Scaling AI Networks with Multicore and Hollow-Core Fiber

Jose M. Castro, Bulent Kose and Paul Huang
Fiber R&D Team
All affiliated with Panduit

IEEE 802.3 New Ethernet Applications Ad Hoc Ethernet for AI Assessment
February 2026

Background

- AI workloads (training and inference) demand increasing computational throughput, which requires faster communication at different network layers: scale-up, scale-out, and scale-across.
- IEEE 802.3 focuses on developing PMDs that are reaching 200G/lane and perhaps even 400G/lane this decade.
- Those data rates currently enable 400G scale-out bandwidth per GPU and are moving toward 800G per port.
- However, challenges related to density, fiber bandwidth, and latency may drive the adoption of “new” fiber types.

Scaling AI Networks: Up, Out and Across

🟩 Scale up switch
🟦 Scale out switch
⬛ Scale across switch

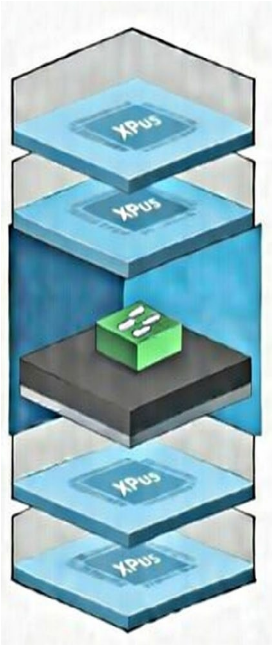

Scale Up

Bandwidth: Very High
Density: Critical
Latency: Extremely Low
Switches: 1 Layer
Media: Copper: *twinaxial*

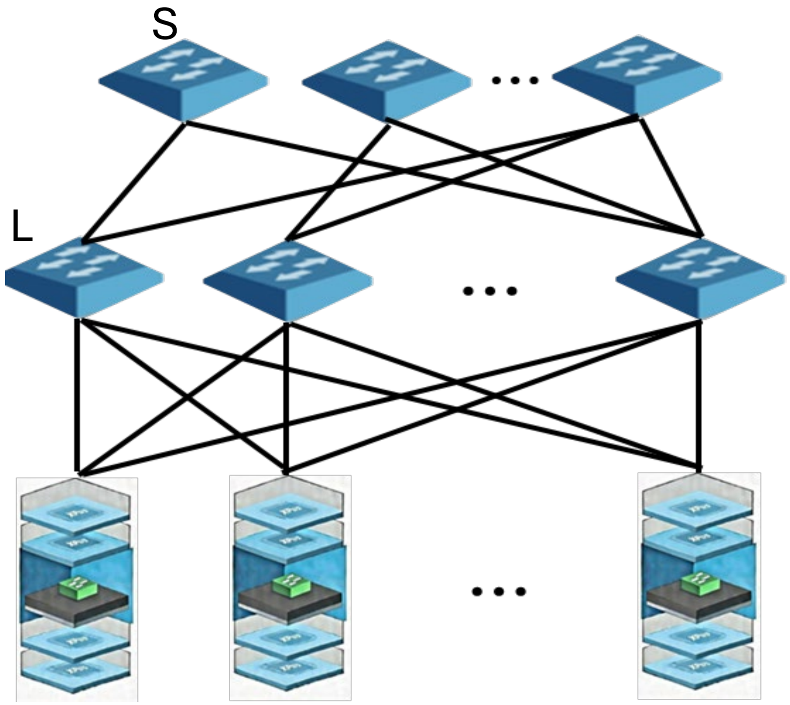

Scale Out

Bandwidth: High
Density: Important
Latency: Low
Switches: 2-3 Layers (shallow buffers)
Single Mode Fiber (SMF) G652D, G657.A1/A2
Multimode Fiber: MMF (OM3/OM4/OM5)

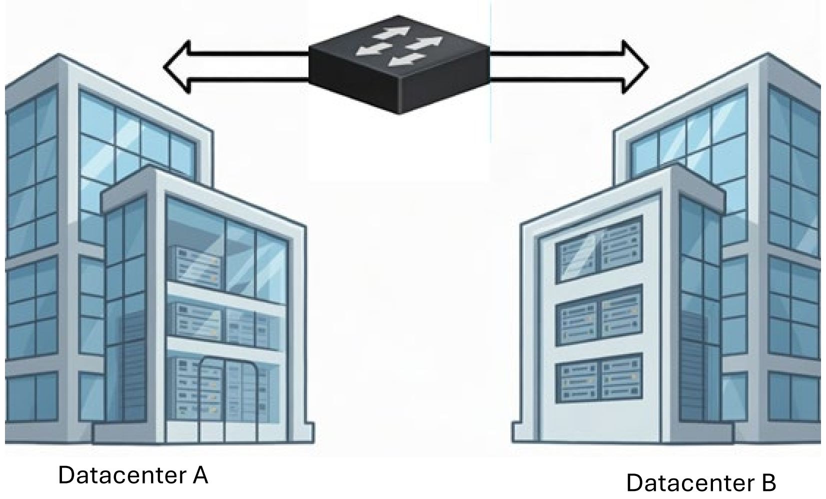

Scale Across

Bandwidth: High
Density: Important
Latency: High
Switches: deep buffers
Single Mode Fiber : G652D, G657.A1/A2

AI Scaling Through Fiber Innovations

Media Challenges

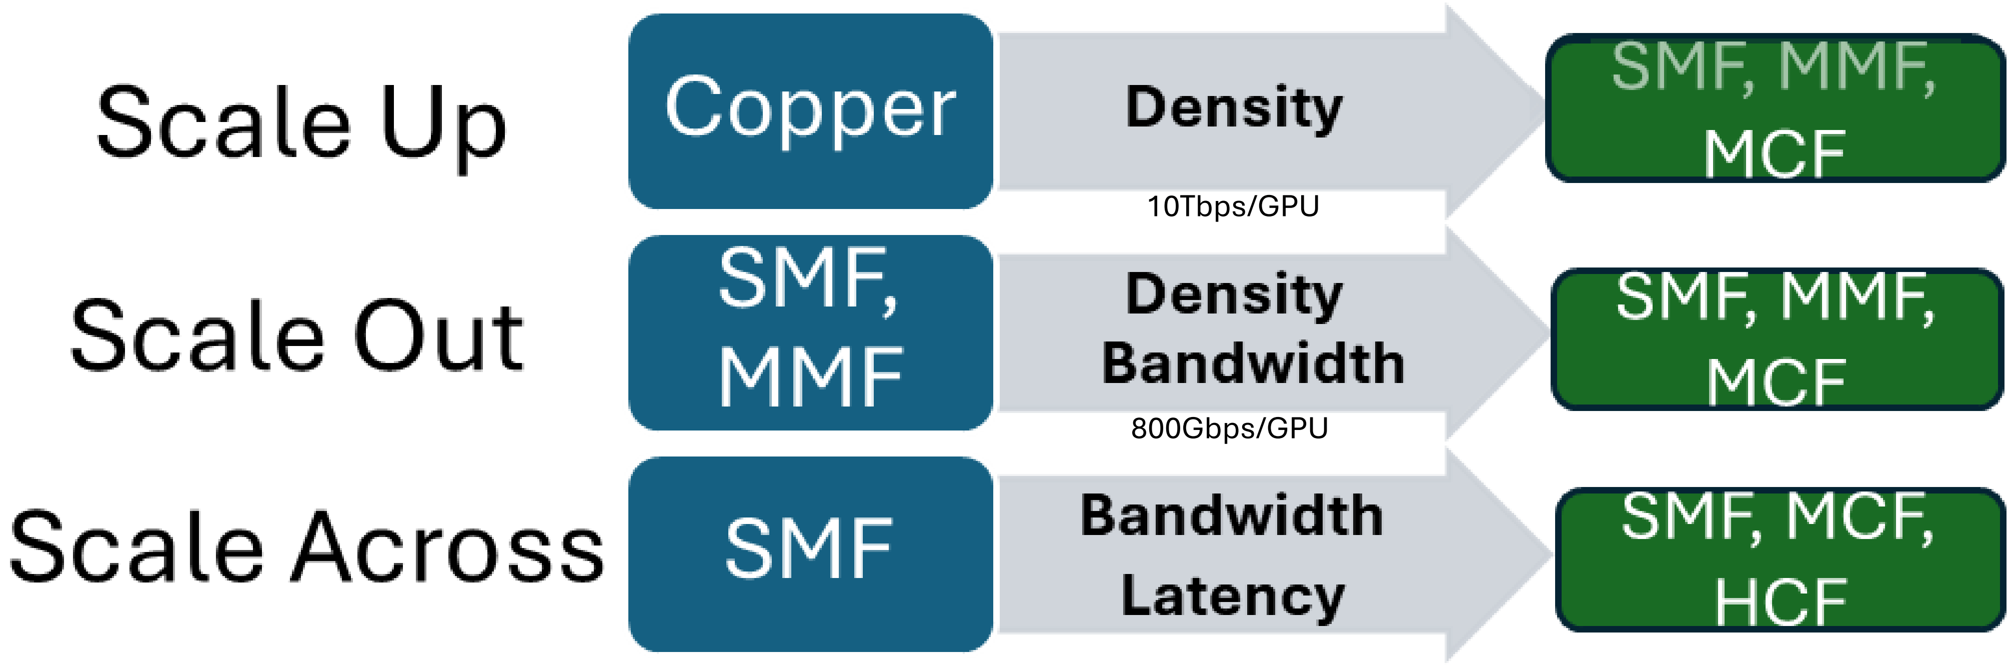

Scale Up	Copper	**Density** 10Tbps/GPU	SMF, MMF, MCF	Transition to optics improves bandwidth, supports disaggregation, reduces memory wall
Scale Out	SMF, MMF	**Density Bandwidth** 800Gbps/GPU	SMF, MMF, MCF	MCF can simplify deployment and cabling management, as well as improve fiber bandwidth
Scale Across	SMF	**Bandwidth Latency**	SMF, MCF, HCF	Interconnecting huge clusters of million-XPUs systems, with high bandwidth and low latency

Note that Multicore fiber can be MMF or SMF MCF

Optics for Scale up

- The number of XPUs in a single scale-up system expected to grow from today’s ~72 devices to many thousands.

- Placing optical engines close to the XPUs and switch ASICs—e.g., co-packaged optics (CPO)—can expand the scale-up network avoiding the need for Megawatt racks.

- MCFs can simplify deployment of those large scale up networks.
  - ≈4K fibers for a 512 port x 800G switch.
  - ≈10K fibers for a NVL72 node .

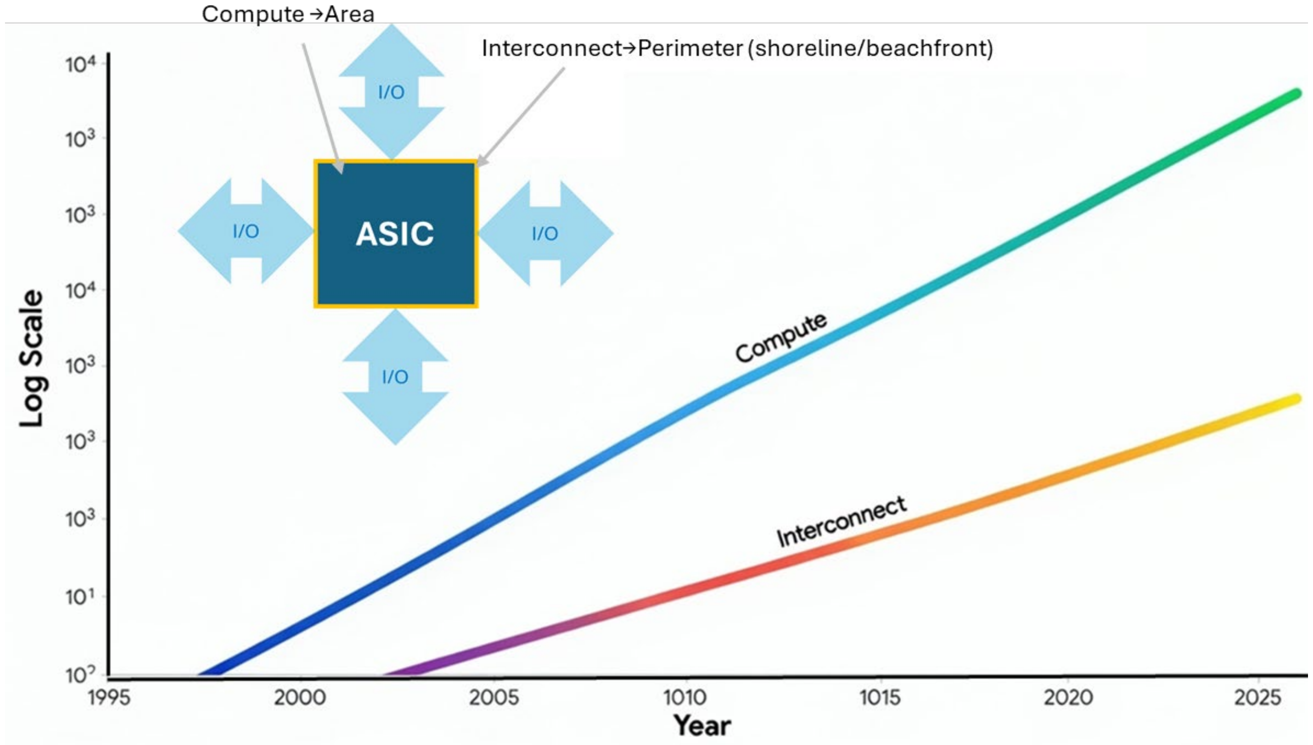

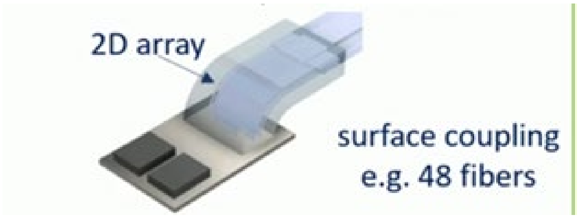

OFC 2025, https://www.youtube.com/watch?v=-PuhRgmTAYc&t=1271s

Multicore fiber types

- Multicore fiber (MCF) can range from 2–4 cores (lowest crosstalk; long-haul/subsea interest) to tens of cores (maximum spatial density)
- For data centers, 4 cores MCFs keep 125 µm cladding and ~200–250 µm coating diameter while limiting crosstalk (XT) which depends on core pitch and core/trench design.
- Core designs that reduce XT must be fully compatible with standard fiber G.652D, G657 to produce chromatic dispersion statistics similar to the ones adopted in 802.3dj models.

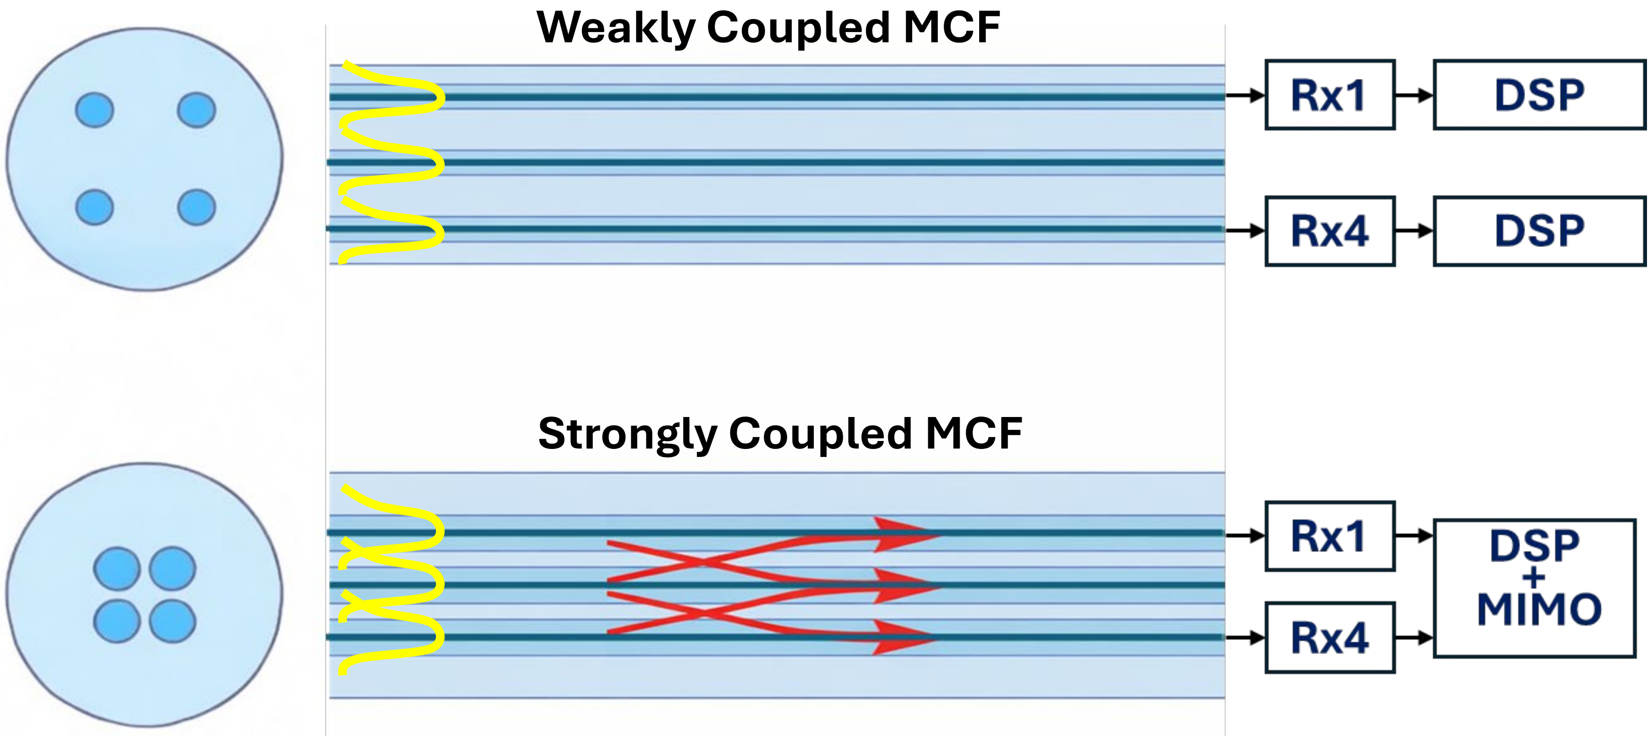

Design Considerations For MCF

Parameter	Being Considered
Similar fiber type	G.657 compliant
Number of cores	4
Cladding diameter	125 µm
Coating diameter	200 µm / 250 µm
Attenuation	< 0.4 dB /km
Mode Field Diameter MFD	8.6 µm- ~9.2 µm at 1310 nm
Chromatic dispersion	G.652.D G.657A1 compliant
Typical crosstalk	≤ −40 dB @1310 nm or @1550 nm (10 km)
Typical fiber loss	≤ 0.35 dB/km @1310 nm ≤ 0.20 dB/km @1550 nm
Macrobending Loss	≈ Conventional G.657.A1
PMD	≈ Conventional G.657.A1

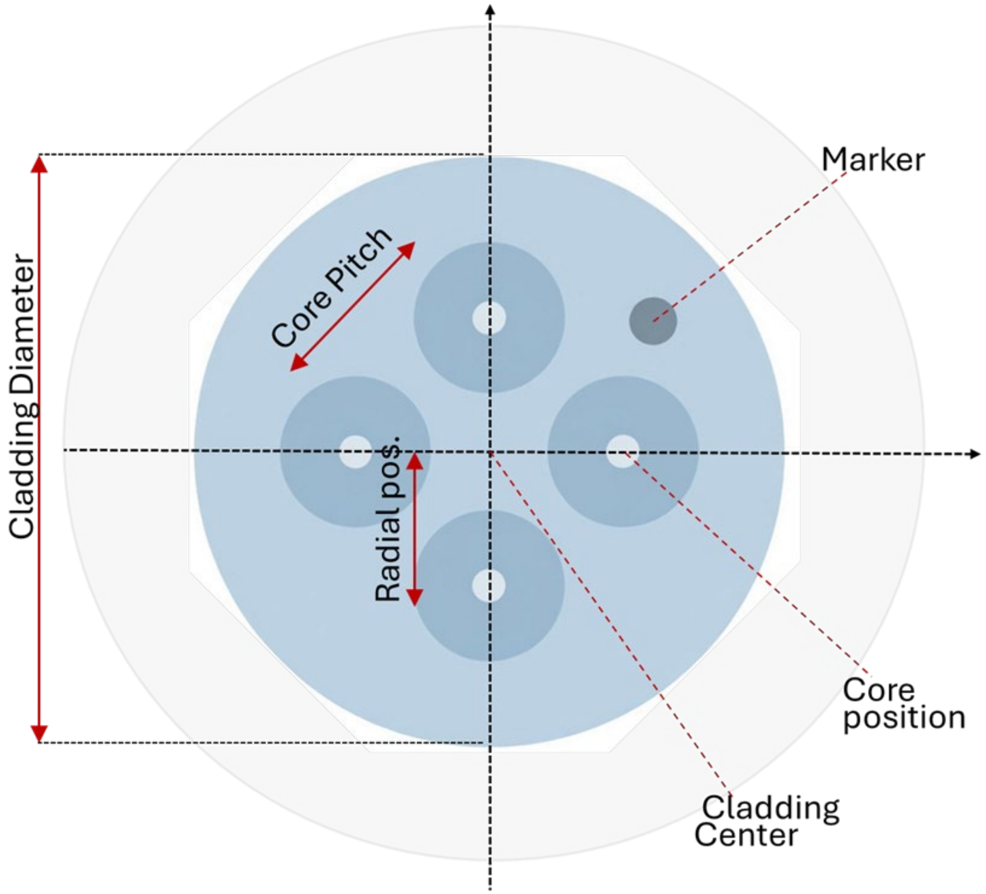

Notes:

- Actual parameters being evaluated in ITU-T G650.x
- Test conditions to be defined, e.g., crosstalk vs Macrobending
- Currently only focus on SMF

MCF for High-Data-Rate Scale-Out AI Networks

- Data center campuses are expanding, and per-lane (per-wavelength) data rates are rising—together, they push more links into a reach-limited regime for multi-wavelength optics.
- Duplex solutions (FR4/LR4) have covered >500 m, but at 400G/λ chromatic dispersion limits practical reach around ~600 m *.
- Covering those paths may require single-wavelength, multi-fiber PMDs (e.g., DR4-2 / DR8-2)—which increases fiber count.
- However, more fibers mean larger trunks and higher deployment complexity.
- Single-wavelength over MCF can potentially extend to ~2 km without increasing cable size

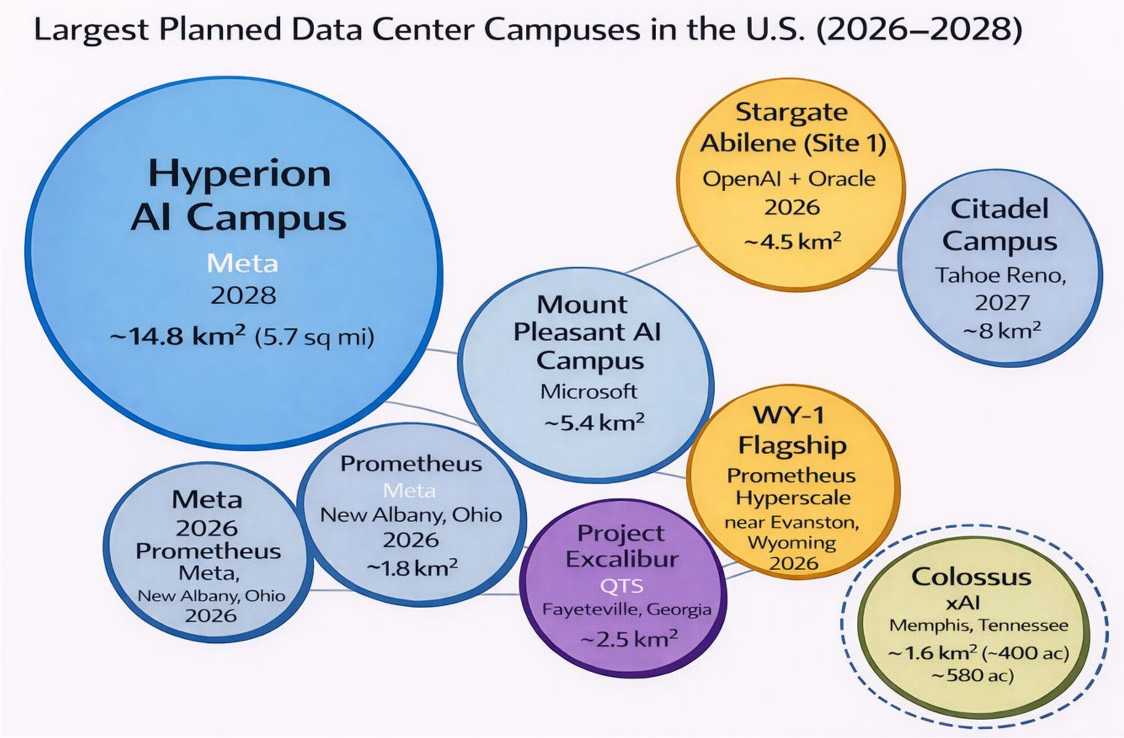

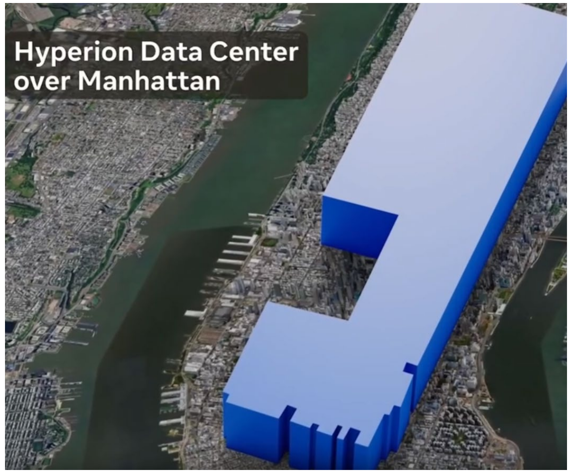

* https://www.ieee802.org/3/dj/public/25_05/yu_3dj_01_2505.pdf

Options for Next-Generation DR MCF Transceivers

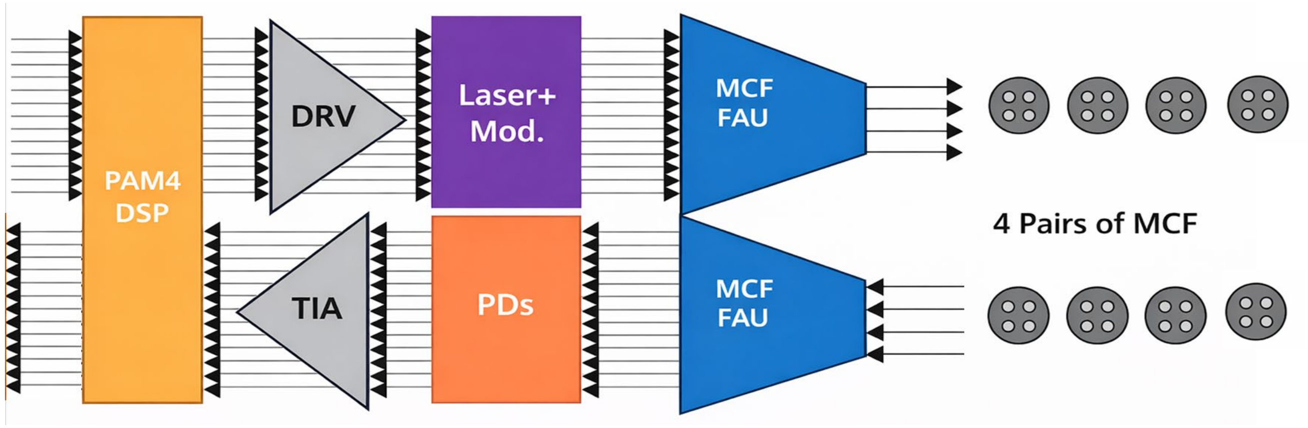

3.2 T (400G per MCF)

Functional Diagram *

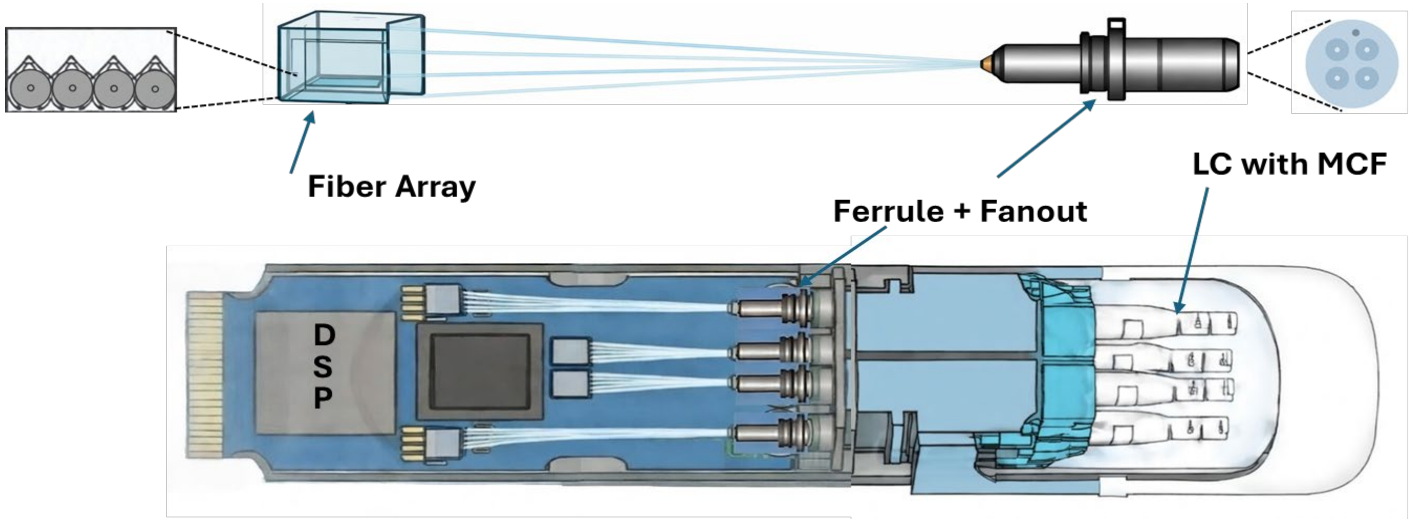

1.6 T (400G per MCF)

Implementation Option **

---

*https://www.ieee802.org/3/dj/public/25_05/yu_3dj_01_2505.pdf

** Image inspired in TeraHop MCF Transceiver Demo (OFC 2025)

MCF Core Fan-In/Fan-Out (FIFO)Technologies

- Links using MCF transceivers can connect directly, without lane breakouts.
- When breakout is required (e.g., shuffle or reroute channels), FIFO modules couple MCF cores to individual SMF or MMF fibers.
  - However, FIFOs increase insertion loss (>0.5 dB)
- Both use cases require new polarity schemes to be standardized within TIA or IEC.

Form Factor Example

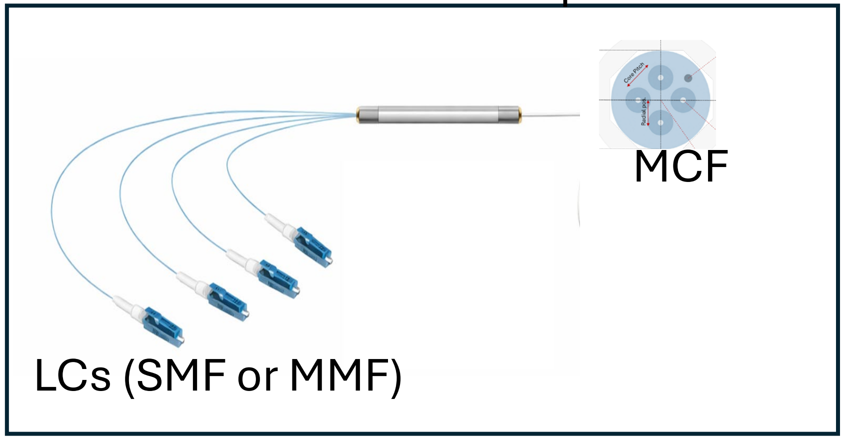

Adiabatic Fused

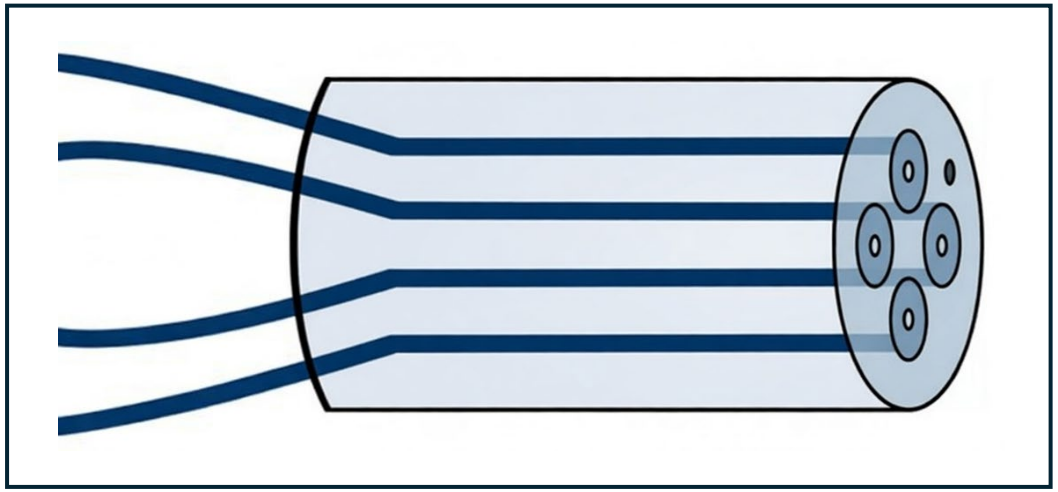

3D Direct Laser Written Waveguide

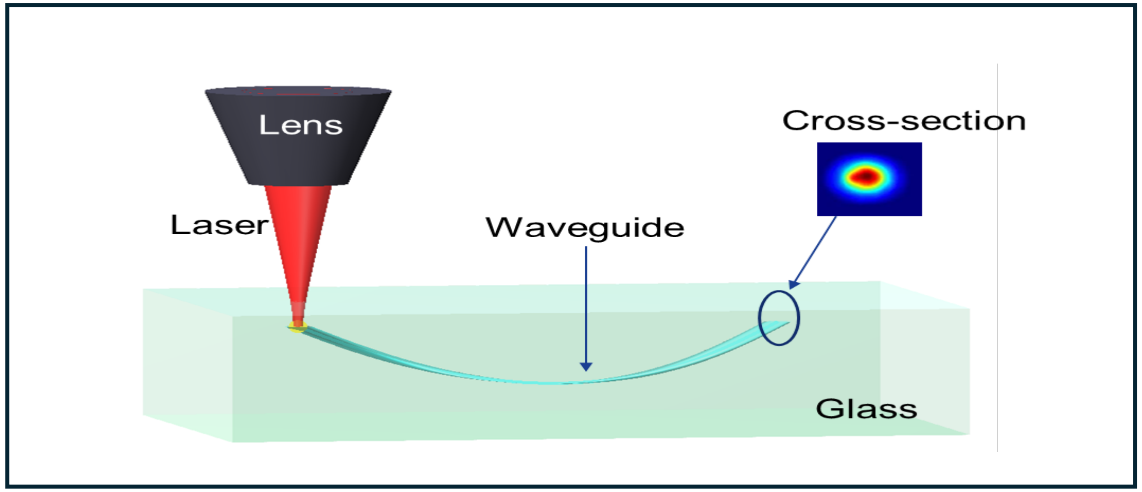

Crosstalk: Wavelength, Macrobend, and Contamination Impact

Red: Macrobend 1 turn 5mm bend radius

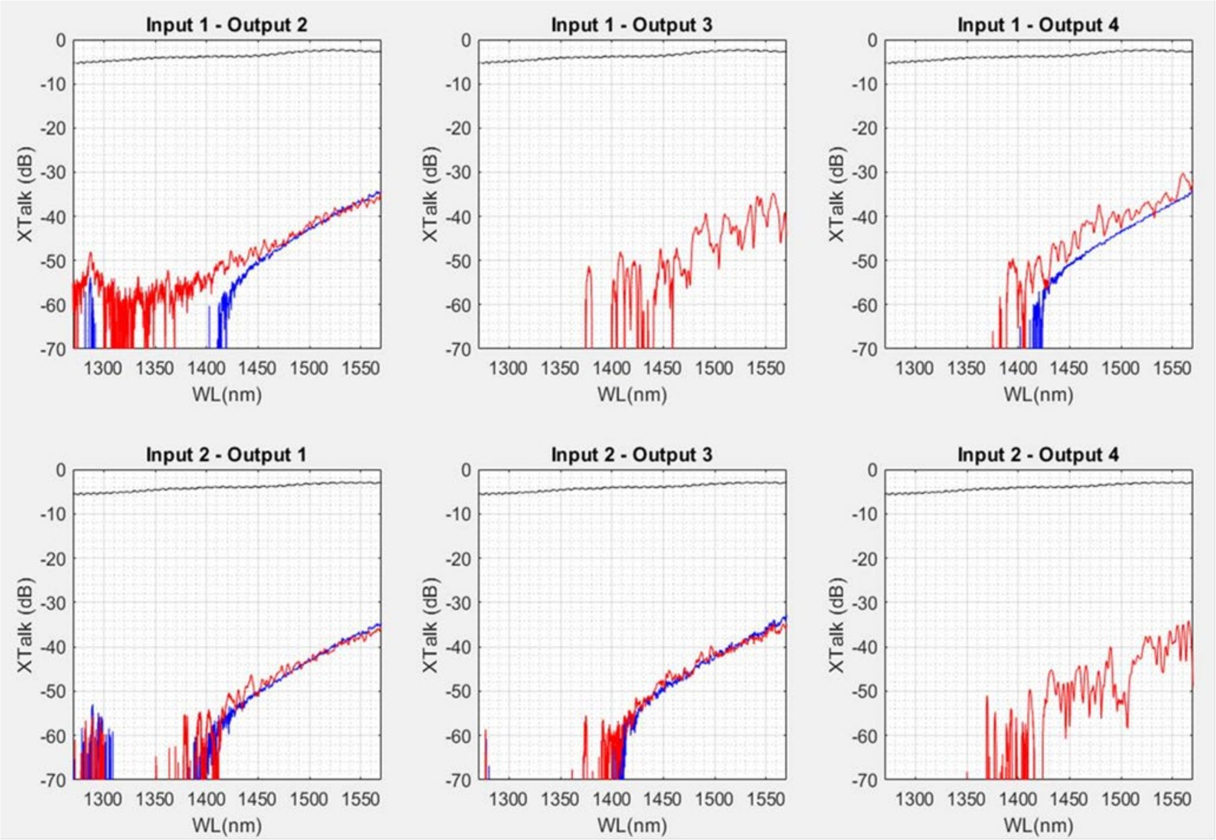

Panduit Fiber R&D Lab measurements of a 4-core MCF (vendor undisclosed)

Contamination can be more impactful in MCF and may require updates to end-face visual inspection standards (e.g., IEC 61300-3-35), since it not only blocks light and increases loss but also scatters or refracts light, increasing crosstalk and multipath interference.

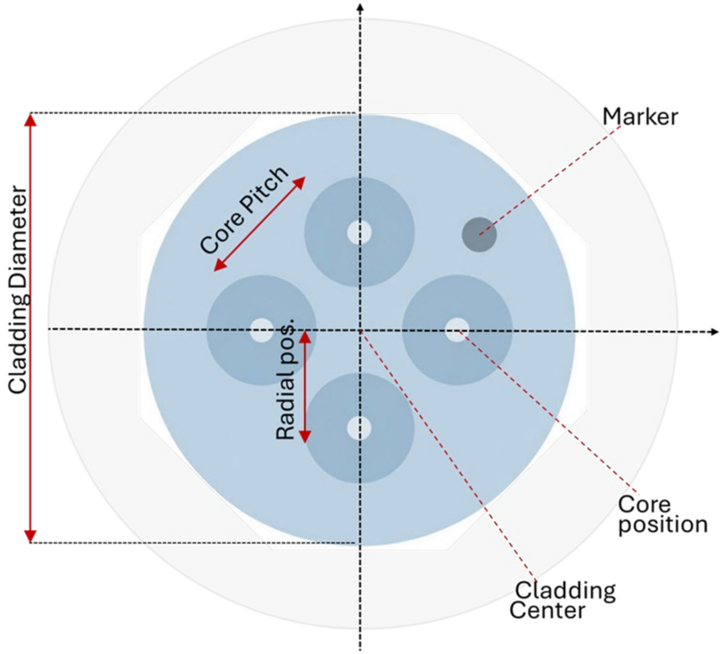

Scale Across with Hollow Core Fiber

Scaling Across: Benefits & Constraints

- Scaling XPUs is critical for AI performance, as it enables complex models and faster computations in training and inference.
- Concentrating 100k+ XPUs in one site risks energy/water constraints, operational hazards, and total investment loss from disasters or outages.
- Distributing data centers across regions allows scaling to 1M XPUs and mitigates single-point failures.
- However, low-latency coherent operation is challenging for training due to propagation latency.

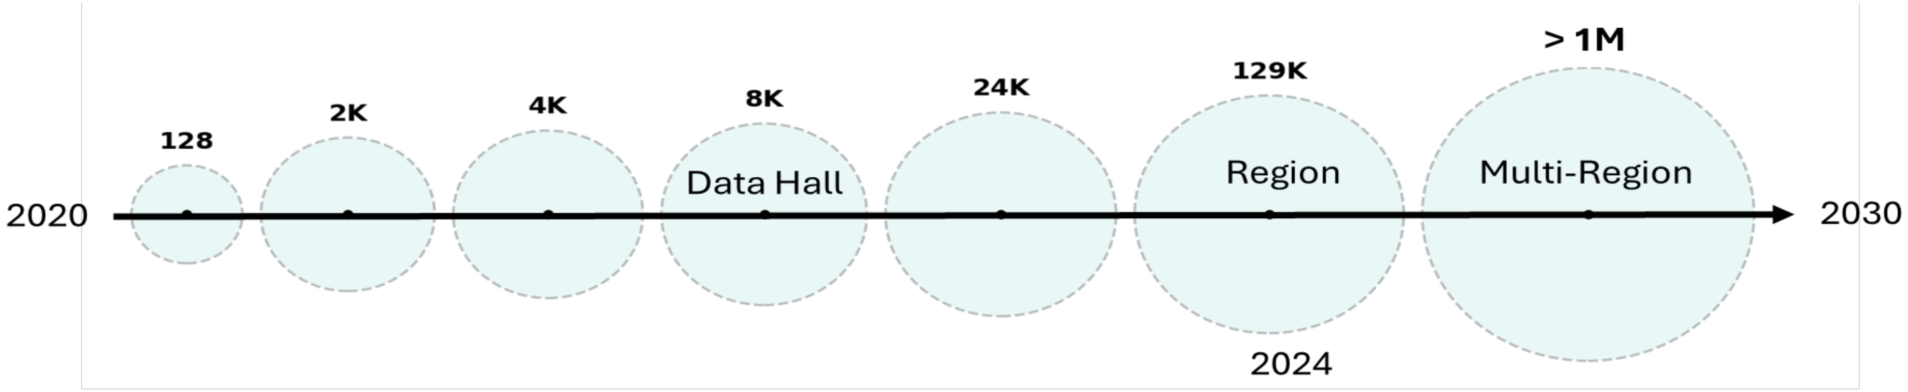

“HCF is absolutely critical for data-center interconnectivity… Microsoft to add ~15,000 km in 24 months”
— Satya Nadella, Ignite 11/24

Figure based on “Scaling the AI Infrastructure to Data Center Regions, presented by Meta https://www.youtube.com/watch?v=P5c-LOQ1pVE&list=PPSV”

What is Hollow Core Fiber ?

- A type of fiber in which the glass core—the region where light propagates—is replaced by air.
  - Light propagates 50% faster in air ($n_{HCF} < 1.001$) than glass ($n_{SMF} = 1.47$) significantly reducing latency (33%)
- Other benefits include less light interaction with glass, which reduces nonlinearities (enabling higher power transmission) and dispersion (enabling higher bandwidth).

Traditional Optical Fiber

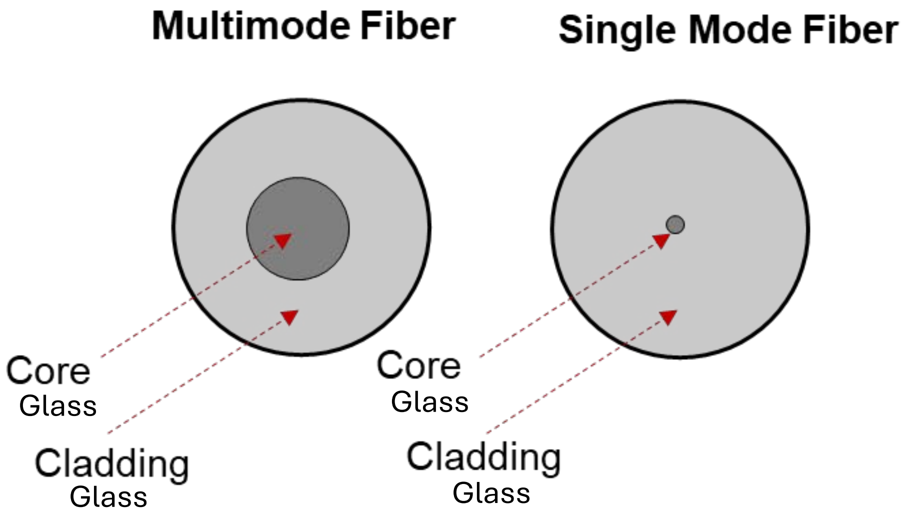

80s-90s

Hollow Core Fiber

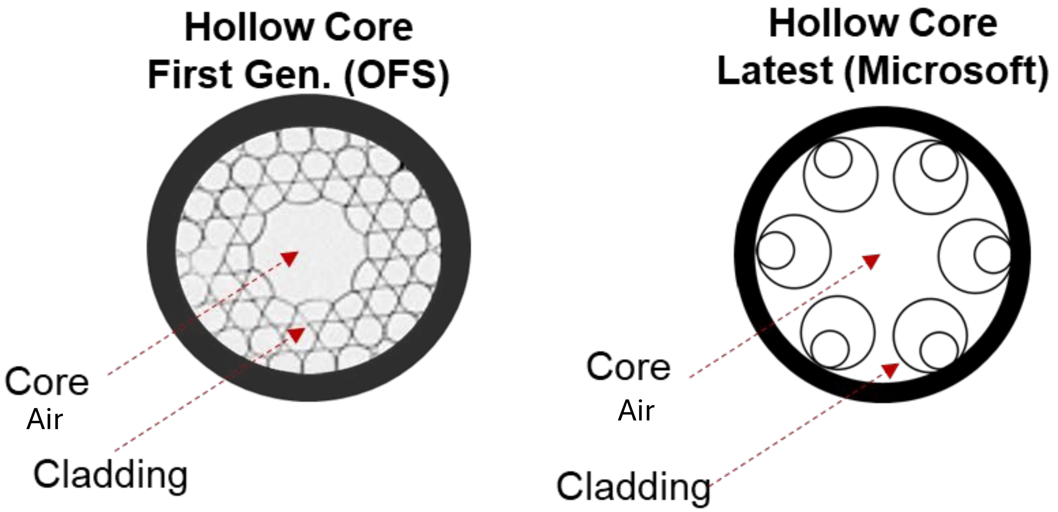

2010s

BandGap

2017

Antiresonant

HCF: Recent Attenuation Milestones

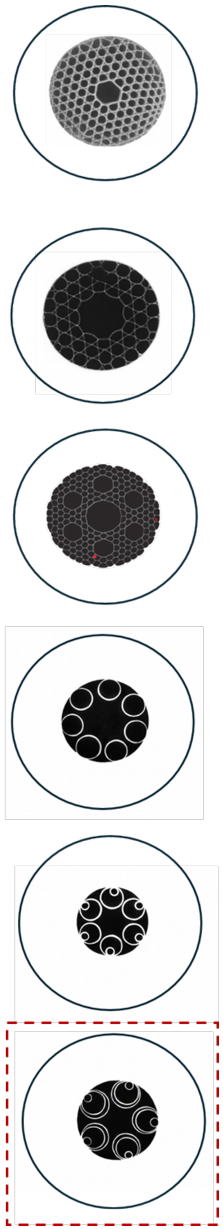

- In our lab we have seen rapid HCF loss improvement:photonic bandgap fibers measured ~6 years ago were >1 dB/km, while antiresonant fibers measured a few years later were ~0.5 dB/km.
- Recent NANF and D-NANF hollow-core fibers from teams like the University of Southampton or YOFC have achieved attenuation approaching conventional SMF.

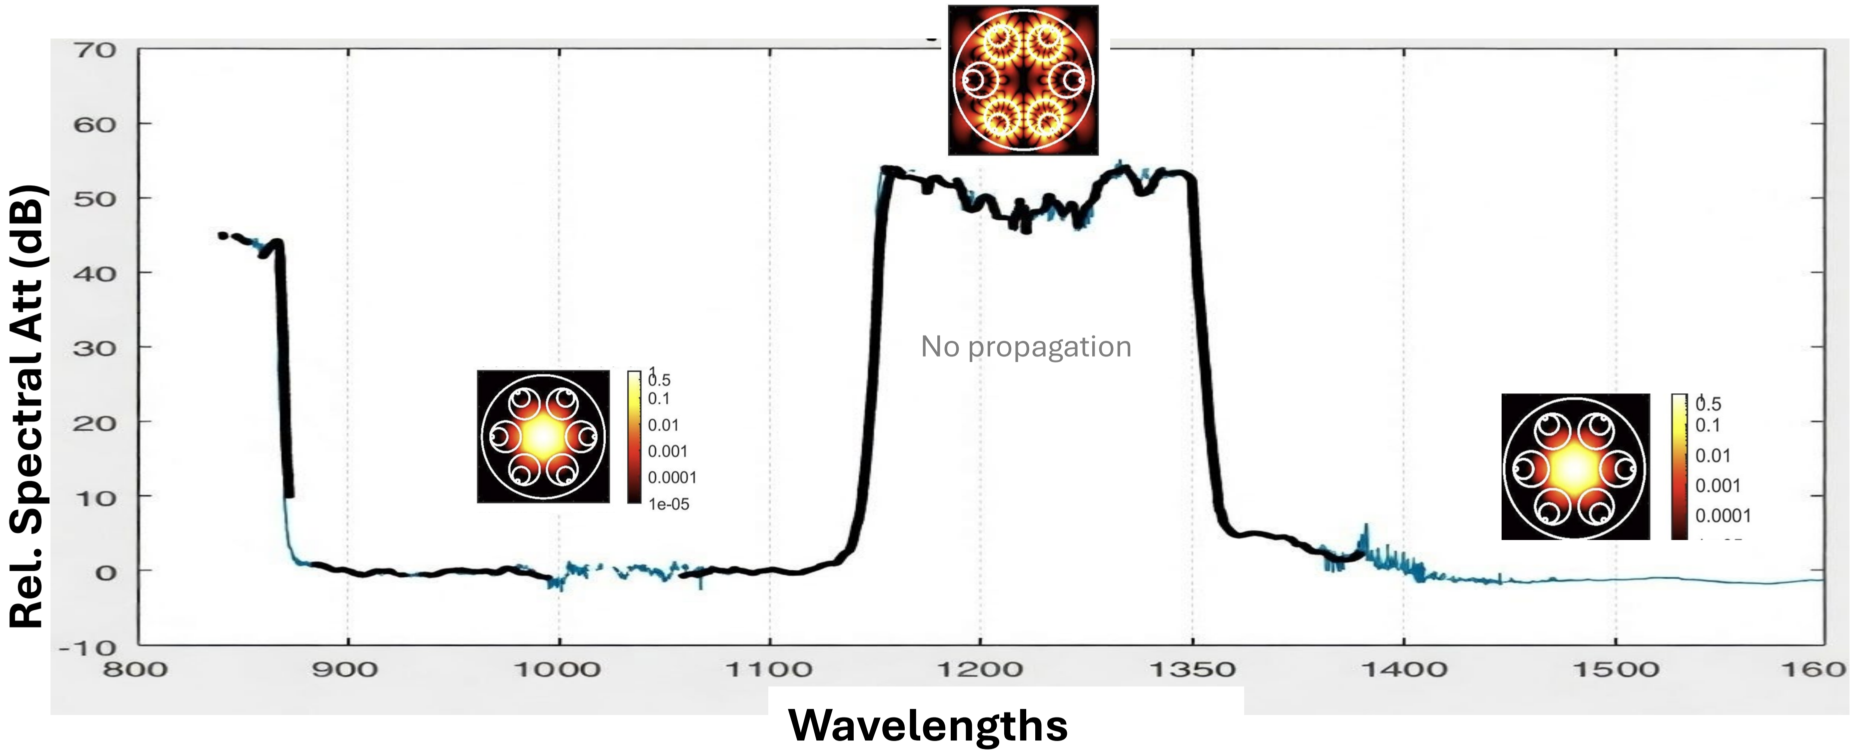

*All figures from Panduit internal modeling and measurements

Chromatic Dispersion Measurements

- The weaker light–glass interaction in HCF enables lower chromatic dispersion. For this specific sample, |D(λ)| is within about ±5 ps/(nm·km) over 1250–1600 nm (vs ~−6 to +20 ps/(nm·km) for G.652D SMF).
  - Note that CD increases sharply near the bandgap edge (~1200–1250 nm).
- To support standard PMDs the fiber must be tuned so the low-dispersion window covers the O and C–L bands.

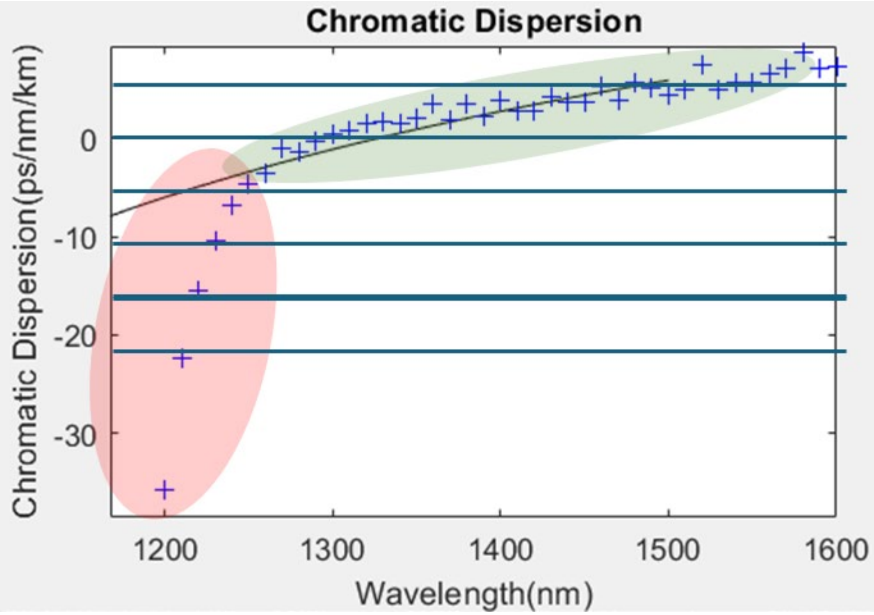

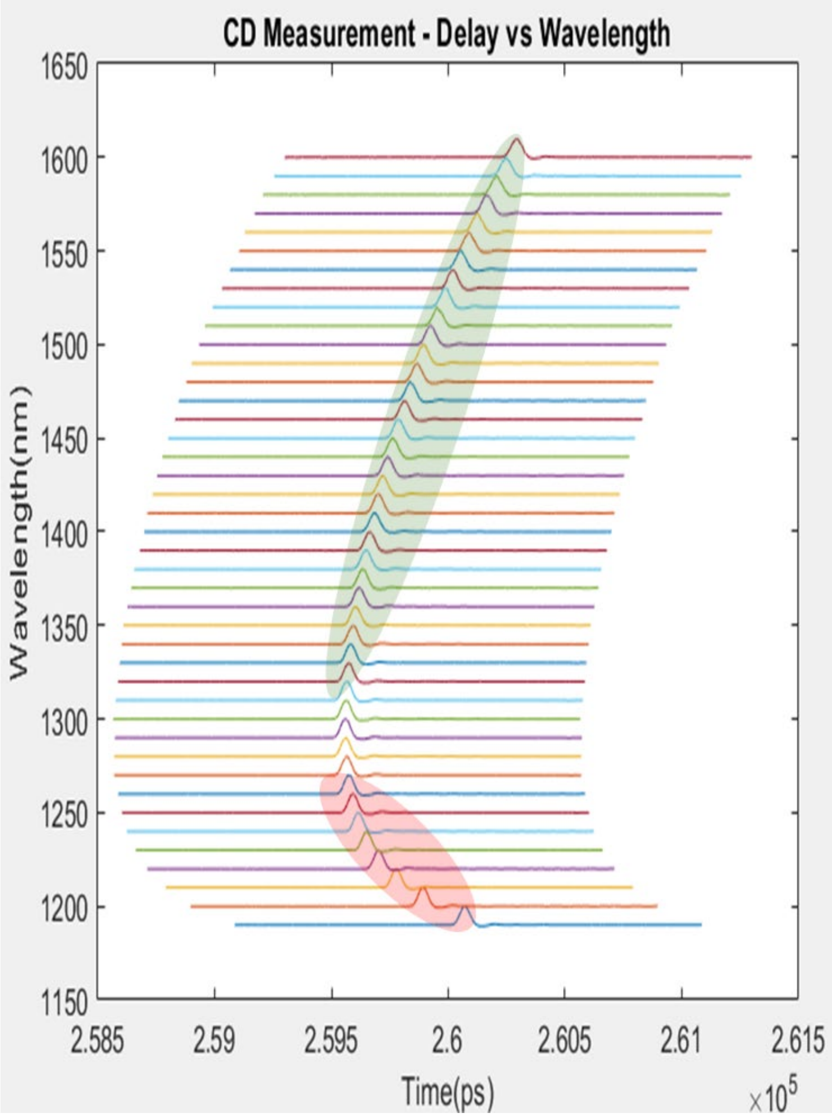

Panduit Fiber Lab Measurements of NANF sample

Scaling Across: Benefits

Deployment Advantages

- Greater flexibility in data center placement for better access to power and water resources, while maintaining relatively low latency.

Network Advantages

- HCF reduces latency (e.g., 0.5 ms / 1.5 = 0.33 ms, using 66% of switch buffer memory))and improve channel traffic control.
- Single fiber supports higher data rates per wavelength due to low chromatic dispersion.
- Expanded optical spectral range possible as nonlinearities (e.g., four-wave mixing) are negligible

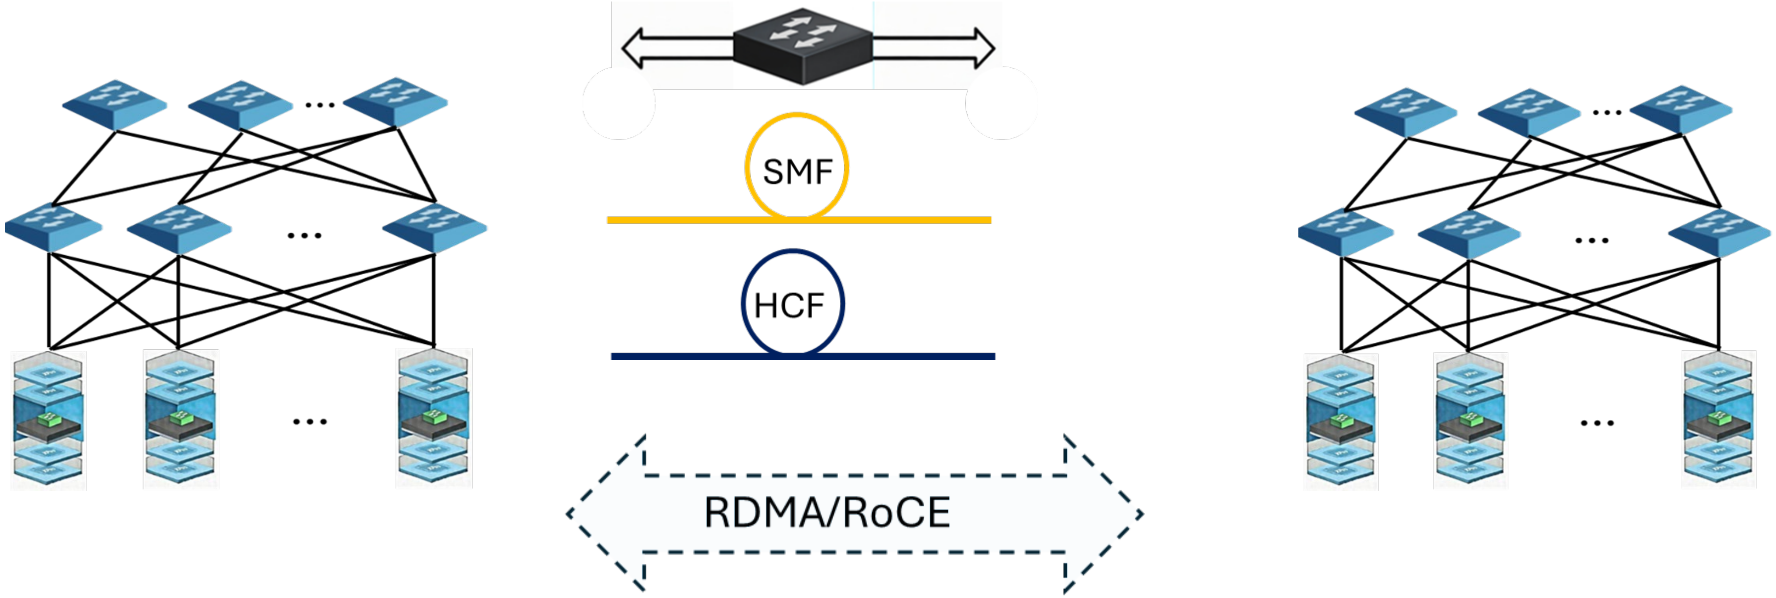

Connectivity Challenges: IL, Reflections, MPI

- A fully hollow-core end-to-end channel is not yet practical: exposing the HCF end face risks contamination and damage, so interfacing to standard SMF remains necessary.
- Managing the strong Fresnel back-reflection at the HCF–glass transition is challenging, because low loss, low reflection, and multipath interference (MPI) are intrinsically linked. Antireflection coatings can help but may restrict the usable spectrum.
  - New approaches such as offset-spliced SMF and GRIN couplers are being explored with promising results*.

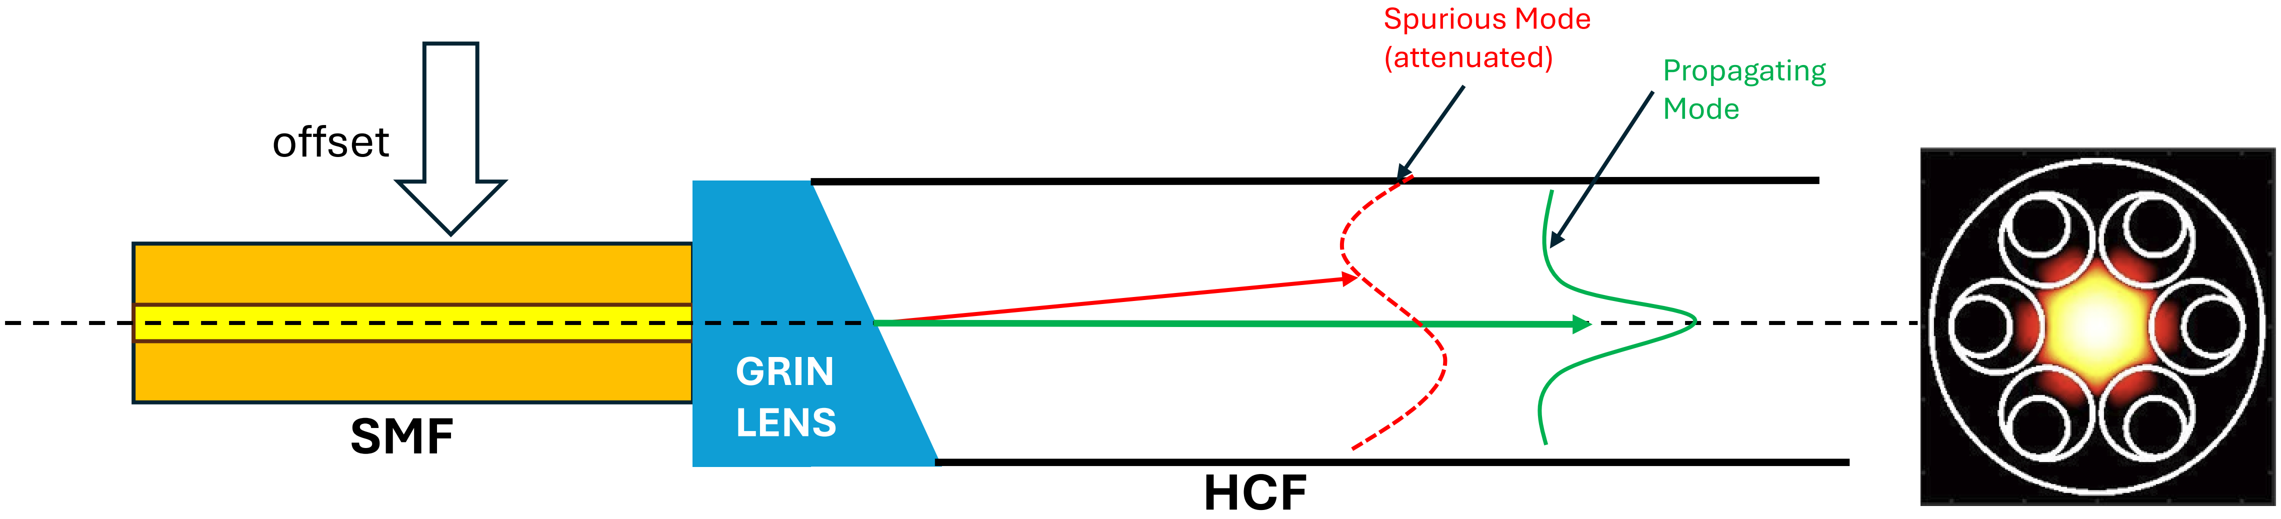

Several reports on advances in improving the coupling quality of HCFs have been published recently by leading teams, such as Radan Slavid's, in conferences/journals including IEEE,OPTICA and ACS.

HCF Splicing Sensitivity

- The challenge with splicing HCF is that when the internal capillary tubes collapse, this causes very high IL.
- HCF-to-HCF splice loss of <0.1 dB is achievable on a regular basis with optimized splice parameters

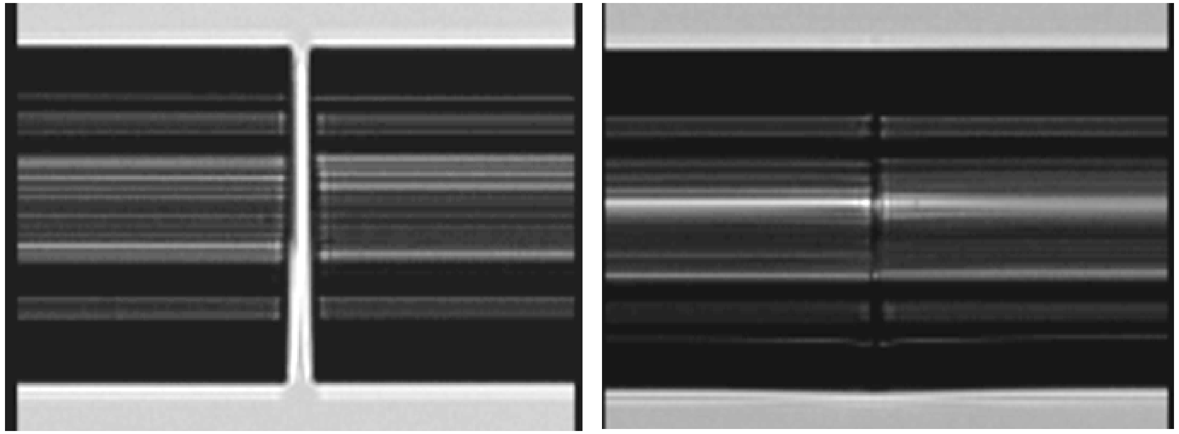

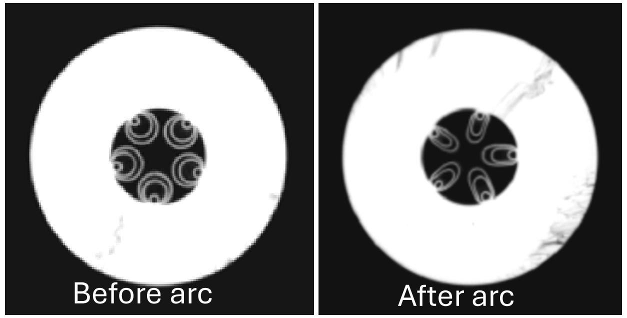

Incorrect arc deforming the capillary tubes (this has to be avoided for a proper splice)

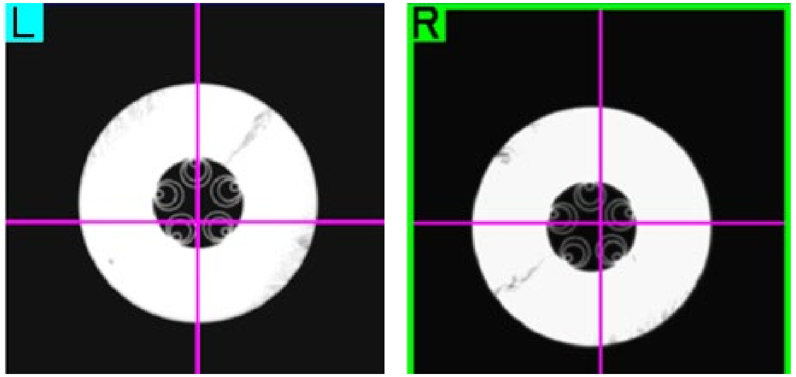

End-view alignment system is needed for proper alignment of internal capillaries

Summary

- AI training and inference require more compute, which depends on faster communication across scale-up, scale-out, and scale-across networks.
- IEEE 802.3 PMDs are moving to 200G per lane—and potentially 400G per lane—this decade, enabling ~400G scale-out per GPU today and providing a path toward 800G ports.
- The advantages of MCF and HCF are clear: improved density (MCF) for scale-out and scale-across, and improved latency (HCF), plus bandwidth headroom for scale-up and campus-scale links.
- However, standardization and interoperability remain open issues—especially for HCF.
- Compatibility, stable designs, and higher-volume production with better yields are expected to reduce cost over time for both fibers, MCF and HCF.

Questions?