As Trillion dollar companies spend hundreds of billions on AI infrastructure, the hardware requirements are becoming ever more problematic.

From a photonics perspective there are fundamental differences between a conventional data center and an AI facility. In a data center, data is moving mostly into and out of the stacks of PCs, referred to as North-South. In an AI facility data also moves between the GPUs, East-West, in addition to North-South. In practice this results in a photonic backplane that is pushed to the limits.

Patrick Kennedy was able to photograph the XAI Colossus Data Center [1]. The sheer volume of optical fiber required highlights both the importance of East-West communications, but also the limits of existing fiber communication. Figure 1 shows the backside of the racks where multiple 400 Gb/s electrical signals are converted to 3.6Tb/s of optical bandwidth per GPU compute server. This approach is clearly not scalable.

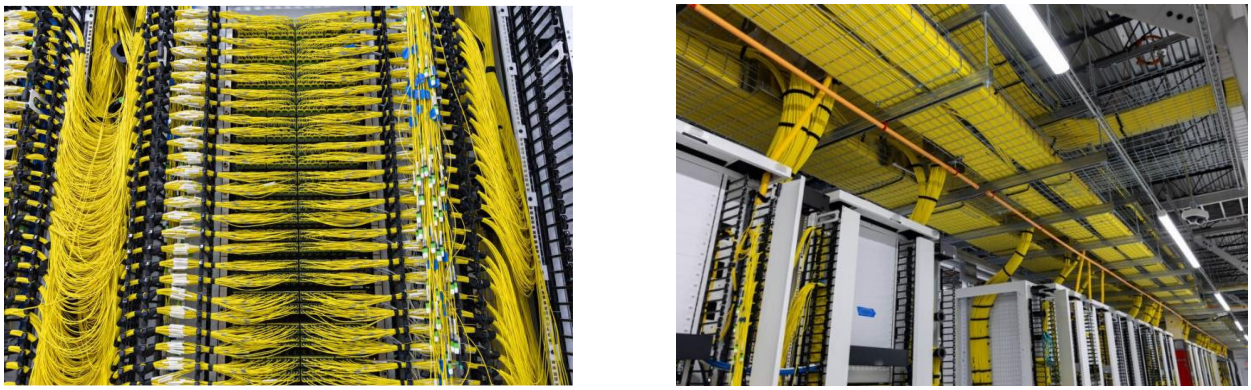

Figure 1. XAI Colossus Data Center

From data center to AI data center

The idea that all the optics required to perform the data transmission function must be integrated on a single silicon substrate has been abandoned by NVIDIA [2]. The reason is likely two-fold. First, the demands on electrical signal propagation, both in terms of signal integrity and energy required result in minimizing the distance the electrical signals must travel before being converted into optical signals. However the result is that so much heat producing circuitry is now co-packaged that liquid cooling is required. This can cause problems for temperature stabilized lasers. The Distributed Feed-Back (DFB) lasers, despite having an embedded grating to control the wavelength, are still sensitive to temperature. This second issue results in the lasers being isolated from the liquid cooled co-packaged optics and electronics.

From a performance standpoint, isolating the lasers is less problematic, as the lasers are CW sources; the high speed modulation is happening in the co-packaged optics. There is a price paid in terms of complexity because each laser is fiber coupled to the co-packaged optics. The architectural differences are shown in Fig. 2.

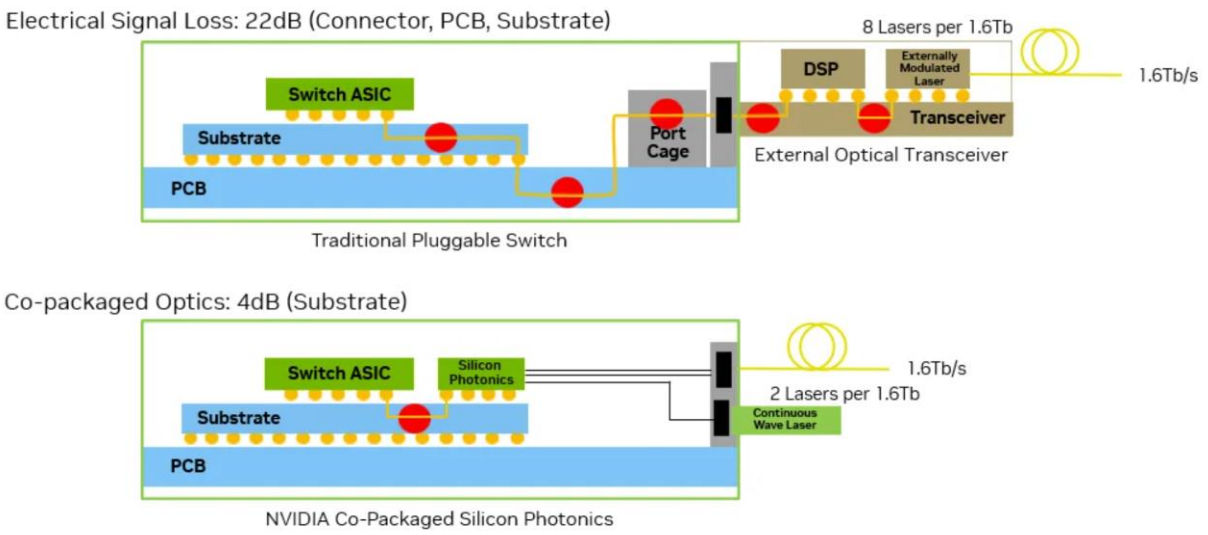

Figure 2. Separation of lasers for other optical functions in NVIDIA optical switch architecture. Red dots indicate electrical path.

NVIDIA Optical Switch

NVIDIA has posted a YouTube video [3] highlighting the various elements of their next generation optical switch. The specs are summarized in Table 1. Clearly, more than 100 Terabytes/sec is an impressive result, but consider that bandwidth management requires 400 billion transistors. The lasers, now moved to separate packages are only 8 lasers per package. Presumably this is limited by low yield from the process of flip-chipping laser die while maintaining optical as well as electrical alignment. Consider what would be required to achieve one Petabyte/sec: 256 lasers operating at 400 Gb/sec [4]. Even for 99% yield for the flip-chip alignment, the final yield is 0.99^{256} = 7.6%. Clearly, an alternative to hybrid packaging will be required.

Total Bandwidth	**115.2 Tbytes/sec**
Transistors	446 Billion
Multiple Push On Fibers	144 (~1000 fibers)
Lasers	144
External Laser Source packages	18
Lasers per External Laser Source	8

Table 1. NVIDIA optical switch specification summary

Nanorods

Nanorods have some unique features that make them desirable candidates for AI. Nanorod LEDs [5], vertically emitting lasers [6] and cleaved edge emitting geometries [7] have all been fabricated. Their small size makes them less susceptible to dislocations with dimensions < 1 μm, the dislocations to relax, allowing the growth of high quality Quantum Wells (QWs) on silicon substrates [8].

However, the nanorod diameter has been shown to be related to gain peak of the QW, allowing devices with a range of emission of nearly a factor of 2 to be fabricated [9]. This means that devices covering the entire telecom band can be fabricated on a single chip. Furthermore, small changes in diameter can be used to compensate for growth variability across a wafer.

Another feature of nanorod fabrication is that by changing the growth conditions, growing a columnar structure can be changed to growing a shell on the column. This allows nanorods to be passivated without breaking vacuum.

Lasers for AI

The fabrication process is shown in Fig. 3. Nanorods are positioned by a mask to be Bragg spaced, both in terms of their spacing and diameter. This results in constructive interference for the entire array, which can include hundreds on nanorods. Note there are no cleaved surfaces or mirror coatings required.

In addition to constructive interference, the nanorods all injection lock one another. This effect will suppress the amplitude of any cavity side modes [10].

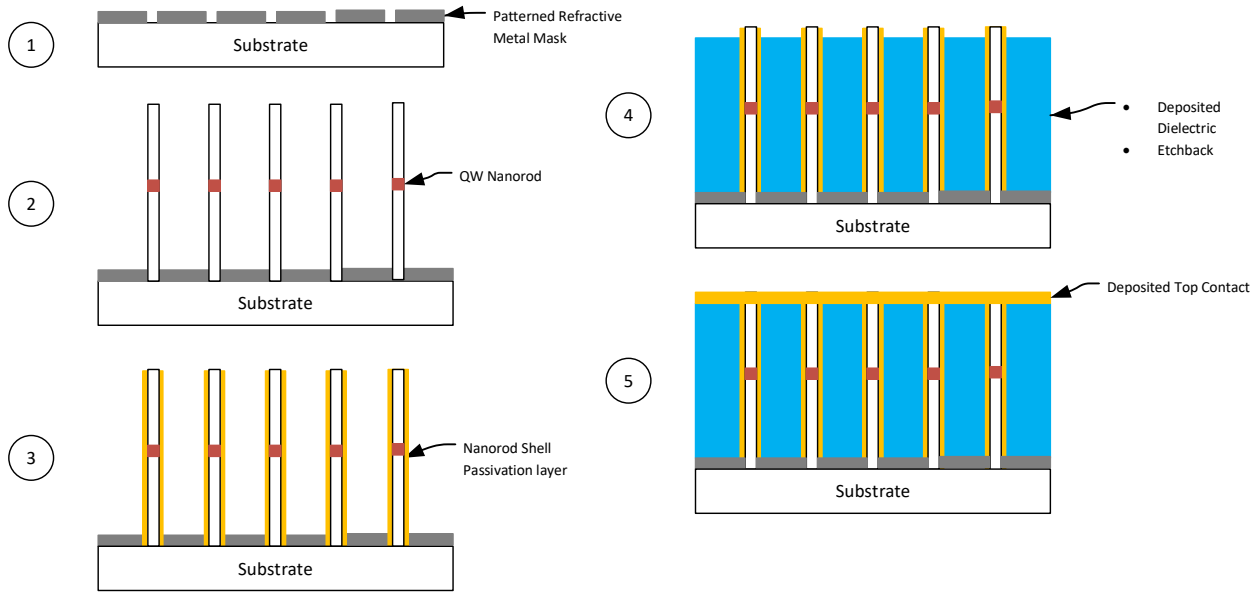

Figure 3. Process steps: (1) Patterned mask deposition. (2) QW nanorod growth. (3) Passivation shell. (4) Dielectric deposition/ and etchback. (5) Top contact deposition.

In conclusion, a straightforward design is presented for a linear array of quantum-well nanorod based in-plane laser using Bragg spaced nanorods that can be fabricated on a silicon substrate using conventional techniques. The superiority of this approach arises from its use of existing technologies and in obviating the need for hybridization, cleaving or mirror coatings. Enhanced performance arises from the injection locking effect not present in other geometries. Its generic applicability to any III-V material system allows for potentially wide application. The approach allows for large numbers of lasers to be fabricated on a single chip.

References

1. supermicro.com/CaseStudies/Success_Story_xAI_Colossus_Cluster.pdf
2 . https://developer.nvidia.com/blog/scaling-ai-factories-with-co-packaged-optics-for-better-power-efficiency/
3. https://www.youtube.com/watch?v=kS8r7UcexJU
4. Yikai Su, Yu He, Ting Wang, Xuhan Guo, Xi Xiao, Guo-Qiang Lo, Appl. Phys. Lett. 125, 230502 (2024).
5. Y. Wu, Y. Xiao, I. Navid, K. Sun, Y. Malhotra, P. Wang, D. Wang, Y. Xu, A. Pandey, M. Reddeppa, W. Shin, J. Liu, J. Min, and Z. Mi, "InGaN micro-light-emitting diodes monolithically grown on Si: Achieving ultrastable operation through polarization and strain engineering," Light, Sci. Appl., vol. 11, no. 1, p. 294, Oct. 2022.
6. S. Arafin, X. Liu, and Z. Mi, "Review of recent progress of III-nitride nanowire lasers," J. Nanophotonics, vol. 7, no. 1, Sep. 2013, Art. no. 074599.
7. E. Stark, T. Frost, S. Jahangir, S. Deshpande, and P. Bhattacharya, "A monolithic electrically injected InGaN/GaN disk-in-nanowire (λ=533nm) laser on (001) silicon," in Proc. IEEE Photon. Conf., San Diego, CA, USA, Oct. 2014, pp. 591–592.
8. S. Deshpande, I. Bhattacharya, G. Malheiros-Silveira, K. W. Ng, F. Schuster, W. Mantei, K. Cook, and C. Chang-Hasnain, "Ultracompact position-controlled InP nanopillar LEDs on silicon with bright electroluminescence at telecommunication wavelengths," ACS Photon., vol. 4, no. 3, pp. 695–702, Mar. 2017.
9. H. Sekiguchi, K. Kishino, and A. Kikuchi, "Emission color control from blue to red with nanocolumn diameter of InGaN/GaN nanocolumn arrays grown on same substrate," Appl. Phys. Lett., vol. 96, no. 23, Jun. 2010, Art. no. 231104.
10. D. R. Dykaar, H. Ghosh, S. Zhao and S. Sivoththaman, "Bragg-Spaced Quantum-Well Nanorod In-Plane Lasers on Silicon," in IEEE Electron Device Letters, vol. 45, no. 5, pp. 789-792, May 2024.