Inter-DC and Intra-DC Networking for AI

OptiNet 2025

Song Jiang

June 2025

AI Use Cases Driving Investment and Technology Transitions

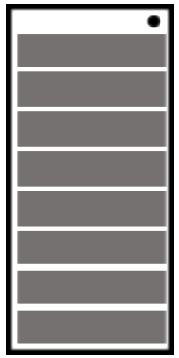

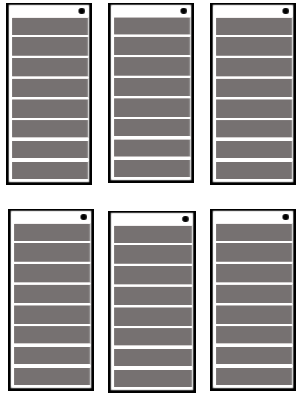

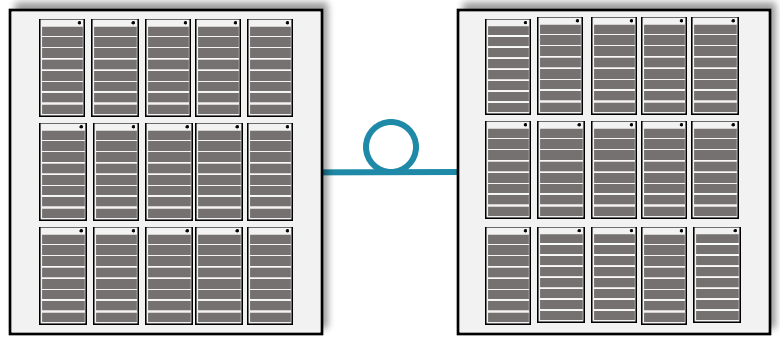

Scale up

- Power efficient optics to replace copper
- Fast vs wide
- LPO/TRO/CPO
- Liquid cooling

Scale out

- Optics need to keep pace with increasing silicon scaling
- High BW material systems
  - TFLN, D-EML, SiPh, InP, etc.
  - Scale speed, not cost
- LPO/TRO/CPO

Multi-DC Training

- Driven by datacenter power limitations
- Campus/Regional Distances
  - Pb/s connectivity
- Coherent-lite
- Power optimized coherent

Acacia

Market Opportunity

- 800G modules will ship in high volume for several years
  - Mostly 100G/lane with some adoption of 200G/lane
- 1.6T modules are early in their lifecycle
- In 2029, the total market for these two speeds is expected to be more than $15B

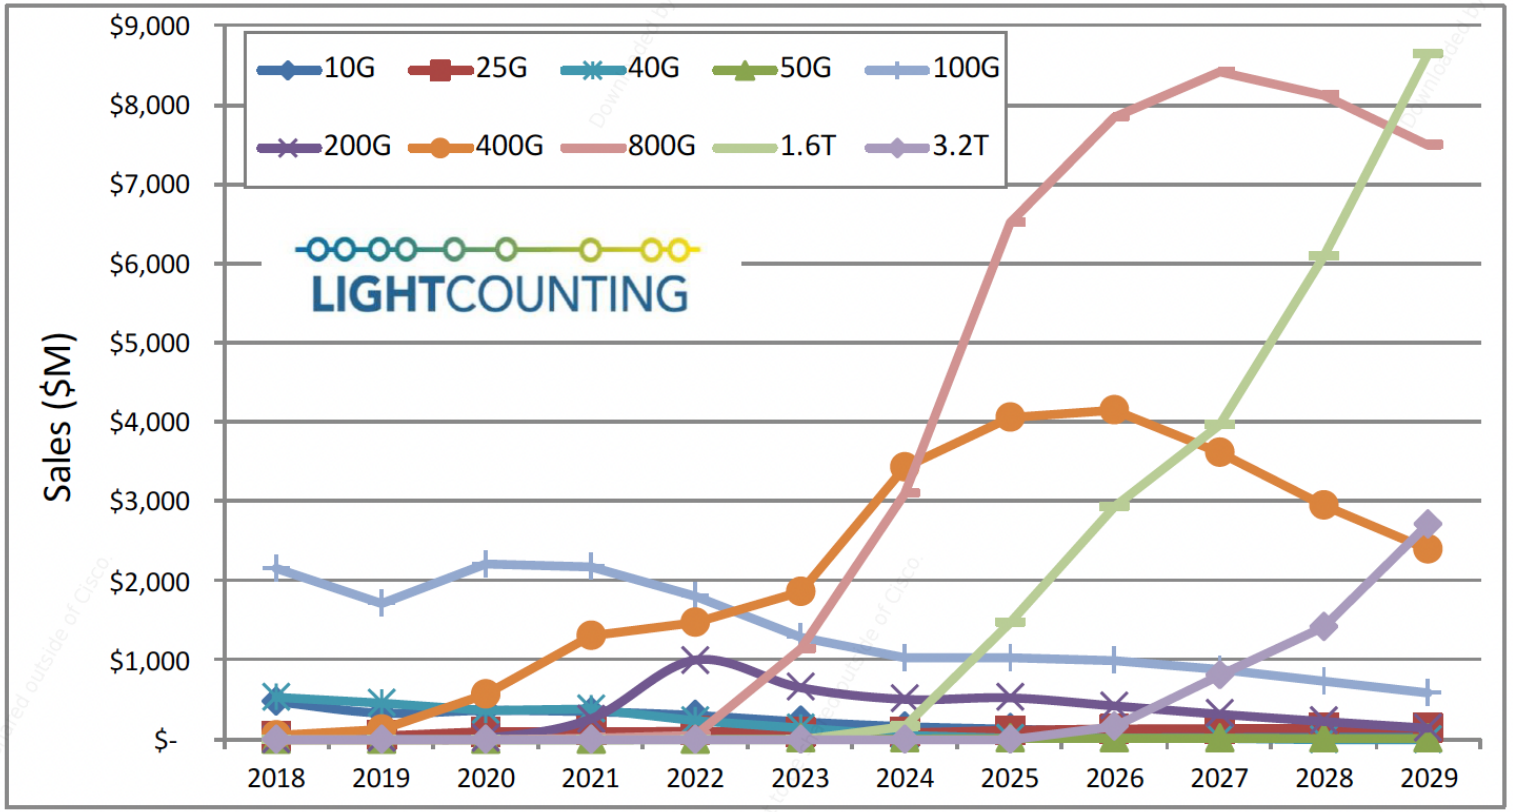

Source: LightCounting

AI Driving Multi-Datacenter Connectivity

- Multi-building campus architectures enable staged build-out
  - New use cases for coherent-lite technology
    - Some require DWDM and others are grey
      - Both O-band and C-band applications
    - Latency is important when fiber runs are a few km
- Training cluster size exceeding power available in single location
  - Synchronous training across geographically diverse datacenters
    - Industry innovation enabling ability to manage latency
  - Multi-Pb/s connectivity may be necessary
    - Significant fiber plant build-outs
    - Power efficient coherent technology is critical
    - Multi-rail line amplification

AI datacenter Examples

Stargate AI Data center, TX	
Total Capacity	1.2GW
Area	900 Acres
GPU	~400K GPU

UAE-US AI Campus, Abu Dhabi	
Total Capacity	5 GW
Area	6400 Acres
GPU	~2.5 million GPU

Coherent Moving to Shorter Reaches

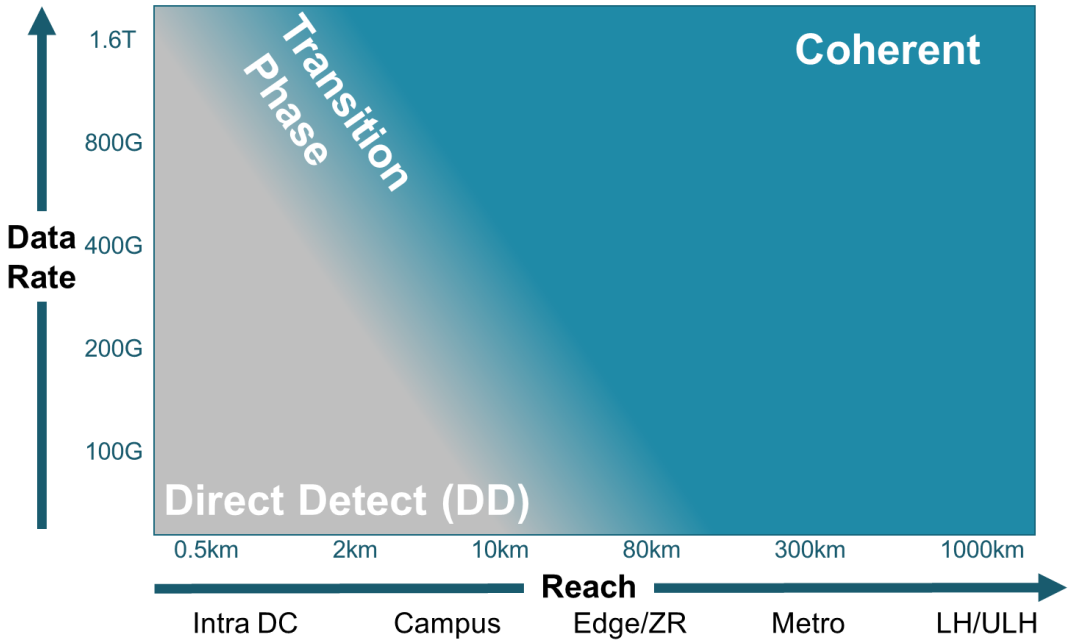

For similar use case, i.e. 800G LR, coherent power is ~20% higher than IMDD

- Primary motivations to move to coherent are chromatic dispersion and WDM
  - PAM4 chromatic dispersion limited reach reduced by a factor of 4 when baud rate is doubled
- IEEE defining 10km interfaces at 800G for both PAM4 and coherent
- At 1.6T, coherent will be needed beyond a few km
  - PAM4 will support the highest volumes (<2km)
- AI build-outs creating new use cases for coherent-lite, but driving unique requirements
  - OIF has a project to define 1.6T coherent-lite with multiple use cases in scope

Coherent-lite Use Cases

Traditional IEEE LR/ER

- LR
  - Service provider use case to overcome patch panel loss
  - Latency less important
- ER
  - Higher loss and chromatic dispersion drives earliest need for coherent
  - Latency not important
- Lower volumes than <2km
  - Margin reduces need to know fiber plant details

Optical Cross-connect

- Short fiber reaches, but higher loss budget (>LR)
- O-band may be required
- Latency sensitive
- Power sensitive
- Creative PAM4 solutions may continue to compete in these use cases
- Potential for high volumes

Campus Interconnect

- AI buildouts moving to multi-data center configurations
- Reaches in the 2-20km range
- Some use cases benefit from WDM capability
- Potential for high volumes

Diverse use cases have led to fragmented coherent-lite implementations despite industry standardization efforts

IMDD vs Coherent in Client Modules

IMDD | IMDD/Coherent | Coherent

Ethernet Rate	Optical Lanes	500m DR	2km DR-2/FR 4dB	10km LR 6.3-8dB	40km ER 12-18dB	80-120km ZR	+500km LH "ZR+"
400G	4x100G	400GBASE-DR4	400GBASE-DR4	400GBASE-LR4	400GBASE-ER4-30		
	1x400G						
800G	8x100G	800GBASE-DR8					
	4x200G	800GBASE-DR4	800GBASE-FR4	800GBASE-LR4	800GBASE-ER1		
	1x800G			800GBASE-LR1			
1.6T	8x200G	1600GBASE-DR8					
	4x400G						
	2x800G						
	1x1600G		Coherent-Lite				

- Low latency
- Low Power
- Simplified DSP
- O-band SiPh

Acacia Expands Client Optics Component Business

New 1.6Tbps Kibo PAM4 DSP

- 8x 200G/lane for host and optical interface
  - Or Gearbox for 800G - 8x100G to 4x200G
- 3nm CMOS node for market leading power efficiency enabling more than 20% lower power compared to existing 1.6T module implementations
- Industry standards compliance enhanced by Acacia algorithms enabling higher performance and lower power
- Transmit Retimed Optics (TRO) configurations with power-efficient support for diagnostic and loopback troubleshooting capabilities
- Support for gearbox and retimer applications
- Designed for 1.6T DR8/2xFR4 modules in OSFP/QSFP-DD form factor

Acacia Expands Client Optics Component Business

New 200G/lane Silicon Photonics Optical Engine (OE) Family

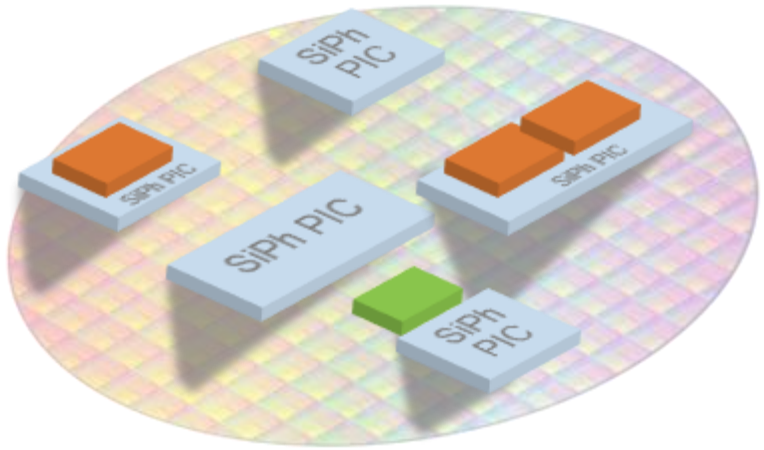

Acacia's Optical Engine Product Family supports a variety of design configurations based on customer requirements

- Builds on success of 100G/lane optical DR4 optical engines
- TX OE for DR4, DR8 and 2xFR4
- Separate transmit and receive components
- Support for 100G/lane and 200G/lane
- Flexible driver configuration support
- High-performance receivers with stacked-die TIA configuration
  - Enable high-performance LPO

Connecting at the speed of light