Dynamic and Efficient Point-to-Point and Point-to-Multipoint Communications by Slicing the Optical Constellation

M. Iqbal[1], M. Ruiz[1], N. Costa[2], A. Napoli[3], J. Pedro[2], and L. Velasco[1*]

[1] Optical Communications Group (GCO), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
[2] Infinera Unipessoal Lda., Carnaxide, Portugal [3] Infinera, London, United Kingdom
e-mail: luis.velasco@upc.edu

Abstract: Optical Constellation Slicing is proposed to convey heterogenous traffic from a source to multiple destinations, while supporting dynamic capacity allocation. Illustrative numerical results reveal the potential of the proposed scheme, while providing significant cost reduction.

1. Introduction

Various advanced applications are much more cost effective when implementing high-bitrate low-latency point-to-multipoint (P2MP) connectivity. In contrast to Point-to-Point (P2P) connections, where a source node sends data to a single destination node, in a P2MP connection the source node sends data to a set of destinations that may be scattered over a geographical area [1]. Just like P2P, P2MP connections can be supported in Wavelength Switched Optical Networks (WSON) (see, e.g., [2]); we denote these connections as *lightpaths* and *light-trees* [3], respectively. Indeed, when the required bandwidth is in the range of a few tens of Gb/s, using dedicated high-capacity optical transceivers to establish independent direct optical connections would be highly inefficient. Alternatively, one single light-tree could be set to support both the P2P and P2MP traffic. In this case, the data signals would reach all destination nodes, which would then filter the relevant data and drop the remaining one. However, this solution entails security considerations (e.g., eavesdropping).

In this paper, we propose to use one single light-tree to transport a combination of P2P and P2MP traffic. This solution could be implemented using digital subcarrier multiplexing (DSCM) technology [1], [4] and dedicate one or more independent subcarriers to support P2P traffic between the source and each of the destinations, as well as P2MP traffic from the source to all or a subset of destinations. However, we explore an alternative solution that consists in partitioning the optical constellation and dedicating a different subset of constellation points for the P2P and P2MP traffic; we call it Optical Constellation Slicing (OCS). We describe the basic principle of OCS and its implementation in current coherent communication networks. Moreover, we also evaluate the throughput and relative cost of OCS.

2. Optical Constellation Slicing (OCS) Concept

Fig. 1a illustrates the concept of optical constellation slicing; an optical signal is generated at the source node and sent to four destinations (Dest-1..4) through a light-tree. Five optical constellation slices are defined (OCS-1..5), where four of them support P2P connectivity between the source node and one of the destinations, whereas OCS-5 supports P2MP connectivity. Fig. 1b highlights the connectivity depicted in Fig. 1a. It is worth noting that the effective bitrate of an OCS can be controlled by selecting the number of constellation points that are assigned to it. In the example in Fig. 1a, the optical signal is modulated using 64-QAM, so 64 constellation points are available. From them, OCS-1 is assigned 32 points, whereas OCS-3 and 4 are assigned just four. Only the constellation points assigned to one OCS can be used for communication to the corresponding destination.

Another important aspect in OCS, particularly in P2MP connections, is data security as all destinations receive the complete constellation. To increase security, an independent lookup table (LUT) is used at the source to encode the transmitted data (*substitution* cipher [5]) which also makes sure that only the destination who knows the corresponding LUT is able to decode the constellation points of the OCSs that have been assigned to it. Note that is in the source where most of the functionalities need to be implemented (including slicing and LUT coding), while destinations perform LUT decoding only.

By enabling a P2MP architecture, OCS enables capital cost reduction by decreasing the number of required transceivers and, in turn, simplifies the interconnection architecture. In the example in Fig. 1a, although only a single configuration making use of a single transmitter (TX) and four receivers (RX) is shown, several different configurations are supported. OCS also simplifies optical connection control, as no light-tree reconfiguration is needed; once the light-tree is created by the network controller, OCS management is implemented by distributing LUTs between the source and the destinations. Moreover, OCSs can be created, modified, and eliminated dynamically, which provides the flexibility of changing the capacity of the connections as per traffic demands, thus reducing operational cost as well. As an example, Fig. 1c presents a different slicing, where two P2MP connections are active among different subsets of destinations. Moreover, the ability of changing the RXs in the light-tree while keeping the same TX makes the system highly flexible.

The research leading to these results has received funding from the European Community through the MSCA REAL-NET project (G.A. 813144), by the AEI through the IBON project (PID2020-114135RB-I00), and by the ICREA institution.

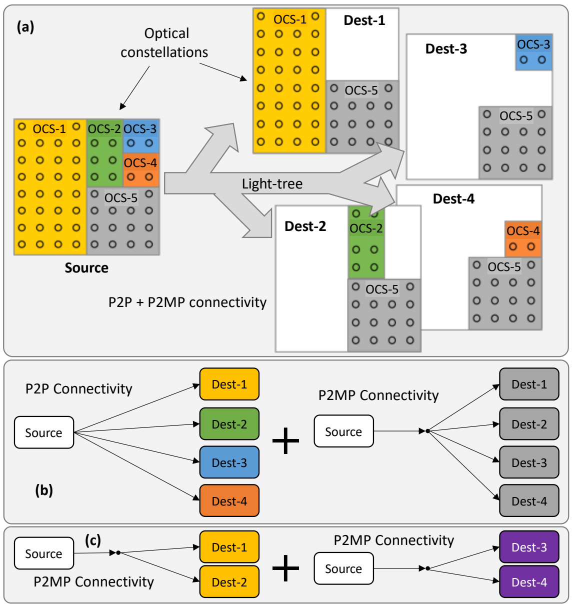

Fig. 1. OCS concept and dynamic configuration.

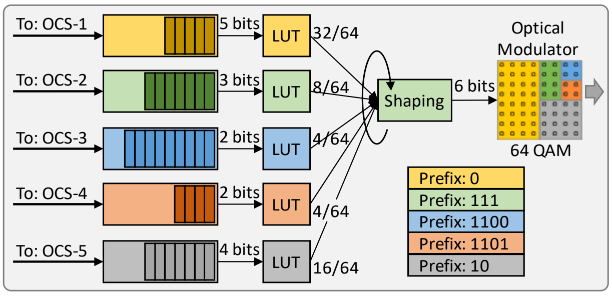

Fig. 2. OCSs at the optical transmitter.

Table 1. Capacity and cost per transceiver

Single-Polarization System	Capacity [Gb/s]	Unitary Cost [m.u.]
64-QAM@64GBaud	384	3.2
32-QAM@32GBaud	160	1.8
16-QAM@32GBaud	128	1
QPSK@32GBaud	64	0.4

3. Optical Slicing

The main step to perform OCS is to select the modulation format (m-QAM), which imposes the maximum throughput and the maximum number of OCSs. Since the minimum number of constellation points that can be assigned to an individual OCS is 2 (i.e., each point would represent 1 single data bit) the maximum number of simultaneous destinations is m/2. In this paper, we assume m=64 for the sake of a wide study.

The next step is to create the OCSs. Fig. 2 illustrates the traffic shaping to be implemented at the transmitter for the example in Fig. 1a. Every OCS is associated with a buffer within the TX, where data streams are temporarily stored. From those buffers, sets of bits of size equal to the number of bits with information (*infobits*) in the OCS are selected and encrypted using the LUT as a substitution cypher. E.g., sets of 5 bits are selected for OCS-1 whereas sets of 3 bits are selected for OCS-2. Next, the shaping block receives one of the encrypted sets from the OCSs at a time and adds the *prefix* that identifies the OCS, i.e., the constellation point. Thus, each RX is served by an OCS formed from a unique <prefix, infobits> pair. Note that prefixes might be of different lengths, so sets of 6 bits are obtained to feed the optical modulator (64-QAM). E.g., it adds prefix 0 to sets of bits from OCS-1 and prefix 111 to sets of bits from OCS-2 (see Fig. 2). The shaping block follows a 64-step cycle, where at every step it receives a set of bits from an OCS; the number of sets of bits selected from each OCS is exactly the number of constellation points assigned to that OCS. E.g., the shaping block receives 32 out of 64 5-bit sets from OCS-1, and 8 out of 64 3-bit sets from OCS-2.

With this arrangement, the throughput of each OCS system can be computed as a function of the number of bits with information, the Symbol Rate (SR), the number of symbols assigned to the OCS, and the modulation format used in the optical transmitter (*m*-QAM).

$$Throughput_i = \frac{infobits_i \times SR \times \#symbols}{m} \quad (1)$$

The throughput of the system is computed as the summation of individual throughputs. In the example, the throughput of OCS-1 and OCS-2 are 160 Gb/s and 24 Gb/s, respectively, assuming SR = 64 GBaud. The total throughput is 264 Gb/s, i.e., 68.8% efficiency with respect to the maximum capacity of the system (384 Gb/s).

4. Illustrative Numerical Results

To evaluate the proposed OCS scheme, we target at serving the traffic between a source and several destinations. Just one single light-tree is set with the OCS scheme, so when the number of destinations increases, the total capacity that might be requested from each destination decreases proportionally. We serve the requested capacity by setting up the needed number of lightpaths and compare the efficiency of the resulting optical system (defined as throughput over capacity), number of transceivers to be installed, and total cost. We assume the capacity in Gb/s and cost in monetary units (m.u.) in Table 1 (based on [4]), where cost is normalized to that of the 16-QAM@32 GBaud transceivers. For the sake of simplicity, we assume single-polarization systems and compute bitrate as the product of bits per symbol and symbol rate. Two approaches are considered: *i*) static capacity allocation, where the amount of traffic to each destination does not change over time, so the capacity of the transceivers can be set to the minimum supporting the requested traffic capacity; and *ii*) dynamic capacity allocation to support traffic dynamicity, while keeping the maximum traffic constant. In this case, the capacity of all transceivers need to support the maximum traffic. Note that OCS allows for dynamic capacity allocation.

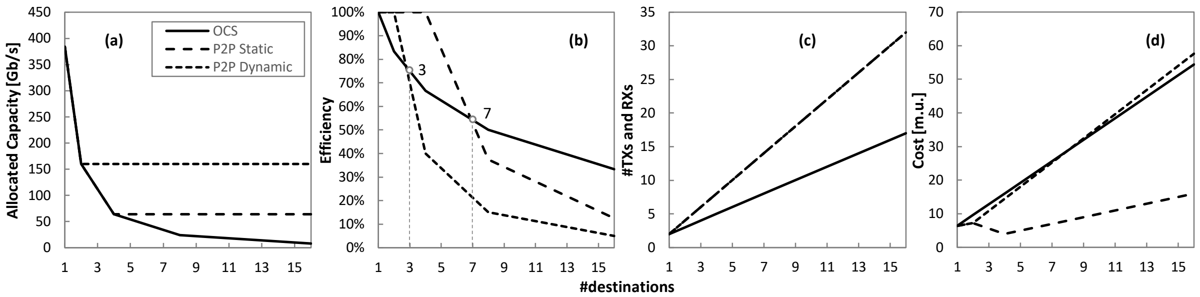

Fig. 3: Comparison for P2P traffic.

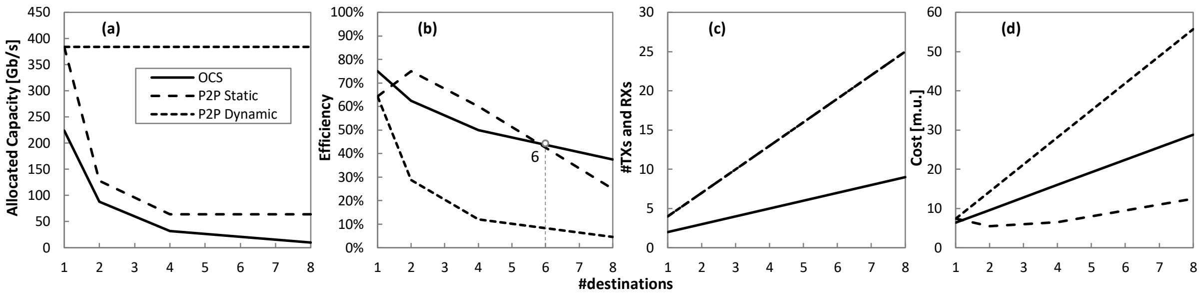

Fig. 4: Comparison for P2P + P2MP traffic.

Fig. 3 presents the results assuming that all the traffic is P2P. Fig. 3a plots the capacity requested by every destination (in continuous line) and the capacity served using lightpaths (discontinuous lines). The maximum capacity that can be requested by each destination is 160 Gb/s, so the dynamic capacity allocation requires transceivers with that capacity. Note that the capacity allocated using lightpaths is higher than the one requested, and therefore, the efficiency of the lightpath solution decreases as shown in Fig. 3b. The efficiency of the OCS approach is related to the number of infobits per symbol, which decreases when the number of assigned constellation points to an OCS decreases. We observe in Fig. 3b that the OCS approach is more efficient than the lightpath solution when the number of destinations is higher than 7, and its efficiency is higher than that of the dynamic capacity allocation when the number of destinations is above 3. As for the number of transceivers needed, the OCS solution largely reduces the number needed with respect to any lightpath solution, as shown in Fig. 3c. Fig. 3d shows that the static capacity allocation minimizes costs. However, in dynamic capacity allocations, the OCS solution reduces the cost w.r.t the lightpath one for an increased number of RXs.

Let us now assume a mix of P2P and P2MP traffic (Fig. 4). To facilitate the analysis, we assume that the capacity for the P2MP traffic is fixed to 64 Gb/s, whereas that of the P2P traffic decreases proportionally to the number of destinations. Fig. 4a plots the capacity allocated for the different solutions. In this case, the maximum capacity that can be requested by each destination is 224 Gb/s, so the dynamic capacity allocation requires lightpaths with at least that capacity. Consequently, its efficiency decreases abruptly with the number of destinations. The static capacity allocation scheme achieves slightly higher efficiency than OCS up to 6 destinations. Fig. 4c shows that the OCS solution serves the traffic with a very limited number or transceivers (Fig. 4c) and, in this case, the cost is closer to that of the static capacity allocation.

To conclude, let us compare the three solutions for the example in Fig. 1. Table 2 presents the maximum throughput and the efficiency of each of the OCSs in the system. The efficiency column is computed as the ratio between the number of infobits with respect to the total number of bits per symbol, whereas the contributed efficiency also considers traffic shaping. Table 3 compares the OCS solution with the lightpaths solutions; it largely reduces the number of transceivers and achieves 32% of cost reduction.

Table 2: Capacity and Efficiency

OCS	Throughput [Gb/s]	Efficiency [%]	Contributed Efficiency [%]
1	160	83.3%	41.7%
2	24	50.0%	6.3%
3	8	33.3%	2.1%
4	8	33.3%	2.1%
5	64	66.7%	16.7%
SUM	**264**	**68.8%**	**68.8%**

Table 3: Required transceivers and costs

Solution	#TX/RXs	Cost [m.u.]	Cost savings [%]
OCS	5	16	32%
Static	13	8	66%
Dynamic	13	23.4	Ref.

References

[1] D. Welch, *et al*. "Point-to-multipoint optical networks using coherent digital subcarriers," IEEE/OSA JLT, 2021.
[2] M. Ruiz and L. Velasco, "Serving Multicast Requests on Single Layer and Multilayer Flexgrid Networks," IEEE/OSA JOCN, 2015.
[3] M. Ruiz and L. Velasco, "Performance Evaluation of Light-tree Schemes in Flexgrid Optical Networks," IEEE Comm. Letters, 2014.
[4] F. Rambach *et al*., "A multilayer cost model for metro/core networks," IEEE/OSA JOCN, 2013.
[5] D. Salomon, *Coding for Data and Computer Communications*, Springer, 2005.