WILL NEXT GENERATION MESH ARCHITECTURES CHANGE SYSTEM AVAILABILITY OR MAINTENANCE TARGETS?

Andy Palmer-Felgate (Facebook), Herve Fevrier (Facebook),
Stephen Fitzpatrick (DRG Undersea Consulting), Elaine Stafford (DRG Undersea Consulting)
Email: estafford@drguc.com
DRG Undersea Consulting, Headquarters Plaza, Morristown, NJ USA 07960

Abstract: Within the last few years, the world's largest technology companies (i.e. Internet Content Providers - ICPs) have jumpstarted the undersea cable market by driving the construction of multiple new cables on the world's biggest routes. Similar to the carriers who led the market historically and spread their capacity over multiple cables serving similar routes, these companies are developing robust mesh networks. But somewhat differently than the carriers who historically invested in consortium cables with an investment share, the next generation of meshes are largely built on individually-owned fiber pairs and/or spectrum. In the event of a cable outage, the owners not only have the ability to individually manage the protection of these fibers and their capacity, but they also have a stronger voice in the maintenance strategies of many cables.

Representatives of these newest cable owners, as well as industry observers, have occasionally remarked that these larger global meshes will mean undersea cables will no longer need to be extraordinarily reliable and MTTR targets can be meaningfully relaxed. Pundits have also suggested that these changes might also drive lower construction and operations cost of cables, i.e., lower individual cable reliability targets may mean lower-cost components, lesser route-engineering/marine costs, and greater reliance on spot-market maintenance services.

This paper examines these hypotheses and presents the results in a series of charts illustrating the impact on mesh-network availability of various parameters, such as cable fill, number of cables, cable MTBF (cable route protection, component reliability or FIT rate), cable MTTR (OA&M options), with the goal of shedding light on the trade-offs between investing in many low-cost cables that are repaired when convenient vs. fewer somewhat more expensive, reliable cables that are repaired as quickly as possible.

1. INTRODUCTION & ASSUMPTION

This paper focuses on how best to economically build <u>and maintain</u> robust transoceanic mesh networks connecting Data Centers which often demand 99.999% network availability.

The authors have modelled the two largest routes, TransAtlantic (TAT) and TransPacific (TPC) and explored those questions by looking at the capex and opex cost of various mesh configurations. Our analysis assumed the following:

Table 1: Assumptions

Point-to-Point	**TAT**	**TPC**
Nominal Capacity/fpr	24 Tbps	18 Tbps
Nom. Fpr Count/Cable	8 fpr	6 fpr
Nom. Capacity/Cable	~200Tbps	~108 Tbps
SDM* Capacity/fpr	20 Tbps	15 Tbps
SDM Fpr Count/Cable	16 fpr	12 fpr
SDM Capacity/Cable	~320Tbps	~180Tbps
Total Repairs/25yrs [1]	9.5	19.5
MTTR	3 wks	5 wks

*Space Division Multiplexing is clarified further below.

a. Modelling was <u>not</u> specific to any individual cable system or owner.

b. Modelling focused on meshes which are sized to provide transport capability comparable to the capacities forecasted by TeleGeography's GBFS, which predicts the need for 2 Pbps of TAT capacity and nearly 1 Pbps of TPC capacity in 2024. GBFS also forecasts that by this time, the vast majority of the demand will come from owners like ICPs who self-provision.

c. Our model characterized individual mesh architectures by their <u>size</u> (number of cables), fill (a percentage reflective of the Working/Protection ratio), and resultant protected capacity. The mesh's fill directly correlates to its efficiency (or economics on a cost/bit basis). Lower fills equate to more stranded cable investment, devoted solely to protection in the quest to achieve five-9's.

d. The model's cable-construction costs are based on recent representative contract prices which remain confidential.

e. The cable fault rates and MTTRs assumed are reasonable approximations of today's market averages based on published International Cable Protection Committee (ICPC) fault history and repair commencement statistics, plus assumed repair durations. Depth related variance in fault frequency is derived from data presented at SubOptic2016 [2]. The MTTR and MTBF were simulated using the methodology described at SubOptic2016 [3] modified using [2]. We have presumed that these ICPC statistics covering the last 10 years characterize the performance of new cables which will utilize today's robust cable-engineering methodologies.

f. Mesh availability is derived from an MTTR which is indicative of the duration of service-outage. The model's MTTR used to explore different fault scenarios presumes the service-outage commences only when the cable ship starts the repair's cutting drive for shunt faults (i.e. it ignores the time to mob, permit and transit to the fault site).

g. Consistent with today's open-system concept, the analysis focused only on the wet-plant reliability (which dominate availability).

h. We assume adoption of high-fpr SDM networks (16-fpr in the Atlantic and 12-fpr in the Pacific). These SDM cables are most likely to be implemented with more efficient cable architectures, which focus on optimizing the total cable capacity, rather than individual fiber pair capacity. SDM has recently started to become somewhat synonymous in the undersea cable market with not only increases in fiber count, but also changes in a number of wet-system design parameters which together result in the cable operating in the more efficient linear regime. This, in turn, results in: (i) more efficient use of the available electrical and optical power; (ii) longer repeater spacing (meaning fewer repeaters); and (iii) lower repeater output power, which together also enable use of a lower-effective area (lower-cost) fiber. Coincident with these SDM architectural changes, system suppliers are increasingly adopting more efficient repeater architectures with greater levels of pump-sharing [4]. The net benefit of these approaches is a more efficient network, compared to more traditional system-design approaches which maximized capacity/fpr.

i. Our reference model assumes component FIT rates for cables which are characteristic of traditional high-reliability undersea system design. Reduced FIT-rate assumptions have also been separately explored to examine their potential impact on mesh MTBF and availability, to study the future possibility that recent pump sharing architectures might provide future opportunity for slight relaxation in individual pump FIT rates, without impacting the overall system FIT rate.

j. Our model assumes that the TAT cables traverse the somewhat risk-prone European shelf and the TPC cables similarly traverse risky areas in Asia, both of which are required to serve major Data Center hubs such as London, Amsterdam, Singapore and Hong Kong.

2. METHODOLOGY

The mathematical model utilized to analyze the availability of various mesh-network configurations assumes that the MTBF and MTTR values follow exponential law. Modelling concepts are typical of those found in [5]. We have calculated the availability of innumerable mesh sizes for each ocean, each capable of transporting a specific amount of protected capacity- with the goal of achieving 99.999% availability with minimum investment for a specific cable architecture. Each system was modelled as an **"m"** *(number of "working" cables required)* out of **"n"** *(number of total cables)* solution. Modelling results are built upon an MTBF derived from: (a) specific deep-water trunk cut/shunt fault assumptions; (b) specific shallow-water shore-end cut/shunt fault assumptions; (c) a shunt/cut ratio; (d) estimated wet-plant MTBF (due to component failure alone) based on typical/representative undersea reliability calculations, and (e) MTTR. TAT and TPC systems use different assumptions for each of the above.

3. CONCLUSIONS

Our conclusions are for the most part not new. The time-tested system design and maintenance approaches which the industry has depended upon for many years continue to apply on the anticipated larger scale of today's and tomorrow's global network needs. However, if meshes continue to grow without fundamental new technology or product breakthroughs, we can imagine that in the distant future, huge meshes may motivate new maintenance methodologies. I.e., extremely large meshes many years from now may mean that both MTBF and MTTR eventually become less critical to the global undersea cable infrastructure.

<u>Conclusion #1</u>: Achieving 99.999% availability is dependent on balancing the number of cables in a mesh and their fill, based on the route's MTBF and MTTR. Larger meshes enable more-efficient, higher-fill cable utilization.

This first conclusion is illustrated in Figure 1, which depicts the availability of a variety of TAT mesh architectures over a range of repair times (MTTRs). For the Atlantic, we have modelled 2- 6-wk MTTRs, the latter of which is unusually high for the Atlantic. We believe a three-week repair time is more realistic. What one can see from is:

a) The "flatter" lines (meshes), which stay above the goal of 99.999% availability for the longest repair/MTTR times, are the most robust. E.g., the lowest-capacity mesh shown (**4 cables @ 0.2 Pbps at 25% fill**) meets the goal for all assumed MTTRs.

b) A similarly-robust 16-cable mesh which transports **ten times** that capacity (but operates much more efficiently **at over 60% fill**) also stays above the five-9 availability target over the full range of MTTRs.

c) None of the other mesh solutions illustrated achieve 99.999% at this unusually-long TAT MTTR, but we have illustrated three examples of meshes performing at 99.999% (each also have 60+% fill) at a more-typical 3-week MTTR.

<u>Figure 1</u>: TAT Fill & Mesh Size Needed to Achieve Five-9's with Increasing Demand

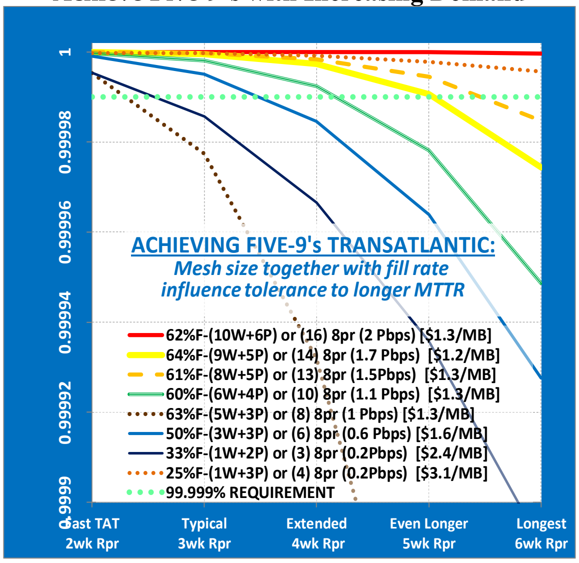

Because TPC routes have both:

a) a shorter MTBF (the Pacific's more hazardous ocean bottom increase the number of cable faults from ship anchors, fishing, or seismic activity), and

b) a longer MTTR (largely due to repair-permit delays and longer transit times),

achieving 99.999% is more difficult for TPC networks than it is for TAT networks. Figure 2 below illustrates 0.4Pbps and 0.6Pbps TPC mesh availability for different protection ratios over a range of MTTRs. Mesh sizes that achieve the goal beyond a typical 5-week repair interval are more inefficient than in the Atlantic, due to the longer MTTR and greater MTBF.

Figure 2: TPC Fill & Mesh Size Needed to Achieve Five-9's with Increasing Demand

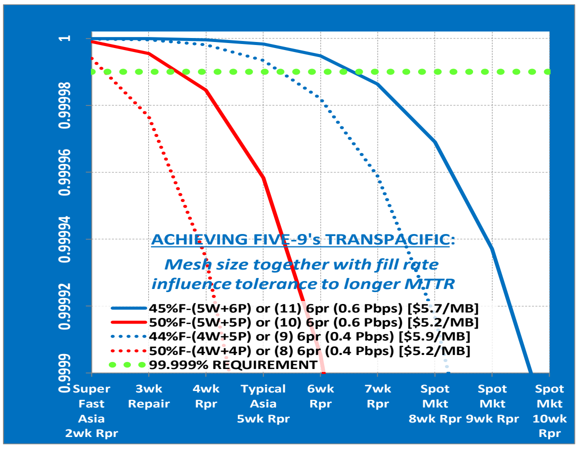

Conclusion #2: Cable faults arising from external aggression are more likely to be the root cause of any cable failure than a fault due to a component failure, and control overall mesh availability.

Mesh networks protect equally well against any type of cable fault, whether the fault is caused by external aggression or the fault caused by a single repeater failure.

- Our model assumes that the number of externally-induced faults over cable's 25-year Design Life is typically on the order of 7.5 and 15.5, respectively, for hypothetical point-to-point TAT and TPC cables.
- The typical forecasted number of faults in 25-yrs resulting from a cable's component failure (based on FIT rates and architecture) is typically somewhere between 1 and 3 (with larger cross-section and longer cables having higher fault predictions).
- TPC cables thus have a double-whammy: more externally-induced faults and a higher likelihood of failure due to component failure.

Because externally-induced (aggression) faults dominate the total number of expected cable faults for most cables, the age-old industry adage that a cable route should be engineered as robustly as economically possible (e.g. adequate armor, burial and other protection) remains true. Investment in armouring and burial should not be scrimped on, as that could significantly decrease the MTBF.

This analysis has also demonstrated that a mesh designed to robustly achieve 99.999% is also reasonably likely to be robust against incremental increases in cable FIT rates. I.e., marginal increases in repeater FIT rates can be tolerated without significant impact to total mesh network availability. In Fig. 3, we demonstrate the availability impact of doubling a TAT mesh's constituent cable's estimated number of cable faults which are not caused by external aggression. As can be seen, meshes which robustly achieve 99.999% across an extended range of MTTR assumptions are also likely to be robust against such increases in repeater FIT rates. Conceivably, recent repeater architectures with the increased levels of pump sharing/redundancy might tolerate incrementally higher FIT-rate pumps without significant impact to cable MTBF. Most other repeater components, which are not protected (redundant), do not have this luxury. Thus, the authors are specifically not advocating a broad relaxation in repeater FIT rate targets to lower the repeater and system cost. Like scrimping on cable protection, significant relaxation in repeater FIT rates could impact the cable overall cable MTBF to the point where it has a significant impact on a mesh's ability to achieve 99.999% unless the mesh becomes less efficient (and more expensive) for comparable transport capacity.

Fig. 3: TAT Network Tolerance to Higher MTBF

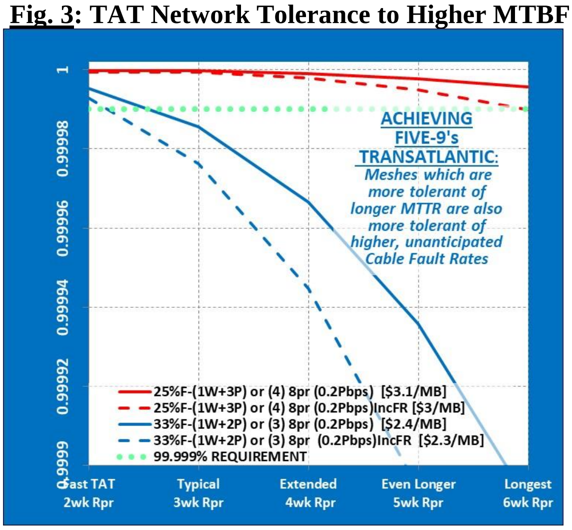

<u>Conclusion #3</u>: Unless mesh networks across a single route are built with more cables than necessary to achieve five-9s, low MTTRs (or reliable repair service) remain critical to achieving 99.999%, especially in the Pacific.

The earlier Figures 1 and 2 illustrate that lower MTTRs correlate with higher mesh availability. Perhaps less immediately obvious is the fact that the more-efficient, more-robust meshes typically have a greater number of cables. The TAT data illustrated shows that, except for the inefficient 25%-fill/utilization (four-cable mesh), the other meshes are increasingly efficient and increasingly robust as more cables comprise the mesh. The reason is that although larger meshes with more cables are more vulnerable to multiple, simultaneous cable faults, the larger meshes also tolerate those circumstances better than smaller meshes.

This means that less investment in extra cables for protection is required if one is confident in shorter repair intervals. This is most-easily illustrated for TPC cables (in Figure 4, which focuses on comparing equivalent-capacity meshes of different architectures across a range of MTTRs).

<u>Figure 4</u>: You Get What You Pay For (TPC)

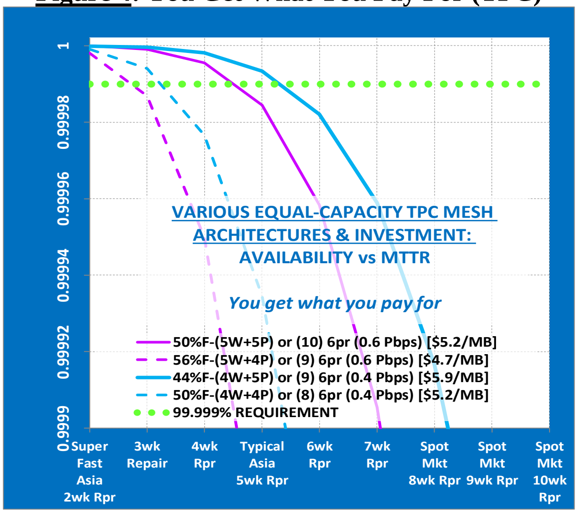

The more-expensive, more-robust meshes (solid lines) cross the five-9 goal line at MTTR/Repair-Times which are approximately 2 weeks longer than the less-expensive, less-robust meshes (dashed lines). Thus, if a mesh owner can build their network with confidence in faster repair times, they can invest in fewer cables (a more efficient or more economic mesh). For each of the cases illustrated, it is just one more cable that provides this two-week improved tolerance against longer repairs. The cost of that extra cable is significant, both in terms of individual cable supply cost (>$250M for a TPC cable), but also the related development costs, landing costs, and both recurring wet and dry OA&M cost. Subscribing to wet-maintenance service (zone or private) often costs millions of dollars per year for a TPC cable (which adds up to a lot over the cable lifetime). Moreover, it is reasonable to presume that the cost of individual repairs will be greater on the spot market than with a zone or private subscription. Avoiding all of these initial and recurring investments for one or more the cables must be weighed against the cost of instead investing repair services which deliver shorter repair intervals. Our math suggests that meshes must get very, very large (beyond sizes envisioned in the next several years) to make this trade worthwhile, especially in the Pacific. Similar conclusions are true for TAT meshes.

We have also tested our thesis that the spot market is a less economic overall maintenance approach by comparing the cost of various meshes which meet the five-9's objective at different capacities. We have done this for both the TAT and TPC routes, the latter of which for reasons explained earlier provides a starker difference. Our calculations estimate that the cost/bit penalty of relying on the spot market, for TPC meshes ranging in transport capacity between 540 Terabit and 1080 Terabits, ranges between 20-30%. This is simply due to the need for more cables in the mesh to reach five-9's when the MTTRs are longer.

<u>Conclusion #4</u>: Single-End Feeding capability is another way of effectively lowering the MTTR and is therefore pivotal to achieving efficient meshes with high availability- especially in the Pacific.

The value of Single End Feeding is often debated, as the electrical-power demands of higher fiber count cables challenge the ability to preserve this age-old system design requirement. If a system can be Single-End Fed, the <u>service impact</u> of any shunt fault can be minimized by powering the faulted cable to the shunt from each shore. This SEF method keeps the faulted cable operational until the arrival of the repair vessel. Thus, the time required for repair-permit approval, vessel mobilization, and transit to the site of the fault do not impact the service. I.e., they do not contribute to the average <u>service</u>-time MTTR in our model.

On average, reports have suggested that cable breaks happen somewhat more often than shunt faults. [7] For simplicity, we assumed 50% and that the MTTR for TAT shunts (in terms of service) is reduced to 1 week (from 3 weeks), resulting in average of 2-wk MTTR. For TPC networks, 5 wks reduces to 3 wks. Fig. 2 is replicated below (now Fig. 5) with overlays illustrating the impact of this MTTR improvement.

<u>Figure 5</u>: The Criticality of Single End Feed -TPC

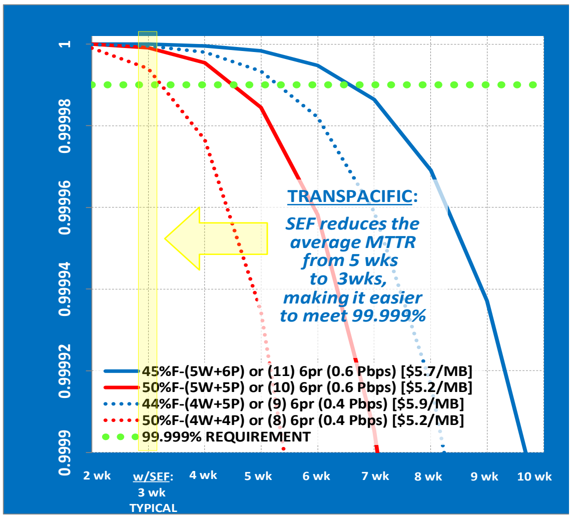

It is intuitively obvious that because larger meshes become more efficient (in terms of a cost/bit), the percentage improvement in cost/bit afforded by SEF lessens also. In other words, the importance of MTTR diminishes as meshes add more and more cables (still with a five-9 availability goal). Our analysis supports that conclusion.

<u>Conclusion #5</u>: A mesh built from trunk/branch cables which utilize fiber-switched branching units (along with terrestrial connectivity between branch landings) effectively also lessens the MTTR by reducing the impact of all types of shallow-water faults (not just shunts) and thereby (depending on terrestrial costs) can more economically achieve 99.999%.

After reading our conclusions to this point in this paper, one might easily presume that the authors advocate building more and more point-to-point cables to create more-capable, more-efficient, more-robust transport across the oceans. Not so. We do believe that more cables will continue to be needed to serve the ever growing transoceanic demand, but we do not believe that single point-to-point cables within a huge mesh are necessarily the best approach. Rather, especially for the TAT route, there is a historically-tested architecture that remains potentially attractive to lessen overall investment cost, while improving availability. Recognizing that the majority of shallow-water faults in TAT Systems are off the European coastline, it is reasonable to contemplate protecting against these shallow-water faults by having multiple landings in Europe, each capable of transporting the entire trunk's capacity. Such cables are more expensive than point-to-point cables, but have lower service MTTR, which allows fewer of them to be built in a mesh targeted at 99.999% availability.

Building a mesh with trunk/branch cables instead of point-to-point cables is most effective if: (a) the branching repeater has fiber switching (so that in the event of branch cable fault, the capacity on <u>all of the trunk</u> fibers can be routed to the unfaulted branch); and (b) the two landing points are connected via an <u>economical</u> terrestrial link that can transport the rerouted demand back to its intended destination- making the restoration path the alternate branch and terrestrial route.

We have modelled both TAT and TPC T/B to Pt-Pt cable meshes, assumed both were capable of SEF (so that the improvement resulted only from the optical switching to the branch). Like other strategies which allowed meshes to become more efficient with lower MTTRs, the cost-advantage diminished as the mesh's capacity grows. The savings is highly dependent upon the cost (length) of the branches and the cost of the terrestrial interconnect. Clearly, if the two together add up to close to the same value as the trunk cable, the advantages are negligible. If they are substantially lower than the trunk price, the advantages are measurable. At lower mesh capacities (<1 Petabit) our estimates suggest a few cables might be saved, with overall cost/bit savings of perhaps 30%. Above those capacities, the savings diminish rapidly as the added cost of the branches drains the benefit.

Building a five-9 mesh using Trunk & Branch cables, instead of point-to-point cables, has the greatest cost/bit benefit:

- for smaller-capacity networks. Refer to the top blue line with lower-capacity.)
- for networks where the branch and terrestrial landing-point interconnection costs are least. (Refer to the slopes below.)

<u>Figure 6</u>: TAT Trunk & Branch Benefits

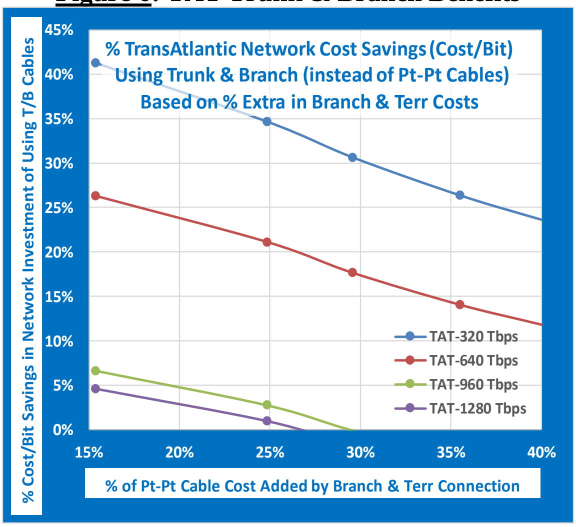

<u>Conclusion #6</u>: Building highly-resilient, large-capacity transoceanic meshes is best most economically and efficiently achieved using SDM cables.

This conclusion may appear counter-intuitive as historically, maximizing a cable's individual fpr capacity was the means to maximize cable capacity. However, today's SDM cables are more efficient/economical overall, as shown in Figures 7 & 8. The first compares the cost/bit of a various-capacity meshes built with "nominal" cables (the red line) vs. SDM cables (the blue line) over a range of transport capacities. The second shows that for equivalent transport-capacity, the TAT cost/bit is lower for meshes using SDM.

<u>Figure 7</u>: 12-fpr SDM Vs 6-fpr TPC Meshes

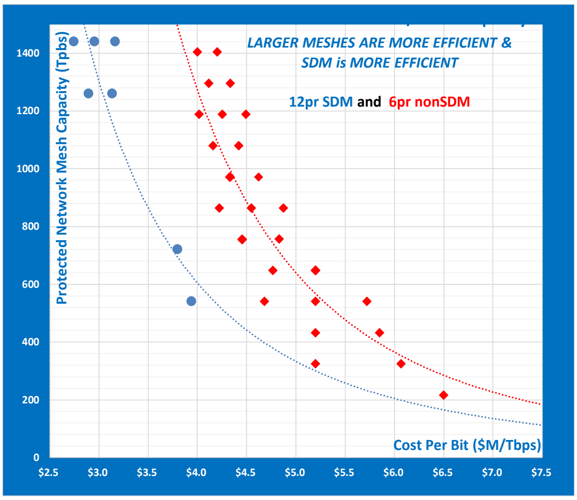

<u>Figure 8</u>: TAT SDM Efficiency

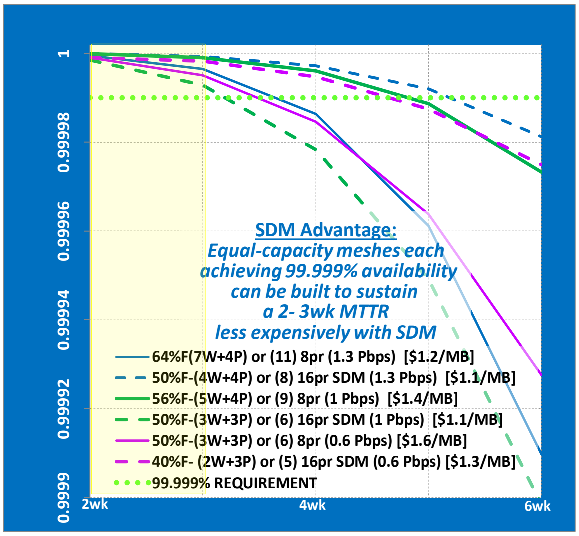

An added benefit of SDM is that the constituent fiber pair cost, when compared to a non-SDM cable's fpr cost, is less. This makes ownership of fiber pairs more affordable, effectively lessening the

Minimum Investment Unit (MIU) to join a cable. With more partners, ICP's with an interest in owning fiber pairs on many diverse cables can spread their capital across more cables, and more cost-effectively grow their mesh to five-9 availability.

4. SUMMARY

The original intent of the authors when proposing this paper to SubOptic was to examine the notion that both MTBF and MTTR make little difference to transoceanic mesh availability and cost and, as a result, that the undersea communications market would evolve towards construction of significantly lower-cost, lower-reliability cables, all of which might eventually be maintained by spot-market wet maintenance services. What we've shown is that for a network availability goal of 99.999%, **both MTBF and MTTR matter**. Moreover,

a. while incrementally-lower individual cable reliability might be tolerable in very large meshes (which we imagine still to be years away), significant increases in MTBF will make achieving five-9's a challenge,

b. MTTR will remain important for capacities forecast in the coming several years– especially in the Pacific. Thus, the authors remain believers in strategies that minimize MTTR, e.g. SEF and maintenance services that provide greater assurance in acceptable MTTRs than the spot market promises.

c. designing highly-reliable, cost-efficient mesh networks starts with architecture. Options for SDM and T/B with fiber switching should both be considered.

The conclusions above generally apply to cable networks across the globe, although meshes are not developing elsewhere at the same rate as they are on the two routes described. Moreover, because the capacity demands along other routes (e.g. Eur-Asia and Africa, but not Intra-Asia) also lag the Atlantic and Pacific, building larger meshes is less economical there, making MTTR all the more important. These other major cable routes often include many more cable landings as part of a single cable, but these landings are not necessarily acting as diverse restoration paths (analogous to the T/B solution described earlier), but rather often represent singular landing vulnerabilities when looked at as individual point-point connectivity. Thus, just like for the TAT and TPC, we cannot see compromising MTTR for those cable routes. For routes which are not be yet connecting large data centers (which demand 99.999%), the lower number of cables on these routes, along with a low MTTR, may reasonably allow a lesser availability to be met.

We can imagine a future where very large meshes make MTTR less critical, but we do not see that within the planning horizon. Compromising MTTR and/or MTBF means less efficient, more expensive meshes, if five-9's remains imperative.

5. REFERENCES

[1] 'Palmer-Felgate, Andy', 'ICPC 2018 Plenary Meeting, Cape Town, South Africa, 25th – 27th April 2018

[2] M. Kordahi, et al. 'Global trends in Submarine Cable System Faults'. SubOptic 2016 Proceedings.

[3] A. Palmer-Felgate, P. Booi How Resilient is the Global Submarine Cable Network? SubOptic 2016 Proceedings.

[4] "Linear and nonlinear submarine networks", E. Mateo, Workshop 7 "Submarine Systems", ECOC 2018 Roma.

[5] 'Jerome/Klion', 'RADC-TR-77-287 "A Redundancy Notebook", dated December 11, 1977'.

[6] H. Fevier, S. Grubb, N. Harrington, A. Palmer-Felgate, E. Rivera-Hartling, T. Stuch, "Future submarine communication needs towards 1 Pb/s cables", ECOC'2018 (September)

[7] Y. Ogasawara, T. Kaida, K. Goto 'Lessons Learned from Past Cases for Achieving Disaster Resilient Routing and Economically Viable Maintenance Operation'. SubOptic 2016 Proceedings.