Technology Highlight

Addressing the I/O Bottleneck in Datacenters with Silicon Photonics

Alex Norton, HPC Research Analyst
October 2019

HYPERION RESEARCH OPINION

This is one of a series of Hyperion Research Technology Highlights looking at the new technologies being developed for HPC, AI, HPDA and other advanced computing areas.

The AI ecosystem is experiencing an influx of startups worldwide, seeking to tackle new AI applications on the premise that they could run more efficiently on hardware specifically designed for AI workloads. From new processors designed for image recognition or data center neural network training to advances in I/O hardware, these start-ups are developing new, innovative ways to solve complex AI problems of today and the future.

Ayar Labs, which grew out of a DARPA-funded research project, is looking to create an optical I/O solution to overcome the copper-based I/O bottleneck that limits high bandwidth data transfer between chips. While processing, compute power and storage have showed improvements over the years and electrical I/O interfaces seem to have plateaued, with 112 Gbps copper-pin I/O appearing to be the end of the line. As the need for bandwidth continues to grow with the rise of large data centers and the increasing deployment of artificial intelligence (AI), as well as other data intensive workloads, new solutions are needed.

Ayar Lab's answer to this is an optical interconnect designed from the ground up to provide up to 2km of distance between chips with up to a 2.56 Tbps link per chiplet using a two-part solution:

- A multi-terabit chiplet that connects to a host device via a wide-low power electrical I/O and outputs an optical signal over fiber. Called TeraPHY, this chiplet is integrated with the target device in package but could also be on board.
- A light supply device called SuperNova, which can support multiple TeraPHY devices.

Ayar Labs believes that such an optical I/O capability could allow for disaggregation of systems, as well as the development of new server/processing architectures, with applicability far beyond the AI space. Such optical I/O solutions could help improve performance on high data-flow based workloads, such as HPC modeling and simulation workloads, as well as support emerging cloud infrastructures running AI workloads with large data transfer requirements, particularly during neural network training jobs.

SITUATION OVERVIEW

Ayar Labs sees a world of ever-increasing bandwidth demands, driven by the explosion of AI applications, the growing complexity of neural networks, and an increase in the processing demands of deep-learning. As the pace of processing capability improvements has slowed down considerably in recent years, the AI community, as well as many in other areas of high performance computing, is exploring new chip architectures to drive stronger performance gains. However, traditional electrical copper-based I/O interfaces are topping out 112 Gbps, creating an additional bottleneck that could significantly decrease the performance capabilities of future systems.

Further complicating copper-based solutions is that as data rates increase, wire trace distances must shorten. As data rates approach the current limit of 112 Gbps, the maximum trace lengths are on the order of inches. This limited reach forces solution developers to tightly pack processing chips, creating both power and thermal issues.

As copper-based I/O solutions were beginning to reach their bandwidth limits, telecommunication companies moved to fiber optic solutions to transfer high bandwidth data across large distances for high-speed networks. This same technology, however, has not been practical for chip-to-chip communication over a short distance for two major reasons:

- Interfaces (optical modules/transceivers) are complex, requiring a mix of materials (silicon, germanium, etc.) and devices that must be integrated in a single package.
- Optical modules are relatively large, power hungry, and difficult to miniaturize.

Ayar Labs' Solution for Increasing Chip-to-Chip Bandwidth

To address these concerns, Ayar Labs' strategy is to move beyond the current approach of the optical module industry and design a silicon solution (chip) from the ground up, focusing on a few key metrics:

- Maximizing silicon area bandwidth density
- Maximizing chip edge bandwidth density (fiber capacity can be increased using dense wavelength division multiplexing (DWDM))[1]
- Maximizing energy efficiency (expressed as energy per bit)
- Creating high reliability at high temperature

Devices

The company's solution set consists of two devices:

- A Terabit transceiver called TeraPHY for in-package optical integration, built on a qualified, high-volume process (the GlobalFoundries RF SOI process), which is the RF version of silicon-on-insulator technology. The initial product has planned availability in 1.28 Tbps and 2.56 Tbps versions.

[1] Ring-resonator based DWDM achieves increased bandwidth density by combining multiple wavelengths of light into a single fiber. Ring-resonators achieve area and energy efficient modulation and detection of densely gridded wavelengths.

- A multiwavelength laser source, referred to as SuperNova light supplies, built on an indium phosphide process. These devices are designed to supply light for up to 256 channels of data, for a total bandwidth of 8.2 Tbps. This device provides the light source for up to multiple TeraPHY devices, depending on configuration.

Ayar Labs reports that this split approach offers a number of advantages over a fully integrated solution:

- Separating the transceiver from the light source allows it to be built on a high-volume commercial CMOS process.
- This makes it easier for IC/system designers to deal with thermal issues: the CMOS transceiver is designed to operate within the harsh conditions of a high power SOC while the laser can be mounted in a less hostile environment.
- It also enables differing topologies where the light sources can be located away from the same PCB as the devices they serve (up to 2 km). For example, light sources for a system could be relocated to the center of the rack to ease cooling load.

System Integration

The TeraPHY is intended to be integrated in a multi-chip module (MCM) with the target device(s), but could be mounted on a PCB close to the target device. Optical connectivity between TeraPHY devices is via standard duplex LC connectors. Light is supplied to the TeraPHY via additional fiber connections.

FIGURE 1

Typical System

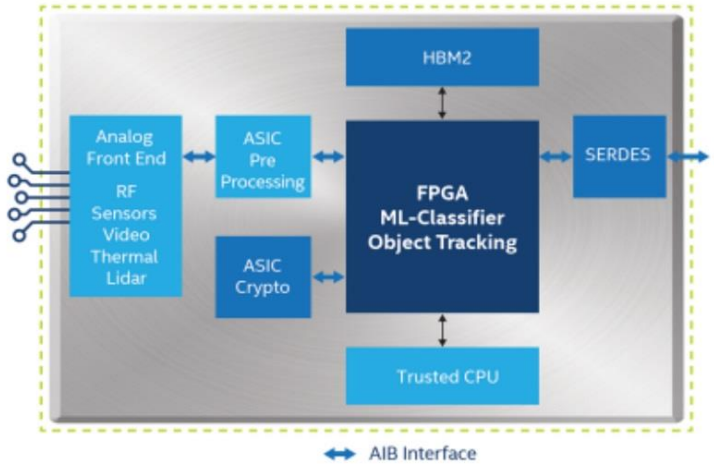

Source: HPCWire and Intel Corp, 2019

AYAR LABS' VISION

The company sees their technology as helping to create a new paradigm for high-performance platforms because an optical interconnect can allow chip-level I/O to continue to scale up past the current bottleneck of 112 Gbps, simplifying system design. Further, because the correlation between bandwidth and distance is broken, new disaggregated architectures are possible. Ayar Labs argues, for example, that chips can communicate at the same speed with chips located 2 km apart as if they were on the same or an adjacent board.

The company plans to have a demonstration of their technology running with one of their partners at Supercomputing 2019 (SC19) in November. Ayar Labs further expects to deliver early units in 2021.

FUTURE OUTLOOK

The AI hardware ecosystem will continue to evolve as new customers and applications enter the community. The desire for more complex models to be trained in shorter times, as well as deployed at the edge, will require new approaches to hardware solutions, and Ayar Labs has tasked itself with redesigning an integral point in the data flow of AI computing. They are part of an initial push to challenge the current hardware on the market by targeting AI from the hardware level, rather than only at a software or application level, to maximize the efficiency and computing power of systems to solve the AI problems of the future.

About Ayar Labs, Inc.

Ayar Labs is disrupting the traditional performance, cost, and efficiency curves of the semiconductor and computing industries by driving a 1000x improvement in interconnect bandwidth density at 10x lower power. Ayar Labs' patented approach uses industry standard cost-effective silicon processing techniques to develop high speed, high density, low power optical based interconnect "chiplets" and lasers to replace traditional electrical based IO. The company was founded in 2015 and is funded by Playground Global and Founders Fund, as well as strategic investment by Intel Capital and GLOBALFOUNDRIES. For more information, please visit http://www.ayarlabs.com.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
