DSP Design for Point-to-Multipoint Transmission

Thomas Duthel[1], Chris R.S. Fludger[1], Bo Liu[1], Antonio Napoli[1], Amir Rashidinejad[2], Stenio Ranzini[1], Sezer Erkilinc[3], Aditya Kakkar[2], Atul Mathur[4], Vince Dominic[4], Parmijit Samra[4], Han Sun[2], Azmina Somani[2], Dave Welch[4]

[1] Infinera GmbH, Nordostpark 12, 90411 Nuremberg, Germany [2] Infinera Inc., Canada [3] Infinera AB, Sweden [4] Infinera Corporation, USA
tduthel@infinera.com

Abstract: Coherent optical transmission systems using digital sub-carriers are ideal for point-to-multi-point applications. Many functional blocks are similar to the ones of single-channel processors. But, several aspects, specific to digital sub-carriers and point-to-multi-point, need consideration in the DSP implementation. © 2022 The Author(s)

1. Introduction

Point-to-multi-point (P2MP) network architectures have the potential to significantly simplify how metro and access networks are designed [1]. Today's P2MP networks – e.g., in the form of passive optical networks (PONs) – are based on intensity-modulated direct detection (IM-DD) transceivers. To support the next generation of wireless network technologies, such as 5G/6G, the backbone optical infrastructure will need to consider some form of coherent technology. First and foremost, to increase the spectral efficiency, but also to achieve better signal-to-noise-ratio (SNR) and sensitivity performance.

Digital signal processing (DSP) for coherent optical communication systems has traditionally been designed for single-channel applications, in which the information is directly modulated onto an optical carrier. In contrast, digital sub-carrier (DSC) based coherent systems split the information into several narrower, intermediate digital carriers before modulating a comb of these onto the optical carrier [1]. This approach occupies a similar bandwidth as in single-carrier modulation (using the same modulation format and same total transport capacity), but offers several advantages, such as better tolerance to non-linear fiber effects and support for spectral water-filling [2]. In addition, DSC systems are especially suitable for a P2MP network architecture. In such a configuration, information is exchanged between several geographically distributed leaf nodes and a central hub node. Whilst the hub sends and receives all DSCs, the individual leaf nodes only modulate and demodulate a subset of them [1]. The network itself can be purely passive with splitters and combiners distributing/aggregating the information. Capacity upgrades can be handled without network reconfiguration, simply by assigning more DSC to a particular leaf location.

Different flavors of a DSC-based DSP can be used to suit different functions inside a P2MP network. A hub DSP would be designed for a large number of DSCs, whilst a leaf node DSP would need to be more power efficient with a fewer number of DSCs, allowing for smaller form factors, e.g., QSFP-DD and QSFP-28 and lower cost modules. A careful decision on a base set of features (e.g., DSC symbol rate, modulation format and FEC coding scheme) allows for a multi-generation design that benefits from power savings over ASIC node sizes while maintaining compatibility to earlier generations. This helps to reduce electronic waste since device upgrades are only needed when capacity limits are reached [1]. In this paper, we describe the architecture of such a DSC-based DSP and detail the special design aspects related to P2MP operation.

2. Digital sub-carrier DSP for point-to-multi-point

A coherent DSP for DSC incorporates the same functional blocks as in a single-channel DSP. The main difference is that most functional blocks are instantiated as one large circuit for a single-carrier DSP, and as multiple smaller circuits for DSC engines. Figure 1 illustrates a high-level block diagram of a transceiver for DSC processing.

On the transmitter side, incoming data streams from client interfaces are sent into forward error correction (FEC) encoding and framing. The TX DSP performs pulse shaping (here root-raised-cosine (RRC)) as well as pre-compensation of e.g., analog bandwidth restrictions and timing skews. Processing blocks that are unique to sub-carrier processors are the gain control blocks (Gain cntrl) per DSC, and the muxing block (SC Mux). The gain control allows the relative DSC power to be changed before the individual streams are modulated onto individual intermediate carriers. After DSC multiplexing, the signal is converted from digital-to-analog (DAC), and finally modulated onto the optical carrier by a dual-polarization optical IQ Mach-Zehnder modulator (E/O) fed by a narrow line-width laser source.

On the RX side, after coherent detection using a local-oscillator, 90° hybrid and balanced photo-detectors (O/E), the signal is digitally sampled by a four channel analog-to-digital converter (ADC). The de-multiplexing into the different sub-carriers is the next processing step, before independent modem engines apply chromatic dispersion

(CD) compensation, polarization de-muxing (POL), carrier and frequency recovery (CR), frame removal, FEC decoding and client-side processing – just like in single-channel [3]. Since all these processing steps take place on a per DSC basis, the mitigation of linear link impairments (e.g., CD, differential group delay, filtering) is straight-forward, even if the different DSCs propagate along different links. But it is worth mentioning here that in P2MP operation the DSP does not average feedback signals from all DSCs.

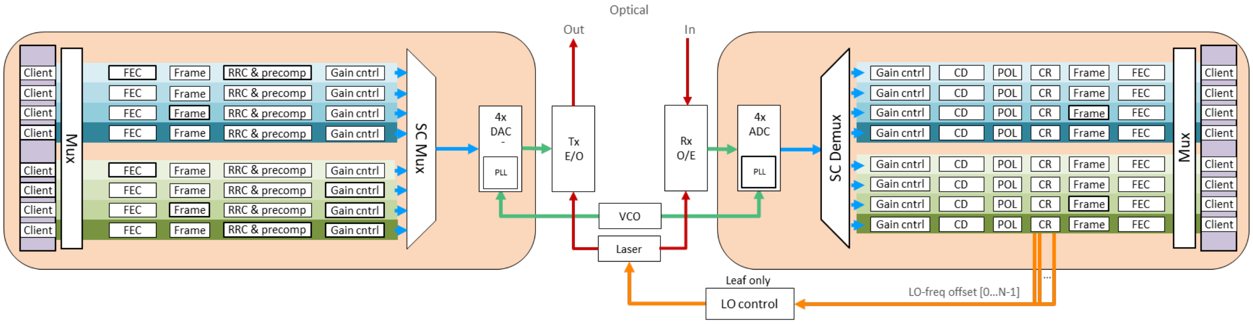

Fig. 1. Digital sub-carrier transceiver module with optics and DSP.

Some critical aspects for P2MP applications are related to timing and carrier frequency recovery. Inaccuracies in reference timing oscillators from each leaf node result in small differences between DSCs which are all sampled by a common ADC. Additionally, the laser frequency recovery is not only limited by the tracking capability of the DSP, but also by the guard band to the next sub-carrier [1]. Collisions between adjacent sub-carriers result in crosstalk and must be avoided to prevent system outages [1]. One solution to both problems is locking both the leaf clock and laser frequency to that of the remote hub. For systems, with lasers shared between TX and RX optics (as shown in Figure 1) the frequency estimate from the leaf RX DSP can be used to enable the leaf-side laser to track the hub laser frequency. In such a scheme, the leaf and hub lasers and reference clocks are to some extent frequency locked. Residual frequency and sampling offsets must be tolerated by the CR circuit [4], or digitally interpolated in the clock recovery circuit [5]

Different link lengths and thus different link losses between leaf sites and the hub is another critical aspect associated with P2MP operation, that leads to power imbalances between (groups of) DSCs received on hub site. Since the received signal sampled by the ADC is the combination of all received DSCs, a RX side per sub-carrier gain control mechanism is required to equalize the power of the DSCs before further DSP processing gets applied.

3. Impairments and operational behavior

This section exemplarily discusses some of the operational aspects of DSC in P2MP applications. In general, DSCs behave differently to impairments than single-channel systems. Whilst single-channel systems show characteristic behavior in the constellation diagram for most of the impairments (IQ skew, quadrature error and IQ power imbalances, see [6]), DSC systems would mainly show a distortion "on the mirror side" of the spectrum [7].

Figure 2 depicts the impact of TX side IQ skew on the spectrum and the related DSC constellation diagrams in a back-to-back configuration. The measurements are taken on a real-time DSC P2MP transceiver, which supports up to 16×4GBaud 16QAM dual polarization DSCs, each transmitting 25Gbit/s data.

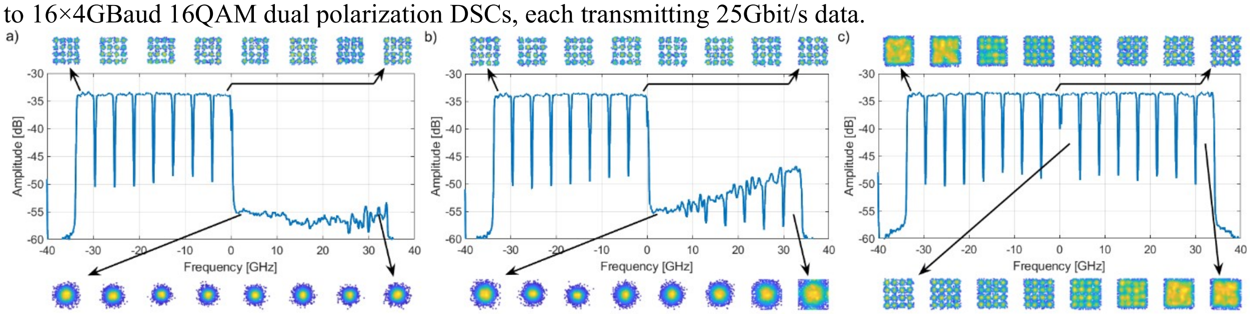

Fig. 2. Measured impact of TX side residual IQ skew on constellations and spectrum:
a) single sideband without IQ skew, b) single sideband with IQ skew, and c) double sideband with IQ skew.

For illustration, the eight right-hand side DSC are turned off in Figure 2a and Figure 2b. Here it can be seen that DSCs are received with clear 16QAM constellation diagrams. It is important to note that the TX side IQ skew added in Figure 2b does not cause any visible distortion in the constellation diagrams but creates mirror images on the other side of the spectrum. The strength of the images depends on the IQ skew value and increases with frequency offset from the optical carrier. The impact from the IQ skew distortion is important when the whole spectrum is filled with DSCs, e.g., when a hub is fully populated with DSCs from different leaf nodes. In this case, the mirror image from the left-hand side DSCs distorts the right-hand side DSCs and vice-versa. This is shown in Figure 2c. Due to the residual/uncompensated IQ skew, DSCs are being more distorted the further they are away from the center of the spectrum. Although, here the impairment is generated at the TX side, residual RX IQ skew causes the same distortion. Other impairments like quadrature error and IQ imbalances have similar effects, and it is clear that those kind of distortions must be calibrated-out during manufacturing. Furthermore, the impact must be minimized or compensated [7] over lifetime and operating conditions.

Power imbalances between DSCs coming from different leaf nodes lead to scenarios, as the ones shown in Figure 3. Here, the setup is simplified to two leaf nodes sending data to a hub site, though up to 16 leaf nodes could be supported. Leaf 1 sends 12 DSCs at the left-hand side of the spectrum, whereas leaf 2 sends 4 DSCs located at the right-hand side of the spectrum. To emulate the impact from different span length, the power of leaf 2 is varied using a variable optical attenuator (VOA). The power of the DSCs coming from leaf 1 is kept constant, and their Q factor shows ~0.5dB increase as the additional 4 DSCs from leaf 2 are attenuated (see Figure 3b). This is due to the slight changes of the RX noise conditions with reducing power from leaf 2. For leaf 2 itself the situation is different, and a power difference results in a Q penalty degradation (also see constellations in Figure 3d) which is mostly caused by the RX noise. But it is worth noting that the gain control mechanism inside the RX DSP allows demodulation of DSCs from leaf 2 even if their power was several dB below the DSCs coming from leaf 1. In real network scenarios, the situation could be further improved by instructing the leaf 2 TX to increase its power, which requires a feedback channel from the hub to leaf.

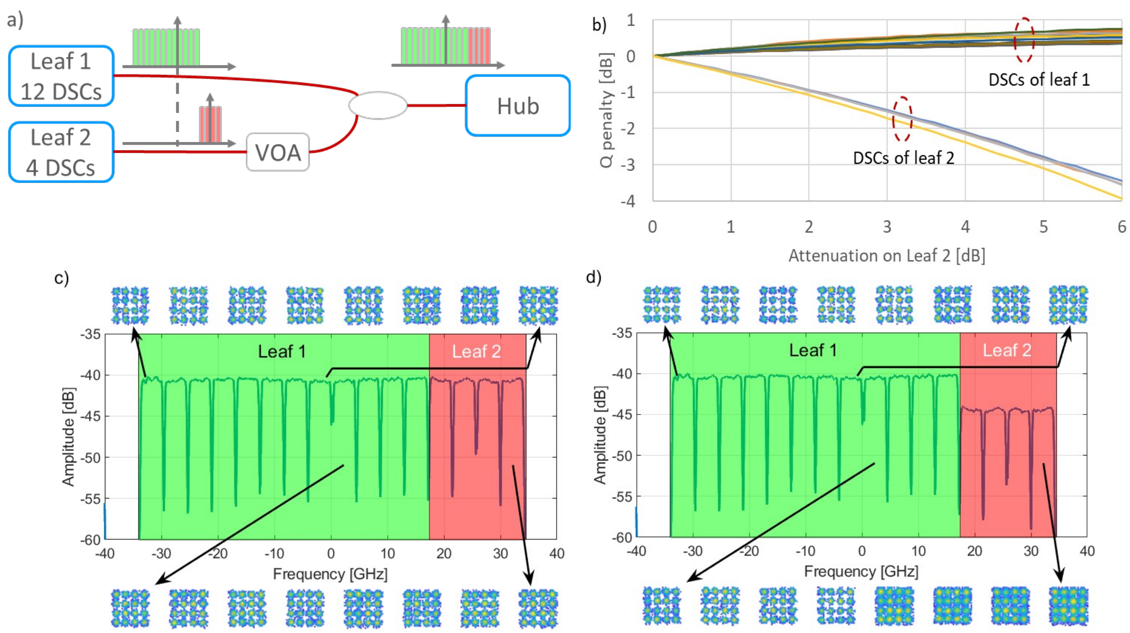

Fig. 3 Measured RX DSP tolerance to power imbalance between sub-carriers, a) setup b) Q penalty vs attenuation on leaf 2 c) spectrum and constellation without extra attenuation d) spectrum and constellation with 4dB attenuation on leaf 2

4. Conclusion

We detailed the architecture of a digital sub-carrier based DSP for coherent optical transceivers and discussed special requirements for point-to-multi-point operation. Experimental results are reported for a 16 sub-carrier 16QAM real-time implementation supporting a total capacity of 400Gbit/s.

5. References

[1] D. Welch et al. "Digital Subcarrier Multiplexing: Enabling Software-Configurable Optical Networks", JLT (October 2022)
[2] H. Sun et al. "800G DSP ASIC Design Using Probabilistic Shaping and Digital Sub-Carrier Multiplexing", JLT 38, no. 17 (Sept. 2020)
[3] S. Savory. "Digital Coherent Optical Receivers: Algorithms and Subsystems.", IEEE J. Sel. Top. Quantum Electron.16, no. 5 (Sept. 2010)[4]
[4] T. Pfau, "Carrier Recovery Algorithms and Real-time DSP Implementation for Coherent Receivers", Paper W4K1, OFC2014
[5] C.R.S. Fludger et al. "Jitter tolerant clock recovery for coherent optical receivers", Paper OTh1F.3, OFC2013
[6] C.R.S. Fludger et al. "Transmitter Impairment Mitigation and Monitoring for High Baud-Rate, High Order Modulation Systems", ECOC2016
[7] G. Bosco et al., "Impact of the Transmitter IQ-Skew in Multi-Subcarrier Coherent Optical Systems", Paper W4A.5, OFC2016