AyarLabs

Technical Brief

Optical I/O Chiplets Eliminate Bottlenecks to Unleash Innovation

Trends Driving In-Package Optical I/O Chiplets

Ayar Labs is the first to deliver monolithic in-package optical I/O (OIO) chiplets, a new universal I/O solution that enables chips to communicate with each other from millimeter to kilometer-scale, at the power, latency, and bandwidth density of in-package interconnect. Several technology trends point to the arrival of OIO chiplets as a critical industry inflection point.

Increasing demand for artificial intelligence, high-performance computing, aerospace data collection and communication, hyper-scale data centers, disaggregated memory and storage has led to the development of a variety of powerful, high-throughput, system-on-chip (SoC) solutions that redefine traditional computing architectures. These new architectures accelerate the data rate of communications between die, sockets, boards, systems, and racks. Electrical I/O has been a bottleneck to scaling processor and SoC performance for the past quarter-century of Moore's Law scaling. It is now a significant barrier to scaling the performance of high-throughput SoC architectures.

Experts agree that electrical SerDes, the most common form of electrical I/O, is hitting a wall. Going beyond 112 gigabits per second (Gbps) is extremely challenging because the large signal losses in copper interconnects at the board level make it hard to transmit data further than a few centimeters at such a high data rate [1]. The next wave of high-performance computing architectures require a new form of universal I/O that eliminate the bottlenecks created by electrical I/O.

Gordon Moore anticipated an eventual slowdown of the scaling benefits created by transistor miniaturization. He envisioned a second phase of Moore's Law where increasingly complex, single-chip solutions give way to an approach that disaggregates common features and enables larger pooling of functional building-blocks connected via high-performance interconnects.

> "It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected."
>
> - Gordon Moore[2]

To deliver the second phase of Moore's Law, the semiconductor industry is moving to chiplets to lower costs, improve yields, enhance reliability, and deliver faster time-to-market for next-generation product designs. Chiplets shorten product cycles by separating the development timelines of different functional building blocks on smaller pieces of silicon instead of putting every function on a single die. This approach allows designers to use advanced process technologies on functionality that requires best available area density and power efficiency, while functionality that does not require the most advanced manufacturing is produced with older, more reliable technology. Chiplet architectures also improve yields by combining smaller pieces of silicon, reducing the risk of manufacturing defects that increase with die size [3]-[6].

Optical communications began replacing electrical cabling in high-performance computing (HPC) and datacenter applications several years ago, changing copper cables between electrical faceplate ports to pluggable optical transceivers with connectorized fiber cables as shown in Figure 1. As faceplates become constrained by mechanical and thermal limits, optical communications are moving from the faceplate to inside the package. In-package optical I/O is a revolutionary approach that integrates silicon-photonic chiplets built with CMOS processes inside a multi-chip package (MCP). OIO chiplets eliminate electrical I/O bottlenecks and transcend process limitations to unleash the next wave of innovation in semiconductor and datacenter design [7]-[9].

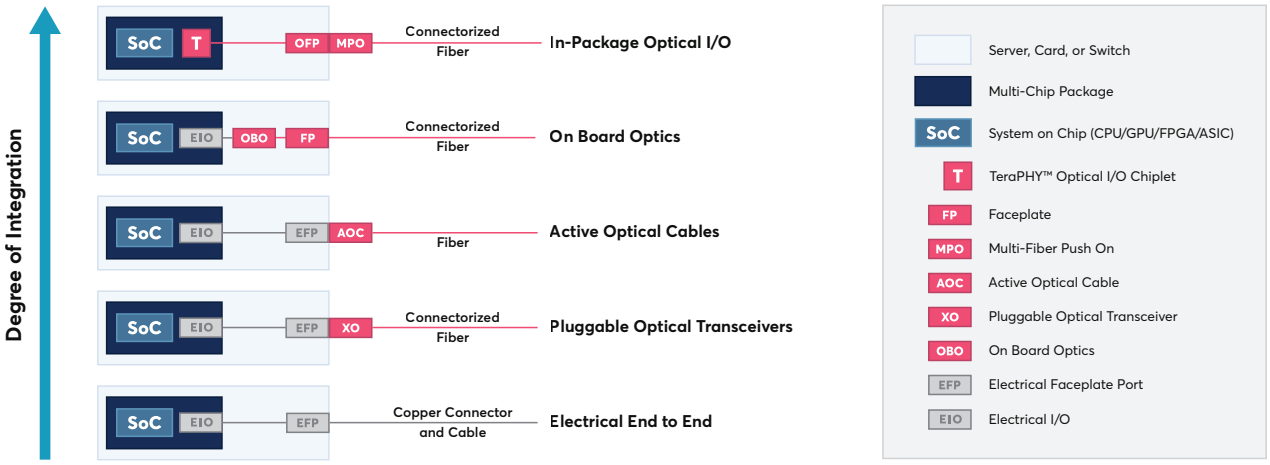

Figure 1 - Industry Approaches to Optical Integration

Electrical I/O Barriers to High-Performance Architectures

When designers need to deliver terabit data rates between chips at distances beyond the faceplate, electrical I/O can't deliver the power efficiency, latency and bandwidth density required for the next wave of high-performance architectures.

Power Efficiency

Power efficiency is a critical limitation when designing electrical systems and datacenters since it directly impacts heat and reliability [10]. Today, the power efficiency of long-reach (LR) electrical I/O at 112 Gbps is 6-to-10 pico Joules per bit (pJ/b). Reaching from the package to the edge of a printed circuit board (PCB) is possible at this data rate but draws a lot of power. Reaching from the package to go across systems, racks and datacenters draws significantly more power, requiring a combination of electrical I/O and pluggable optics.

Latency

Latency is another critical design factor, limiting the size and number of components interconnected into a system. On-board and off-board electrical I/O at data rates above 50 Gbps require forward error correction (FEC) coding, which introduces added latency of ~100 ns. Such latency, while tolerated in networking applications, is not tolerated in distributed computing systems (memory semantic fabrics) such as those used for machine learning training, inference, and other high-performance computing applications.

Bandwidth Density

Today electrical I/O provides bandwidth density around 200 Gbps/mm, supporting 25.6-Tbps (terabits per second) Ethernet switch chips. Next-generation 51.2-Tbps Ethernet switch chips will require bandwidth density around 500 Gbps/mm. Future 102.4-Tbps Ethernet switch chips will require bandwidth density around Tbps/mm. While the latest generation of 112-Gbps long-reach electrical SerDes solution can deliver bandwidth density at 200-500 Gbps/mm, there is no roadmap for SerDes technology to achieve bandwidth densities beyond this limit. Long-reach electrical SerDes at high data rates also suffer from increased package complexity, cooling requirements, and costs.

Reach

Electrical reach is the distance between bumps inside an MCP or ball-to-ball across the PCB. Table 1 shows reach segmentation by length, loss, and applications. Electrical I/O is suitable for applications within a package but does not scale for connections outside the package due to signal loss, the need for repeaters, and the amount of error correction required. Correcting for these inherent electrical limitations quickly drives power curves beyond sustainable levels.

Length		Loss	Application
< 10 mm/0.4 in	USR	1.5dB@14GHz 3dB@28GHz	Bump-to-bump Inside MCP or 3D Stack
< 50 mm/2.0 in	XSR LPPI	4dB@14GHz 8dB@28GHz	Ball-to-Ball Across PCB
< 200 mm/7.9 in	VSR C2M	10dB@14GHz 20dB@28GHz	Ball-to-Ball
< 500 mm/19.7 in	MR C2C	20dB@14GHz 40dB@28GHz	Ball-to-Ball
< 1000 mm/39.4 in	LR C2F	35dB@14GHz	Ball-to-Ball

Table 1 - Electrical Reach Length, Loss and Applications [10]

Ayar Labs In-Package Optical I/O Chiplet

Ayar Labs' solution combines TeraPHY™, an in-package optical I/O chiplet, with SuperNova™, a multi-wavelength optical source, to eliminate I/O bottlenecks, transcend process limitations, and unleash innovative architectures. TeraPHY chiplets disrupt the traditional performance, cost, and efficiency curves of the semiconductor and computing industries by combining silicon photonics with standard CMOS manufacturing processes to deliver up to 1000x bandwidth density improvement at 1/10th the power compared to electrical I/O [12].

Developed in a high-volume GlobalFoundries 45 nanometer process, TeraPHY chiplets integrate millions of transistors with hundreds of photonic devices to drive tens of Tbps of bandwidth up to 2 km out of the package with unmatched power efficiency of less than 5pJ/b. Latency is only 10ns + 5ns/m, point-to-point, with no need for repeaters or FEC, allowing designers to create logically connected, physically distributed compute architectures that scale across racks. TeraPHY delivers bandwidth density in excess of 200 Gbps/mm today with a roadmap to Tbps/mm for future generations of high-performance architectures.

Optical I/O Requirement	Ayar Labs TeraPHY™ Performance
Power efficiency	<5pJ/b, roadmap to <2pJ/b
Latency	< 2 x 5ns + TOF
Bandwidth density	In excess of 200 Gbps/mm, roadmap to Tbps/mm
Reach	Package to package connections from mm up to 2 km

Table 2 - Summary of Optical I/O Requirements and TeraPHY Performance

Figure 2 shows an example MCP assembly hosting two TeraPHY chiplets, each delivering up to 2 Tbps of optical I/O bandwidth.

Figure 3 shows an opened-lid MCP containing four TeraPHY chiplets integrated into the package with an SoC. MCP technologies (Embedded Interconnect Bridge (EMIB), Silicon-interposer, High-density fanout) provide many advantages over large monolithic dies, including improved yield, support for multiple process nodes in one package, and increased power efficiency.

Figure 2 - Example Multi-Chip Package Assembly with TeraPHY™ Optical I/O

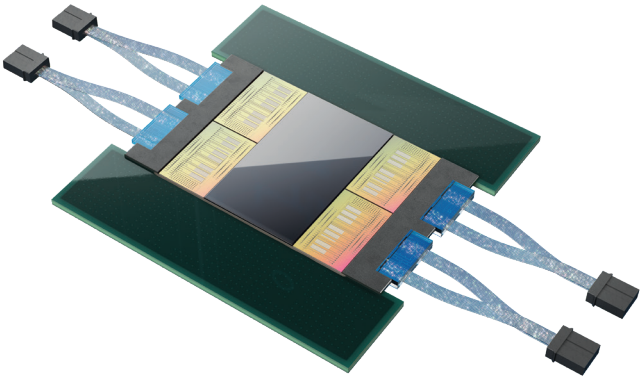

Figure 3 - TeraPHY™ Optical I/O Multi-Chip Package Integration Example

Current pluggable optical solutions undergo full multi-chip packaging (driver chips, silicon photonic chip, laser diodes, etc.) for every optical port (100-400 Gbps). As shown in Figure 4, TeraPHY OIO chiplets (each containing up to eight 256-Gbps optical ports) are flip-chip attached (e.g. C4 or Cu-pillar/µBump) to simplify in-package integration of many optical ports and automate assembly.

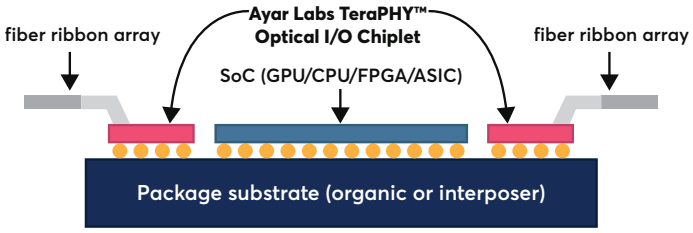

Figure 4 - Flip-Chip Attached for High Volume Manufacturing

Figure 5 presents an example OIO application with a TeraPHY chiplet and SuperNova optical source hosted on the same PCB. One of the key scalability ingredients of the OIO solution is moving the heat-sensitive laser diodes into a single optical supply module situated outside of the hot MCP package. In the example, two systems (servers, cards, or switches) are connected by single mode fiber optic cables across distances up to 2 km.

TeraPHY chiplets support up to 8 ports, each operating at up to 256 Gbps, enabling 2 Tbps of optical I/O bandwidth per chiplet. Each port connects to three fibers, one for transmit (Tx), one for receive (Rx), and one for the laser light source. The light originates at the Ayar Labs SuperNova laser (multi-wavelength optical source), supplying eight O-band wavelengths to each TeraPHY optical port.

Many other configurations are possible, including multiple TeraPHY chiplets in a single package, hosting multiple MCPs on the same board, and hosting SuperNova modules on a separate board or rack to supply light to TeraPHY chiplets on boards distributed across multiple datacenter racks.

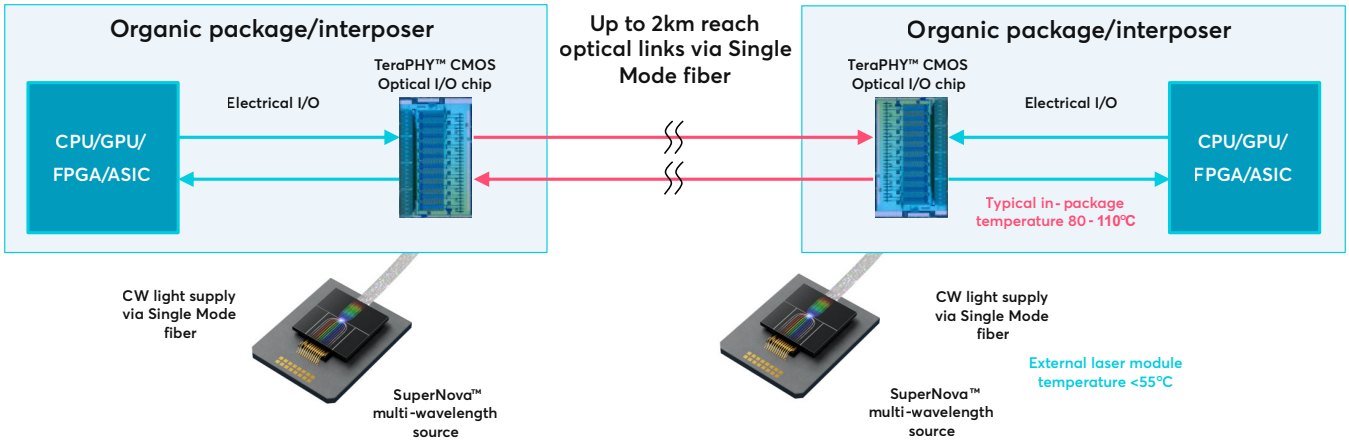

Figure 5 - Example Application of TeraPHY™ Optical I/O

Figure 6 zooms in to show the floorplan for the TeraPHY OIO chiplet. The Advanced Interface Bus (AIB) provides an electrical interconnect to the SoC, operating at an aggregated bandwidth of 2 Tbps [12]. This interface requires 2.5D package integration, which can be realized with a range of technologies such as Embedded Multi-Die Interconnect Bridge (EMIB) [13], silicon interposer, Chip-on-Wafer-on-Substrate (CoWoS), InFO, and redistribution layer (RDL) [14]. Future TeraPHY chiplets can support other parallel interfaces such as OpenHBI/BoW/proprietary or high-speed serial interfaces such as PCIe/JESD/XSR.

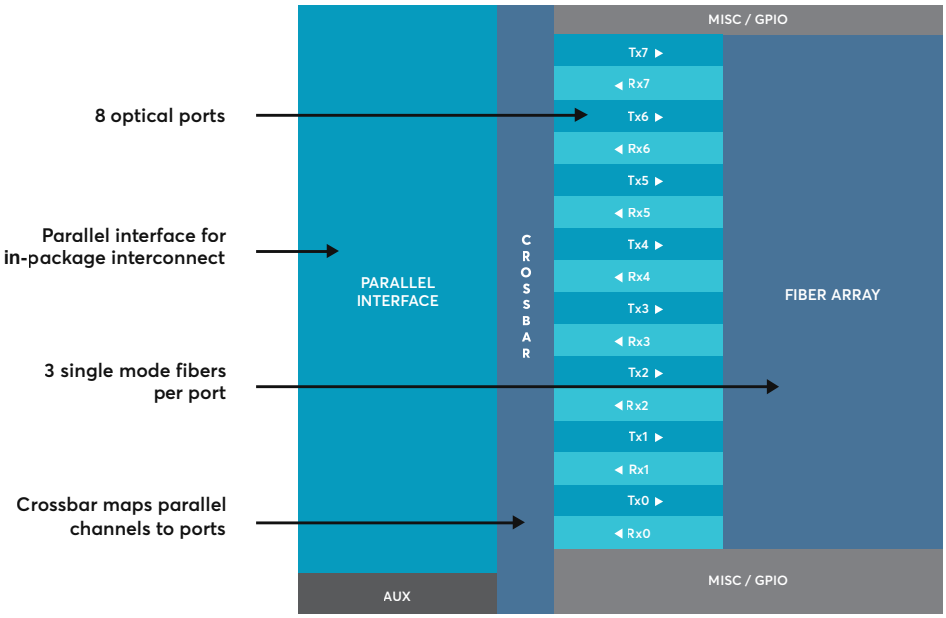

Figure 6 - Ayar Labs TeraPHY™ OIO Chiplet Floorplan

Figure 7 abstracts a single TeraPHY Tx channel and Rx channel, showing microring resonators encoding data from the AIB onto the wavelengths supplied by SuperNova. The encoded light is transmitted up to 2 km to another TeraPHY chiplet, where the corresponding microring resonator on the Rx port converts the light back into an electrical signal. This approach does not require additional protocols, repeaters, or FEC – minimizing both latency and power [9].

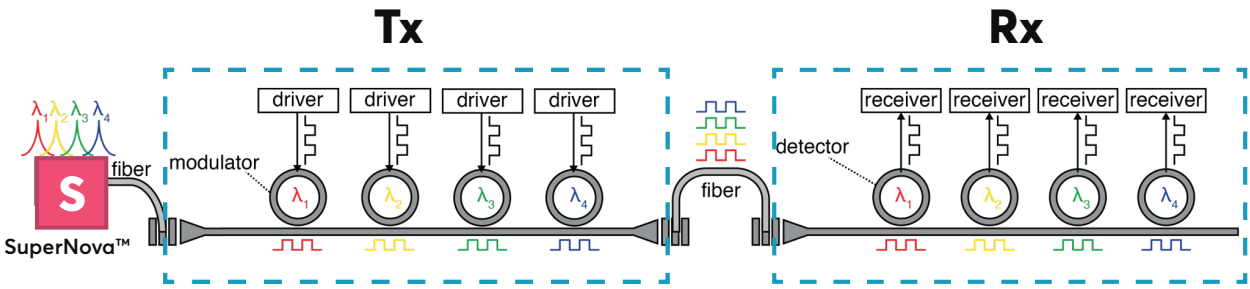

Figure 7 - TeraPHY™ Tx and Rx Channel Abstraction

Port capacity scales with the number of wavelengths per port and data rate per wavelength. Data may be encoded at 16 Gbps/λ, 25 Gbps/λ or 32 Gbps/λ to achieve up to 256 Gbps per optical port and 2 Tbps per chiplet as shown in Table 3. Ayar Labs has demonstrated technology components providing up to 100 Gbps/λ, paving the way for a long-term bandwidth scaling roadmap.

No. of λ per Optical Port	Gbps per λ	Gbps per Optical Port	Tbps per Chiplet	Bandwidth Density Gbps / mm
8	16	128	1.024	114
8	25	200	1.6	178
8	32	256	2.048	228

Table 3 - TeraPHY™ Data Rates and Bandwidth Density

Ayar Labs has developed SuperNova, the industry's first multi-wavelength, multi-port optical source with 64 wavelengths and the first optical source designed to be compliant with the CW-WDM MSA specification released in 2021. The Ayar Labs SuperNova remote light source (Figure 8) can be deployed across a wide range of applications including high-speed I/O, artificial intelligence, optical computing, and high density, co-packaged optics.

As the backbone of Ayar Labs' OIO solution, SuperNova provides up to 16 wavelengths of light, powering up to 16 ports, and is capable of supplying light for 256 channels of data (or 8.192 Tbps at full capacity). The solution provides up to 1000x the bandwidth at 1/10th of the power compared to electrical I/O alternatives – effectively eliminating I/O bottlenecks and transcending process limitations.

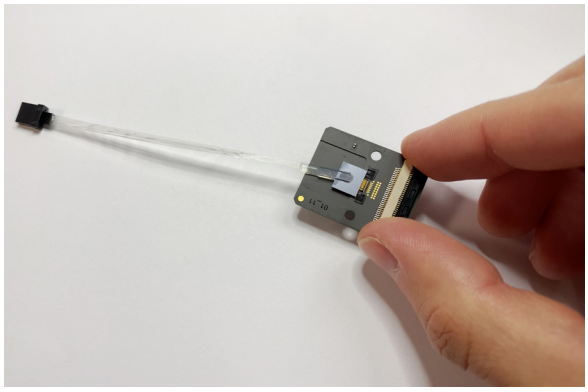

Figure 8 – Ayar Labs SuperNova™ Optical Source

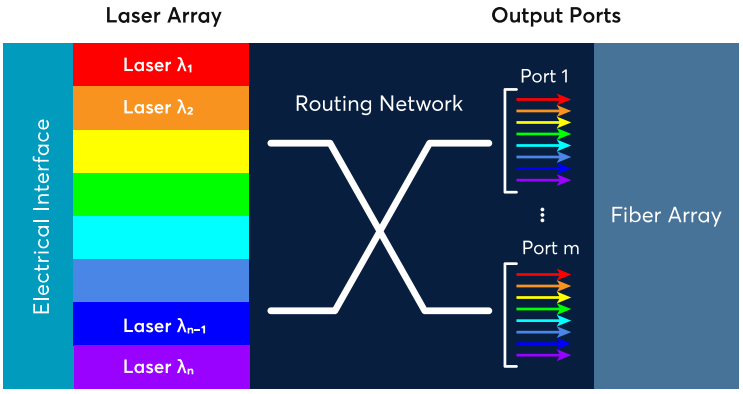

Figure 9 - SuperNova™ Optical Source Floorplan

Applications of Optical I/O

By eliminating the constraints imposed by electrical I/O, TeraPHY OIO enables new system architectures for applications across several markets. These markets include cloud computing, HPC and artificial intelligence (AI), telecommunications and an increasing array of intelligent edge applications.

Cloud Computing

Cloud applications include rack-scale architectures and resource pooling, scaling the power of systems that were previously constrained within a server to extend across rows of racks in a datacenter. The emergence of OIO presents new opportunities for much more powerful data centers, providing technological avenues that were previously impossible, or cost-prohibitive. With the ability to deliver up to a 1000x improvement in interconnect bandwidth density at one-tenth of the power, Ayar Labs is disrupting the traditional performance, cost, and efficiency curves and lowering the bar to higher performing architectures. OIO is helping to democratize the access to more powerful computing, creating opportunities and avenues for innovation in virtually every industry that accesses cloud computing. With OIO, the computing industry is being transformed, providing the performance and scalability needed for advanced applications and creating more opportunity for innovation.

HPC and AI

HPC and AI are converging as rapidly evolving algorithms drive the need for greater processing power, high-bandwidth and low-latency access to higher memory capacity, and low-latency between computational elements. OIO is effectively overcoming the limits of transistor density, which has stalled Moore's law, and limiting the ability of CPUs and HPC systems to run at their full potential. By bringing optics to the CPU and memory, OIO is able to dramatically improve performance, while enabling truly disaggregated and distributed architectures that support more flexible systems – all while requiring only a 10th of the power required for copper-based data movement. With the ability to provide high-bandwidth density, new architectures are being explored that enable every element in the system to communicate with low latency, and effectively eliminating many of the traditional impediments to scaling. OIO is emerging as the best solution to drive the next phase of Moore's Law and help HPC centers accelerate the slope of compute progress needed to tackle ever-growing scientific problem sizes and HPC/AI convergence.

Connectivity

Bandwidth demands are scaling to astronomical levels as mmWave technologies such as 5G (and eventually 6G) put new pressures on telecommunications systems. As these new networking technologies drive faster data rates over existing infrastructures, bottlenecks at interconnection points create network congestion and unacceptable points for failure. Replacing copper cables, OIO technology provides interconnection that enables high-bandwidth and low-power connectivity between antenna/sensing elements and the digital signal processing infrastructure – all within a compact package that can be used in field where space and weight constraints are critical factors for success. With the ability to deploy new architectures in places where low-latency, high-bandwidth, and low-power capabilities are paramount, OIO technology has the potential to completely revamp previous-generation components with more nimble and capable equipment and architectures. As the edge-computing ecosystem being built around 5G networking evolves and matures, OIO technology will increasingly deliver fast, high-bandwidth connectivity required to process and distribute data at the edge.

Intelligent Edge

The capabilities of OIO, particularly the unique characteristics of the TeraPHY chiplet, which efficiently couples the electrical and optical domains, opens a vast realm of opportunities for intelligent edge use cases, such as those being employed in autonomous vehicles and in aerospace and defense organizations. The capabilities of OIO enable the development of new sensing and computer architectures that disaggregate units within the system – providing a range of benefits for improving communications, strengthening defensive measures, and enabling improvements in real-time decision-making. With capabilities for high-bandwidth, low-latency, and low-power interconnection resistant to electromagnetic interference, OIO is being used to create improvements in systems of all types. Applications include improvements in air traffic control systems, ground-to-air communications with orbital vehicles, phased array radar and communications systems, unmanned aircrafts system, communications with interconnected satellites and more. OIO effectively solves the key challenges of accelerating compute performance while reducing size, weight, and power requirements, making it an ideal choice for the growing number of intelligent edge applications.

Conclusion

Semiconductor designers are creating new, chiplet based architectures that deliver the second phase of Moore's Law to lower costs, improve yields, enhance reliability, and deliver faster time-to-market for next-generation designs. At the same time, the limitations of electrical I/O are driving the search for a new form of universal I/O that quickly delivers terabit data rates from die-to-die and across data centers with less power. TeraPHY OIO chiplets eliminate electrical I/O bottlenecks and transcend process limitations to unleash the next wave of innovation in semiconductor and datacenter design.

Ayar Labs is the first to deliver monolithic in-package optical I/O chiplets, a new universal I/O solution that replaces traditional electrical I/O and enables chips to communicate with each other from millimeters to kilometers, to deliver orders of magnitude improvements in latency, bandwidth density, and power consumption.

Contact us at www.ayarlabs.com to learn more about how TeraPHY OIO can accelerate your design performance today.

References

[1] A. Steffora Mutschler, "Wrestling With High-Speed SerDes," Semiconductor Engineering, 2019. [Online]. Available: https://semiengineering.com/wrestling-with-high-speed-serdes/.

[2] G. E. Moore, "Cramming more components onto integrated circuits. In: Electronics," Electronics, 1965.

[3] L. T. Su, S. Naffziger, and M. Papermaster, "Multi-chip technologies to unleash computing performance gains over the next decade," in Technical Digest - International Electron Devices Meeting, IEDM, 2018.

[4] T. Simonite, "To Keep Pace With Moore's Law, Chipmakers Turn to 'Chiplets,'" Wired, 2018. [Online]. Available: https://www.wired.com/story/keep-pace-moores-law-chipmakers-turn-chiplets/.

[5] T. Coughlin, "Chiplets for All," Forbes, 2019. [Online]. Available: https://www.forbes.com/sites/tomcoughlin/2019/05/11/chiplets-for-all.

[6] E. Sperling, "Chiplets, Faster Interconnects, More Efficiency," Semiconductor Engineering, 2019. [Online]. Available: https://semiengineering.com/chiplets-faster-interconnects-and-more-efficiency/.

[7] C. Sun et al., "Single-chip microprocessor that communicates directly using light," Nature, 2015.

[8] V. Stojanović et al., "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes [Invited]," Opt. Express, vol. 26, no. 10, p. 13106, 2018.

[9] M. Wade, M. Davenport, M. Rust, and F. Sedgwick, "TeraPHY : A Chiplet Technology for Low-Power , High-Bandwidth In-Package Optical I / O," Hot Chips Conference, 2019. [Online]. Available: https://ayarlabs.com/teraphy-a-chiplet-technology-for-low-power-high-bandwidth-in-package-optical-i-o/.

[10] R. Danilak, "Why Energy Is A Big And Rapidly Growing Problem For Data Centers," Forbes, 2017. [Online]. Available: https://www.forbes.com/sites/forbestechcouncil/2017/12/15/why-energy-is-a-big-and-rapidly-growing-problem-for-data-centers.

[11] J. Goergen, C. Systems, and V. Parthasarathy, "Spanning SERDES Across Reaches - Finding the Best Modulation Approach," no. November, 2014.

[12] R. Meade et al., "TeraPHY: A High-Density Electronic-Photonic Chiplet for Optical I/O from a Multi-Chip Module," 2019 Opt. Fiber Commun. Conf. Exhib. OFC 2019 - Proc., no. Mcm, pp. 5–7, 2019.

[13] D. Kehlet, "Accelerating Innovation Through A Standard Chiplet Interface : The Advanced Interface Bus (AIB)," 2019.

[14] R. Mahajan et al., "Embedded Multi-die Interconnect Bridge (EMIB)-A High Density, High Bandwidth Packaging Interconnect," in Proceedings - Electronic Components and Technology Conference, 2016.

[15] M. Lapedus, "Bridges Vs. Interposers," Semiconductor Engineering, 2018. [Online]. Available: https://semiengineering.com/using-silicon-bridges-in-packages/.

AyarLabs

For general information see our website at www.ayarlabs.com

For inquiries, please email us at info@ayarlabs.com

Emeryville Office
6460 Hollis St, Ste A
Emeryville, CA 94608

Silicon Valley Office
3351 Olcott St
Santa Clara, CA 95054