



Chipmakers Getting Serious About Integrated Photonics – July 12, 2021

This technology could enable the next wave of Moore's Law. What's needed to make that happen?

JULY 12TH, 2021 - BY: [BRIAN BAILEY](#)

Integrating photonics into semiconductors is gaining traction, particularly in heterogeneous multi-die packages, as chipmakers search for new ways to overcome power limitations and deal with increasing volumes of data.

Power has been a growing concern since the end of Dennard scaling, which happened somewhere around the 90nm node. There are more transistors per mm², and the wires are thinner, which increases resistance and capacitance and generates heat. Alongside of that, the amount of data that needs to be processed and moved continues to grow, so various processing elements, memories and I/Os are being utilized more intensively than in the past. This makes it harder to move data, to deliver enough power where it is needed, and to dissipate the heat.

Photonics offers a potential solution. In fact, it could provide a step-function improvement, opening the door to new applications that today are limited by fixed power budgets and copper interconnects.

“The cost of communications using copper is starting to become prohibitive,” says James Pond, director of product management at [Ansys](#). “The challenge with electrical interconnect is that as you go up in performance, or up in reach, your power costs go up. The industry is getting to the point where electrical interconnects will consume your power budget in its entirety, and that’s going to happen within a few years.”

Until recently, this was prohibited by cost. “It gets to be really interesting when the cost of photonics crosses below the cost of copper,” says Gilles Lamant, distinguished engineer at [Cadence](#).

Are we close to that point? To answer that question, it is necessary to look at other developments within the industry.

The reticle limit defines the maximum size that a lithography machine can etch. For 193nm immersion steppers, which are used to produce a large proportion of chips today, that limit is 33 x 26 — a little over 800mm². At the same time, Moore's Law is slowing in terms of cost effectiveness for many companies or design types. That means the number of transistors that can be economically fabricated on a single piece of silicon is reaching its limit.

The next phase of Moore's Law requires that devices start becoming assemblies of multiple dies, and the industry has been investing in this technology. Today, proprietary versions of it are found in most high-end CPUs and GPUs, FPGAs, and AI processors. As the industry increases adoption, problems are being solved, costs are coming down, and that is enabling larger potential markets.

Heterogeneous integration of multiple die points to the notion of chiplets. These are pre-designed and fabricated pieces of functionality that can be assembled in a package in much the same way as chips are assembled on a PCB today. There are several challenges that have to be overcome in the industry before this becomes a broadly adopted reality. Some of them are technical, others financial or legal, and some of them relate to the creation of standards that can ensure a large enough market for these standardized pieces.

Chiplet technology also will help heterogeneous integration of things other than electronics (see figure 1). We have already seen video sensor devices that are a combination of an optical layer and an electronic layer being fused together using wafer bonding. The volumes are enormous, and the technology well proven since 2016.

LiDAR

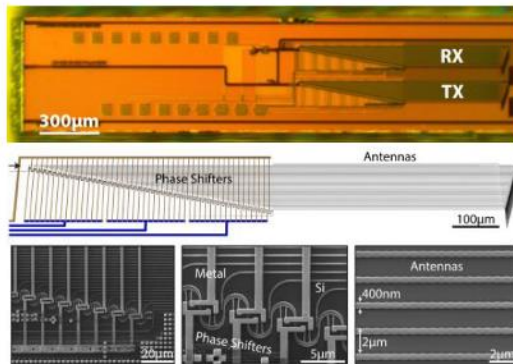




Image: Christopher V. Poulton <https://spectrum.ieee.org/tech-talk/semiconductors/optoelectronics/mit-lidar-on-a-chip>

Biosensing



Genalyte Maverick Detection System <https://www.genalyte.com/>

Artificial Intelligence



Y. Shen et al. Deep learning with coherent nanophotonic circuits, Nature Photonics, <https://doi.org/10.1038/nphoton.2017.93>

Quantum Computing




Image: Xiaogang Qiang <https://spectrum.ieee.org/tech-talk/computing/hardware/building-quantum-computers-with-photons>

Fig. 1: Emerging applications for heterogeneous integration. Source: Ansys

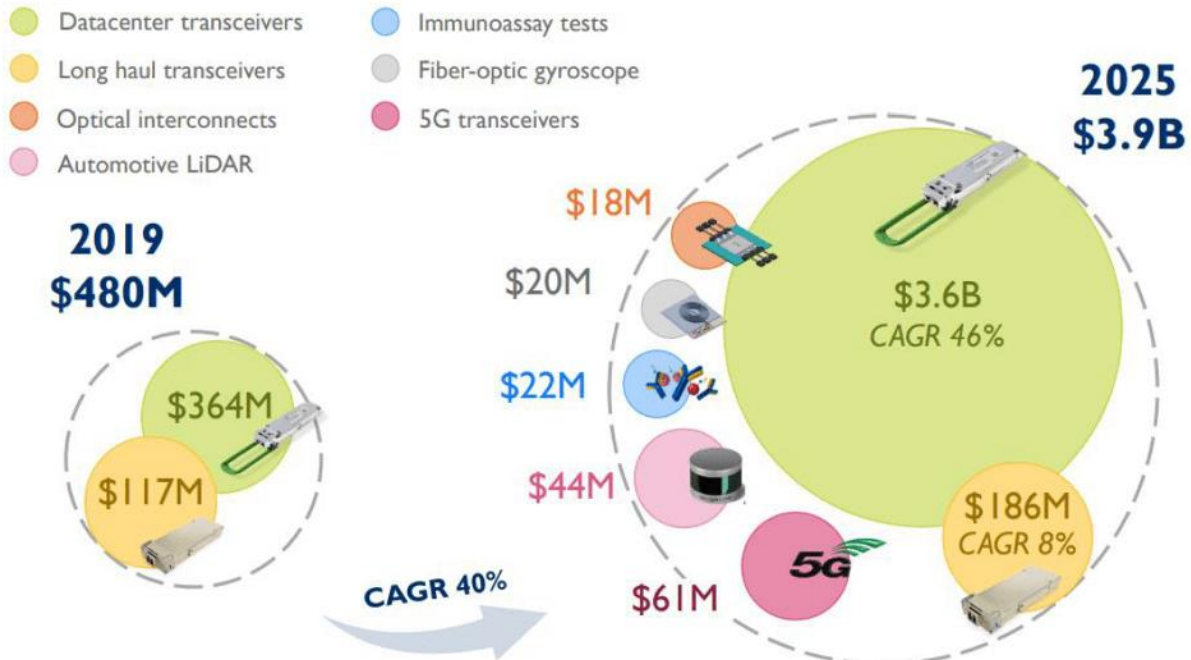
One technique used to integrate chiplets is to utilize an interposer. This is the equivalent of the PCB in this type of packaged system. The interposer serves as the interconnect between the various chiplets. This approach has been successfully used with high-bandwidth memory (HBM). The problem with an interposer is that it is a separate “chip” that has to be fabricated, which adds to total cost. Other companies are looking at stacking chiplets directly onto the primary die or using small pieces of interposer that bridge between various die.

We are seeing the photonics market start to make use of heterogeneous integration and it could reimagine many aspects of the ways die, chips and systems communicate.

Growing markets

The cost of photonic components has been high because there was a lack of volume products. “The promise of low-cost, high-speed transceivers using integrated photonics — and especially [silicon photonics](#) — has been realized,” says Ansys’ Pond. “We have seen millions of units shipped, and they are being used in data centers. That’s expected to continue to grow at an annual growth rate above 40% until 2025 at least (see figure 2). Co-packaged optics is the next frontier. It is required to solve the data

bandwidth problem that we have. Optical interconnects and co-packaged optics is the only way to do that.”



Yole Développement, 2020 http://www.yole.fr/Si_Photonics_Datacom_Sensing.aspx

Fig 2. Growth of transceiver market. Source: Ansys.

New markets will have different sets of requirements. “Telecom applications have been driven by a full width at half maximum (FWHM) approach so they could have a lot of different wavelengths very close to each other,” says Martin Eibelhuber, deputy head of business development for EV Group. “This typically led them to integrate edge-emitting lasers at different wavelengths closely together. New applications, in particular, lidar based on photonic integration, are becoming more popular, and that will be a bigger market in the near future.”

Lidar has different requirements. “For lidar, it’s not so important to have different wavelengths, but the beam shape is much more important,” says Eibelhuber. “What people are aiming for when you go to optical phased arrays is more about steering and output power. You want to modulate the waves that go outside. They also have different constraints in terms of size and how to build them. For example, datacoms is less sensitive to physical dimensions or weight, compared to something that goes into an automobile. Lidar in automotive would also be a mission-critical system.”

The biggest part of the step function comes from rethinking system-level issues rather than just a technology replacement. “In high-performance computing between two SerDes, you have retiming,” says Cadence’s Lamant. “Because there is variation on the copper line, the device needs to have the ability to self-adapt to each condition. That has been included in most high performance SerDes, and in all of the high-performance communication standards. While they don’t ask for it specifically, the requirements on the quality of the signals makes it absolutely impossible to work without those devices.”

“With a photonics solution, you can completely get rid of that part of the circuit,” adds Lamant. “This is an excellent illustration of what is to be won by integrated photonics. We are not talking about a small 10% gain. We are talking about 40% gain. The retimer is an analog circuit, and one that is very complicated to

get right. With photonics, you don't have those impedance matching problems and losses in the fiber. This is the type of change that integrated photonics can make happen, and they are step functions."

Getting the light

It is unlikely that lasers will ever be built on silicon, just because physics is not conducive to it. "You have to decide whether the light source should be on chip or out of chip," says Lamant. "To have photonics, you need to have light, and then you need to do something with it. Some companies are building and growing, or welding, wafer bonding, a laser right on top of a silicon substrate. Building it right on top of the silicon substrate is the holy grail type of thing because it would be fully within a single fab. Other companies are doing it with chipllets that are wafer bonded on top of their big silicon wafer."

Even on-chip there are different approaches. "You can consider wafer level integration schemes for heterogenous integration of different parts," says EV Group's Eibelhuber. "If we think about integrating the laser chipllets, compound semiconductors on silicon, things like that, that is basically happening already today."

Size mismatch

Another reason why photonics has been separated from electronics is a mismatch in the geometries used for each piece. "Silicon photonic devices are always going to be large and can typically be made at foundries that don't require the same advanced node technologies, just because of the size of the device," says Pond. "The sizes of the photonic devices are limited by the physics, so you can't just keep shrinking them down as people have been doing with the electrical ICs. There is a quite a bit of control electronics that is required for silicon photonics to keep everything operating properly, to maintain the thermal heaters and feedback loops to keep everything going. But when it comes to advanced nodes, it's just two different technologies and it probably makes sense to keep them separate."

But you can build useful amounts of electronics in older nodes. "There is at least one foundry offering an integrated monolithic solution today," says Lamant. "The transistors are the last generation of planar transistor, and that's because of all the compatibility requirements with the sizes that you need for photonics. There are challenges. You need a balance because they cannot really go much smaller than 45nm because that would create too many challenges for the photonics."

The opportunity

In a chipllet technology, an interposer may be used as the substrate onto which the chipllets are mounted. This is also at a larger geometry than is used for intensive electronic functionality. "The chipllet and 3D IC model is really good and you can have an integrated photonic interposer on which you build up that 3D IC," says Pond. "That can serve as the communication layer, both between those chipllets as well as with the outside world. Plus, when you start thinking about moving to every 3D IC having an integrated photonic interposer layer for communications, you're talking volumes that are orders of magnitude larger than data center transceivers. This is an important stage for integrated photonics that moves from, what could be compared to electronics is definitely sort of a niche size market to something that's more substantial."

While chipllets may be enabling photonics, photonics may also impact chipllets. "Today, the industry thinks about having electronic connections between the chipllets," says Lamant. "This allows two chipllets to have a vertical connection between themselves. But instead of having a bump that carries electrical signals, why not use optics? They call that optical bumping. The challenges here are alignment, but those are challenges for chipllets in general."

This now takes what was an expensive piece of the system — the interposer — and uses it to create a step change in the communications between the heterogeneous pieces. "The photonic layer should be the interposer and provide the communication layer, both between chipllets and to the outside world," says Pond.

Integration issues

One problem with all 3D IC systems is heat, and photonics adds an extra dimension to this. “How the whole system behaves in a critical environment is important and thermal management is a key part of that,” says Eibelhuber. “They need active measures for temperature control. Temperature management is a general problem for the whole system in the 3D package world, including for the electronics.”

Lamant agrees. “The basic engines for analyzing thermal are known. This is not just for photonics, you see the same complexity appearing on the chiplets. Chiplets are wonderful, but once you start putting them on a 2.5D or a 3D IC, they’re very close to each other, and you have to do analysis across the boundaries of each of them. It is not specific to photonics.”

Photonics does have some unique problems, though. “Electronics generates a lot of heat and they are going to be modifying the temperature of that integrated photonic interposer,” says Pond. “Some photonic devices are incredibly sensitive to temperature changes. The modeling challenge is to be able to understand what is happening to the temperature of the entire 3D IC and, from a simulation perspective, to be able to calculate the operating temperature in different configurations. What is the local temperature of every individual photonic component? Then we can study the impact of that on the performance of the photonic circuit.”

A typical chip can go through a wide range of temperatures based on operating environment or workload. “Even a fraction of degree change in a photonic demodulator can have a big impact on the performance,” adds Pond. “It can drive it completely out of its operating points. A one-degree change could destroy the performance of the device. You might ask how could this ever function? That is the role of thermal heaters and control loops that are maintaining a correct operating point for the devices. It is not the absolute temperature that matters, it is the differential temperature between different arms of an interferometer. As long as you have the control electronics and heaters maintaining the device at the right operating point, then you can deal with larger global temperature ranges. The challenge is that if you’re having to overcompensate, you can start to burn up a lot of power in heaters just to keep everything operating at the right temperature.”

Standards

The biggest challenge is not technical but associated with industry cooperation and standards. “There is a significant drive within the industry for standards,” says Eibelhuber. “They want to collaborate, and to have standards that would enable them to access similar processes at reasonable costs and not have to develop everything on their own. There are some co-packaging solutions on the way that try to keep as much on the wafer-level as possible. There are good solutions on the chip level, and you can basically take the best of both worlds. However, not everything that works in electronics can be copied for the photonics industry, and that’s the gap that has to be closed.”

Reliable transfer of data between dies requires more than just a PHY. “Instead of the very low-level interface standards, higher-level standards have to be implemented,” says Andy Heinig, group leader for advanced system integration and department head for efficient electronics at [Fraunhofer IIS’ Engineering of Adaptive Systems Division](#). “Such higher-level protocols are likely to be application-oriented. They will be different between an analog-digital application, such as might be found in an optical front end, or digital accelerators such as would be found in data centers for AI application.”

Standards enable innovation. “The reason why pluggables have been so successful is because of standards,” says Pond. “They have a standardized form factor and communication standards. That meant a whole lot of innovative companies could go off and design things like integrated photonics transceivers, and push the performance and speeds up. Then someone building a data center could purchase the pluggable transceivers and, as long as they meet the standards, combine them from different vendors. The challenge now, when talking about bringing optics onto a chip or into a co-packaged system, is that there are no real standards about how to connect those fibers and exactly what the communication standards are going to be. Once that is achieved, we’ll see a similar flurry of innovation.”

This is no different than the standards needed for electronics where there has to be some common agreement on what communication standards are going to be used.

Conclusion

The need for chiplets has caused the industry to address certain problems. They were also problems for integrated optics, but that market was too small to make a difference. As we near solutions for 3D ICs, some of the barriers for integrated optics have been removed.

But that is only the first part of the story. With those problems solved, optics can change the way we think about communications within the package and move from an electrical interconnect interposer to an optical interposer. What was seen as an expensive overhead now becomes a faster and lower-power solution that was not been possible before. When those two technologies come together, new system-level solutions become possible. Power will drop considerably, and total product cost may decline eventually.

<https://semiengineering.com/chipmakers-getting-serious-about-integrated-photonics/>