Open Undersea Cable Systems for Cloud Scale Operation

Jamie Gaudette and Tim Stuch

Microsoft, Redmond, WA; jagaudet@microsoft.com; tstuch@microsoft.com;

Abstract: A true open cable system is designed specifically to operate in a disaggregated, vendor agnostic manner. We outline Microsoft's approach to open cable systems and discuss the technical challenges.

OCIS codes: (060.2330) Fiber optics communications; (060.1660) Coherent communications; (060.4250) Networks

1. Introduction

Cloud service providers are faced with accommodating exponential growth while maintaining near-perfect service availability. To meet the challenge, we focus on simplicity and efficiency. To continually maintain best-in-class hardware efficiency, cloud providers have pushed the Open Line System (OLS) concept, where the line system components (amplifiers, ROADMs, etc.) are disaggregated from optical transceivers associated with that line system [1]. For simplicity, centralized SDN controllers originally focused on layer-3 service orchestration [2,3], have been extended to control the OLS, with the goal of complete end-to-end service automation.

Interestingly, some of the principle concepts of the OLS have been in practice for many years on legacy undersea cables. The very early adoption of disaggregation in the undersea community was driven by a need to reduce upgrade costs and maximize spectral efficiency while operating over the extreme case of a long-term line system, the 25-year lifetime cable. In almost all cases, legacy systems were purchased closed, with the cable supplier's transponders and subsea line terminating equipment (SLTE) attached at start-of-life. These systems were then later converted to notionally open systems through SLTE upgrade. The conversion of a closed cable system to open does have significant drawbacks from the perspective of a cloud provider. First, the true performance of the cable system is abstracted through the original supplier's proprietary hardware, and results are confidential. This makes it difficult for other SLTE vendors to estimate performance without lengthy field trials, slowing the upgrade cycle. Next, since the cable system hardware and software was not designed to be open, the operators are left without the option to integrate cable monitoring software into SDN controllers. Instead, operators combine multiple network management systems to run the end-to-end service, adding complexity and risk.

A true OLS or open cable system (OCS) is one that is designed specifically to operate in a disaggregated, vendor agnostic manner, amenable to integration into centralized SDN controllers. In this paper, we outline how Microsoft extends our approach to OLS to define and operate OCSs.

2.0 Specification and Measurement of the OCS

The OCS concept is illustrated in Figure 1a. Like the OLS, the OCS includes open, programmable hardware with vendor agnostic APIs, running directly on the hardware, offering streaming telemetry rich in data. Today, a good implementation is represented with REST APIs for proactive data collection, efficient UDP-based alerting and alarming for reactive events, and simple CLIs for configuration. However, to obtain simplicity, the OLS is not concerned with squeezing the last ten-or-twenty percent of capacity out of the system. Spectrum on an undersea cable can cost 2-orders of magnitude more than spectrum on a terrestrial long-haul segment and therefore the OCS must allow the extraction of all efficiencies. To this end, we see two necessary additions to the OCS over and above the OLS: (1) comprehensive specification of all performance impacting attributes and (2) the ability to measure these specifications at acceptance and continually measure throughout the lifetime of the cable while separating any impact from SLTE.

For specification, the table in Figure 1b is used. Before wet-plant procurement, the table from all wet-plant proposals have vendor information removed and are shared with potential SLTE providers for 3rd party analysis of the

M2E.1.pdf

best technical solution. After procurement, both the average and worst case values are verified at acceptance. Once in-service, the worst case values are actively measured and used as alarm thresholds. Active measurement is discussed later in this section. Many of the parameters in the specification table are intuitive, such as power and OSNR, however, we pay particular attention to the shape of the spectrum using slope of tilt and gain deviation parameters, illustrated in Figure 1c. We have found that tilt and gain deviation can dominate transmission performance if allowed to become excessive, and can impair capacity as the cable system ages. It should be noted that end-of-life and repair specifications are also required, but have been left out of this example for simplicity.



Figure 1. (a) Open Cable Concept, (b) Open Cable Specification Table, and (c) Illustration of Gain Deviation and Slope of Tilt



Figure 2. Relationship of Components in OCS system and Measurements of the Specification Table

To verify the specification table at acceptance, and continually throughout the lifetime of the OCS, we introduced the concept of the Open Cable Interface (OCI). The OCI combines coupling to the wet-plant, high resolution spectral measurement, and fast power monitoring to resolve many of the target specification in real-time. The simplest implementation of the tracking engine includes an Optical Spectrum Analyzer (OSA) and fast Power Meter (PM). Co-ordination across both ends of the OCS is required to remove SLTE contribution to the specifications, so the OCI must be connected to the central SDN controller for correlation with the far-end.

M2E.1.pdf

As shown in Figure 2, the OCI combined with SLTE and wet-plant supervisory now provides direct measurement of all performance impacting parameters in the specification table, except fiber core size, and allows for cable-specific SLA tracking while eliminating lengthy field trial cycles. Our preference is the ability to verify fiber core through measurement, a work around is to obtain fiber manufacturing data with core size verification as part of system acceptance. To simplify the co-ordination of the three separate systems, a common set of APIs accessed directly by the SDN controller are necessary.

3. Challenges

Disaggregation has challenges, and the OCS is no exception. In an OCS, fault segregation to the cable alone is imperative. The goal is to have the OCI provide that function by monitoring the output of the cable and comparing to far-end input from SLTE. Simple UDP based reactive alerting, or proactive data collection with trending, can trigger notification of spectrum degradation or power fluctuation, in real-time, at the input and output of the cable separately, and should be extended to include alarming on OSNR degradation and end-to-end tilt degradation. Alarm thresholds should be set to the worst-case column in the specification tables.

Monitoring points provided by the OCI are limited to the end points of the OCS. For parameters that can vary per span, such as span loss and slope of tilt, there is a desire to track the specification on a per span basis. While traditional cable supervisory can track many of these parameters, the slope of tilt cannot be tracked today. We prefer this functionality reside in the cable supervisory system, and are investigating advances in C-OTDR based line monitoring equipment that can accurately resolve the slope of tilt on a per-span basis.

Cable systems are often shared amongst many operators. In an open cable system, each operator can select their own SLTE and OCI vendors. If each operator has their own fiber pair, segmentation of the OCI is straight forward, with a separate OCI allocated to each fiber pair or operator. If multiple operators share spectrum on a fiber pair, the OCI must have the ability to partition separate secure access for each operator on the shared fiber. In a shared spectrum use-case, the OCI has the added benefit of providing unbiased 3rd party monitoring of all parties on the fiber pair. Unlike the OCI, all operators will be tied to the same cable supervisory system. Therefore, it is imperative that supervisory systems also allow secure, partitioned access to the supervisory system for all operators.

Specification and acceptance based on the cable's performance elevates the importance of the optical parameters and necessitates more scrutiny when designing and laying the cable. We have found that this can increase system costs or slow cable deployment if the specifications are treated as ridged boundaries. In practice, we have found that close partnership with our wet-plant suppliers to grant exceptions on a case-by-case basis can allow speedy, efficient cable lay without significant compromise on performance or intent.

Lastly, acceptance on the OCS is more complicated than a closed system. Acceptance should be performed on all fiber pairs, and requires three-way co-ordination between SLTE supplier, wet-plant supplier, and the operator to verify all parameters in the optical specification table. Ideally, all purchasers share acceptance outcomes with one-another and collaboratively agree on overall contract acceptance.

4. Acknowledgement

We would like to acknowledge Dr. Georg Mohs at TE Subcom, JJ Crawford and Stephen Grubb at Facebook, and David Mackey and Tim Alfaro at Microsoft for their contribution to Open Cable Systems.

5. References

[1] M. Filer et al, Elastic Optical Networking in the Microsoft Cloud, JOCN, Vol. 8, Issue 7, pp. A45-A54 (2016)],

[2] P. Sun et al, A Network-State Management Service, Sigcomm 2014

[3] Jain et al, B4: Experience with a Globally-Deployed Software Defined WAN, Sigcomm 2013