Waiting for Chiplet Standards

The need and desire for chiplets is increasing, but for most companies that shift will happen slowly until proven standards are in place.

Interoperability and compatibility depend on many layers and segments of the supply chain coming to agreement. Unfortunately, fragmented industry requirements may lead to a plethora of solutions.

Standards always have enabled increasing specialization. In the early days of the chip industry, a company had to design, implement, and fabricate everything. For most companies, fabrication became separated from design by well-defined interfaces and models, such as PDKs, BSIM models, and libraries.

The emergence of the IP industry enabled companies to concentrate on the design of the system and the pieces that provided them with differentiation, but they still had to do most of the implementation themselves. Standards for interconnect and the models that were transferred between IP provider and consumer made this possible.

Today, we are on the cusp of another level of specialization, where a company will only design the system — and design and implement the pieces of the system that provide differentiation — without having to worry about the implementation or fabrication of commodity parts of the design. These would be available in the form of chiplets, which are fully implemented and fabricated pieces that can be assembled with custom silicon to form a system. To get there, the industry needs some new standards.

The pioneering work has been done by large systems companies, which own both the system and the chiplets. (See figure 1.) This enables them to make larger, or more modular offerings, and along the way to iron out many of the kinks. They have developed proprietary means to make these systems.

Unsurprisingly, there is significant variation in those solutions. "The industry is divided up into a range of offerings, simply because there was a need by ASIC companies," says Ketan Mehta, senior director of SoC IP product marketing at OpenFive. "They are building custom silicon, and they want solutions right away. They don't want to wait for the standards to be developed and evolve. So proprietary implementations are being developed and proven in all these companies."

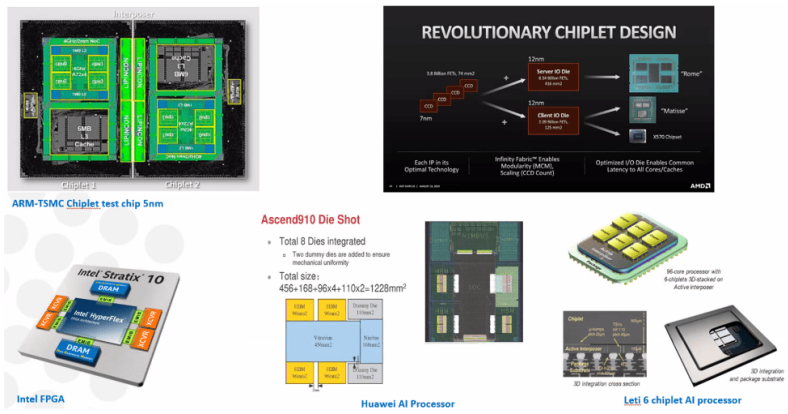


Fig. 1: Some early pioneers in 2.5D integration. Source: OpenFive

The first partially open systems, where IP and system are developed by different companies, has been seen with high bandwidth memory (HBM). Here, the DRAM is

provided by one company, utilized in a system designed by another, and packaged by a third. This provides a solution for a restricted application, and there are other fabrication/packaging technologies that also focus on bringing memory closer to logic.

A bigger gain will come when logic can be connected to other logic using chiplets that are available off-the-shelf. This will break the restriction where everything has to be fabricated in the same technology node. While reliability issues remain, due to issues like differential expansion and warpage, it is likely these will get resolved over time. Also, a viable business model for this has not yet been developed.

The demand for chiplets is coming from several directions. "This is a real opportunity to help alleviate a lot of the challenges that companies see in this space," says Rob Mains, executive director for the CHIPS Alliance. "It requires a standardized interface. It needs a standardized PHY. It has to be instantiated for a particular chip process technology, or packaging technologies. And then it requires an EDA ecosystem that goes along with that. DARPA's vision is right, and it is a matter of getting the education level set with design teams globally. This will bring an understanding of the benefits and provide a level of assurance that it will yield effective results."

That's not where the industry is today. It's closer to computer scientist Andrew Tanenbaum's observation, "The good thing about standards is there are so many to choose from." However, that is beginning to change as an increasing number of players attempt to consolidate the field and deal with issues that bind fabrication and packaging with electrical standards. Protocols are required to ensure data integrity across a system. Beyond that, a whole host of other issues need to be addressed, such as physical layout, power delivery networks, test, debug, monitoring, and many others. Some of these are being investigated today.

Previous articles have looked at the overall push toward chiplets, and the impacts on the development flow. The focus of this article is the evolving standards that may enable a market to develop, although it is by no means a complete account of where everyone stands, or the relationships between them.

The physical layer

Moving from separate chips that are packaged and placed on a board, to a package that integrates multiple dies, dramatically changes interconnects. "A traditional ASIC has large I/O drivers necessary to drive signals through the package, board and external interfaces," says Tony Mastroianni, advanced packaging solutions director for Siemens EDA. "This could range from tens of millimeters to several meters. 2.5D die-to-die interfaces deploy smaller I/O drivers that are only required to drive horizontal connections to adjacent die through the interposer, which may be on the order of tens to hundreds of microns. 3D die-to-die interfaces deploy even smaller I/O drivers, which are only required to drive vertical connections directly to the die stacked above or below. Those may be on the order of a few to hundreds of nanometers. The reduced drive strength and shorter trace lengths inherent in the 2.5 and 3D approaches enables dramatic reductions of power and increased I/O bandwidth, which offers orders of magnitude of improved energy efficiency (pJ/bit)."

There are several options available at this point. "One method of chiplet integration avoids the use of fine-geometry interconnect altogether," says Brian Holden, vice president of standards for Kandou. "With this method, the interconnect between

the chiplets is simply the organic package substrate. This avoids a complex manufacturing processes, and the extra cost and yield loss associated with silicon interposers. Low-power ultra-short reach (USR) SerDes are used to enable high-speed interconnect between the chiplets."

The physical interface suggests a solution. "When you disaggregate your die into multiple dies, you can put it either on a substrate or put it on an interposer," says OpenFive's Mehta. "That creates a big distinction. With an interposer you can do thousands of signals, whereas with the substrate you can only do a few hundred at most. If the customer is implementing a large die that also has HBM, for example, they have no choice but to implement it on an interposer. That leads you toward a parallel interface, because the interposer will accommodate thousands of signals."

Intel has developed its own chiplet strategy around its Embedded Multi-die Interconnect Bridge (EMIB). Instead of using a large silicon interposer typically found in 2.5D approaches, EMIB uses a very small bridge with multiple routing layers. This bridge is embedded as part of their substrate fabrication process.

Parallel or serial?

The debate between parallel and serial will probably continue for a long time, and it is unlikely that there will ever be a single solution. Each evolving standard is a tradeoff between many different factors.

"What customers really care about are lowest possible latency, lowest possible power, bandwidth for beachfront, performance in terms of reach, and then cost, which is basically the yield," explains Manmeet Walia, senior product manager for high-speed SerDes at Synopsys.

	Monolithic SoC	Chiplet-Based
Cost	High	
Effort	High	
Risk	High	
Power		Acceptable?
Performance		Acceptable?
Area		Acceptable?

Fig. 2: Defining acceptable interfaces. Source: Cadence

Standards are heading in several directions to optimize the various design factors. "Serial connections use very lightweight SerDes," says Walia. "They have minimalistic PHYs, and you don't need any decision feedback equalization — simply DLL clock-based forwarded approach."

The serial standards are being driven by the Optical Internetworking Forum (OIF). "This is referred to as 112G USR, or extra short reach (XSR) links," he says. "These should be ratified in 2021 timeframe. But remember that activities don't happen

based on standard ratifications. They happen based on the drafts that get made available for the standards. OIF drafts are available now."

On the parallel side, there are several standards efforts. First is Open High Bandwidth Interconnect (OpenHBI). This is an effort led by the Open Compute Project's (OCP) Open Domain-Specific Architecture (ODSA) subproject. Ratification for that is targeted to be the middle of this year.

Intel has developed the Advanced Interface Bus (AIB). "The specification for AIB 2.0 is already in the CHIPS Alliance GitHub," says Jose Alvarez, senior director in the CTO Office for the Programmable Solutions Group at Intel. "It is work in progress, and very close to being released. Our goal is 4 gigabits per second per wire, a total of about 7.6 terabits per second of bandwidth per interface. But it's not just about bandwidth itself. It's about energy efficiency. Today we are at 0.85 picojoules per bit of energy utilization. We went to 0.5 picojoules per bit, and the DARPA PIPES program wants to push this to a 0.1 pico Joule per bit. That's a much longer horizon, but we are leading toward that."

Many companies have deployed an approach that is called Bunch Of Wires (BOW). A November 2020 press release from GUC showed some performance numbers for that interface and demonstrates some of the performance tradeoffs. It quotes error-free communication between dies with full duplex 0.7 Tbps traffic per 1 mm of beachfront, consuming 0.25 pJ/bit. GUCs believe the next generation will support 1.3 Tbps error-free full duplex traffic per 1 mm of beachfront, with the same 0.25 pJ/bit power consumption using TSMC 5nm process.

How does that compare to a serial connection? GUC says that power consumption for a parallel connection is 6 to 10X lower than alternative solutions using ultra-short reach SerDes-based communication through package substrate.

Protocols

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 in the future," 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."

Productivity and reusability come with abstraction. "The next layer of interconnected is in terms of communication structure, protocols, busses, networks," says Michael Frank, fellow and system architect at Arteris IP. "CCIX and CXL are coming. People are building to them, but I do not see a standard that allows you to build a system of a handful of chiplets that all talk to each other."

Compute Express Link (CXL) is a cache-coherent interconnect for processors, memory expansion and accelerators. The 2.0 specification was released in November 2020. The goal is to maintain memory coherency between the CPU memory space and memory on attached devices, which allows resource sharing and reduced software stack complexity.

Similarly, Cache Coherent Interconnect for Accelerators (CCIX) has been migrating from in-system application to in-package. "With momentum towards 2.5D and chiplets, you essentially get rid of these longer latency and high power SerDes, or

interfaces and have parallel interfaces or very low latency XSR or short reach SerDes," says Millind Mittal, chairman of the Technical Steering Committee of CCIX Consortium and technical lead for CCIX, CXL and ODSA consortiums at Xilinx. "CCIX leverages the datalink layer of PCIe, but after that it separates out into optimized paths. We are defining our next version, and it is adapting to new transports. For 2.0, we are looking at adapting to in-package integration options." (See figure 3)

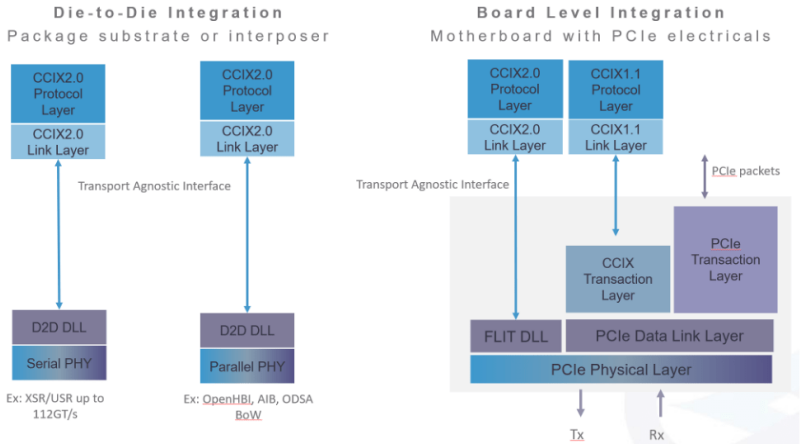


Fig. 3: CCIX 2.0 integration options. Source CCIX Consortium

There is also a standard from Arm. "This is part of the fabric where they have what is referred to as their Coherent Mesh Network (CMN) fabric," says Walia. "If you have two compute chips talking to each other on a die-to-die interface, the fabric-to-fabric has to look like a single fabric. And that's where zero latency is very important."

More than signals

Getting standards for signal interfaces is important, but more is required to make reusable chiplets. "We have to customize the IP today," says Walia. "This may mean removing the standard C4 bumps, replace them with micro-bumps. We have to work very closely in an iterative manner. There are often three or four iterations that go back and forth between us and the customer and their package provider."

Some of these issues are being addressed. "ODSA and AIB have proposed a bump maps," says Mehta. "This defines how the SerDes are going to be laid out, or how the parallel wires are going to be laid out. When both devices belong to the same customer, they have a little bit of flexibility. But standards are required for many things, like power and thermal, if it is not a closed-loop system."

Power is a big issue. "How do you bring 100 watts up through these tiny micro bumps?" asks Marc Swinnen, director of product marketing at Ansys. "You need to have a separate power distribution connectivity, physical connectivity scheme, thick TSVs, or something that can carry the power up through the chip. The technique that's most often used today is to aggregate a whole collection of these micro-bumps into bump farms that act as a single connection. So you take 100 of these, and they are all V_{ss} or V_{dd}, and they all work in concert, with the current being divided amongst them. Now you have to do very careful analysis to see that none of these contact points overheat and cause local melting."

Additional models are required. "I need a power model for this, a thermal model for that," says John Park, product management group director for IC packaging and cross-platform solutions at Cadence. "What is the pin pitch standard? There is a

checklist of things that people go through when they start thinking about standards for chiplet-to-chiplet interfaces. I believe there cannot be one standard for this. There will probably be a half dozen, a dozen, maybe even more. There are so many different types of packages that no standard is going to work for everything. And then, of course, there's reach. There may be dozens, and potentially hundreds, of chiplets in a big design, and if you design with laminate, you can get very big. So how far does the signal need to travel?"

As the big issues get addressed, new ones will bubble up. "With AIB 2.0, within the CHIPS Alliance, we are adding other concepts to chiplets like security," says Intel's Alvarez. "We are also looking at other ways of handling the interface, protocols, etc. We want to provide a more complete framework of hardware for chiplet development."

Conclusion

The semiconductor industry is transitioning proprietary chiplets into one based on standards. Many of the proprietary solutions are being placed into the hands of standards bodies today. The industry is coming together to consolidate those solutions, but only a certain level of consolidation is possible or perhaps even desirable.

Use cases will drive the rate of adoption of the proposals, and if initial success is had, many more use cases will look to move in this direction. But they all may require variants of the standards. Flexibility and optimization are always tricky to balance.