the very beginning of the ecosystem, multi-die systems are changing how companies engage with chip IP cores—the specific ways in which companies design, patent, and sell integrated circuit layouts. “Today, it’s possible to buy IPs from a number of companies, but many don’t want to give away their ‘secret sauce,’” mansategular modefrom a8.346 33.5 and s-5 (lita)10.35 (e)Moorhea

these very large workloads, we actually need to be able to put more silicon in the package than would fit in a single monolithic die—you just physically could not print it in a single reticle,” explains Talbot. (The reticle defines the largest chip that can be manufactured with current technology.)

Frank also sees chiplet-based design as a vehicle to improve the viability of large language models (LLMs) and other generative AI engines like ChatGPT—technologies he says are currently seeing “exponential growth”—and enabling them to be deployed at scale. Beta testing is one thing, but these tools “require a lot of hardware, and they are actually hardware-limited today,” he notes.

Chiplet technology may also allow for better and more nuanced customization. With multi-die design, it’s possible to put some circuits on less advanced nodes while dedicating other critical tasks to newer nodes. “You can mix and match and optimize for a better performance-at-cost solution,” says Frank. Zhang explains that, for example, one piece of a chiplet can do computation-intensive work while another focuses on I/O interface and yet another is dedicated to memory.

François Piednoël, distinguished chief mSoC (multiple system on chip) architect at Mercedes-Benz, sees the power of this custom functionality in the auto industry. Chiplets, he explains, can facilitate complex power management in autonomous driving, which is defined by specific levels. “If a user wants to drive with Level Two, where they still need to be in control of the car and pay attention but with some lane-keeping assistance, chiplets can provide this capability without consuming the power of an entire chip,” he says.

Finally, there are potential cost benefits to multi-die design, particularly when it comes to manufacturing. “Chiplets are very economical, and they allow whoever

is designing them to avoid having to pay for very large dies—especially for a very complex machine,” says Piednoël. “If you were to try to do this in a monolithic way, the top-end, data-centered chips would be too expensive.”

Chiplet challenges

Despite the many advantages of chiplets, multi-die systems haven’t yet seen widespread adoption—there are still ecosystem innovations necessary.

Specifically, challenges can be found with integration technologies, power limits, and testing. “You’re dealing with physics when you have to connect between chips. Spent physical size—ther (ol o)0 -2.889pano(t)5 (er) .

The road to standardization

Universal Chiplet Interconnect Express (UCIe)—a comprehensive specification being co-developed by a number of semiconductor industry heavy hitters—is rapidly gaining popularity as a standard for die-to-die connectivity in chiplet-based design.


In short, UCIe is an open specification for die-to-die interconnect between chiplets. It specifies factors like the physical layer, protocol stack, software model, and compatible compliance testing procedures. “It makes it possible to take an off-the-shelf chiplet and connect to it,” says Uri Frank, vice president of engineering at Google.

Embracing the future

The dynamic nature of semiconductor technology is what draws many to work in the field to begin with. Though the shift to multi-die systems presents some hurdles, it also offers a wealth of opportunities for those who enjoy the satisfaction of complex problem-solving. “The fact that you’re always learning new things is one of the exciting things about this,” says...

He... including... “We’re not talk... to test the water—it’s... adding that Synopsys is co... 100 unique systems that are be... multi-die. “The need is there, the push... is absolutely there. Now what makes it ex... the innovation we’re seeing across the chain—... architecture all the way to manufacturing—and he... we’re collaborating to optimize the whole technology stack so multi-die systems can reach scale across global markets.”

“What makes
the innovation
across the chain
architecture



About MIT Technology Review Insights

MIT Technology Review Insights is the custom publishing division of *MIT Technology Review*, the world's longest-running technology magazine, backed by the world's foremost technology institution—producing live events and research on the leading technology and business challenges of the day. Insights conducts qualitative and quantitative research and analysis in the U.S. and abroad and publishes a wide variety of content,
