

Feature | Computing | Hardware

24 Sep 2019 | 15:00 GMT

Goodbye, Motherboard. Hello, Silicon-Interconnect Fabric

Bare chiplets on a silicon-interconnect fabric will make computers smaller and more powerful.

By Puneet Gupta and Subramanian S. Iyer




Illustration: Harry Campbell

The need to make some hardware systems tinier and tinier and others bigger and bigger has been driving innovations in electronics for a long time. The former can be seen in the progression from laptops to smartphones to smart watches to hearables and other “invisible” electronics. The latter defines today’s commercial data centers—megawatt-devouring monsters that fill purpose-built warehouses around the world. Interestingly, the same technology is limiting progress in both arenas, though for different reasons.

The culprit, we contend, is the printed circuit board. And the solution is to get rid of it.

Our research shows that the printed circuit board could be replaced with the same material that makes up the chips that are attached to it, namely silicon. Such a move would lead to smaller, lighter-weight systems for wearables and other size-constrained gadgets, and also to incredibly powerful high-performance computers that would pack dozens of servers' worth of computing capability onto a dinner-plate-size wafer of silicon.

This all-silicon technology, which we call silicon-interconnect fabric, allows bare chips to be connected directly to wiring on a separate piece of silicon. Unlike connections on a printed circuit board, the wiring between chips on our fabric is just as small as wiring within a chip. Many more chip-to-chip connections are thus possible, and those connections are able to transmit data faster while using less energy.

Silicon-interconnect fabric, or Si-IF, offers an added bonus. It's an excellent path toward the dissolution of the f(requently) big, complicated, and difficult-to-manufacture systems-on-chips that currently run everything from smartphones to supercomputers. In place of SoCs, system designers could use a conglomeration of smaller, simpler-to-design, and easier-to-manufacture chiplets tightly interconnected on an Si-IF. This chiplet revolution is already well under way, with AMD, Intel, Nvidia, and others offering chiplets assembled inside of advanced packages. Silicon-interconnect fabric expands that vision, breaking the system out of the package to include the entire computer.

To understand the value of eliminating the printed circuit board, consider what happens with a typical SoC. Thanks to Moore's Law, a 1-square-centimeter piece of silicon can pack pretty much everything needed to drive a smartphone. Unfortunately, for a variety of reasons that mostly begin and end with the printed circuit board, this sliver of silicon is then put inside a (usually) plastic package that can be as much as 20 times as large as the chip itself.

The size difference between chip and package creates at least two problems. First, the volume and weight of the packaged chip are much greater than those of the original piece of silicon. Obviously, that's a problem for all things that need to be small, thin, and light. Second, if the final hardware requires multiple chips that talk to one another (and most systems do), then the distance that signals need to travel increases by more than a factor of 10. That distance is a speed and energy bottleneck, especially if the chips exchange a lot of data. This choke point is perhaps the biggest problem for data-intensive applications such as graphics, machine learning, and search. To make matters worse, packaged chips are difficult to keep cool. Indeed, heat removal has been a limiting factor in computer systems for decades.

If these packages are such a problem, why not just remove them? Because of the printed circuit board.

The purpose of the printed circuit board is, of course, to connect chips, passive components, and other devices into a working system. But it's not an ideal technology. PCBs are difficult to make perfectly flat and are prone to warpage. Chip packages usually connect to the PCB via a set of solder bumps, which are melted and resolidified during the manufacturing process. The limitations of solder technology combined with surface warpage mean these solder bumps can be no less than 0.5 millimeters apart. In other words, you can pack no more than 400 connections per square centimeter of chip area. For many applications, that's far too few connections to deliver power to the chip and get signals in and out. For example, the small area taken up by one of the Intel Atom processor's dies has only enough room for a hundred 0.5-mm connections, falling short of what it needs by 300. Designers use the chip package to make the connection-per-unit-area math work. The package takes tiny input/output connections on the silicon chip—ranging from 1 to 50 micrometers wide—and fans them out to the PCB's 500- μm scale.

Recently, the semiconductor industry has tried to limit the problems of printed circuit boards by developing advanced packaging, such as silicon interposer technology. An interposer is a thin layer of silicon on which a small number of bare silicon chips are mounted and linked to each other with a larger number of connections than could be made between two packaged chips. But the interposer and its chips must still be packaged and mounted on a PCB, so this arrangement adds complexity without solving any of the other issues. Moreover, interposers are necessarily thin, fragile, and limited in size, which means it is difficult to construct large systems on them.

We believe that a better solution is to get rid of packages and PCBs altogether and instead bond the chips onto a relatively thick (500- μm to 1-mm) silicon wafer. Processors, memory dies, analog and RF chiplets, voltage-regulator modules, and even passive components such as inductors and capacitors can be bonded directly to the silicon. Compared with the usual PCB material—a fiberglass and epoxy composite called FR-4—a silicon wafer is rigid and can be polished to near perfect flatness, so warping is no longer an issue. What's more, because the chips and the silicon substrate expand and contract at the same rate as they heat and cool, you no longer need a large, flexible link like a solder bump between the chip and the substrate.

Solder bumps can be replaced with micrometer-scale copper pillars built onto the silicon substrate. Using thermal compression—which basically is precisely applied heat and force—the chip's copper I/O ports can then be directly bonded to the pillars. Careful optimization of the thermal-compression bonding can produce copper-to-copper bonds that are far more reliable than soldered bonds, with fewer materials involved.

Eliminating the PCB and its weaknesses means the chip's I/O ports can be spaced as little as 10 μm apart instead of 500 μm . We can therefore pack 2,500 times as many I/O ports on the silicon die without needing the package as a space transformer.

Even better, we can leverage standard semiconductor manufacturing processes to make multiple layers of wiring on the Si-IF. These traces can be much finer than those on a printed circuit board. They can be less than 2 μm apart, compared with a PCB's 500 μm . The technology can even achieve chip-to-chip spacing of less than 100 μm , compared with 1 mm or more using a PCB. The result is that an Si-IF system saves space and power and cuts down on the time it takes signals to reach their destinations.

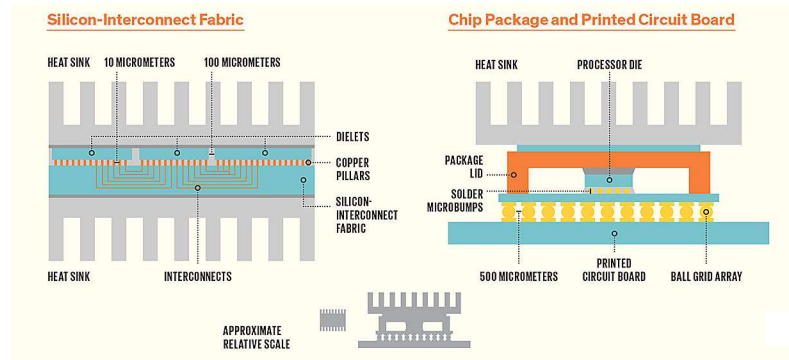
Furthermore, unlike PCB and chip-package materials, silicon is a reasonably good conductor of heat. Heat sinks can be mounted on both sides of the Si-IF to extract more heat—our estimates suggest up to 70 percent more. Removing more heat lets processors run faster.

Although silicon has very good tensile strength and stiffness, it is somewhat brittle. Fortunately, the semiconductor industry has developed methods over the decades for handling large silicon wafers without breaking them. And when Si-IF-based systems are properly anchored and processed, we expect them to meet or exceed most reliability tests, including resistance to shock, thermal cycling, and environmental stresses.

There's no getting around the fact that the material cost of crystalline silicon is higher than that of FR-4. Although there are many factors that contribute to cost, the cost per square millimeter of an 8-layer PCB can be about one-tenth that of a 4-layer Si-IF wafer. However, our analysis indicates that when you remove the cost of packaging and complex circuit-board construction and factor in the space savings of Si-IF, the difference in cost is negligible, and in many cases Si-IF comes out ahead.

Silicon-Interconnect Fabric vs. Chip Package and Printed Circuit Board

Diagram: Brandon Palacio



A silicon-interconnect fabric links chips with denser, higher-bandwidth connections; it allows them to pack together more tightly and removes heat better.

Let's look at a few examples of how Si-IF integration can benefit a computer system. In one study of server designs, we found that using packageless processors based on Si-IF can double the performance of conventional processors because of the higher connectivity and better heat dissipation. Even better, the size of the silicon “circuit board” (for want of a better term) can be reduced from 1,000 cm² to 400 cm². Shrinking the system that much has real implications for data-center real estate and the amount of cooling infrastructure needed. At the other extreme, we looked at a small Internet of Things system based on an Arm microcontroller. Using Si-IF here not only shrinks the size of the board by 70 percent but also reduces its weight from 20 grams to 8 grams.

Apart from shrinking existing systems and boosting their performance, Si-IF should let system designers create computers that would otherwise be impossible, or at least extremely impractical.

A typical high-performance server contains two to four processors on a PCB. But some high-performance computing applications need multiple servers. Communication latency and bandwidth bottlenecks arise when data needs to move across different processors and PCBs. But what if all the processors were on the same wafer of silicon? These processors could be integrated nearly as tightly as if the whole system were one big processor.



This concept was first proposed by Gene Amdahl at his company Trilogy Systems. Trilogy failed because manufacturing processes couldn't yield enough working systems. There is always the chance of a defect when you're making a chip, and the likelihood of a defect increases exponentially with the chip's area. If your chip is the size of a dinner plate, you're almost guaranteed to have a system-killing flaw somewhere on it.

But with silicon-interconnect fabric, you can start with chiplets, which we already know can be manufactured without flaws, and then link them to form a single system. A group of us at the University of California, Los Angeles, and the University of Illinois at Urbana-Champaign architected such a wafer-scale system comprising 40 GPUs. In simulations, it sped calculations more than fivefold *and* cut energy consumption by 80 percent when compared with an equivalently sized 40-GPU system built using state-of-the-art multichip packages and printed circuit boards.

These are compelling results, but the task wasn't easy. We had to take a number of constraints into account, including how much heat could be removed from the wafer, how the GPUs could most quickly communicate with one another, and how to deliver power across the entire wafer.

Photos: UCLA Chips

Chiplets Ahoy! Test dielets, or chiplets, link up on a silicon-interconnect fabric built on a 100-millimeter-wide wafer. Unlike chips on a printed circuit board, the dielets can be placed a mere 100 micrometers apart, speeding signals and reducing energy consumption.



Power turned out to be a major constraint. At a chip's standard 1-volt supply, the wafer's narrow wiring would consume a full 2 kilowatts. Instead, we chose to up the supply voltage to 12 V, reducing the amount of current needed and therefore the power consumed. That solution required spreading voltage regulators and signal-conditioning capacitors all around the wafer, taking up space that might have gone to more GPU modules. Encouraged by the early results, we are now building a prototype wafer-scale computing system, which we hope to complete by the end of 2020.

Silicon-interconnect fabric could play a role in an important trend in the computer industry: the dissolution of the system-on-chip (SoC) into integrated collections of dielets, or chiplets. (We prefer the term *dielets* to *chiplets* because it emphasizes the nature of a bare silicon die, its small size, and the possibility that it might not be fully functional without other dielets on the Si-IF.) Over the past two decades, a push toward better performance and cost reduction compelled designers to replace whole sets of chips with ever larger integrated SoCs. Despite their benefits (especially for high-volume systems), SoCs have plenty of downsides.

For one, an SoC is a single large chip, and as already mentioned, ensuring good yield for a large chip is very difficult, especially when state-of-the-art semiconductor manufacturing processes are involved. (Recall that chip yield drops roughly exponentially as the chip area grows.) Another drawback of SoCs is their high one-time design and manufacturing costs, such as the US \$2 million or more for the photolithography masks, which can make SoCs basically unaffordable for most designs. What's more, any change in the design or upgrade of the manufacturing process, even a small one, requires significant redesign of the *entire* SoC. Finally, the SoC approach tries to force-fit all of the subsystem designs into a single manufacturing process, even if some of those subsystems would perform better if made using a different process. As a result, nothing within the SoC achieves its peak performance or efficiency.


The packageless Si-IF integration approach avoids all of these problems while retaining the SoC's small size and performance benefits and providing design and cost benefits, too. It breaks up the SoC into its component systems and re-creates it as a system-on-wafer or system-on-Si-IF (SoIF).

Such a system is composed of independently fabricated *small* dielets, which are connected on the Si-IF. The minimum separation between the dielets (a few tens of micrometers) is comparable to that between two functional blocks within an SoC. The wiring on the Si-IF is the same as that used within the upper levels of an SoC and therefore the interconnect density is comparable as well.


Source: Puneet Gupta & Subramanian S. Iyer

Bandwidth, Latency, Energy: Waferscale integration using silicon interconnect fabric boosts bandwidth, reduces latency, and uses less energy than a conventional system on a printed circuit board.


Bandwidth (gigabytes/second)



Energy (femtojoules/byte)



Latency (picoseconds)



The advantages of the SoIF approach over SoCs stem from the size of the dielet. Small dielets are less expensive to make than a large SoC because, as we mentioned before, you get a higher yield of working chips when the chips are smaller. The only thing that's large about the SoIF is the silicon substrate itself. The substrate is unlikely to have a yield issue because it's made up of just a few easy-to-fabricate layers. Most yield loss in chipmaking comes from defects in the transistor layers or in the ultradense lower metal layers, and a silicon-interconnect fabric has neither.

Beyond that, an SoIF would have all the advantages that industry is looking for by moving to chiplets. For example, upgrading an SoIF to a new manufacturing node should be cheaper and easier. Each dielet can have its own manufacturing technology, and only the dielets that are worth upgrading would need to be changed. Those dielets that won't get much benefit from a new node's smaller transistors won't need a redesign. This heterogeneous integration allows you to build a completely new class of systems that mix and match dielets of various generations and of technologies that aren't usually compatible with CMOS. For example, our group recently demonstrated the attachment of an indium phosphide die to an SoIF for potential use in high-frequency circuits.

Because the dielets would be fabricated and tested before being connected to the SoIF, they could be used in different systems, amortizing their cost significantly. As a result, the overall cost to design and manufacture an SoIF can be as much as 70 percent less than for an SoC, by our estimate. This is especially true for large, low-volume systems like those for the aerospace and defense industries, where the demand is for only a few hundred to a few thousand units. Custom systems are also easier to make as SoIFs, because both design costs and time shrink.

We think the effect on system cost and diversity has the potential to usher in a new era of innovation where novel hardware is affordable and accessible to a much larger community of designers, startups, and universities.

Over the last few years, we've made significant progress on Si-IF integration technology, but a lot remains to be done. First and foremost is the demonstration of a commercially viable, high-yield Si-IF manufacturing process. Patterning wafer-scale Si-IF may require innovations in "maskless" lithography. Most lithography systems used today can make patterns only about 33 by 24 mm in size. Ultimately, we'll need something that can cast a pattern onto a 300-mm-diameter wafer.

We'll also need mechanisms to test bare dielets as well as unpopulated Si-IFs. The industry is already making steady progress in bare die testing as chipmakers begin to move toward chiplets in advanced packages and 3D integration.

Next, we'll need new heat sinks or other thermal-dissipation strategies that take advantage of silicon's good thermal conductivity. With our colleagues at UCLA, we have been developing an integrated wafer-scale cooling and power-delivery solution called PowerTherm.

In addition, the chassis, mounts, connectors, and cabling for silicon wafers need to be engineered to enable complete systems.

We'll also need to make several changes to design methodology to deliver on the promise of SoIFs. Si-IF is a passive substrate—it's just conductors, with no switches—and therefore the interdielet connections need to be short. For longer connections that might have to link distant dielets on a wafer-scale system, we'll need intermediate dielets to help carry data further. Design algorithms that do layout and pin assignments will need an overhaul in order to take advantage of this style of integration. And we'll need to develop new ways of exploring different system architectures that leverage the heterogeneity and upgradability of SoIFs.

We also need to consider system reliability. If a dielet is found to be faulty after bonding or fails during operation, it will be very difficult to replace. Therefore, SoIFs, especially large ones, need to have fault tolerance built in. Fault tolerance could be implemented at the network level or at the dielet level. At the network level, interdielet routing will need to be able to bypass faulty dielets. At the dielet level, we can consider physical redundancy tricks like using multiple copper pillars for each I/O port.

Of course, the benefit of dielet assembly depends heavily on having useful dielets to integrate into new systems. At this stage, the industry is still figuring out which dielets to make. You can't simply make a dielet for every subsystem of an SoC, because some of the individual dielets would be too tiny to handle. One promising approach is to use statistical mining of existing SoC and PCB designs to identify which functions "like" to be physically close to each other. If these functions involve the same manufacturing technologies and follow similar upgrade cycles as well, then they should remain integrated on the same dielet.

This might seem like a long list of issues to solve, but researchers are already dealing with some of them through the Defense Advanced Research Projects Agency's [Common Heterogeneous Integration and IP Reuse Strategies \(CHIPS\)](#) program as well as through industry consortia. And if we can solve these problems, it will go a long way toward continuing the smaller, faster, and cheaper legacy of Moore's Law.

About the Authors

[Puneet Gupta](#) and [Subramanian S. Iyer](#) are both members of the electrical engineering department at the University of California at Los Angeles. Gupta is professor and vice chair of computer engineering, and Iyer is Distinguished Professor and the Charles P. Reames Endowed Chair.