

Genetic Algorithm Optimization of Multi Core Fibre Transmission Links based on Silicon Photonic Transceivers

A. Ottino¹, A. Saljoghei¹, T. Hayashi², T. Nakanishi², C. Kochis³, P. De Dobbelaere³, G. Zervas¹

¹Optical Networks Group, University College London, United Kingdom

²Optical Communications Laboratory, Sumitomo Electric Industries, Ltd., 1, Taya-cho, Sakae-ku, Yokohama 244-8588, Japan

³Luxtera Inc., Carlsbad CA 92011, USA,

g.zervas@ucl.ac.uk

Abstract: We demonstrate a genetic algorithm based system that can optimize optical interconnects using silicon photonic multi-core fibre coupled transceiver. The GA selects 48 parameters to deliver a minimum 6.9×10^{-16} BER on channels with diverse losses.

OCIS codes: (060.0060) Fiber Optics and optical communication; (060.2360) Fiber optics links and subsystems.

1. Introduction

It has been envisioned that the speed of the I/Os on semiconductor chips in computing systems will increase to reach capacities beyond 1 Pb/s by 2030 [1]. A front-runner application for fulfilling this purpose is a board-detachable optical transceiver in the form of a mid-board optic (MBO), which offers high bandwidth density and energy efficiency. In a recent work [2] a multi-processor system on chip (MPSoC) based memory disaggregated data center network (DCN) was reported using an optical circuit switched network using MBOs. In these disaggregated systems, high bandwidth density and high capacity optical transceivers are required to support dynamic all-optically routed communication between processors and remote high bandwidth memory (HBM) modules that can support over 1 Tb/s bandwidth.

Moreover, to reduce the front panel port and bandwidth density in DCNs, multi-core fibre (MCF) coupled transceivers have been recently explored [3]. MCF based data centre networks (DCNs) have also shown to outperform wave division multiplexing (WDM) systems in terms of performance blocking, cost and energy efficiency [4]. Thus, it can be anticipated that the integration of MCFs with MBOs in future DCNs can lead to substantial performance gains. However, opto-electronic transceivers embedded on disaggregated CPUs and memory modules offer a multitude of control parameters. In practice, transceiver channels might offer different signal quality and experience various levels of degradation and attenuation throughout the network. Thus, the optimum selection of these parameters can maximize system power budget, potentially leading to forward error correction (FEC) free operation, which is essential in low-cost, low-complexity and low-latency DCNs.

In this paper, we develop a purpose-made genetic algorithm (GA) and use it in real-time to optimize 48 equalization and amplification parameters for an 8 channels MCF-MBO driven by a Xilinx MPSoC with each channel operating at 10 Gb/s over an optical channel with diversely losses. The process took 13 hours instead of 1.44×10^{33} hours required for a brute force search method. Results suggest significant performance enactment in terms of bit error rate (BER).

2. System Optimization

Each channel of the transceiver has a total of 6 parameters to be optimized (making up 48 global parameters for all 8 channels), each with 5 to 32 possible values. The best approaches to handle such a problem are learning techniques or heuristic methods. Compared with heuristic methods, genetic algorithms (GAs) are known to better handle problems with high dimensionalities [5]. GAs are a subset of the evolutionary algorithms which are key for the optimization of non-linear problems by making use of principles such as natural selection, survival of the fittest and mutation. In optical communications, GAs have been studied in the context of various optical network optimization problems [5].

In this work the control parameters relate to the voltage drive levels and the equalization subroutines inherent to opto-electronic and electrical transceivers used. The GA monitors in real-time the bit error rate (BER), the bit error count and bit count performance for each individual optical channel for a given set of employed optimization parameters and changes them as the system evolves. The GA in the system proposed here starts with a given population n of random combination of control parameters (*individuals*), where the performance of each combination is evaluated in real-time. The best performing individuals (tournament size m) are selected, cloned, mutated and crossed over in order to create another population. The process is repeated for a finite number of iterations called generations. The population size n and tournament size m are the two parameters that control the diversity of the population and largely affect the performance of a GA.

3. System setup

The experimental setup used in this work is presented in Fig. 1 (a) and (b). A Xilinx Ultrascale+ MPSoC connects 8 of its bi-directional electrical transceivers via 16 pairs of differential RF links passing through two stages of connection conversion steps to a MBO that is coupled to two LC connectorized 2x4 core layout trench-assisted MCFs. It should be noted that the coupling losses associated with the integration of the MCF interface on the MBO account for 0.6-1.8 dB of loss. The MBO is supporting on-off-keying (OOK) up to 25 Gb/s [3]. The MBO's optical connections are looped back either directly or after fan-in-fan-outs (FIFOs) and passing through in-line attenuators to emulate practical link losses and reductions in optical signal to noise ratio (OSNR). Each channel of the electrical transceiver has 3 controlled parameters, one defines the differential drive level that is delivered to the RF links and the other two determine the transmitter side de-emphasis (pre- and post-) equalization (Fig. 1(a)). The opto-electronic transceiver supports three control parameters per channel, one for the transmitter side continuous time linear equalizer (CTLE), one for the receiver side pre-emphasis equalizer and one for the output driver.

The GA is implemented in Python running on a PC and it directly communicates in real time with the FPGA and the MBO. The procedure by which the GA operates in this work is outlined in Fig. 1 (c) and is allowed to run for a total of N generations (N_GEN). It should be noted that due to the dependencies between different channels, each optical transceiver in the MBO cannot be optimized individually, thus each individual contains control parameters for all 8 channels. At the first generation the system starts by first creating a population of n random combinations of 48 control settings (individuals). For each set of control settings, the performance of all 8 optical channels is measured in terms of BER and bit error counts, this measurement continues for three seconds. If no errors are encountered at 3 seconds, the measurement is continued until a bit error is encountered after which the measurement moves to the next combination of settings. Once all elements in the first generation are measured for, the m (tournament size) best combination of parameters in the population, which led to the lowest BERs, are selected. These are crossed over and mutated to create the n combinations that will be used on the next generation.

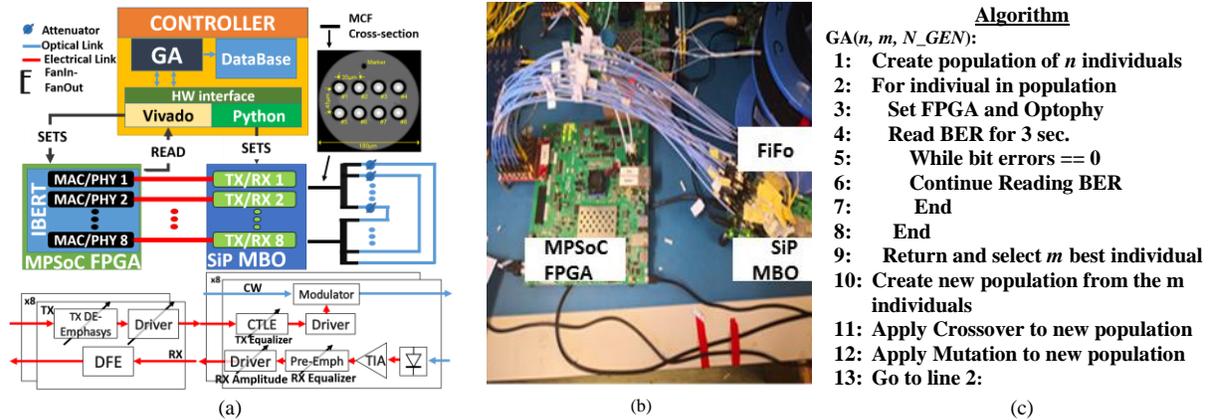


Fig. 1. a) Logical system (boxes with arrows signify the blocks controlled by the GA), b) experimental setup and c) GA algorithm used

4. Results and discussions

The performance of the GA algorithm highly depends on the initial population size n and the tournament size m . For this purpose, the performance of the system described in Fig. 1 with only FIFOs was evaluated for various population sizes of $n=100, 1000$ and tournament sizes of $m=15, 30, 150, 300$. The results in Fig. 2(a) and Fig. 2(b) were evaluated under an average optical loss of 4.8 dB and variance of 0.73 dB (combined loss of FIFOs and coupling loss of MCF interfaces). These results were evaluated for four different combinations of populations and tournament sizes, each repeated four times. It should be noted that the BER results shown in Fig. 2 (a) are for combination of control parameters which led to the minimum BERs in each run. Moreover, each BER value is derived based on the total number of errors received and bit transmitted from all eight channels (i.e. a single link with 80 Gb/s is assumed). Fig. 2(a) showcases that the higher population size $n1000$ (specially with lower tournament size) offers the best performance, where both cases of 150 and 300 tournament sizes achieve median BER rates (red lines) of 1×10^{-13} . This is due to the fact that a larger population size introduces more diversity into the system, increasing the possibility of crossovers and mutations. This increases the probability of the algorithm to find optimum solutions in least possible generations. Furthermore, it can be seen the $n1000$ - $m150$ case shown to arrive to BERs values below 5.5×10^{-15} . Fig. 2(a) also presents the performance of the system where GA is only applied to electrical transceivers (FPGA only) saturating the performance to a BER of approximately 3.1×10^{-12} . Fig. 2(b) shows the evolution of the GA over various generations in terms of the median BER achieved in each generation for four independent runs. It's evident that the

n1000-m150 and n1000-m300 cases arrive to a lowest possible BER in the least number of generations. In terms of time taken to arrive to the lowest BER, high population cases take longer, since each generation has to analyze more elements. n1000-m150 and n1000-m300 run for total of 13 hours and 5.7 hours respectively in order to achieve the lowest BER. However, n100-m15 and n100-m30 reach their lowest BERs in 2.9 and 0.8 hours respectively. As it can be further seen in Fig. 2(a) and (b), the variability in the measurements has a dependency on the tournament size where increasing it leads to lower variance in results; this is due to the fact that a larger m increases the diversity in subsequent population. Finally, to assess the performance of the proposed system over various channel conditions the GA is further evaluated for a) back to back configuration (loop back of MCF connectors) achieving mean loss of 1.1 dB with variance of 0.3 dB across all channels, b) FIFO (mean attenuation 4.8 dB and variance of 0.73 dB), c) attenuation 1 (mean att. 10.24dB, var. 2.3dB), d) attenuation 2 (mean att. 12.9dB, var. 9.67dB) while running the test case GAs n1000-m150 and n1000-m300 twice each time. As it can be seen in Fig. 2(c) with Att. 1, the min BER achieved is approximately at 3.88×10^{-12} and for Att. 2 its approximately 5.79×10^{-11} while for back to back scenario the min BER achieved is 6.94×10^{-16} . For comparison, a semi-brute force scheme not considering channel dependencies is used to optimize the electrical transceiver (took over 19 hours) after which the MBO control parameters were manually optimized. The results for this are plotted in Fig. 2(c) with Att. 1 and back to back scenarios. The semi brute-force & manual tuning approach led to a BER reading above 1×10^{-4} and 9.8×10^{-13} for Att. 2 and back to back scenarios respectively whereas the GA reached BER below 1×10^{-10} and 6.94×10^{-16} (limited to measurement time) for the respective scenarios. If a brute-force method was used to fully analyze all the various combinations of control settings (5.03×10^{36} combinations), it would have taken the system over 1.4×10^{33} hours. Single channel analysis is shown in Fig. 2(d) in terms of BER vs received optical power after optimization, where it can be noted that the BER values for each channel are approximately constant irrespective of loss (apart in Att. 2 due to the high variance in optical loss).

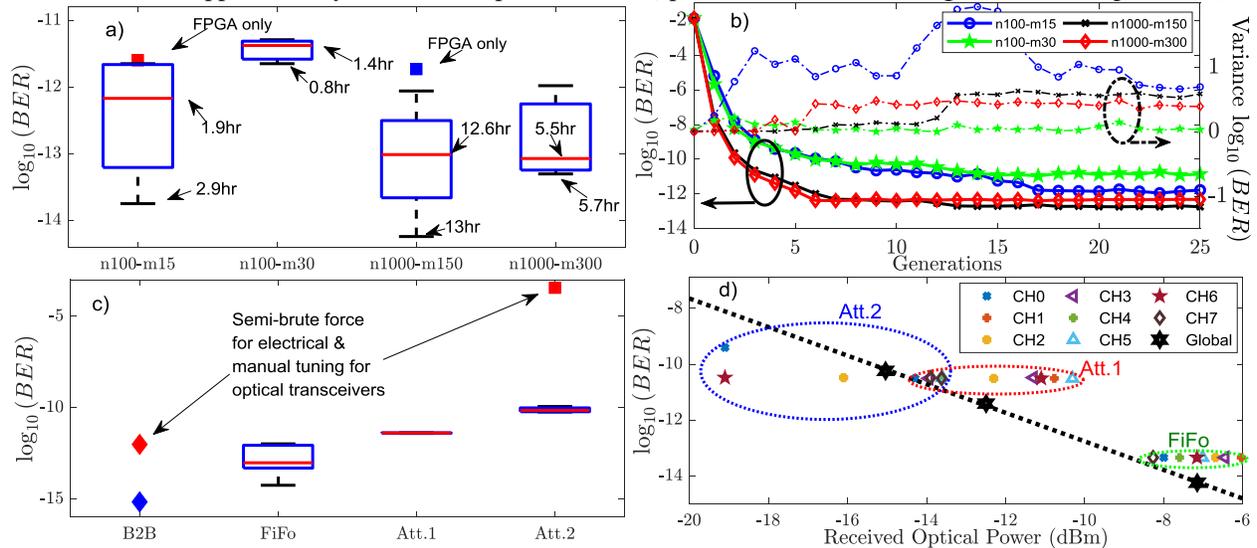


Figure 2(a) Box plot of minimum BER found on FIFO configuration for various tests along the time taken to achieve mean and min BER, (n100-m15: population size 100, tournament size 15). (b) The mean BER versus generations vs variance under FIFO configuration. (c) Minimum BER (measured in condition with n=1000, m=150,300) for various loss scenarios, (d) BER vs. received optical power for various combinations

5. Conclusion

In this paper, genetic algorithms were used to identify a set of 48 optimum parameters used to control an 8 channel MCF based transceiver for data center network interconnects which led to the achievement of BER levels below 6.94×10^{-16} . It was further shown that GA outperformed the combination of manual and semi-brute-force optimization. The GA reaches the global optimal BER in 32-orders of magnitude faster than a brute force sweep method. Crucially, the GA can deliver uniform BER across all channels even if they experience optical loss variations up to 9.76dB. A tradeoff between various choices of population and tournament size suggested that higher population sizes lead to best solutions.

Acknowledgement This work was supported by EC H2020 dReDBox and EPSRC TRANSNET EP/R035342/1.

References

- [1] G. Zervas *et al.*, IEEE JOCN, vol.10, no.2, p270-285, 2018.
- [2] A. Saljoghei *et al.*, paper W1C.1, OFC 2018,
- [3] T. Hayashi *et al.*, paper Th.2.A.4, ECOC 2017
- [4] H. Yuan *et al.*, IEEE JOCN, vol.10, no.4, pp.272-288, 2018.
- [5] J. Mata, *et al.*, Optical Switching and Networking, vol.28, pp.43-57, April 2018