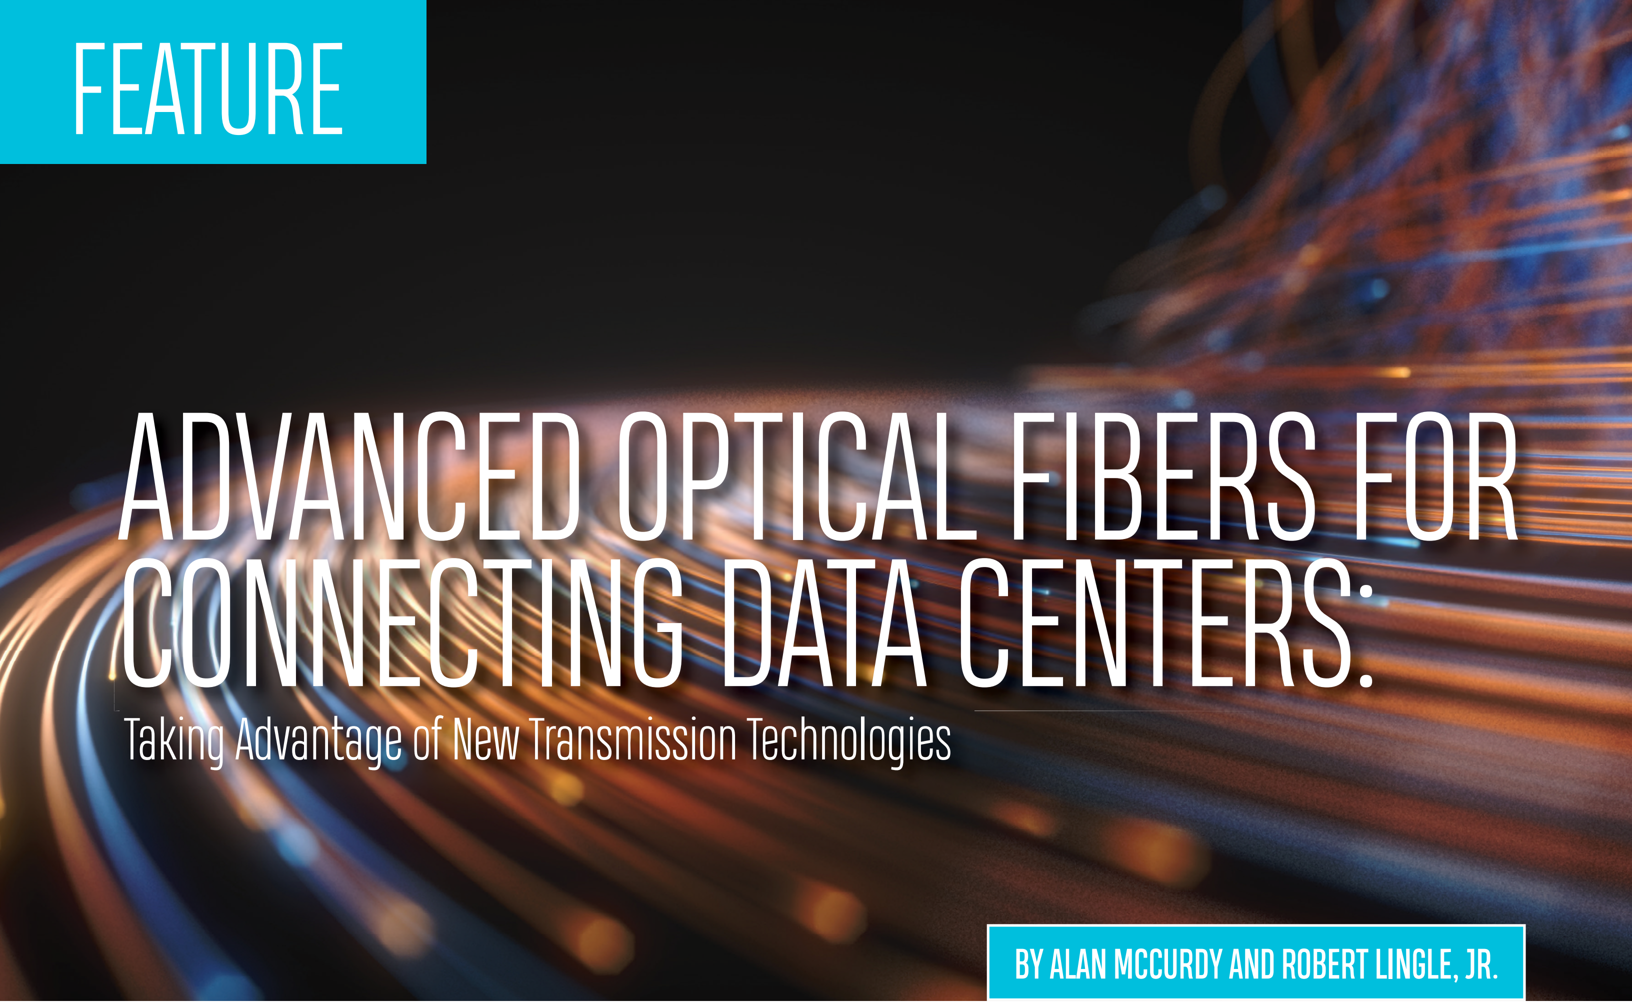

ADVANCED OPTICAL FIBERS FOR CONNECTING DATA CENTERS:

Taking Advantage of New Transmission Technologies

BY ALAN MCCURDY AND ROBERT LINGLE, JR.

Connections between large cloud data centers can be divided into four regimes conceptually. The shortest are metro links up to (perhaps) 60 to 80km, where the link latency is low enough such that several DC's comprise one large distributed switching fabric. The second is data center connectivity over the regional network in the range of 500 to 600km with pluggable 400G optics over amplified spans, loosely termed "400G-ZR+," similar to the range of the Open ROADM project. Massive cross-continent long-haul data flows also occur point-to-point between large data center campuses. Finally, subsea traffic is a critical aspect of the DC network, as shown by the fact that Google & Facebook are two of the largest funders of subsea networks in the last five years, with Microsoft and AWS also participating. Large effective area, ultra-low loss (ULL) fibers provide key benefits in network design in each of these spaces. In particular, it is possible to construct a seamless network of ULL optical fibers with 125 mm^2 mode field area (ITU category G.654.E) across the data center interconnect networks.

ADVANCED OPTICAL FIBERS

Modern optical fiber design and manufacturing excellence has enabled competitively priced fibers with properties which were unattainable just a few years ago. The properties which are of increasing value today when using new transmission technologies are: attenuation, optical mode field size and nonlinear index of the fiber core. Optical fibers with low attenuation (~ 0.15 dB/km for submarine fibers and ~ 0.17 dB/km for terrestrial fibers at 1550 nm wavelength) serve the purpose of preserving the signal strength, so that less gain is required at the amplifier locations (every 60 – 100 km depending on the link design). Since each amplifier adds noise, a low attenuation fiber offers a way to preserve high signal-to-noise ratio (OSNR). The fiber mode field size and nonlinear index both affect the distortions and intermodulation (generally called nonlinear impairments) which occurs due to interactions between the transmitted light and the glass in the fiber core. Even though these impairments are very small, they can add up to problems in two scenarios: 1) transmission over long distances (such a transoceanic links which are many thousands of kilometers) and 2) transmission using high-order modulation formats (16QAM and higher) where the signal constellations are very sensitive to noise. Spreading the optical power over a larger cross-sectional area lowers the nonlinear interaction with the glass. Large mode field fibers require special design to preserve the signal through the inevitable fiber bends which occur during cabling and deployment. The most advanced optical fibers have light-guiding modes which occupy roughly twice the cross-sectional areas as conventional fibers (150 mm^2 versus 80 mm²). The improved nonlinear index is a function of getting away from the typical germanium doping, used to raise the refractive index of conventional optical fiber cores, and moving to pure silica core fibers. This requires lowering the refractive index of the surrounding cladding in order to confine the light to the core. Hence fluorine doping is generally used in the cladding glass. Though the effect of the nonlinear index reduction is not as large as changing the mode field size, every little bit helps!

TERRESTRIAL CONNECTION BETWEEN DATA CENTERS

Over the past ten years, enormous "hyperscale" data centers have been multiplying around the world. But data center planners must balance the cost savings of centralization of functionality with the performance advantages (especially latency) of geographical distribution (as close as possible to end users). The challenge of connecting these data centers (through both terrestrial and submarine links) has driven new technologies. On the terrestrial side many of these "DCI" connections are in the range of ~ 100 km in length. This has allowed the use of very high bit-rate wavelengths over metro fiber without the need for inline amplifiers, where DCI links operate with wavelength speeds > 100 Gbps. An example of this is the recent demonstration by Furukawa Electric Company and Fujitsu of 600 Gbps wavelengths over 100 km reach with commercial gear and installed ribbon cable [4], using the Fujitsu flex-rate 1FINITY™ T600 transponder. This test compared the performance of ITU-T category G.652.D (OFS AllWave®) fiber against G.654.E (OFS TeraWave® ULL) fiber and found that the reach could be extended by 25% with the advanced G.654.E fiber. This is shown in Figure 1 where the OSNR margin (excess OSNR beyond that required to hit the desired bit error rate) is shown at various link reaches. The G.654.E fiber has > 2 dB OSNR margin across the range of link distances tested and hits 125 km maximum distance compared to 100 km for the G.652.D fiber. Other tests (demonstrated at the OFC 2019 conference this year) pitting G.654.E fiber against installed-base G.652 fiber on spools have showed a 50% reach advantage for G.654.E.

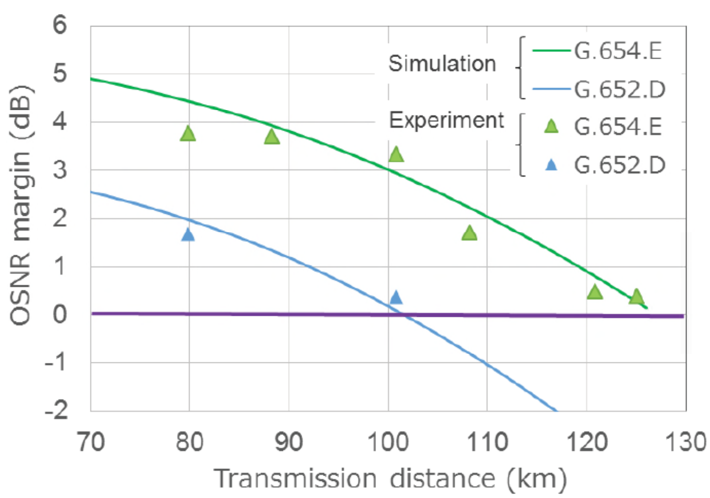

Figure 1. Enhanced OSNR margin using G.654.E fiber. The advanced fiber allows 2 dB more optical power to be launched, thereby spreading out the 64 QAM constellation and preserving performance.

The value of ultra-low fiber loss in un-amplified applications is obvious, as is well-known in unrepeatered submarine links. But the role of the large mode field of G.654 in such short distances may not seem obvious at first glance. Again, the large mode field combats fiber/light nonlinear interactions, which are typically thought to gradually build up over very long distances (thousands of kilometers). To see nonlinear effects at 100 km reach may seem surprising. However high bit rates on single wavelengths are obtained through increasing the complexity of the modulation format (far beyond conventional QPSK used for 100 Gbps). The transponder used in these demonstrations utilizes 64QAM modulation which is quite sensitive to noise of any type, so nonlinearity quickly becomes an issue over conventional fibers.

The 400G ZR optical module should be a very popular form factor for metro DCI. The Optical Internetworking Forum (OIF) has been working for several years on the 400G ZR agreement and is now considering an 800G version. Wavelengths carrying data at these speeds will certainly benefit from the improved OSNR performance of advanced G.654.E fibers. This is particularly true as these agreements intend to use less costly, interoperable coherent optics which will likely benefit from lower impairments over G.654 fiber than over conventional fiber.

High-speed links over regional distances require in-line amplifiers. Our colleagues at Fujitsu also applied engineering rules to predict performance of the 1FINITY™ T600 transponder with amplified spans operating at 400 Gbps. The results are shown in Table 1, where a comparison is shown between transmission over G.652.D and G.654.E fiber. There is a substantial reach and cost benefit to using the G.654.E fiber. This benefit accelerates for larger amplifier hut spacing (cost savings through long spans). The implication of a longer maximum reach is the avoidance of signal regeneration, which is exceedingly expensive.

Maximum Unregenerated Reach (km)			
Fiber Type	Span length (km)		
	80	100	120
AllWave (G.652.D)	560	400	240
TeraWave ULL (G.654.E)	720	700	600
Reach Improvement	**29%**	**75%**	**150%**

Table 1: 400 Gbps link reach for different fiber types using Fujitsu design rules with 1FINITY™ T600 transmission gear.

DATA CENTER TO THE SHORELINE FOR SUBSEA TRANSPORT

Terrestrial wavelengths are historically terminated at shoreline terminal equipment located at a cable landing station, where the traffic is groomed and re-transmitted onto more highly efficient submarine transmission systems. However, companies providing colocation services, such as Equinix, Inc., are now suggesting a model where the typical shoreline termination is bypassed and traffic flows directly to the inland data center [2]. Power feed equipment, of course, typically remains consolidated for all the cable's users at the shoreline. Among the several advantages of the new approach is the preservation of the high-performance submarine link all the way to the end user (of which there can be several, in diverse geographical locales). In some cases, submarine fibers can be cabled in the terrestrial network and included in the design rules governing the overall submarine link. Optimizing end-to-end links containing both undersea and terrestrial segments obtains a lower cost per unit capacity (cost-per-bit, CPB) than separately optimizing the two segments, as demonstrated in the study in [3]. This optimization includes the choice of fiber type, span lengths, etc. Such an analysis shows that there is no point in designing (and paying the high cost of) a submarine segment to deliver high OSNR when the terrestrial portion is not capable of supporting a similar OSNR. This immediately shows the value of upgrading the terrestrial connecting segments if possible. A secondary result of this work is that use of an advanced optical fiber on land incurs a smaller CPB penalty than conventional fiber when the repeater noise figure (NF), shoreline terminal equipment, or span lengths are suboptimal (often terrestrial segments must work with brownfield infrastructure with fixed hut spacing). This can be understood from a consideration of the Shannon Law scaling of spectral efficiency (how much information can be carried in a unit of bandwidth, bit/sec/Hz) with signal-to-noise ratio, as shown in Figure 2. This scaling is not linear but is logarithmic. Hence the slope of this curve is lower at the operating point of an advanced G.654 fiber than for a conventional G.652 fiber. If any network impairment occurs, the spectral efficiency will degrade for both fiber types, but more so for the G.652 fiber because of the steeper slope at its operating point.

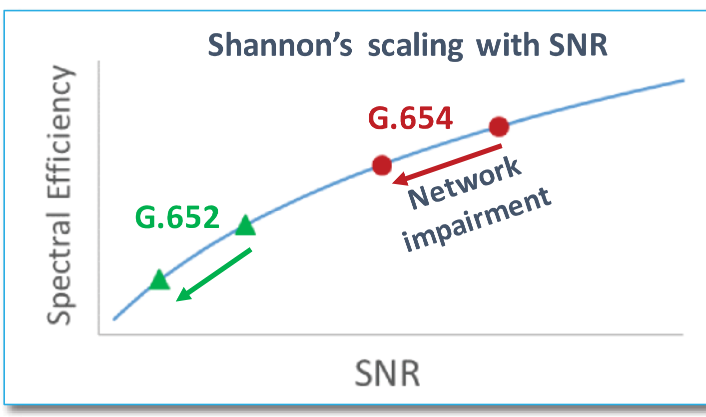

Figure 2. Because advanced (ultra-low loss G.654) fibers operate at higher spectral efficiency, they suffer lower penalty due to shortcomings in the network (because of the shallow slope of the Shannon curve at their operating point) than conventional G.652 fibers.

In the ideal scenario, the transoceanic link can be given continuity by transitioning to an advanced terrestrial optical fiber type at the shoreline, so that maximum capacity can be maintained on each transmission wavelength. Terrestrial G.654.E fibers don't have quite the performance of the best submarine fibers but can be cost effective and sometimes easier to cable. Appropriate care should generally be taken in cabling G.654 fibers, as the larger mode fields make the fibers somewhat more susceptible to enhanced loss through fiber bending (either in the cable construction or installation). For these reasons, OFS specifically designed the TeraWave® SCUBA 125 subsea fiber for use in both ocean and terrestrial cable constructions, for network designs requiring highest performance. Loose tube cable constructions are typically preferred, though ribbon cables are possible, as described earlier in the joint Furukawa/Fujitsu work. An additional advantage of high-performance fibers on shore is that the distribution of traffic out of the submarine fibers is simplified.

THE TRANS-OCEANIC LINK BETWEEN DATA CENTERS

Transmission under the ocean is fundamentally different than on land because of the difficulty of powering repeaters far from shore. Carrying the electrical power for these repeaters requires a sizeable conductor in the cable which adds size, weight and expense. Because of these conductor constraints, terminal voltages are typically quite limited, which constrains the amount of electrical power available for each repeater (optical amplifier). A new technology which is receiving a lot of attention lately is space-division multiplexing (SDM). The idea here is that the valuable repeater power can be most effectively used, in the sense of providing lowest CPB or highest throughput, by lowering the signal power in each fiber span and adding more fibers to the cable. These extra fibers have the added benefit of being a convenient unit of capacity trading (as opposed to optical wavelengths) which can be pairwise marketed to customers.

Until recently, cables with four to six fiber pairs were the norm, which enabled high power wavelengths to be used, thereby favoring optical fibers with highest possible mode field. 150 mm^2 is a favored mode field area which delivers unprecedented spectral efficiency on a single fiber. With an SDM strategy it is awkward to significantly increase the fiber count of the 150 mm^2 fiber because of bending loss concerns (assuming minimal increase in the cable size, the fibers will be more tightly packed in the SDM scenario). So even though the large mode area fibers might perform better, even at lower signal power, packaging concerns give reason to lower the fiber area. How low should one go? This now becomes a techno-economic question where a trade-off is made between the cost (added fibers and repeaters) and the capacity (the maximum information rate the entire cable carries) and packaging effects in the cable. Operating at lower signal power with longer spans lowers the spectral efficiency so more fibers are required to make up the difference. It is always possible to raise the capacity of a cable by adding more fiber pairs, but it is not always cost-effective to do so.

Recent studies have examined this cost trade-off for an array of fibers types [4], taking into account the technical and economic factors impacting the cost and capacity of an SDM cable. Table 2 shows optimum fiber type choices for different cable fiber counts for transoceanic distances (trans-Pacific is roughly twice the trans-Atlantic reach). Note that it is assumed here that the cable design does not need to be changed (made more expensive) as fibers pairs are added. In reality, there will be added cable cost for the very high fiber-count cases.

Trans-oceanic SDM cable systems currently in planning, as well as under construction, use 12 to 16 fiber pairs, which can be considered first-generation SDM, including sharing pump lasers across amplifiers in the repeater to save cost and increase reliability. For example, Google has announced a trans-Atlantic cable project named Dunant which will employ 12 fiber pairs, including power-optimized repeaters, with a design capacity of 250 Tb/s of data [5]. Google also announced the Equiano SDM cable from Portugal to South Arica. In the future, a second generation of SDM cables will deploy 24 to 32 fiber pairs.

Figure 3 shows the results of calculations from [4] that

Optimum # fiber pairs (fp) for given fiber effective area	Fiber mode effective area		
	150	125	80
Trans-Pacific	< 10 fp	10-24 fp	> 24 fp
Trans-Atlantic	< 16 fp	16-32 fp	> 32 fp

Table 2: Fiber Pair (fp) count regimes where advanced fibers with different core sizes are preferred to achieve lowest cost-per-bit transport, following Table II of Ref [4]. Note that all fibers are assumed to have 0.155 dB/km attenuation at 1550 nm, but the 80 mm^2 variant is still in R&D.

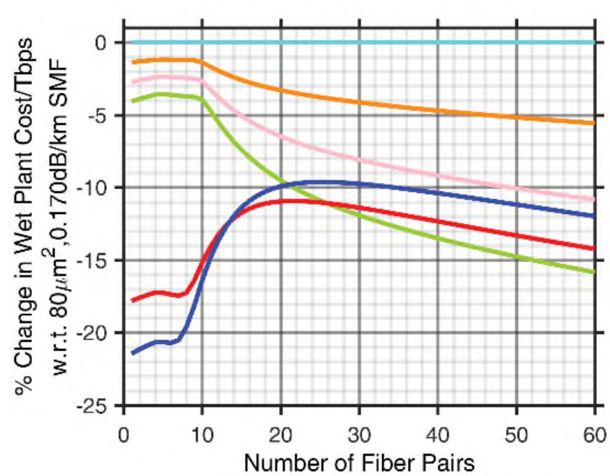

Figure 3. Minimizing CPB as a function of fiber pair count, using different fiber mode area choices, in a representative 11,000 km trans-Pacific SDM cable system example [4]. Note that an 80 mm^2 fiber with 0.155 dB/km loss (green curve) is still in R&D.

estimate the "wet-plant cost-per-bit" benefit from using an advanced, ULL fiber design, as a function of fiber pair count in the cable, referenced to a standard terrestrial-grade ULL fiber (0.17 dB/km attenuation) with standard mode area (80 mm^2) in light blue. Wet plant CPB results from dividing the estimated cost of the wet-plant assets by the calculated total capacity of the cable, as a function of number of fiber pairs. We exclude the dry plant costs on the shore because that equipment is independent of the fiber design chosen, whereas the optimum spacing between repeaters, as one example, depends critically on the fiber properties. The 150 mm^2 fiber shown in dark blue has been the workhorse of the past few years for non-SDM cable with 4 to 8 fiber pairs; however, it will be too bend sensitive to deploy in high fiber count SDM cables. On the other hand, the fiber shown in red, with 125 mm^2 mode field area, (such as OFS TeraWave® SCUBA 125 fiber) offers the optimum cost-performance trade-off in the range of 10 to 24 fiber pairs, giving good performance in higher fiber count cables. The importance of the 125 mm^2 is highlighted by the fact that the fiber represented by the green curve – a standard core size fiber with 0.155 dB/km attenuation at 1550 nm – is not a manufacturable fiber at this time. The fiber choice represented by the pink curve becomes the fiber of choice > 24 fp, for the trans-Pacific case, in accord with Table 2 even with the higher cost of the 0.160 dB/km case. This is because the bend loss in the tightly packed cable may preclude use of the higher mode area fiber.

CONCLUSIONS

Both the number of data centers and traffic between data centers continue to increase. That is giving opportunity for new technologies to help optimize optical transport for DCI. Trends that have been noted include wavelengths carrying ever higher data rates over terrestrial segments, SDM submarine cables with cost savings through more efficient use of limited electrical power, and migration of submarine terminal points inland from the shore. In all these cases, advanced G.654 fibers can provide enhanced DWDM capacity and cost savings. These fibers have been a de facto standard in submarine networks for some years now and are now gaining traction in terrestrial data center interconnects. The combination of ultra-low fiber attenuation and a range of large mode field sizes seem to ideally position G.654 fibers for new applications in the foreseeable future. STF

Trends that have been noted include wavelengths carrying ever higher data rates over terrestrial segments, SDM submarine cables with cost savings through more efficient use of limited electrical power, and migration of submarine terminal points inland from the shore.

ALAN MCCURDY *is a Distinguished Member of Technical Staff at OFS. He does technical business case development and support for new fiber products, marketing of the same, and, when time allows, works on advanced noise and fiber measurement problems in optical communications. Alan has worked in telecommunications since joining the Enterprise Networks Group at Lucent Technologies twenty years ago. Prior to that, he spent nine years on the Electrical Engineering faculty of the University of Southern California. He earned B.S. degrees in Chemical Engineering and Physics from Carnegie Mellon University, and a Ph.D. from Yale University.*

ROBERT LINGLE, JR. *is Sr. Director of Strategic Marketing at OFS in Norcross, GA, as well as Adjunct Professor of Electrical and Computer Engineering at the Georgia Institute of Technology. He has a research background in short pulse lasers and their application to fundamental processes in liquids and interfaces, with a Ph.D. in physics from LSU and a postdoc in surface physics at UC Berkeley. At Bell Labs and now OFS, he first worked in sol-gel materials chemistry and then managed the development and commercialization of many new optical fiber products. He currently works with customers and technology partners to help OFS understand and influence the technology trends in the market.*

REFERENCES

[1] T. Gonohe *et al.*, "Transmission performance of 600-Gbps per wavelength signals over installed slotted-core cable using G.654.E Optical Fibers," International Wire and Cable Symposium 2019, 9-1.

[2] F. Salley *et al.*, "We landed the cable; now what?," SubOptic 2019, OP8-2.

[3] A. McCurdy *et al.*, "Techno-economic study of optical links containing both submarine and terrestrial segments," SubOptic 2019, OP8-3. A. McCurdy, Submarine Telecoms Forum Issue 103, November 2018, pp 20 – 23 "Cost Effective Intercontinental Data Center Connections", https://issuu.com/subtelforum/docs/subtelforumissue103

[4] K. Balemarthy *et al.*, "Optimum fiber properties and pair counts for submarine space division multiplexing," SubOptic 2019, OP18-2.

[5] https://cloud.google.com/blog/products/infrastructure/a-quick-hop-across-the-pond-supercharging-the-dunant-subsea-cable-with-sdm-technology