VIRTUALIZING THE SUBMARINE NETWORK REQUIREMENTS, ARCHITECTURE, AND USE CASES

Philippe Perrier (Globenet), Mark Hinds (Ciena), Bogdan Cotruta (Ciena).

Email: philippe.perrier@globenet.net

GlobeNet, 200 East Las Olas Blvd, Suite 1700, Fort Lauderdale, FL 33301, US

Abstract: Most submarine cable systems being deployed today have been designed as Open Cable Systems (OCS), where the line system (mostly the wet plant) is “disaggregated” from the transponders. At the same time, advances in transmission technology (advanced modulation formats, coherent detection, digital signal processing…) can deliver massive capacity per fibre pair, not only on the newer OCS but also on legacy (closed) submarine cable systems.

Both transformative shifts have given cable owners the opportunity to introduce a new business model: Spectrum Partitioning. In this model, a cable / fibre owner can offer parts of the total cable / fibre bandwidth to prospective customers. Depending on the customer needs, spectrum partitioning can take several forms, from the degenerate case of a dedicated fibre, to managed services, managed spectrum, and finally shared spectrum. In this last option, the spectrum customer not only owns the transmission equipment (SLTE) but also controls its spectrum. No matter the spectrum partitioning option, the entire spectrum must be managed correctly, avoiding conflicts and minimizing impacts of actions of one spectrum owner on another. To this end, a way of orchestrating these actions and responding to fault events is required in order to guarantee flexibility and security to all spectrum customers.

1. INTRODUCTION

The motivation for spectrum partitioning from a cable operator (owner) point of view may include: selling reservations and physical access to unused spectrum when available. This allows an owner to monetize part of an existing fibre pair or an unlit fibre pair by hosting multiple customers with needs for only a subset of the spectrum (and related capacity). The owner still retains overall management responsibility for the fibre pair, including power management when upstream customer faults occur.

The motivation for spectrum partitioning from a customer point of view is to invest in immediate needs without having to take ownership of an entire fibre pair. Customers will have the choice to select the modem technology, independent of the owner options. Some customers will further be motivated by the autonomy offered in some spectrum partitioning options, where they can choose when new channels are added or modem capacity upgraded, without having to co-ordinate with the cable owner. At the limit of autonomy, some customers would setup their own SLTE working on a virtual submarine fibre pair provided by the owner solution.

This paper will discuss the motivation for spectrum partitioning, will present the system architectures of the different spectrum partitioning options, and the key functions of each architecture will be described. In addition, the benefits of each spectrum partitioning architecture will be compared from both an owner and a customer perspective.

Attributes to ensure successful spectrum partitioning (i.e., that network integrity is maintained when offering spectrum services) will also be discussed, including considerations of spectrum partitioning on legacy versus new cables. A management architecture using open APIs is proposed upon which orchestration functionality can be built.

2. FUNCTIONS and ARCHITECTURES

In developing spectrum partitioning options, the following functionalities are identified as key requirements for a robust and scalable architecture; [1]

- Spectrum Security (in add direction),
- Spectrum Privacy (in drop direction),
- Spectrum Monitoring, and
- Power Management.

The security and privacy requirements shape the choice of filter technology used to provide Add / Drop direction barriers between customers in this shared environment. Signal integrity for all customer channels is achieved by maintaining spectral isolation between customer allocations and limiting noise addition penalties (or superposition).

The spectrum monitoring and power management requirements determine the ancillary equipment and functionality needed to manage the system in both normal and faulted conditions. The spectrum monitoring function must measure spectral power with enough resolution to accurately measure per channel power, thus allowing verification of the agreed upon launch powers and spectral density at the customer / owner interfaces and wet plant interfaces. The monitor function should assist in isolating faults through detection of power changes vs set baselines (also referred to as Threshold Crossing Alerts) for per channel power and total power at each interface point. Another recommended feature is an output monitor tap for an external OSA connection, which can be used for new customer insertion and fault isolation activity.

Power management involves several functions including: active per channel power level adjustment and idler channel insertion in unused spectrum or in place of faulted channels or idlers. The decision to replace a customer channel or allocation with idlers is a joint responsibility shared by the customer and owner. This shared responsibility also applies to the decision to switch back to traffic when the fault clears.

Different spectrum partitioning architectures allow multiple customers to share a fibre pair including "owner-managed" configurations and "shared spectrum" arrangements.

The owner "managed" configurations include managed services and managed spectrum.

Managed services or client port leasing [2], illustrated in Figure 1, is the most common method for multiple customers to share a fibre pair, but does not really fall under the spectrum partitioning concept.

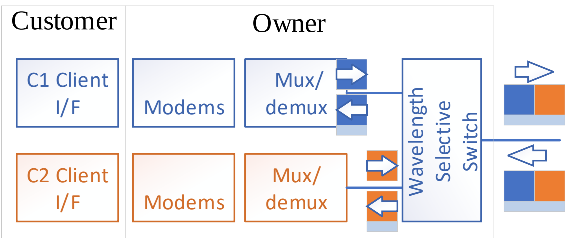

Figure 1: Managed Client Services

Managed spectrum or Line Mux/Demux port leasing allows an owner to dedicate a filtered mux/demux structure for each customer. This achieves the privacy and security requirements, as illustrated in Figure.

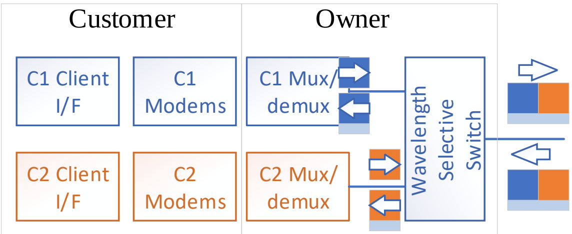

Figure 2: Managed Spectrum Ports

The "Shared Spectrum" configurations that have been deployed and discussed recently include:

- Broad Band Coupler/Splitter,
- Fixed Band Filter mux/demux,
- Flexible Filter Add only, and
- Flexible Filter Add /Drop.

A broad band splitter/ coupler configuration, depicted in Figure 3, has been the traditional method to support shared spectrum (in cable system upgrades in particular) at low cost, high reliability and with full flexibility. This option exhibits several critical drawbacks: no spectrum security or privacy and no common power management scheme. Note that due diligence is needed to ensure that work on one SLTE does not affect traffic from the other SLTE. Furthermore, this configuration does not block noise addition across the full band which introduces an additional Tx OSNR penalty [3]. For these reasons, this configuration is not well suited to spectrum partitioning and should be reserved for a single owner network supporting different SLTEs from successive upgrades.

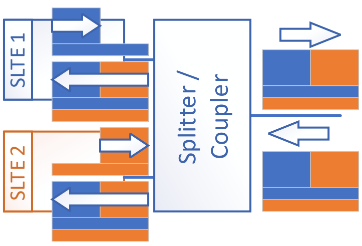

Figure 2: Broadband Coupler/Splitter

Fixed Band OADM filters can be used in place of the traditional coupler / splitter approach and provides spectrum security and privacy, but this configuration is inflexible to changes in allocated bandwidth, with filter skirts (dead band) wasting some spectrum. The Fixed filter configuration, as illustrated in ***Figure 4***, also encompasses the C & L band application, where the two bands could be different customers and monetized separately.

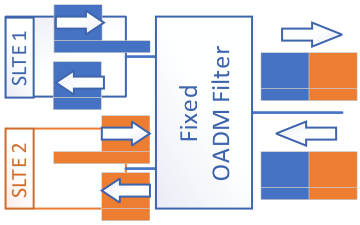

Figure 3: Fixed OADM Filter

Wavelength Selective Switch technology allows for a Select on Add Tx direction, which provides the spectrum security of Tx filters mentioned above but has the added benefit of flexible filter width and much smaller filter dead bands. The down side of unidirectional ADD Tx filtering is that it uses splitters to broadcast in the Drop / Rx direction which does not provide spectrum privacy, as illustrated in ***Figure 5***.

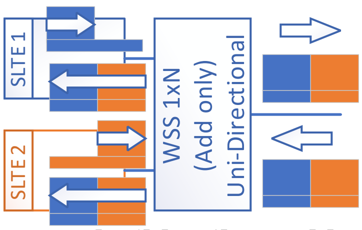

Figure 4: Flexible Filter Add Only

The use of Wavelength Selective Switch technology in both Add Tx and Drop Rx direction provides both spectrum security and spectrum privacy, while maintaining the added benefits of flexible filter width and much smaller filter dead bands. The added privacy feature is depicted in ***Figure 6***.

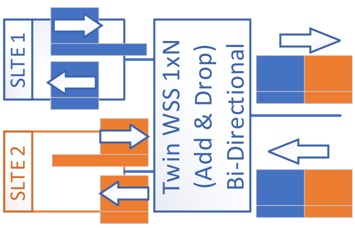

Figure 5: Flexible Filter Add/Drop

The flexible filter add/drop configuration is the preferred option for a Spectrum Partitioning solution. A comparison of the features supported by these different configurations is provided in **Error! Reference source not found.**.

The proposed flexible filter configurations support a wide range of owner/customer deployment models. For example, some customers would like to reserve a spectral range on day 1 and deploy on the full allocation day 1. Other customers may prefer to reserve a spectral range on day 1 but only do a partial deployment day 1 and ensure the reserved spectrum is available on request. The reserved block would be deployed in stages as needed. Other options are available limited only by the creativity of the customer / owner negotiations.

Partitioning Option	Requirements			Usage
	Security	Noise Addition	Privacy	
Managed services	n/a	n/a	n/a	Not spectrum partitioning
Managed spectrum	Y	N	TBD	For small channel count
Coupler /splitter	N	Y	N	For upgrades
Fixed OADM Filter	Y	N	Y	Not recommended
Flexible Add Filter	Y	N	N	For new D+ cables
Flexible Add /Drop Filter	Y	N	Y	For new D+ cables

Table 1: Functional Comparison of Spectrum Partitioning Options

3. PERFORMANCE VARIATION

In the spectrum negotiation process an important consideration is that the performance across the band is not uniform (depending on the cable system vintage), impacting both capacity and power management requirements and thus influence the way in which spectral ranges are selected and monetized.

Legacy dispersion compensated cables have highly variable performance across the band which means that not all spectral regions provide the same performance resulting in a much smaller net capacity. The effort to partition the spectrum into virtual bands, i.e. spectrum sharing option, is much less attractive and less cost effective. The design of legacy cables was originally focussed on IMDD 10G technology resulting in higher per channel output power and strong non-linear interaction between co-propagating carriers, all of which result in more complex channel add procedures and specialty power management solutions. Because of these complexities, legacy cables require a single entity managing the full spectrum which is not compatible with a full shared spectrum solution. However, a managed spectrum solution would be recommended in this case.

New uncompensated cables present the most practical spectrum sharing opportunity given the relatively uniform performance across the band and given the much higher net capacity available with new coherent modems. There are still some variations in linear and non-linear propagation impairments across the band for new D+ cables [4], such that the spectrum in different areas of the band may support different net capacity, which may factor into the negotiation process and other commercial considerations. Some of the variation in performance can be determined via modelling tools and refined with initial deployment measurements.

4. NETWORK MANAGEMENT

Once hardware is in place for one or more spectrum partitioning options, both cable owner and customers will enrol the hardware in their respective management and control layers to support day-to-day operations including additional bandwidth allocation, line rate changes, fault isolation, power management, and fault recovery.

The interface between customer hardware and owner hardware is of importance, where both customer and owner will need to independently monitor the ingress and egress interfaces regarding total power with fast update rate to detect transients as well as spectral content (i.e., power spectral density).

In the event of faults, detected by either party, the party responsible for the power management will replace the faulty traffic with the appropriate content of power management signals, until such time that the fault is repaired, and traffic channels can be re-added to the system. As an example, in spectrum sharing configurations, each customer is responsible for maintaining the agreed spectral density across their allocated band, but the owner has the responsibility of maintaining the working conditions of the fibre pair for all customers. Hence, when the owner detects a spectral density violation deemed to be detrimental to other customers on the fibre pair, they have the duty to replace the violating customer spectrum with ASE as a temporary measure until the condition clears. The assumption is that the "fault" conditions and the "fault recovery" conditions are stipulated in a contract between owner and customer.

Additional bandwidth and/or capacity increases may require actions from both the owner and the customer network management systems (NMSs). For additional bandwidth allocation in the case of shared spectrum, it will be the responsibility of the owner NMS to slowly increase the size of the partition while the customer NMS fills the additional partition with the appropriate signal content. For line rate increases in the case of the managed spectrum, it is the responsibility of the customer NMS to change the transponder line rate, while the owner NMS will be required to provide the higher line rate signal with a differential OSNR boost (margin bias) compared to other signals in the link to enable performance at the higher line rate.

5. ORCHESTRATION

With the advent of open networking, the support of northbound interfaces using ReST for provisioning and management and web sockets for alarm and event notification, are commonplace in the industry. Moreover, standardized interfaces between different management and control systems are being developed (e.g. ONF TAPI) and an Orchestrator at the application layer may therefore automate the dynamic bandwidth allocation and capacity increases (see section 4).

For example, in the context of managed spectrum, in ***Figure 6***, the green customer's (customer 2) user privileges give him access to a larger partition of the owner's orchestrator, allowing for automated wave – line synchronization for new channel additions as well as automated OSNR margin boost for line rate increases.

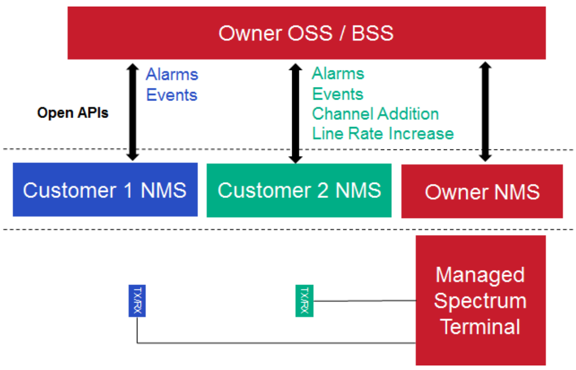

Figure 6Managed Spectrum Orchestration

In the case of shared spectrum, in ***Figure 7***, the green customer's (customer 2) user privileges give him access to a larger partition of the cable operator orchestrator, allowing for automated bandwidth increase.

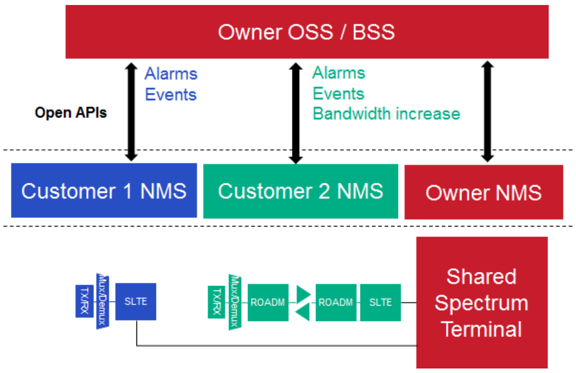

Figure 7 Shared Spectrum Orchestration

For all spectrum partitioning architectures, the alarms and events open APIs will be used to inform all parties if any line monitoring activities or spectral loading changes are being performed on the cable so that all participants can start fast monitoring their transponders' performance during these cable maintenance states.

The BSS may additionally be equipped with automatic pricing for all bandwidth on demand operations.

6. VISUALIZATION

Giving customers access to the spectral view of their channels into the wet plant may help with the joint owner-customer fault diagnostic and recovery.

Additionally, the owner may allow customers to visualize the allocated or "in use" spectrum partitions as well as the unallocated or reserved spectrum (see Figure 9) such that a customer may plan future expansion in those regions. If dynamic bandwidth expansion using the owner orchestrator is available to the customer (see section 5), it can make the bandwidth on demand operation easy to implement and monetize. Included in the visualization may be information regarding the available performance of the cable in that region of the spectrum (see section 3) so that the customer may decide in which region to expand depending on the available performance vs cost of that region.

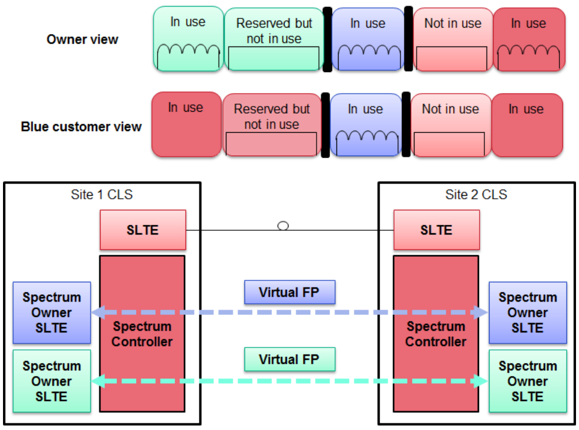

Figure 8 Shared Spectrum Visualization

7. COMMERCIAL CONSIDERATIONS

From the previous sections, it should be apparent that the implementation of a spectrum partitioning comes at a cost (CAPEX and OPEX) for both the owner and the customer. Customers wanting a Spectrum Partitioning solution need to be aware of these added costs, responsibilities, potential operational challenges, and recognize that there are limited assurances regarding performance. The implementation of a shared spectrum option will typically require more activity from both the customer and the owner. On one hand, the customer will need to select a vendor, develop the link performance budget, provide, and operate its own equipment. On the other hand, the owner will need to interact with the customer and their equipment supplier, define performance parameters, assist with equipment installation, monitor, and manage optical parameters. Also, it should be clear that the owner cannot provide any performance guarantees other than OSNR.

There are customers for whom the benefits of shared spectrum outweigh its cost. Usually, those customers already own and / or operate submarine cable capacity and are looking to purchase many waves. Customers who need a few waves (single digit) should be steered towards a capacity solution for lower risk, faster delivery, overall simpler implementation. Customers needing a moderate number of waves would benefit from a Managed Spectrum solution.

From the owner perspective, the spectrum must be priced properly to compensate for foregone capacity, including the loss of spectrum due to any guard bands. However, different business models can be envisioned that would benefit both the customer and the owner. For instance, instead of charging the customer for its full allocated spectrum, the owner could charge only for the spectrum used by the customer provided the right to use the unused spectrum until the customer needs it.

Finally, as discussed in section 3, the cable performance is not uniform across the bandwidth, with some parts being more valuable than others (the difference in performance is most pronounced in legacy cables but it also exists in newer cables). Methods for clearly measuring the available performance and assigning a dollar value is of utmost importance [4].

Tables 2 and 3 attempt to compare, in terms of cost (CAPEX and OPEX), upgradability, and application space, the various spectrum partitioning options discussed in this paper from both a customer and owner perspectives.

The OPEX cost implies a certain level of knowledge required from the customer in order to achieve the trustworthiness necessary to operate spectrum partitioning cables in co-operation with the owner. As Management and Orchestration features mature, there is a reasonable expectation of improved user experience and reduced OPEX.

Partitioning Option	Considerations assuming a fixed spectrum allocation			
	CAPEX	OPEX	Future Modems	Legacy cables
Managed services	Lowest	Lowest	No	Yes
Managed spectrum	Medium	Low	Yes	Yes
Broadband coupler/splitter	High	Higher	Yes	Not easy
Fixed OADM Filter	High	Medium	Yes	Not easy
Flexible Add Filter	Higher	High	Yes	Not easy
Flexible Add/Drop Filter	Higher	High	Yes	Not easy

Table 2: Options from Customer view point

Partitioning Option	Considerations assuming a fixed spectrum allocation			
	CAPEX	OPEX	Future Modems	Legacy cables
Managed services	High	Highest	Yes	Yes
Managed spectrum	Low	Medium	No	Yes

Broadband coupler/splitter	Low	Highest	No	Not easy
Fixed OADM Filter	Low	Low	No	Not easy
Flexible Add Filter	Medium	Medium	No	Not easy
Flexible Add/Drop Filter	Medium	Medium	No	Not easy

Table 3: Options from Owner view point

8. CONCLUSION

In this paper we have presented managed and shared spectrum options for partitioning the spectrum and evaluated those options technically and commercially from both a customer and owner perspective. The paper also discussed management and orchestration features for operating spectrum partitioning options.

Several other points came out of this evaluation:

For legacy, dispersion compensated cable systems, which have non-uniform performance across the spectrum and strong interaction between the signals sharing the spectrum, shared spectrum solutions may be too costly to implement and operate. Instead, a managed spectrum solution, where all the cable optimization falls with one party and one controller, is the recommended option.

In comparison, D+, uncompensated cables, which have more uniform performance across the spectrum with less interaction between co-propagating signals and much higher capacity, are more suited to the flexible add / drop shared spectrum options. The flexible add / drop solution lends itself to automation (as discussed in the Orchestration section)

Although the hardware required for the proposed spectrum partitioning options is mature, the related management and orchestration software layers are only starting to be developed. The richer features in software are intended to improve the user experience and ultimately reduce OPEX.

On the commercial side, the ongoing challenge is the negotiation of spectrum reservation, allocation, pricing and policing rules under fault conditions. As proposed in the paper, the development of orchestration tools may assist with all these challenges.

One final note for future consideration, as new cable systems evolve to Spatial Division Multiplexing (SDM) designs, the per fibre pair capacity may decrease to a point which may reduce the need for spectrum partitioning solutions.

9. REFERENCES

[1] M. Hubbard, "Spectrum Sharing: Options for Accessing Capacity," TNC 2018 Workshop on Spectrum Sharing .

[2] E. West, "The Impact of Convergence on Submarine Cable System Deployment," Proceedings SubOptic 2013, 10/TuB1.5.

[3] D. Evans, "Spectrum Sharing in a Multi-Vendor Environment," Proceeding SubOptic 2016.

[4] J. K. Perin, "Importance of Amplifier physics in Maximizing the Capacity of Submarine Links," *arXiv:1803.07905v2 [physics.app-ph] 5Jun 2018.*

