



Multi-channel collision-free reception for optical interconnects

Downloaded from: <https://research.chalmers.se>, 2025-12-05 00:12 UTC

Citation for the original published paper (version of record):

Lu, Y., Agrell, E., Pang, X. et al (2018). Multi-channel collision-free reception for optical interconnects. Optics Express, 26(10): 13214-13222. <http://dx.doi.org/10.1364/OE.26.013214>

N.B. When citing this work, cite the original published paper.

MULTI-CHANNEL COLLISION-FREE RECEPTION FOR OPTICAL INTERCONNECTS

YANG LU,¹ ERIK AGRELL,² XIAODAN PANG,^{3, 4} OSKARS OZOLINS,⁴ XUEZHI HONG,⁵ RUI LIN,³ YUXIN CHENG,³ ALEKSEJS UDALCOVS,⁴ SERGEI POPOV,³ GUNNAR JACOBSEN,⁴ AND JIAJIA CHEN^{3,5*}

¹College of Communication Engineering, Hangzhou Dianzi University, Hangzhou, China

²Department of Electrical Engineering, Chalmers University of Technology, Göteborg, Sweden

³KTH Royal Institute of Technology, Stockholm, Sweden

⁴Networking and Transmission Laboratory, RISE Acreo AB, Kista, Sweden

⁵ZJU-SCNU Joint Research Center of Photonics, South China Normal University, Guangzhou 510006, China

* jiajiac@kth.se

Abstract: A multi-channel reception scheme that allows each node to receive an arbitrary set of wavelengths simultaneously (i.e., collision-free) is proposed for optical interconnects. The proposed scheme only needs to use a few receivers and fixed-wavelength filters that are designed based on error-control coding theory. Experiments with up to four channel collision-free reception units are carried out to demonstrate the feasibility of the proposed scheme.

© 2018 Optical Society of America

OCIS codes: (060.2330) Fiber optics communications; (060.2360) Fiber optics links and subsystems; (060.4264) wavelength assignment; (060.4510) Optical communications.

References and links

1. C. Kachris and I. Tomkos, "A survey on optical interconnects for data centers," *IEEE Communications Surveys & Tutorials*, **14**(4), 1021–1036 (2012).
2. C. Kachris, K. Kanonakis, and I. Tomkos, "Optical interconnection networks in data centers: Recent trends and future challenges," *IEEE Commun. Mag.*, **51**(9), 39–45(2013).
3. M. Fiorani, S. Aleksic, M. Casoni, and L. Wosinska, "Energy-efficient elastic optical interconnect architecture for data centers," *IEEE Commun. Lett.*, **18**(9), 1531–1534(2014).
4. Y. Cheng, M. Fiorani, L. Wosinska, and J. Chen "Reliable and cost efficient passive optical interconnects for data centers," *IEEE Commun. Lett.*, **19**(11), 1913–1916(2015).
5. Y. Gong, X. Hong, Y. Lu, S. He, and J. Chen, "Passive optical interconnects at top of the rack: Offering high energy efficiency for datacenters," *Opt. Express*, **23**(6), 7957–7970(2015).
6. X. Hong, Y. Gong, Y. Yang, and J. Chen, "AWG based passive optical interconnects for datacenters," in *Proceedings of Photonic Networks and Devices*, (2016), NeTu1C.3.
7. T. Ye, T. Lee, M. Ge, and W. Hu, "Modular AWG-based interconnection for large-scale data center networks," *IEEE Trans. Cloud Comput.*, **PP**(99), 1-1(2017).
8. R. Lin, X. Pang, O. Ozolins, Z. Feng, A. Djupsjoebacka, U. Westergren, R. Schatz, G. Jacobsen, M. Tang, S. Fu, D. Liu, S. Popov, and J. Chen, "Experimental Validation of Scalability Improvement for Passive Optical Interconnect by Implementing Digital Equalization," in *Proceedings of ECOC*, (2016), VDE.
9. R. Murano and M. J. L. Cahill, "Low Cost Tunable Receivers for Wavelength Agile PONs," in *Proceedings of ECOC*, (2012).
10. Y. Lu, E. Agrell, X. Pang, O. Ozolins, X. Hong, R. Lin, Y. Cheng, A. Udalcovs, S. Popov, G. Jacobsen, and J. Chen, "Matrix Receiving Scheme Supporting Arbitrary Multiple-Wavelength Reception for Optical Interconnects," in *Proceedings of ECOC*, (2017).
11. G. C. Clark and J. B. Cain, "Error-Correction Coding for Digital Communications". New York: Springer(1981)
12. E. Agrell, "Tables of binary block codes," online: <http://codes.se/bounds>.
13. E. Agrell, A. Vardy, and K. Zeger, "A table of upper bounds for binary codes," *IEEE Transactions on Information Theory*, **47**(7), 3004–3006(2001).
14. Grassl, Markus, "Bounds on the minimum distance of linear codes and quantum codes." Online available at <http://www.codetables.de>.
15. Juniper Networks, "Understanding Pre-FEC BER Monitoring and BER Thresholds," https://www.juniper.net/documentation/en_US/junos/topics/concept/otn-signal-degrade-monitoring-understanding.html

1. Introduction

In recent years, optical interconnects have been introduced to sustain the ever-increasing traffic demand in data center networks (DCNs) [1, 2]. To provide high power efficiency, passive optical interconnects (POIs), which use only passive optical components for servers' interconnections have been proposed [3–7]. It has been demonstrated in [3] that the power consumption can be reduced by a factor of ten when the commodity switches are replaced by POIs. Besides, the POIs also bring other advantages, such as low cost, high reliability, easy maintenance and transparency to the data rate and modulation format. Most of the existing POIs use either couplers [3, 4] or arrayed waveguide gratings (AWGs) [5–7] for interconnection. The coupler-based solutions suffer from a high insertion loss and cannot afford large-size interconnects [8]. The issue of high insertion loss can be avoided by the AWG-based solutions because the insertion loss of a cyclic AWG is relatively low and independent on the size of the interconnects.




Fig. 1. A possible collision that could occur at reception in an AWG-based POI. WTT: wavelength tunable transmitter; Rx: receiver.

Without loss of generality, we take a typical AWG-based POI architecture presented in [5] as an example to illustrate its operation principle (see Fig. 1). Each node is usually equipped with a wavelength tunable transmitter (WTT) and a receiver (Rx) for data transmission and reception. Due to the cyclic wavelength routing property of the $N \times N$ AWG, each node can transmit a signal to any other nodes by tuning the wavelength. Node i is able to send data to Node j via $\lambda_{(j-i) \bmod N+1}$. For example, Node 1 sends data to Node 2 (Node N) via λ_2 (λ_N). The routing is achieved by using WTTs and an AWG, so the Rx in each node is not necessarily wavelength-selective. It helps to reduce the complexity compared to the coupler-based POIs. However, in such an AWG-based POI, a reception collision, which means that several nodes send data to the same destination simultaneously, still exists. Figure 1 depicts a possible reception collision, where Node 2 receives the signals transmitted from Node 1 and Node N at the same time.

Considering unicast traffic (one-to-one connection), where all N nodes have an equal opportunity to choose any other nodes as the destination, the probability that one node receives the signals from M other nodes at the same time is $\binom{N-1}{M} \frac{(N-2)^{N-1-M}}{(N-1)^{N-1}}$. Although the probability of simultaneous arrivals is a function of N , the difference caused by changing the value of N is minor. In the limit of large N , the probability of M signals colliding at a given receiver is $1/(eM!)$. If there is only one receiver at each node, the probability of reception collision (i.e., $M \geq 2$) is more than 25 % (exactly $1-2/e = 26.4$ % for large N). Thus, if no specific scheme is employed to handle collisions, the blocking probability exceeds 25 %, which is not acceptable. This probability for M simultaneous signals are, if N is large, 18.4, 6.1, and 1.5 % for $M = 2, 3$, and 4, respectively, and only 0.4 % for $M \geq 5$. It means the blocking at the receiver is expected to be of very low possibility if one can receive multi-channel (e.g., up to 4) signal without collision.

One solution to mitigate the risk of reception collisions is to make each node capable of receiving multiple channels simultaneously. A straightforward way is to have a demultiplexer and $N-1$ receivers equipped at each node. However, such a method is not cost-effective, especially for large-scale POIs, and has quite low utilization of the receivers because the signals that simultaneously arrive are usually from no more than 4 nodes (the possibility of receiving

combined signals from more than 4 nodes is only 0.4 %). An alternative way is to equip each node with a limited number of receivers, each of which needs a wavelength-tunable filter (WTF) to dynamically select the wavelength. However, it is very challenging to have WTFs with a high tuning speed and a large tuning range [9], which satisfies the switching requirements in the optical interconnects.

We propose a multi-channel reception scheme that only uses a limited number of wavelength-fixed filters and receivers, allowing each node connected with the AWG-based interconnects to receive multiple wavelengths simultaneously. As an extension of our previous work [10], this paper includes the following new contributions: (i) we present a general theory for the proposed multi-channel reception scheme, not just one special case as in [10]; (ii) we perform an extensive analysis, including the calculation of the probability of simultaneous arrivals, a thorough treatment of signal recovery, more case studies for different parameters and an example of our improved filter matrix design; (iii) besides the experimental demonstration of the feasibility of our scheme, we also improve the transmission performance by equalization to meet a lower FEC limit requirement or even eliminate the FEC requirements for some cases.

The remainder of the paper is organized as follows. We present the introduced matrix reception scheme and illustrate its operation principle in Section II. In Section III, we propose filter matrix design strategies that minimize the number of required wavelength-fixed filters with any given set of system parameters, e.g., the total number of wavelengths and the number of channels that are able to receive multiple signals simultaneously. Section IV presents the experimental results, where up to 4-channel reception is carried out. The bit error rate (BER) performance can achieve the second generation of forward error correction (FEC) limit without any specific equalization schemes, and can be further enhanced to reach the first generation of FEC limit (i.e., G.709 GFEC) by implementing equalization. Finally, conclusions are drawn in Section V.

2. Matrix Receiver




Fig. 2. The proposed multi-wavelength reception scheme.

The proposed multi-wavelength reception scheme is illustrated in Fig. 2. Our scheme consists of a $1 \times F$ splitter, F filters, and F receivers. Each filter is a multibandpass filter which allows one or a few channels to pass. Each receiver follows one filter, and its received wavelength channels are constrained by the filter. To achieve optimized system configuration, the passbands of the filters, collectively represented as a filter matrix, need to be carefully designed (see the next section). The control plane should collect all communication requests from all nodes and avoid that more than M nodes transmit to the same destination simultaneously. The signals to be sent to this node are synchronized on the bit level and adjusted to the same power. The control plane informs the receiver about the wavelengths of the arrived signals. The arrived

signals (carried by different wavelengths) are first split and sent to all filters. For each branch, the wavelengths that are in passbands pass through the filter and are received by the following receiver. At reception, there may be null, a single signal, or a combination of multiple signals carried by different wavelengths that are received. Then all received signals are sent to a processor to recover the origin signals by a matrix operation.

The recovery process of multi-channel signals can be explained as follow. The received signals carried by different wavelengths $\lambda_1, \dots, \lambda_N$ is denoted by a data vector $\mathbf{d} = [d_1, \dots, d_N]^T$, where each element d_j is the signal carried by the wavelength λ_j . If wavelength λ_j does not arrive, d_j is zero. The passbands of the filters is represented by an $F \times N$ filter matrix \mathbf{A} , where F represents the number of filters. Each row of \mathbf{A} corresponds to a filter and each column corresponds to a wavelength. Each element has the value of “1” or “0”, meaning that the wavelength can or cannot pass the corresponding filter, respectively. The received signal vector $\mathbf{r} = [r_1, \dots, r_F]^T$ is obtained as $\mathbf{r} = \mathbf{A}\mathbf{d}$. Thus, F linear equations are obtained: $r_i = \sum_{j=1}^N a_{ij}d_j$ ($i = 1, \dots, F$). If these F equations are solvable, then the signals d_j from the F filters can be recovered. Solving a linear system of F equations needs additional digital signal processing (DSP), while the number of optical receivers can be reduced from N to F . Thereby the proposed receiver trades optical complexity for electronic complexity, which is attractive since the cost of high-speed DSP has dropped dramatically in recent years.

As discussed in the previous section, it is of low probability that signals carried by more than four wavelengths arrive at one node simultaneously. It means most of elements in \mathbf{d} are zero. Assuming that signals carried by M wavelengths arrive at one time, there are M non-zero elements in \mathbf{d} . We denote the positions of these nonzero elements by k_1, \dots, k_M and define $\mathbf{d}' = [d_{k_1}, \dots, d_{k_M}]^T$. Similarly, we define the receiver matrix \mathbf{A}' composed of columns k_1, k_2, \dots, k_M of \mathbf{A} . Now the equations are simplified to $r_i = \sum_{j=1}^M a_{ik_j} \times d_{k_j}$ ($i = 1, \dots, F$), or, equivalently, $\mathbf{r} = \mathbf{A}'\mathbf{d}'$. If the rank of \mathbf{A}' is M (in other words, all M columns in \mathbf{A}' are linearly independent), then the equations $r_i = \sum_{j=1}^M a_{ik_j} \times d_{k_j}$ ($i = 1, \dots, F$) are solvable, meaning that M simultaneously arrived signals can be recovered by using only F receivers.

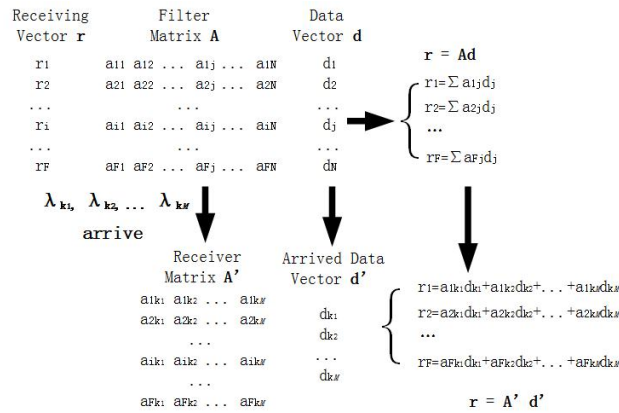


Fig. 3. Matrix recovery for M simultaneously arrived signals.

Figure 4 shows one example of our scheme. Four filters are used in each node to support in total 15 wavelengths for interconnection. In the filter matrix, any two columns are linearly independent, so the rank of any two-column receiver matrix in this scheme is always two. Such a scheme could recover two simultaneously arrived signals carried by any two wavelengths. For example, d_3 and d_{11} carried by wavelengths λ_3 and λ_{11} arrive at the same time. The receiver matrix now is composed of the third and eleventh columns of the filter matrix. By transforming the receiver matrix, we can solve the receiver equations recover d_3 and d_{11} from r_1 and r_2 .

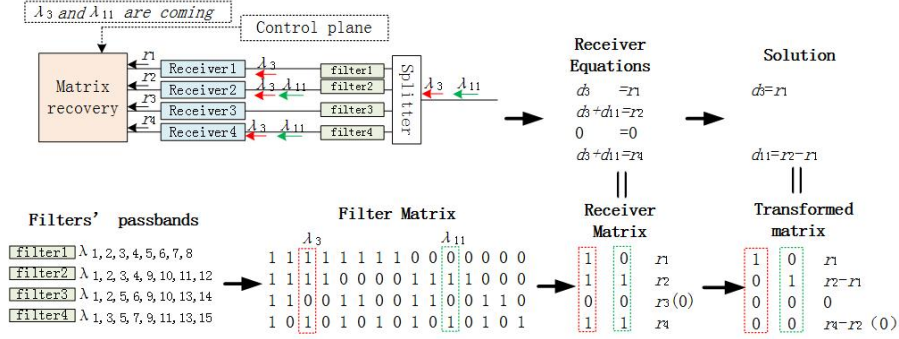


Fig. 4. One example: four filters for 15 wavelengths' interconnection.

The scheme in Fig. 4 is applicable to receive any two channels because any two columns in the filter matrix are linearly independent. If more than two channels arrive simultaneously, the feasibility of collision-free recovery depends on whether or not the corresponding columns in the filter matrix are linearly independent. If the three simultaneously arrived signals are d_3 , d_7 , and d_{11} carried by the wavelengths λ_3 , λ_7 , and λ_{11} (see Fig. 5 (a)), then columns 3, 7, and 11 of the filter matrix are extracted into a receiver matrix. The three columns are linearly independent, so the signals (d_3 , d_7 , and d_{11}) can be successfully recovered. On the other hand, if d_7 is changed to d_8 , recovery cannot be achieved because columns 3, 8, and 11 of the filter matrix are no longer linearly independent (see Fig. 5 (b)).

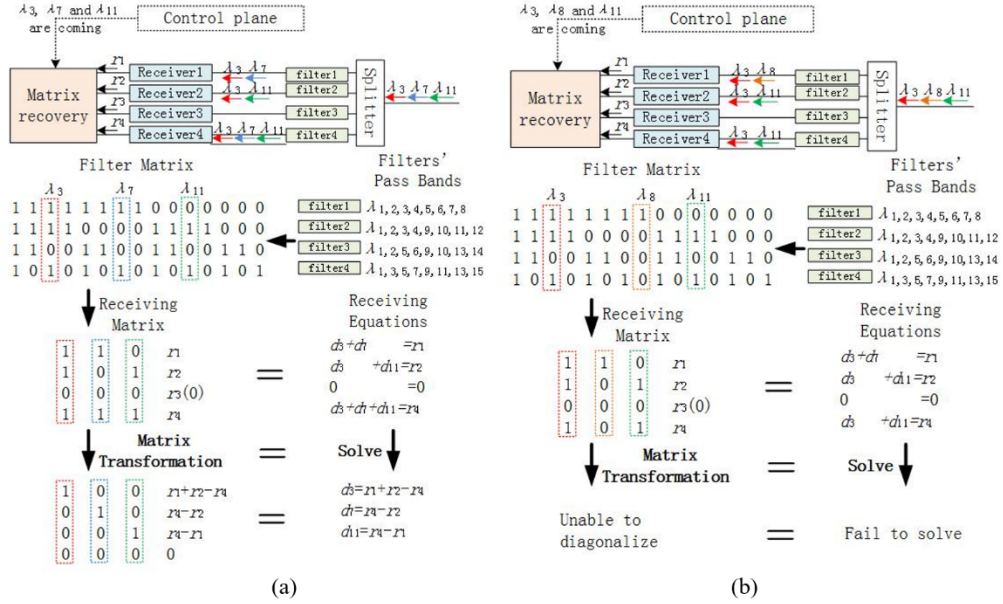


Fig. 5. Cases with three signals arrived simultaneously: (a) A successful case: d_3 , d_7 , and d_{11} arrive at the same time; (b) A failed case: d_3 , d_8 , and d_{11} arrive at the same time.

3. Filter Matrix Design

As illustrated above, the signals carried by different wavelengths can be recovered if and only if the columns of the receiver matrix are linearly independent. To ensure that the data carried by any M wavelengths can be recovered, any M columns in the filter matrix should be linearly independent. In general, a filter matrix should be designed to maximize the number of

wavelengths M that can be simultaneously received, for a given total number of wavelengths N and a given total number of filters and receivers F in each node, or to minimize F for given N and M . An optimized filter matrix is an $F \times N$ binary matrix, in which any M columns are linearly independent.

We propose to use error-control coding theory to design filter matrices. A binary block code with parameters (n, k, d) is a set of 2^k binary codewords with a length of n , such that every pair of distinct codewords differ in at least d positions. If the code is linear, it is fully specified by a parity-check matrix \mathbf{H} . \mathbf{H} is an $(n-k) \times n$ binary matrix in which any $d-1$ columns are linearly independent [11, Sec. 2.1.2]. With parameters $N = n$, $F = n-k$, and $M = d-1$, any parity-check matrix \mathbf{H} for a binary block code (n, k, d) can be used as a filter matrix (N, F, M) . There is a rich literature on binary block code design and extensive tables of the best known number of codewords of a binary block code for given n and d are available [12-14].

According to the relation among n , k and d in [14], we are able to determine the required number of filters F as a function of N and M (see Fig. 6). It can be seen that the minimum F is much smaller than N . For a given N , F increases with M . Furthermore, M could be increased from two to three or from four to five at the cost of just one more filter (and Rx), because the optimum k for any pair (n, d) with d even is the same as for $(n-1, d-1)$.




Fig. 6. Relations from coding theory among the numbers of required filters F , wavelengths N , and simultaneously received signals M .

In the error-control coding theory, linear independence is defined in the binary Galois field $\text{GF}(2)$. For $M = 2$, the condition of linear independence in $\text{GF}(2)$ is equivalent to linear independence in the real-number field. The binary block codes with parameters $(n, k, 3)$ always provide the optimal filter matrix design. Such codes are referred to as Hamming codes, for which $k = n - \lceil \log_2(n+1) \rceil$ ($\lceil \cdot \rceil$ is the ceiling function). Thus, F filters are able to support $\lfloor 2F-1 \rfloor$ ($\lfloor \cdot \rfloor$ is the floor function) wavelengths. However, when M is larger than 2, linear independence in $\text{GF}(2)$ is a sufficient but not necessary condition for linear independence in the real-number field. Thus, optimal binary block codes do not always provide the optimal solution to the filter-design problem. We take the matrix design with $M = 3$ as an example, which has the highest possibility when a reception collision occurs with $M > 2$, to illustrate the possibility of further reducing the number of required filters. We construct a filter matrix where the number of “1” elements in any column ranges from p to $2p-1$, for any integer p . In this way, the sum of any two columns has no less than $2p$ “1” elements, which cannot match any third column. Therefore, any three columns in the filter matrix are linearly independent. There are in total $\sum_{i=p}^{2p-1} \binom{F}{i}$ available column candidates for such a matrix. To get the largest N , the optimum p should be $\lfloor (F+1)/3 \rfloor$. Compared with the matrices provided by binary block codes, the improved ones could save one filter (and one Rx) in some cases when $M = 3$ (see Fig. 7).




Fig. 7. F vs. N for $M = 2$ and 3.

One example of the improved matrix ($F=4$, $N=10$, and $M=3$) and the corresponding filters is shown in Fig. 8. The four filters in Fig. 8 can support three-channel collision-free reception over ten wavelengths.

Receivers Matrix	Filter Passbands
1 1 1 0 1 1 1 0 0 0	filter1 $\lambda_{1, 2, 3, 5, 6, 7}$
1 1 0 1 1 0 0 1 1 0	filter2 $\lambda_{1, 2, 4, 5, 8, 9}$
1 0 1 1 0 1 0 1 0 1	filter3 $\lambda_{1, 3, 4, 6, 8, 10}$
0 1 1 1 0 0 1 0 1 1	filter4 $\lambda_{2, 3, 4, 7, 9, 10}$

Fig. 8. An improved filter matrix for $F = 4$, $N = 10$ and $M = 3$.

4. Multi-Level Signal Reception Experiments

Proof-of-concept experiments were carried out to verify the capability of multi-level reception. Four 10 Gb/s on-off-keying (OOK) signals were generated. To cover both wide and narrow channel spacing cases, the wavelengths were chosen as 1530, 1550, 1555, and 1556 nm. After bit synchronization and power adjustment, the signals were combined and then detected by a photo-detector. Measured eye diagrams and BERs for the received signals with data carried by two, three, and four wavelengths are shown in Fig. 9. The BER of the combined signals are estimated based on the Q factors of the eye diagrams. As each bit of the signal carried by one wavelength already brings some noise, the combination of more wavelengths shows worse performance.

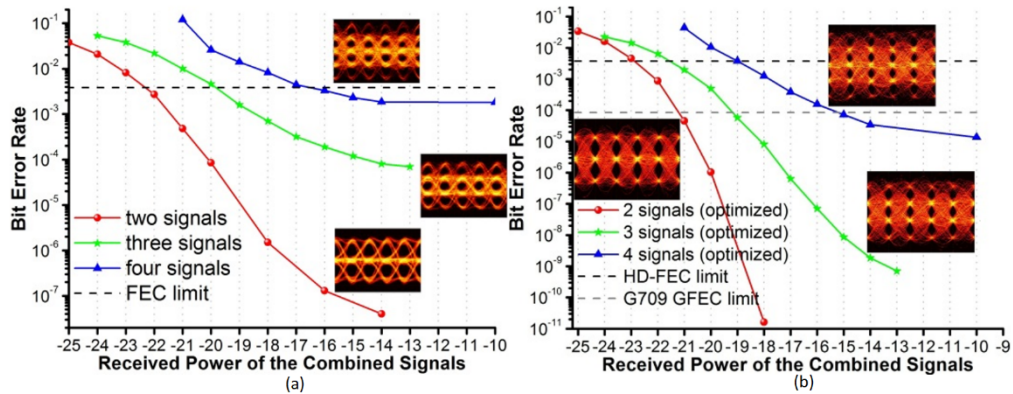


Fig. 9. BER vs. received optical power: (a) without any equalization and (b) optimized.

As reported in our ECOC paper [10], a 4-channel reception can reach the second generation hard-decision FEC limit at 7% overhead, i.e., $BER = 3.8 \times 10^{-3}$. Considering this FEC limit, the

receiver sensitivities are -22.3 , -19.8 , and -16.5 dBm for the combinations of two, three, and four wavelengths, respectively (see Fig. 9(a)). The sensitivities for the combinations of two, three and four wavelengths have 3, 5, and 6 dB difference compared with the receiver sensitivity for each signal carried by one wavelength. This is because the measured received power is the sum of all the received wavelengths and the signals for each wavelength are of the same power. Thus, the receiver sensitivity of each wavelength channel is -25.3 , -24.8 , and -22.5 dBm, respectively. The total link loss consists of an AWG loss (typically 5 dB), a power splitter loss (typically 6–12 dB for $F = 4$ –16) and a filter loss (typically 1 dB). The required transmitting power can be kept below 0 dBm, which implies that commercially available transceivers are sufficient for the proposed scheme.

In this work, we implemented electrical equalization (7 feed-forward taps and 3 feed-back taps) to further improve the transmission performance. The results are shown in Fig. 9(b). Compared to the cases without equalization, implementing equalization improves the sensitivities (under the second generation FEC limit) for the combinations of two, three, and four wavelengths to -22.9 , -21.6 , and -19 dBm, respectively, while the receiver sensitivity of each wavelength channel is -25.9 , -26.6 , and -25 dBm, respectively. The electrical equalization could provide another 2.5 dB power budget margin for 4-channel reception. It should be also noted that the 4-channel reception after equalization can reach the G.709 GFEC limit (i.e., $\text{BER} = 8 \times 10^{-5}$) [15]. G.709 GFEC is known as the first-generation FEC limit, which greatly reduces the complexity of FEC compared to the second-generation FEC limit. The sensitivity under the G.709 GFEC limit is -15.2 dBm for the combination of four wavelengths, while the sensitivity for each channel is -21.2 dBm. The required transmitting power can still be kept below 0 dBm, for which commercially available transceivers are also sufficient. For the cases $M = 2$ and $M = 3$, 8 filters are enough to support more than 100 wavelengths, so the splitter loss will be no more than 9 dB. The BERs reach 1×10^{-9} at -18.8 dBm ($M = 2$) and -13.5 dBm ($M = 3$). If the receiver is designed only for 2-channel or 3-channel collision-free reception and the transmitter power is above 1.5 dBm, FEC is no longer required. For modern AWGs with hundreds of channels, the power budget needs to be increased and the trade-off between performance and power considered. As illustrated in Figs. 7–8, the required number of filters, and hence the splitting loss, grows relatively slowly with the number of wavelengths.

5. Conclusions

A matrix reception scheme, which allows for receiving multiple wavelengths simultaneously, was proposed to avoid reception collision in optical interconnects. With only a few wavelength-fixed filters, the proposed multi-channel reception scheme is able to recover all signals, no matter by which wavelength they are carried. The optimized filter matrix design was achieved using error-control coding theory. For receivers that are able to receive more than two wavelengths, an improved filter matrix design was introduced. Experimental demonstrations verified the capability of receiving up to four wavelengths with a line rate of 10 Gb/s, which can be realized by commercially available transceivers. Implementing electrical equalization can further improve transmission performance, providing another 2.5 dB power margin for 4-channel reception.

For future optical interconnects operating beyond 10 Gb/s, transmission performance is expected to be degraded. To maintain acceptable quality, either the number of signals that can be simultaneously received needs to be reduced or a more complicated equalization scheme is required, which calls for future work.

Funding

This work was supported by the Swedish Foundation for Strategic Research, the Göran Gustafsson Stiftelse, the Swedish Research Council, the Swedish ICT-TNG, the National

Natural Science Foundation of China (NSFC) (61605047, 61671212, 61550110240), and the Natural Science Foundation of the Zhejiang Province (LQ17F050003).