CORNING

Deploying Robust and Scalable Co-Packaged Optics Fiber Infrastructure

Introduction

Co-packaged optics (CPO) is a much-anticipated revolution in the architecture of high-bandwidth switches and distributed-computing hardware used in data centers (DCs). The prevailing technology involves electrical connections inside the boxes, with optical/electronic transceivers that plug into the face plates ("pluggables") and connect to optical cables that link switches, servers, distributed memory, and processors together. As link bandwidths increase and demand grows for better power efficiency, nearly lossless fiber will enter the box to replace lossy copper traces and allow the transceivers (TRXs) to move much closer to the core components such as switch ASICs or GPUs. The box will now contain a large number – hundreds and perhaps thousands – of optical fiber links between the tightly coupled Electronic Integrated Circuits (EICs) and Photonic Integrated Circuits (PICs) that form the densely integrated TRXs, and the face plate.

While the introduction of this CPO optical infrastructure can substantially reduce the power required to drive data transfer between different locations in the DC, this infrastructure may be unfamiliar to manufacturers of these boxes. Moreover, at least to some extent, the combination of very large amounts of fiber cabling in close proximity to electrical components and fans — and at a consistently elevated temperature — is unprecedented.

The use of pluggables enables the "hot-swapping" of failed units, but with the advent of CPO any link failure may necessitate the replacement of an entire switch box. Thus, high reliability of the CPO box is critical, and designers and manufacturers must ensure that the deployment of the optical infrastructure and the assembly process does not compromise the reliability of the feedstock subassemblies.

The requirements for CPO reliability are not yet well defined, but reasonable working targets for reliability engineers are a lifetime expectation of at least 5 years and a failure rate of <0.1% per annum for the entire CPO infrastructure. This is in line with current pre-CPO experience where the all-cause rate of switch failure requiring replacement of the entire assembly is around 1% p.a. [1, 2].

This white paper will explain the design and handling practices that have been developed over half a century to deliver high and consistent levels of reliability wherever glass optical fibers are used in critical applications. Insights in fiber deployment and handling will be provided to maintain high levels of reliability in the context of emerging CPO designs.

Scope

The paper will discuss the reliability of silica-based glass optical fibers deployed in a CPO switching environment which may involve complex paths with some tight bends. The fibers in scope include single-mode fibers used in Tx/Rx signal paths and polarization-maintaining fibers carrying power from external laser sources.

General overview of fiber reliability

Fiber failure mechanisms include breaks, darkening and the "fiber fuse," but given the wavelengths and power ranges anticipated for CPO, the only one that we consider here is fiber breaks. The mechanical reliability of silica-based glass optical fibers is long-studied and well understood qualitatively and quantitatively [3, 4, 5, 6]. Here we provide an overview of work on fiber reliability and the practical consequences, based on the extensive literature on this subject.

A simple model of fiber breaks is as follows:

1. The surface of a glass fiber is covered in microscopic flaws or cracks that act as stress concentrators.
2. Under tension, these flaws grow through a water-driven stress-corrosion process at a velocity that is dependent on a fatigue parameter that itself depends on the surface glass composition.
3. As the flaws grow, the stress intensity at the flaw tip increases and the corresponding velocity increases exponentially.
4. The stress intensity eventually exceeds the fracture toughness of the glass, at which point the flaw grows at an extremely high speed and breaks the fiber.

Given data about the fiber strength distribution, which is determined by the manufacturing process and the level of proof-testing, the failure probability for a given applied tensile stress and time to failure can be predicted in a statistical sense, i.e., what percentage of a population of fibers deployed in that way will have failed after a given time. Note that the reliability of the fiber is determined only by the population of extrinsic flaws and applied tensile stresses – the optical properties of the fiber are irrelevant, and all fibers made to the same quality in terms of surface flaws will behave the same. The population of extrinsic flaws is controlled by a proof test, establishing a minimum threshold in the fiber strength distribution. In proof-testing, a defined stress is applied at every point along the fiber and all flaws that are below a target strength level will grow to failure during the test. The surviving fiber has a truncated population of flaws with a lower strength limit roughly equal to the proof test and determined by the unloading rate in the proof test.

The most common way that stress is applied to a fiber is by bending it. This produces a tensile stress on the outside of the bend proportional to the fiber diameter and inversely proportional to the bend radius. Thus, CPO designers deploying fibers in bent configurations must consider the impact on reliability as well as on optical loss.

The tension-induced crack growth in bullet point 2 above is an exponential process, so initial crack velocity and correspondingly failure rates can increase rapidly as the fiber is bent more tightly. Because the rate of crack growth scales with stress intensity with an exponent in the 20–40 range, a short exposure to high stresses can result in significantly more breaks.

The crack growth process is such that the failure rate increases logarithmically with time, so the difference between the 5–10 year operational lifetime required in a typical DC application and the 25 years specified for cable plant does not markedly relax deployment and handling constraints.

Strength distribution data applies only to flaws created at the point of measurement, and proof testing protects only against flaws that existed when the proof stress was applied. Care needs to be taken in subsequent handling steps to prevent damage to the coating or glass that can create new flaws.

Even the slightest contact with the bare glass can also induce lifetime-limiting flaws, so tough polymer coatings are used to isolate the glass surface [7]. Damage and compromised reliability can occur where the fiber coating is stripped, for example, for connectorization or splicing.

Optical fiber reliability can be enhanced by reducing the cladding diameter, by using a modified surface glass composition – for example, titania-doped silica – that has a higher fatigue constant, or by increasing the proof-test level [8].

Fiber deployment for reliability

As noted above, the key failure mechanism for coated fiber is crack growth under tension at the glass surface. Tension can be induced by bending the fiber or twisting it about its axis. Stretching the fiber also causes tension but it is unlikely that the fiber in CPO deployments will be subjected to such a deformation.

For example, a fiber with a glass diameter of 125 µm bent with a radius of 15 mm, will (in a simple beam-bending model, where the neutral axis coincides with the fiber axis) have a tensile strain on the outside of the bend of about 0.4%. For comparison, a twist of 90° in 25 mm, such as might occur when rotating a fiber ribbon to align it to a PIC, will also produce a shear strain of 0.4% but, over the entire surface area of the glass fiber, equivalent to 0.17% tensile strain (Figure 1).

Torsional Stress: $\sigma_{tor} = G\,\varepsilon_{tor}$

Where:

G = Shear Modulus = 31 GPa (4500 kpsi)

and

$$\varepsilon_{tor} = \frac{r\,\theta}{L}$$

Example: 1 meter twisted 360°

$$\varepsilon_{tor} = \frac{0.0000625 \text{ x } 2\pi}{1}$$

$$\sigma_{tor} = G \text{ x } 0.00039$$

$$\sigma_{tor} = \underline{12 \text{ MPa}}$$

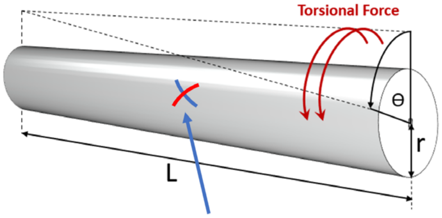

Principal stress axes are at 45° to fiber axis – blue axis is under tension, red under compression. For an isotropic material such as silica, the principal stresses under the deformation shown here have the same magnitude as the torsional stress. The equivalent extensional strain for the same stress under uniaxial tension is given by $G\varepsilon_{tor}/E$ where E is Young's modulus.

Figure 1: Stresses in twisted fiber.

Corning has carried out extensive testing of fibers to determine strength distributions. The fibers are tested in uniaxial tension, and the resulting strength distribution can be statistically scaled to account for the non-uniform stresses applied on the fiber surface in bending. Corning's reliability models incorporate both material behavior and the flaw strength distribution for Corning fibers. It is important to note that the conclusions drawn may not be applicable to fibers from other manufacturers.

As an example, consider a coil 30 mm in length of fiber with a glass diameter of 125 µm. For a five-year failure rate of no more than 1 x 10^{-6}, we see from the lifetime design diagram in Figure 2 below [9] that the minimum bend radius is about 5.5 mm. Now consider the case where the same fiber is bent in a 360° loop with a radius of 3.3 mm, which corresponds to a length of about 21 mm of fiber. The lifetime design diagram indicates that 30 mm of fiber bent at that radius has a five-year failure probability of roughly 7 x 10^{-6}. Such low probabilities will scale with fiber length, so the failure rate for one turn of bend radius 3.3 mm will be (21/30) x 7 x 10^{-6}, or about 5 x 10^{-6}.

The data for these calculations is collected at room temperature, but the behavior will be very similar at the temperatures and humidities in a typical CPO application.

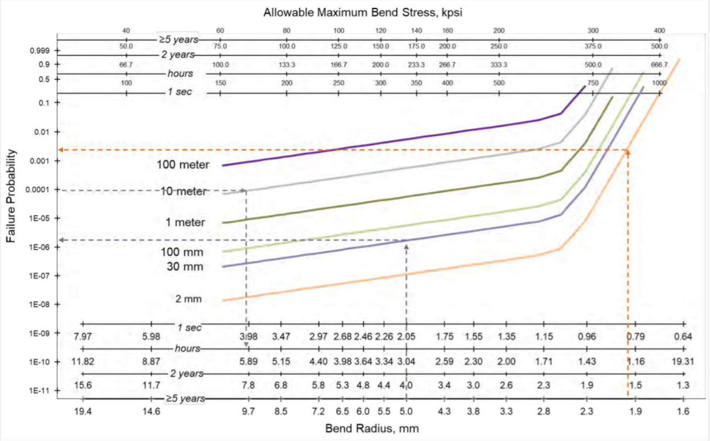

Figure 2: Lifetime design diagram for fiber lengths in bending (figure 14 of reference [9])

The failure probability *F* vs. time of a piece of fiber deployed in a specified path can be similarly modeled in detail, but it can also be estimated as the sum of the failure probabilities f_i of the uniformly bent arcs and uniformly twisted segments *i* that together approximate the path.

$$F = \sum_i f_i$$

Technically we should write:

$$F = 1 - S = 1 - \prod_i s_i = 1 - \prod_i (1 - f_i)$$

Here *S* and s_i stand for the survival probabilities, but for *S* and s_i very close to 1 – i.e., very small failure probabilities for each element and for the total path – the initial formulation is adequate.

Finally, for an ensemble of the population of links found in a CPO switch box, the probability S_{box} for the box to survive to the specified time without a fiber break is computed in the same way from the survival probabilities S_l for each link *l*.

$$S_{box} = \prod_l S_l$$

The failure rate rises rapidly as the bend radius is reduced, so the tightest bends on the fiber path dominate the overall performance.

Figure 3 below shows a ribbon-management device designed for CPO. The flanges impose a minimum radius on the ribbon bend, limiting the reliability risk from bending stress. As the same set of bent segments are present for different ribbon lengths stored in the cavity the overall reliability is independent of the stored length.

Figure 3: Concept for ribbon deployments in CPO, with capacity to manage slack resulting from manufacturing variations or from connecting different paths with a single ribbon jumper length.

The reliability data and models have been used to calculate the five-year failure probability for different numbers of 12-fiber ribbons deployed in such an accumulator where all the bends are at the minimum bend radius set by the flanges. The results are shown in Figure 4 below. Clearly, the larger the number of fibers in the box, the higher the probability of having a fiber break becomes.

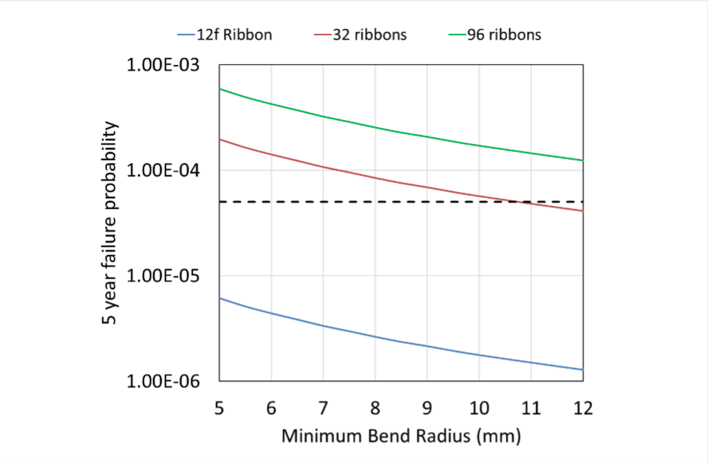

Figure 4: Modeled failure probabilities for various ribbon counts and different accumulator design radii (failure is defined as at least one broken fiber). The dashed line shows a 1-FIT probability level.

Similar computations can be carried out for any CPO fiber deployment, and the target reliability performance may constrain design rules such as the minimum bend radius. Note that crossovers of individual fibers that are placed under pressure can be points of small bend radius, so the layout design should at least minimize and ideally eliminate this possibility.

The design of the fiber layout in the switch box must consider not only the fiber fatigue stress limits, but also the thermal environment of the fiber coatings, because if these are compromised by thermal degradation, the glass surface of the fiber may be damaged and break before the design lifetime. Standard fiber coatings are capable of long-term operation at temperatures well above those expected in an operating CPO assembly [10], and specialized higher-temperature coatings are available, but potential hot spots must be avoided, and the insulating properties of fiber and fiber bundles must be considered to avoid creating hot spots. If necessary, spacers or routing hardware should be used to keep the fiber infrastructure in suitable microenvironments.

The mechanical environment of the coatings is also an important consideration, since loss of the coating integrity may lead to glass surface damage and thus a shortened fiber lifetime. Pinching, abrasion, and nicking must be avoided to prevent coating failure. Thus, fibers should not be deployed against rough or sharp surfaces or routed through narrow openings, which over time and combined with mechanical shocks and vibrations may lead to damage. In particular, strain-relief structures for optical connections must observe manufacturers' limits on stress applied to the fibers and their packages – cables, ribbons etc.

Fiber handling for reliability

The previous section dealt with designing the fiber deployment for acceptable reliability performance. This section covers fiber handling during assembly, rework or repair to ensure that the reliability design intent is realized.

We assume that the optical infrastructure input to the assembly process is a set of jumpers or harnesses, that is to assemblages of fiber packaged in some way – for example in loose tubes or ribbons, possibly with furcations and with the fibers terminated at both ends in some form of connector.

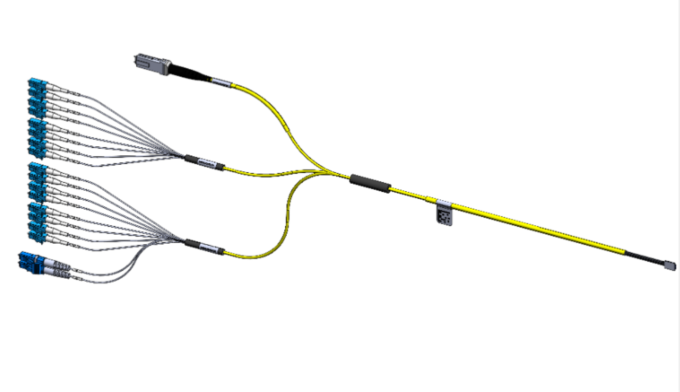

Figure 5: Example of fiber harness design that protects the fiber surface during handling.

During any manipulation of the fiber, three factors that may compromise its reliability must be avoided:

1. Damage to the glass surface
2. Application of large tensile stresses
3. Damage to the integrity of the coating

Since any contact with the glass may cause new microscopic surface flaws that can result in fiber failure, there should be no exposed glass in the subassembly being installed in the box. Moreover, care must be taken not to damage the integrity of the fiber coating. To avoid nicks, cuts, abrasion and punctures of the coating, sharp edges must be minimized in the design of the box and in the assembly environment wherever fiber may be placed. The fiber must be routed away from any remaining ones or contact must be prevented by shields or routing hardware. Excessive pressure from pinching or clamping fiber can also damage coatings by causing delamination, and exposure to organic solvents can weaken the coating material. Particulates, particularly of hard materials, should be avoided, and electrostatic forces controlled, to avoid the failure mode where particles are attracted to the fiber coating surface and then driven through the coating by external pressure, ultimately damaging and weakening the glass surface. In general, unbuffered fiber should be protected (e.g., with tubing) when not being worked on.
If any fiber break occurs during assembly and the fiber ends are exposed, careful cleaning must be carried out to eliminate all shards of glass that might cause coating or glass damage.

It is likely that fibers with reduced diameter coatings will be deployed in CPO applications to reduce the volume of the fiber infrastructure and to improve mechanical compliance through reduced stiffness. Assembly tools and procedures must be controlled to ensure that reduced coating thickness does not lead to tighter bends or glass surface damage from abrasions or punctures.

Excessive tensile stresses that pose a significantly higher probability of causing failures can arise if there are snags or tight bends, even for a brief period. A minimum bend radius for the fiber must be specified and respected, along with slack management strategies that avoid fiber kinks and buckling.

Further details can be found in reference [11], which includes an audit checklist for any process handling fibers.

CPO equipment will likely be built by incorporating optical subassemblies into the equipment enclosure, with the subassemblies being built separately and forming part of the final assembly kit.

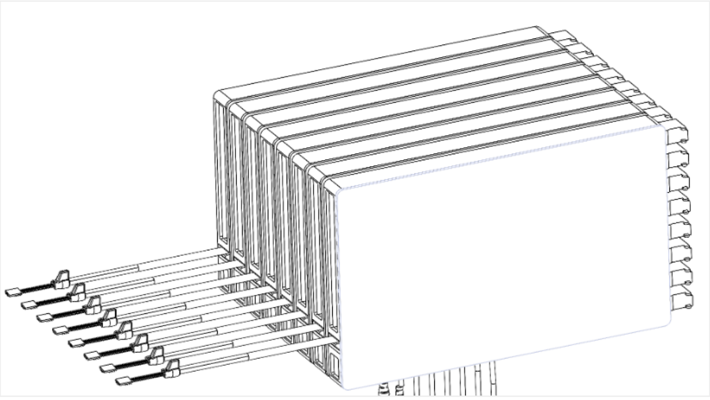

Figure 6: Example of fiber sub assembly concept: the fibers are supplied as “known good fiber” pre-routed in cassettes that reduce handling at the assembly house.

The subassembly manufacturing process will necessarily be carried out in a manufacturing operation with extensive experience and strong capabilities in fiber handling and connector assembly to deliver a reliable product. This facility will be well-placed to ensure that the subassemblies are free of process-induced damage that can cause failures during transport to final assembly, in final assembly itself, or through the service life of the equipment. Additionally, the subassemblies will undergo a series of screening tests to determine “known good fiber.” The tests will encompass return loss (RL), Optical Coherence Component Reflectometry (OCCR), insertion loss (IL) and visual inspection. Thus, the final assembly process will receive reliable subassemblies (“known good fiber”) with a very high probability.

The final assembly operation will work with a much broader range of components, and the level of skill in fiber handling may well be significantly lower than in the subassembly operation. It is therefore desirable that fiber handling be minimized. The interventions required to manage slack arising from faceplate-to-PIC path length differences or manufacturing variations in optical fiber lengths may be minimized by utilizing cassettes that store or dispense fiber or fiber bundles with a simple mechanism providing a function much like a domestic cable storage device. With this aid, final assembly operators will be able to deploy the fiber outside the cassette according to a reproducible and deterministic pre-designed scheme. This will help to avoid issues such as buckling and tight bends that might compromise reliability.

Designing the fiber layout to facilitate reliability assurance in assembly and servicing

In addition to ensuring long-term mechanical reliability, the design of the fiber layout should support good reliability practices during assembly or servicing of the CPO module. Good design practices include:

1. Avoiding the interweaving or entanglement of different optical subassemblies wherever possible.
2. Keeping each optical subassembly well-groomed so that individual cables or fibers can be handled without interference.
3. Providing a clear path for working on the infrastructure that minimizes the risks of snagging or pinching fibers; in particular, module components that overlie the optical infrastructure should be moveable to allow access in the case that servicing is required.

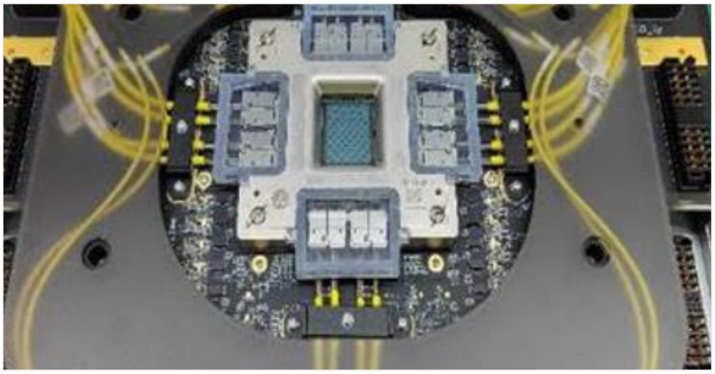

Figure 7: Example of fiber routing inside the box without interweaving (Image courtesy of Micas Networks Inc.)

Precedents for reliable dense fiber deployment

Fibers have been deployed with tight bends inside confined spaces for years. For example, Corning’s EDGE™ solutions for datacenters include 1RU housings that can accommodate 576 fibers, with the limitation being the connector density at the front panel, not the number of fibers in the box. Corning’s various splice management enclosures also pack many fibers into tight spaces, and in many cases these enclosures protect

PON splices – which require fiber stripping and handling that creates reliability risks – made by technicians out in the field and subsequently exposed to extremes of cold and heat. The ambient temperature for CPO boxes, while likely to be stable, is probably warmer than for typical passive module installations inside a datacenter but may not reach the extremes for PON equipment in some locations around the world.

Of course, with CPO the level of congestion can be expected to be significantly higher due to electronic components, powers supplies and fans as well as the fiber infrastructure, but the deployment principles are the same and will, if observed, deliver a reliable product.

Conclusions

Fiber infrastructures for co-packaged optics can be designed and assembled to achieve high reliability at scale. Failure mechanisms of optical fibers are well understood and following proper layout designs as well as handling guidelines will ensure that the overall reliability of the optical connections will meet the desired levels of reliability. The fiber harnesses that interconnect the photonic ICs to the equipment's faceplate are designed to protect the mechanical integrity of the fiber surfaces, which is the key factor in reducing failures over the lifetime of the system. The layout of the fibers within the CPO box is also an important consideration to ensure not only long-term mechanical reliability but also ease of assembly and serviceability; this is particularly important in denser footprints and higher fiber count configurations. Fibers that can accommodate smaller bend radii and fiber ribbon accumulators are examples of designs aimed at easing fiber layout and assembly. Use of pre-assembled fiber cassettes is another approach to reduce chances of fiber mishandling by downstream box-level assemblers who may not have extensive fiber handling experience.

References

[1] Intel, "IT@Intel: Green Computing at Scale," 2021. [Online]. Available: https://www.intel.com/content/www/us/en/it-management/intel-it-best-practices/green-computing-at-scale-paper.html.

[2] R. Singh, M. Mukhtar, A. Krishna and A. Parkhi, "Surviving switch failures in cloud datacenters," [Online]. Available: https://www.racheesingh.com/papers/sigcomm-ccr-final465-with-open-review.pdf.

[3] Corning Inc., "WP8002: Optical Fiber Mechanical Reliability," 2017. [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP8002.pdf.

[4] Corning Inc., "WP5053: Mechanical Reliability: Applied Stress Design Guidelines," [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP5053.pdf.

[5] Corning Inc., "WP3690: The Mechanical Reliability of Optical Fiber in Bending," [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP3690.pdf.

[6] Corning Inc., "WP1282: The Mechanical Reliability of Corning Optical Fiber in Small Bend Scenarios," [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP1282.pdf.

[7] Corning Inc., "AN2000: Optical and Mechanical Resilience of Reduced Coating Diameter Fibers," 2021. [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/application-notes/AN2000.pdf.

[8] Corning Inc., "AN7080: Determining Appropriate Proof Test Level for G.657 Fiber (designed for tight bending)," 2020. [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/application-notes/AN7080.pdf.

[9] G. S. Glaesemann and Y. Shu, "Optical Fiber Mechanical Reliability Modeling Extended to Small Bend Radii," in Proceedings: IWCS Cable & Connectivity Industry Forum, 2021.

[10] Corning Inc., "WP4250: Beyond 85°C: Thermal Aging Impact on Optical Fiber with Corning® CPC® Coatings," 2017. [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP4250.pdf.

[11] Corning Inc., "WP3267: Suggested Guidelines for the Handling of Optical Fiber," 2015. [Online]. Available: https://www.corning.com/media/worldwide/coc/documents/Fiber/white-paper/WP3267.pdf.

Corning Optical Communications LLC • 4200 Corning Place • Charlotte, NC 28216 USA
800-743-2675 • FAX: 828-325-5060 • International: +1-828-901-5000 • www.corning.com/opcomm