


Multi-die systems define the future of semiconductors



Preface

“Multi-die systems define the future of semiconductors” is an MIT Technology Review Insights report sponsored by Synopsys. The report was produced through interviews with technologists, industry analysts, and experts worldwide, as well as a cross-industry poll of executives. Stephanie Walden was the writer for this report, Teresa Elsey was the editor, and Nicola Crepaldi was the publisher. The research is editorially independent, and the views expressed are those of MIT Technology Review Insights.

This report draws on a poll of the MIT Technology Review Global Insights Panel, as well as a series of interviews with experts specializing in the semiconductor industry and chip design and manufacturing. Interviews occurred between December 2022 and February 2023.

Special thanks to the following individuals for their insights:

Uri Frank, Vice President of Engineering, Google

Sassine Ghazi, President and Chief Operating Officer, Synopsys

Patrick Moorhead, Founder, CEO, and Chief Analyst, Moor Insights & Strategy

François Piednoël, Distinguished Chief mSoC Architect, Mercedes-Benz Research & Development North America

Gerry Talbot, Corporate Fellow, AMD

Kevin Zhang, Senior Vice President of Business Development, TSMC

Poll methodology

In January and February 2023, MIT Technology Review Insights polled its global panel of executives about their understanding and adoption of multi-die systems. Of the 302 respondents, about 48% were C-suite executives or directors. They represented more than a dozen industries, from retailing to transportation, with the largest share of respondents from IT and telecommunications (22%) and manufacturing (13%). Respondents came from all geographies, representing North America (27%), Europe (22%), Asia-Pacific (24%), Latin America (13%), and the Middle East and Africa (14%). Their organizations varied in size: 28% of respondents' companies had total revenue of more than \$5 billion in 2022, 40% had revenue between \$100 million and \$5 billion, and 32% had revenue under \$100 million.

To say that semiconductor technology is part of the fabric of modern society is not an overstatement – it underpins everything from our cars to our phones to our home appliances. In 2021, the semiconductor industry shipped a record 1.15 trillion chips, and sales topped half a trillion dollars worldwide, while thousands of new chip designs entered the market.

A new semiconductor chip architecture, termed “multi-die system” or “chiplet-based design,” will be instrumental in meeting this decade’s burgeoning demand for processing power. Because this new approach will pose technical challenges throughout the semiconductor ecosystem – remaking how products are imagined, designed, and fabricated – opportunities for innovators across the value chain will emerge from this shift. Business leaders across industries who identify use cases for these advanced chips will benefit from their ability to power unique and customized customer experiences.

Few business leaders, however, are keeping pace with the latest developments in this arena. Multi-die technology is still an enigma to many executives. A recent poll by MIT Technology Review Insights asked business leaders about their awareness of this design strategy – and found that 62% of respondents are either uninterested, unaware, or only somewhat aware of this technology’s capabilities.

A few chip-reliant industries obviously need to keep a close eye on advancements in semiconductor tech: automotive companies, artificial intelligence firms, hyperscale data processing organizations, and smart device manufacturers, to name a few. But because advanced semiconductors are foundational to modern-day business operations, even executives whose functions don’t directly touch technology should care about chip design trends – including those that will define the sector’s next chapter.

Figure 1: Few business leaders are savvy about multi-die

How familiar are you with multi-die chip designs?

VERY AWARE AND COMMITTED TO DESIGN IN	5%
VERY AWARE AND TAKING ACTIVE STEPS	5%
AWARE AND EXPLORING	28%
SOMEWHAT AWARE	35%
NOT AT ALL AWARE, OR NOT INTERESTED	27%

Source: MIT Technology Review Insights poll, 2023.

Why semiconductors matter

While the global semiconductor shortage that began in 2020 had its proximate causes in natural disasters and geopolitics, its effects drew widespread attention to the fact that just about every industry relies on chips. And pandemic-related ripple effects aside, the silicon status quo has been in flux for some time. New technologies like artificial intelligence and machine learning (AI/ML), which require greater computing efficiency and performance, have strained traditional systems in recent years.




Figure 2: Companies are building more advanced smart products

How would you classify your organization's participation in the Internet of Things or connected devices ecosystem?



- Beginning:** Adding basic intelligence to previously non-smart devices.
- Intermediate:** Increasing the capability and features of existing smart products.
- Advanced:** Adding specific AI or machine learning abilities to products.
- Cutting-edge:** Adding advanced intelligence allowing multiple workloads to run in parallel on a single device.
- Non-participant:** We do not make this type of product or device.


Source: MIT Technology Review Insights poll, 2023.

With the rise of the Internet of Things (IoT), customers have also come to expect intelligence in everything from refrigerators to lightbulbs. Innovators are responding accordingly. Our poll found that nearly one-third (31%) of business executives plan to improve upon their companies' existing smart products, and almost another third (29%) intend to add AI/ML capabilities to their products soon. Only 9% of respondents said they were not producing IoT or connected devices.

This type of technology, however, necessitates robust edge computing and on-device processing, which requires greater and more efficient hardware performance. Complicating matters, the cloud data centers powering this compute shift are also voracious energy consumers. This is another area where traditional silicon is stagnating: sustainability. The

cost of producing superfluous silicon is not just bad for business – it has an environmental impact. And while there's an ongoing push toward net-zero carbon emissions within the semiconductor supply chain, the industry isn't yet on track to meet the emissions standards set forth in the UN 2016 Paris Agreement.

An industry shift toward multi-die design could be part of the solution to these challenges. Instead of a single monolithic chip ("system on chip"), multi-die designs consist of a collection of chips (chiplets or dies) linked in a sophisticated package ("systems of chips"), which can include stacking blocks in a 3D configuration for greater density. Multi-die system designs are capable of supporting the rollout of AI/ML at scale, and they can improve silicon yields, reducing waste during chip manufacturing.



“Chiplets enable smaller companies with smaller pocketbooks to use semiconductors for unique competitive advantage.”

Patrick Moorhead, Founder, CEO, and Chief Analyst, Moor Insights & Strategy

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.



“The broader sense of Moore’s law is about predictability for our industry continuing to drive improvement in terms of performance, power, functionality, and systems-level costs.”

Kevin Zhang, Senior Vice President of Business Development, TSMC

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 for companies that have based their business models on outdated expectations of what next-generation computing will cost. “If you have a business where you need to serve more clients every year, and you’re still assuming you can offset those client growth costs with better computing performance, you’re going to have a problem,” he says.

Our poll suggests that many companies may need to check their assumptions. The vast majority of respondents say they expect their organization’s demand for advanced chips to grow during the next two years. One-quarter (24%) say they anticipate demand will greatly increase, while another 46% are expecting it to increase somewhat.

This is not to say Moore’s law has become entirely irrelevant. “A lot of people tie Moore’s law to simple, linear transistor scaling. But that’s not exactly what Moore’s law states,” explains Kevin Zhang, senior vice president of business development at semiconductor manufacturing company TSMC. “The broader sense of Moore’s law is about predictability for our industry continuing to drive improvement in terms of performance, power, functionality, and systems-level costs.”

Figure 3: Demand for advanced chips set to grow

How much do you expect your organization’s demand for advanced silicon chips to change in the next two years?




Source: MIT Technology Review Insights poll, 2023.

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,’” explains Frank. “But they are willing to sell you a chiplet. In that way, you can get access to some IPs that you can’t get in a regular mode of work.”

Moorhead projects that we’ll see an “explosion” of companies creating IP content in the coming months and years, in a “democratization of IP.” “With chiplets,” he says, “innovation isn’t nearly as capital intensive or as resource dependent, and it doesn’t take five years to get to market. You’re shaving off at least a year, maybe two years, if you have a chiplet implementation because you’re not spending time hardening the monolithic design.”

And while multi-die isn’t a new concept, it is gaining traction across the ecosystem. According to Synopsys, if late 2022 was an inflection point for multi-die systems, 2023 looks to be the year when these architectures really take off. Zhang agrees that chiplet technology is on the cusp of a breakthrough. “In the future, the chiplet is going to become a prevailing way to bring more function and more capabilities together to achieve better system-level performance and power efficiency,” he says.

Chiplet opportunities

Multi-die systems have the potential to solve some of the semiconductor industry’s biggest pain points. For one, they’re a much better fit for advanced workloads like those required by AI/ML. “To be able to tackle



“To be able to tackle 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.”

Gerry Talbot, Corporate Fellow, AMD

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. The physical size of components, interconnect, layer transitions, the size of the package – all of these things limit you from being able to scale the bandwidth of the interface between chiplets,” explains Talbot.

There are also thermal limitations when chips are packed together, notes Frank. “All of a sudden, instead of touching the die and cooling it immediately, you have multiple layers of dies. Cooling that is more problematic,” he says.



TSMC is developing a family of chip stacking and packaging technologies, called TSMC 3DFabric, to address some common challenges. “At TSMC, we’ve developed a broad portfolio of different ways to integrate multiple chips together,” says Zhang. With 3D stacking, he says, “you can create a system with high-bandwidth interconnects between the chips while maintaining very small form factor.”

Another hurdle to making multi-die systems mainstream is a lack of industry standards. “If you want to use a chiplet in a machine, you need to define how you power it on, how power management works, how you check the chip,” says Frank. “There’s a lot of commonality needed that is not ready today as a general solution.”

Piednoël agrees that standards – like Universal Chiplet Interconnect Express (UCIe), which debuted in 2022 – will likely be a critical part of enabling chiplet innovation. “You have to think about how to leverage

standardization that already exists so that nobody has to reinvent the wheel every time they change or they improve the chiplet,” he says.

Finally, while manufacturing costs are often lower for chiplet-based designs, overall costs are less clear. Taking a large chip and dividing it into a series of smaller ones improves metrics like yield, but it’s also a complex process. More design, architecture, verification, and testing work goes into it – all of which translates into greater cost.

Moorhead sees this upfront cost as a nuanced component of overall ROI. “I liken it to the public cloud versus on-prem,” he says. “Everybody knows that you’re not going to save any money going to the public cloud. You do it because it’s faster. You can get in; you can get out. That’s the way I view chiplets. You have to buy into the notion that you’re going to do more designs for more people, and it’s going to have an ROI.”



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.

Standardization for the entire chiplet ecosystem is further off, however, particularly when it comes to supporting real-world use cases, such as AI-enabled automobiles. “There are many layers of standardization needed,” explains François Piednoël, distinguished chief mSoC architect at Mercedes-Benz. “Chiplet is one, software is another... if you want to uniformize power management of chiplets inside your circuits, your chiplets need to be able to negotiate with each other over how much power they’re using.”

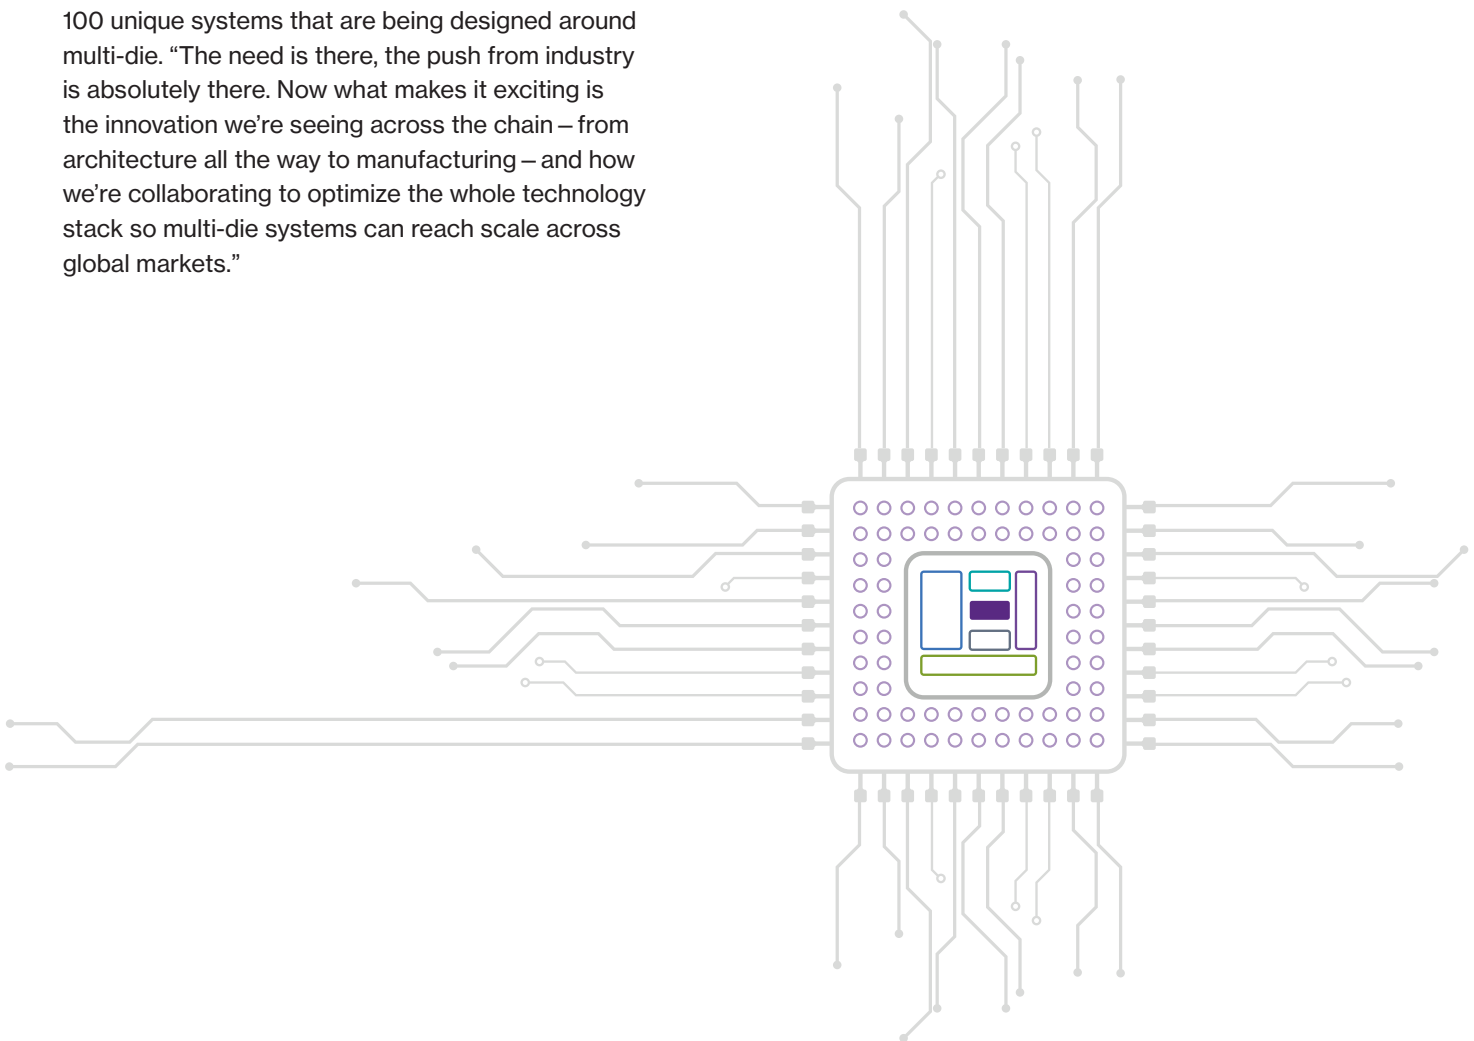
Embracing the future

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 profession – it is constantly evolving, and not just from the point of view of the actual chip itself, but all of the tools and infrastructure and methodologies we use,” says Talbot.

Ghazi, too, sees this moment of “discontinuity” as an opportunity for the technology industry. He’s optimistic about what’s in store for multi-die, including the cascading innovations it may prompt. “We’re not talking about the industry just trying to test the water – it’s well beyond that,” he says, adding that Synopsys is currently tracking more than 100 unique systems that are being designed around multi-die. “The need is there, the push from industry is absolutely there. Now what makes it exciting is the innovation we’re seeing across the chain – from architecture all the way to manufacturing – and how we’re collaborating to optimize the whole technology stack so multi-die systems can reach scale across global markets.”

“What makes it exciting is the innovation we’re seeing across the chain – from architecture all the way to manufacturing – and how we’re collaborating to optimize the whole technology stack.”

Sassine Ghazi, President and Chief Operating Officer, Synopsys



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, including articles, reports, infographics, videos, and podcasts. And through its growing MIT Technology Review [Global Insights Panel](#), Insights has unparalleled access to senior-level executives, innovators, and entrepreneurs worldwide for surveys and in-depth interviews.

About Synopsys

Synopsys, Inc. (Nasdaq: SNPS) is the Silicon to Software™ partner for innovative companies developing the electronic products and software applications we rely on every day. As an S&P 500 company, Synopsys has a long history of being a global leader in electronic design automation (EDA) and semiconductor IP and offers the industry's broadest portfolio of application security testing tools and services. Whether you're a system-on-chip (SoC) designer creating advanced semiconductors, or a software developer writing more secure, high-quality code, Synopsys has the solutions needed to deliver innovative products. Learn more at www.synopsys.com.


The Synopsys logo, featuring the word "SYNOPSYS" in a bold, purple, sans-serif font, with a registered trademark symbol (®) to the upper right of the "S".

Illustrations

Illustrations assembled by Tim Huxford with elements from Shutterstock.

While every effort has been taken to verify the accuracy of this information, MIT Technology Review Insights cannot accept any responsibility or liability for reliance by any person in this report or any of the information, opinions, or conclusions set out in this report.

© Copyright MIT Technology Review Insights, 2023. All rights reserved.



MIT Technology Review Insights

www.technologyreview.com

insights@technologyreview.com