Tunable Lasers and DSPs in the Age of AI

The use of generative artificial intelligence (AI) models is transforming several industries, and data centers are no exception. AI models are computationally heavy, and their increasing complexity will require faster and more efficient interconnections than ever between GPUs, nodes, server racks, and data center campuses. These interconnects will have a major impact on the ability of data center architectures to scale and handle the demands of AI models in a sustainable way.

As we discussed in a previous article, transceivers that fit this new AI era must be fast, smart, and adapt to multiple use cases and conditions. However, what impact will that have on the tunable lasers and digital signal processors (DSPs) inside these transceivers? This article will review a couple of trends in lasers and DSPs to adapt to this new era.

The Power of Laser Arrays

In 2022, Intel Labs demonstrated an eight-wavelength laser array fully integrated on a silicon wafer. These milestones are essential for optical transceivers because the laser arrays can allow for multi-channel transceivers that are more cost-effective when scaling up to higher speeds.

Let's say we need an intra-DCI link with 1.6 Terabits/s of capacity. There are three ways we could implement it:

- **Four modules of 400G:** This solution uses existing off-the-shelf modules but has the largest footprint. It requires four slots in the router faceplate and an external multiplexer to merge these into a single 1.6T channel.
- **One module of 1.6T:** This solution will not require the external multiplexer and occupies just one plug slot on the router faceplate. However, making a single-channel 1.6T device has the highest complexity and cost.
- **One module with four internal channels of 400G:** A module with an array of four lasers (and thus four different 400G channels) will only require one plug slot on the faceplate while avoiding the complexity and cost of the single-channel 1.6T approach.

	4x400G	1x1.6T	1x4x400G
# of modules	4	1	1
Cost/complexity	Medium	High	Low
Faceplate density	Low	High	High
External MUX	Yes	No	No

Same spectral efficiency

Table 1: Comparison of different ways to implement a 1.6 Terabit/s link: four modules of 400G, a single-channel 1.6T module, and a module with four channels of 400G.

Multi-laser array and multi-channel solutions will become increasingly necessary to increase link capacity in coherent systems. They will not need more slots in the router faceplate while simultaneously avoiding the higher cost and complexity of increasing the speed with just a single channel.

Co-Designing DSP and Optical Engine

Transceiver developers often source their DSP, laser, and optical engine from different suppliers, so all these chips are designed separately from each other. In such cases, the DSP is like a Swiss army knife: a jack of all trades designed for different kinds of optical engines but a master of none.

For example, current DSPs are designed to be agnostic to the material platform of the photonic integrated circuit (PIC) they are connected to, which can be Indium Phosphide (InP) or Silicon. Thus, they do not exploit the intrinsic advantages of these material platforms. Co-designing the DSP chip alongside the PIC can lead to a much better fit between these components.

A PIC and a standard *platform-agnostic* DSP typically operate with signals of differing intensities, so they need some RF analog electronic components to "talk" to each other. This signal power conversion overhead can constitute up to 2 Watts or about 10-15% of transceiver power consumption.

However, the modulator of an InP PIC can run at a lower voltage than a silicon modulator. If this InP PIC and the DSP are designed and optimized together instead of using a standard DSP, the PIC could be designed to run at a voltage compatible with the DSP's signal output. This way, the optimized DSP could drive the PIC directly without needing the RF analog driver, doing away with most of the power conversion overhead we discussed previously.

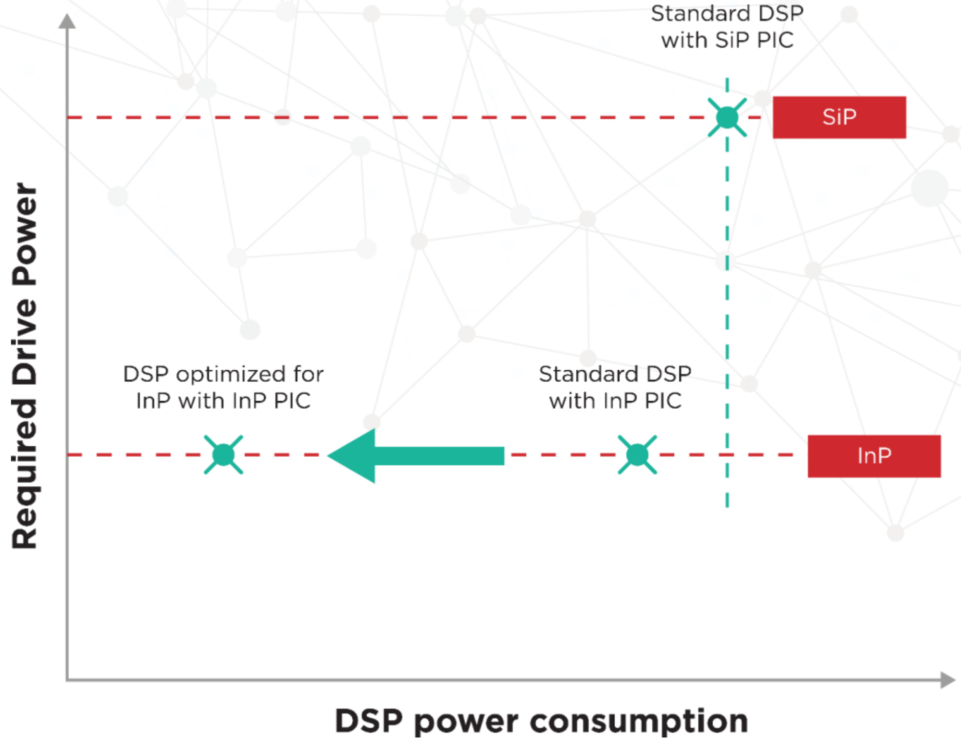

Figure 1: Comparison between the drive power and power consumption of three different PIC + DSP pairings: standard DSP with a silicon PIC, standard DSP with an InP PIC, and an optimized DSP with an InP.

Smart Devices Built In Scale

Making coherent optical transceivers more affordable is a matter of volume production. As discussed in a previous article, if PIC production volumes can increase from a few thousand chips per year to a few million, the price per optical chip can decrease from thousands of Euros to mere tens of Euros. Achieving this production goal requires photonics manufacturing chains to learn from electronics and leverage existing electronics manufacturing processes and ecosystems.

While vertically-integrated PIC development has its strengths, a fabless model in which developers outsource their PIC manufacturing to a large-scale foundry is the simplest way to scale to production volumes of millions of units. Fabless PIC developers can remain flexible and lean, relying on trusted large-scale manufacturing partners to guarantee a secure and high-volume supply of chips. Furthermore, the fabless model allows photonics developers to concentrate their R&D resources on their end market and designs instead of costly fabrication facilities.

Further progress must also be made in the packaging, assembly, and testing of photonic chips. While these processes are only a small part of the cost of electronic systems, the reverse happens with photonics. To become more accessible and affordable, the photonics manufacturing chain must become more automated and standardized. It must move towards proven and scalable packaging methods that are common in the electronics industry.

If you want to know more about how photonics developers can leverage electronic ecosystems and methods, we recommend you read our in-depth piece on the subject.

Takeaways

In conclusion, tunable lasers and DSPs are adapting to meet the rising demands of AI-driven data center infrastructure. The integration of multi-wavelength laser arrays and the co-design of DSPs and optical engines are crucial steps towards creating more efficient, scalable, and cost-effective optical transceivers. These devices will be smart and provide some telemetry data, and must shift towards volume production and the adoption of electronics industry methodologies. These innovations not only promise to enhance the capacity and efficiency of data center interconnects but also pave the way for a more sustainable growth trajectory in the face of AI's computational demands.

More Articles from EFFECT Photonics