Hollow-Core Fibre and Multi-Core Fibre: Use Cases vs. Traditional Single Solid-Core SMF

Introduction

For over four decades, the optical communications industry has relied almost exclusively on single-mode, solid-core silica fibre (SMF) as its transmission medium of choice. The physics are well understood, the manufacturing is mature, and the ecosystem of transceivers, amplifiers, and connectors is vast. Yet two emerging fibre technologies — hollow-core fibre (HCF) and multi-core fibre (MCF) — are beginning to find genuine commercial footholds, each solving a different set of problems that standard SMF cannot address. This article examines where each technology excels, where it falls short, and how the industry is beginning to deploy them.

A Brief Recap of Standard SMF

Conventional single-mode fibre guides light through a solid germanium-doped silica core, typically 8–10 µm in diameter, surrounded by a lower-refractive-index cladding. Light is confined by total internal reflection, propagating at roughly 68% of the speed of light in vacuum — a consequence of the silica core's refractive index of approximately 1.467 at 1550 nm. This introduces an irreducible latency floor of around 4.9 µs per kilometre. SMF also suffers from chromatic dispersion, non-linear optical effects (self-phase modulation, four-wave mixing, cross-phase modulation), and polarisation mode dispersion — all of which become significant constraints as data rates and spectral efficiency demands increase.

Hollow-Core Fibre (HCF)

How It Works

HCF guides light through an air (or vacuum) core rather than a solid glass one. Two principal designs have emerged: photonic bandgap fibre (PBF), which confines light using Bragg-like interference in a periodic cladding microstructure, and the more commercially promising anti-resonant fibre (ARF) — particularly the nested anti-resonant nodeless fibre (NANF) design. NANF has been further improved with DNANF (double-nested) and TNANF (triple-nested) designs which achieve robust bend loss for SMF like mode field diameters. In ARF designs, thin-walled silica tubes surround the hollow core, and light is confined by anti-resonant reflection rather than total internal reflection.

Key Advantages Over SMF

Latency. With the modal field propagating predominantly in air, HCF achieves a refractive index close to 1.0, raising propagation speed to around 99.7% of the speed of light in vacuum. This translates to a latency reduction of approximately 31% compared to SMF — roughly 1.5 µs/km saved. On a 1,000 km route, that is approximately 1.5 ms round-trip improvement, which is transformative in latency-sensitive applications.

Non-linearity. Silica's non-linear refractive index (n_2) is 2 to $3 \times 10^{-20}\ \mathrm{m^2/W}$. With the optical field confined largely to air, where n_2 is roughly 1,000 times lower, HCF has an effective non-linearity order of magnitude below that of SMF. This fundamentally relaxes the optical power limits that constrain long-haul amplifier spacing and modulation format choices.

Chromatic dispersion. HCF can exhibit very low and flat dispersion across broad wavelength ranges, which simplifies or eliminates the need for dispersion-compensating fibre or complex digital signal processing.

Radiation hardness. Because the signal propagates in air rather than glass, HCF is significantly more resistant to ionising radiation — important in space, nuclear, and certain defence environments.

Use Cases for HCF

Financial trading and low-latency interconnects. The single most commercially compelling application today is arbitrage and algorithmic trading, where every nanosecond of reduced latency represents directly monetisable competitive advantage. Routes between co-location facilities — Chicago to New York, London to Frankfurt — are the primary targets. Vendors including Lumenisity (now part of Microsoft Azure) and others have been deploying commercial HCF on these corridors.

Data centre interconnects (DCI) requiring ultra-low latency. Within and between hyperscale campuses, latency symmetry between compute nodes is critical for distributed AI training and HPC (High Performance Compute) workloads. HCF allows operators to reduce the effective electrical length of a fibre run without physically shortening the route.

Future coherent long-haul. As capacity demands push towards higher-order modulation formats (64QAM, 256QAM), the non-linear Shannon limit of SMF becomes the binding constraint rather than amplifier noise. HCF's near-zero non-linearity could substantially extend the theoretical capacity ceiling of a single fibre, though long-haul deployment is still constrained by several practical challenges noted below.

Sensing and defence. Gyroscopes, hydrophones, and electromagnetic-interference-immune signal links benefit from HCF's isolation of the optical field from the glass medium. In military environments, radiation hardness is an additional driver.

Challenges

While many existing hollow-core fibre (HCF) designs show higher attenuation than the lowest-loss single-mode fibre (SMF), advances in design and fabrication have demonstrated that HCF can reach comparable or even lower loss. This has been achieved at key telecommunications wavelengths, with reported values below 0.25 dB/km at 1310 nm and 1550 nm. The more pressing manufacturing challenge is consistency and yield: producing HCF in the continuous, unjointed lengths required for long-haul applications without degradation in the microstructured cladding geometry remains difficult, and yield limitations drive up effective deployed cost significantly. The anti-resonant microstructure that enables HCF's properties also necessitates a larger outer diameter than standard 125 μm SMF. This increases the minimum bend radius considerably — a practical constraint in cable designs, duct installations, and termination enclosures where tight bends are unavoidable, and one that limits how closely HCF can substitute for SMF in existing physical infrastructure without modification.

A further loss mechanism that has no equivalent in solid-core fibre is the loss incurred at the transition between HCF and conventional SMF. Mode-field mismatch between the air-guided mode in HCF and the glass-guided mode in SMF means that every HCF-to-SMF interface introduces additional loss that must be budgeted independently of the fibre's intrinsic attenuation. Critically, this transition is not a field splice — it is a factory-made component supplied by the manufacturer, either as a pigtail (where the installer splices HCF-to-HCF at one end and SMF-to-SMF at the other) or as a short transition length, typically around one metre, that serves the same purpose. This approach keeps both field joints like-for-like, avoiding the need for an installer to attempt a direct HCF-to-SMF splice in the field — which would be technically demanding and introduce uncontrolled loss. It's worth noting, when compared to a SMF-to-SMF fusion splice, there is additional reflection loss due to the air-to-glass interface (HCF-to-SMF), from the difference of the two refractive index mediums. This can be compensated for with anti-reflection coatings on the SMF – another reason this transition would be more challenging in the field. Nonetheless, the transition loss itself is inherent to the technology and present at each end of every HCF span, meaning a system using HCF as a segment within a predominantly SMF network must account for these losses at both terminations.

The absence of in-fibre Raman amplification also complicates long-haul system design. In SMF, the glass itself acts as the gain medium when pumped, enabling distributed Raman amplification across a span. With the optical field confined to air, this mechanism is unavailable in HCF, removing a useful tool from the system engineer's kit and potentially requiring shorter amplifier spacing or different launch power strategies compared to an equivalent SMF system. This does not impose a hard distance limit — long routes such as New York to Chicago can still be served using discrete EDFA amplification at repeater sites in the conventional way — but it does reduce design flexibility and may increase infrastructure cost for extended deployments. The more binding near-term distance constraints remain the manufacturing yield limitations on continuous fibre lengths and the accumulation of transition losses across multiple spliced segments.

Multi-Core Fibre (MCF)

How It Works

MCF packages multiple individual cores — each capable of supporting one or more spatial modes — within a single fibre cladding. Designs typically feature 4, 7, 8, 12, or 19 cores arranged in various geometrical patterns. Each core in a single-mode MCF behaves analogously to a separate strand of SMF, with inter-core crosstalk (XT) being the primary design challenge. Trench-assisted designs and careful core pitch optimisation suppress crosstalk to acceptable levels (typically better than −30 dB over transmission distances of interest).

The 8-core count is particularly significant for data centre applications. Parallel optic standards such as 400GBASE-SR4 and 800GBASE-SR4.2 use 4 transmit and 4 receive lanes — 8 lanes in total — making 8-core MCF a natural fit that allows a single fibre to replace an entire parallel fibre ribbon without any change to the transceiver interface. The 8-core geometry is however more technically demanding than the geometrically natural 7-core arrangement: fitting 8 cores within a standard 125 µm cladding while maintaining sufficient core pitch to control crosstalk is challenging, and some designs increase the cladding diameter beyond 125 µm to provide adequate core separation — a trade-off that must be weighed against compatibility with standard handling equipment and splice machines.

Key Advantages Over SMF

Spatial efficiency. A 7-core MCF in a standard 125 µm outer diameter cladding delivers 7× the spatial channels of a single SMF in the same physical cross-section. In environments where duct space, sub-sea cable diameter, or data centre tray density is the binding constraint, MCF multiplies capacity without multiplying cable plant.

Reduced cost-per-bit per unit of physical infrastructure. When duct fill is the constraint — as it increasingly is in dense urban environments and transoceanic cable systems — MCF offers a route to capacity growth without new civil works.

Compatibility with space-division multiplexing (SDM). MCF is a natural platform for SDM, the next major paradigm shift in optical capacity scaling after wavelength-division multiplexing (WDM). SDM combines spatial channels (cores or modes) with WDM to break through the non-linear capacity limits of single SMF spans.

Use Cases for MCF

Submarine cable systems. This is the leading commercial deployment environment. NEC, SubCom, and Alcatel Submarine Networks have all been developing or trialling MCF-based cable designs. A 12-core MCF in a sub-sea cable of the same outer diameter as a conventional 2-fibre pair cable increases capacity by 6×, dramatically improving the economics of transoceanic capacity. The physical volume of cable and repeater housings is the primary constraint driving MCF adoption here.

High-density data centre cabling. Within data centres, tray and conduit space between Top-of-Rack switches, spine switches, and patch panels is at a premium. MCF combined with fan-out assemblies (splitting each core to individual connectors at endpoints) allows a single cable run to carry 4 or 7 times the traffic of a conventional fibre. This is particularly relevant in hyperscale and AI-focused data centres where port densities are climbing rapidly.

Metro aggregation and urban fibre. In cities where ducts are fully occupied and new civil engineering is prohibitively expensive, MCF offers a capacity upgrade path by replacing existing fibre with MCF using the same duct cross-section. The economics depend heavily on the cost of fan-in/fan-out components and MCF-compatible amplifiers.

Research and future backbone networks. Laboratory demonstrations using SDM over MCF have achieved petabit-per-second aggregate capacity over single fibre segments. As component costs decline and SDM transponders mature, MCF backbone deployment becomes more plausible.

Challenges

MCF requires dedicated fan-in/fan-out (FI/FO) components at every termination point, multiplying connector hardware. Inter-core crosstalk management becomes critical over long distances, particularly for higher core counts. MCF-compatible EDFA (erbium-doped fibre amplifier) designs require multi-core doped fibres or spatial-multiplexed pump configurations. Connectorisation demands precision alignment across all cores simultaneously, and the ecosystem of splicing equipment, test instruments (OTDR sources and detectors that can address individual cores), and compatible transceivers is still maturing.

Comparative Summary

Property	Traditional SMF	Hollow-Core Fibre (HCF)	Multi-Core Fibre (MCF)
Propagation speed	~68% c	~99.7% c	~68% c (per core)
Latency	Baseline	~31% lower	Same as SMF per core
Non-linearity	Moderate	Near-zero	Same as SMF per core
Attenuation	~0.15 dB/km	<0.25 dB/km (best-in-class)	Similar to SMF
Chromatic dispersion @ 1310 nm	~0 ps/nm.km (zero-dispersion wavelength)	Very low	Same as SMF per core
Chromatic dispersion @ 1550 nm	~17 ps/nm.km	Very low (similar to 1310 nm)	Same as SMF per core
Cladding diameter	125 µm (standard)	>125 µm (design dependent)	125 µm (some 8-core designs larger)
Bend radius	Standard	Larger than SMF	Same as SMF
Additional losses	None beyond intrinsic	SCF-to-HCF transition loss at each span end	Fan-in/fan-out insertion loss at terminations
Spatial efficiency	1 channel/fibre	1 channel/fibre	N channels/fibre
Ecosystem maturity	Very high	Early commercial	Early commercial
Primary use case	Everything today	Low-latency critical links	Capacity-constrained plant

Standardisation and Supply Constraints

A consideration applicable to both HCF and MCF is the current absence of formal standardisation across either technology. Unlike SMF, where ITU-T G.652/G.654 and IEC standards define fibre geometry, attenuation, dispersion, and interconnect compatibility in precise terms, neither HCF nor MCF has yet attracted equivalent international standards coverage. This creates interoperability uncertainty between vendors and complicates the procurement, testing, and long-term maintenance of installed systems.

Compounding this, supply of many HCF and MCF products remains limited due to unresolved intellectual property and patent disputes between manufacturers. Key aspects of fibre design, preform fabrication, and draw processes are subject to competing claims, and in some cases these disputes have constrained the number of suppliers able to bring product to market or export freely across jurisdictions. Prospective deployers should be aware that vendor choice may be restricted not solely by technical capability or price, but by the evolving IP landscape — and that this situation is likely to develop materially over the coming years as patent positions are tested or cross-licensing agreements are reached.

The Hybrid Future

It is worth noting that HCF and MCF are not mutually exclusive — research into incorporating hollow cores within multi-core fibre structures is already underway, with the goal of combining low latency and near-zero non-linearity with the spatial multiplexing advantages of MCF. Such fibres could eventually address both the latency-sensitive and capacity-constrained segments of a network simultaneously, though the manufacturing complexity is considerable even relative to each technology in isolation.

More practically, the near-term picture is one of segmented deployment: SMF continues to dominate the vast majority of network infrastructure where its performance is adequate; HCF carves out premium, latency-critical point-to-point links where the cost premium is justified by latency value; and MCF addresses capacity bottlenecks in physical-layer-constrained environments — submarine cables first, dense urban ducts next, and high-port-density data centres in parallel.

Conclusion

Traditional single solid-core SMF remains the dominant transmission medium and will continue to be for the foreseeable future. However, it faces genuine physical limits that neither additional WDM channels nor improved DSP can fully overcome. HCF addresses the latency and non-linearity constraints, opening new possibilities for time-sensitive applications and potentially extending the capacity ceiling of long-haul systems. MCF addresses the spatial efficiency constraint, enabling capacity growth where the physical cable plant — not the spectral bandwidth — is the bottleneck. Together, they represent complementary and increasingly credible additions to the optical networking toolkit, rather than wholesale replacements for the infrastructure on which the internet already runs.