Proceedings of the ASME 2021 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems
InterPACK2021
October 26-28, 2021, Virtual, Online

IPACK2021-73253

BER ILLUSION METHODOLOGY: A NOVEL, OPEN SOURCED AND SCALABLE APPROACH TO TROUBLESHOOTING HIGH RADIX PHOTONICS INTERCONNECTS IN A MODERN HYPERSCALE DATACENTER

Shreyas Rao
Facebook, Inc.
Menlo Park, USA

Abhijit Chakravarty
Facebook, Inc.
Menlo Park, USA

Dharmesh Jani
Facebook, Inc.
Menlo Park, USA

ABSTRACT

Facebook Datacenter consists of a large number of servers that run diverse Facebook services aggregated to serve any given user request. To allow this aggregation, servers have to interact with each other via different traffic flows which are managed by networking fabric. The underlying connection powering this fabric consists of a large number of pluggable optical interconnects and On Board Optical (OBO) modules carrying production data. This connectivity at scale requires fast and reliable detection of the link failures to ensure resolution. In the first generation of the deployments, detection of the link failure was sequential and a slow process. The troubleshoot process was equally tedious as the available tools required characterizing one optical transceiver at a time. Further, the failure analysis also presented a majority of resolution with no failed optics as a root cause resulting in high No Trouble Found (NTF) rate.

In this paper we introduce a novel link failure detection and resolution method that improves on the previous method across three dimensions: faster resolution, reliable troubleshooting and scalable implementation. We introduce BER Illusion Methodology (BIM) that is a highly scalable and resource efficient solution that significantly reduces the time taken to troubleshoot pluggable optical interconnects. This is also scalable to next-gen OBO modules at Facebook datacenters aiming to lower the NTF rate and optimally utilizing the available resources. BIM, which is based on Open Compute Platform (OCP) network switches, can be used to troubleshoot 128 QSFP28, 64 QSFP56 or 32 OBO modules simultaneously in under 30 minutes. The tool is easy to implement and capable of also reporting diagnostics on the transceiver such as Transmitter Power, Transmitter Bias Current, Receiver Power, Case Temperature, Bit Error Rate result per channel, Vendor information and Manufacturing part number. This additional test data report along with true failure indication helps optic suppliers gain confidence and build customer credibility. The open-source nature and the universal applicability of this tool offers possibility for other users to adopt and further customize it for their networking needs.

Keywords: Optical Communication; Optical High-Speed Interconnects; Post-Production Troubleshoot Process

NOMENCLATURE

OCP	Open Compute Platform
BER	Bit Error Rate
OBO	On-Board Optics
NTF	No Trouble Found
QSFP	Quad Small Form-factor Pluggable
BOM	Bill of Materials
UUT	Unit under test
MTTR	Mean time to resolve
ASIC	Application Specific Integrated Circuit
DOM	Digital Optical Monitoring
FPGA	Field programmable gate array
CLI	Command Line Interface

1. INTRODUCTION

Facebook services are utilized by over a billion people on a regular basis, relying on the seamless and highly available performance. This network infrastructure needs to constantly evolve and keep up with the application needs. The next-generation Facebook data center building is a disaggregated system that replaces the large cluster devices into smaller high-performance units that provides connectivity among the server racks within the data center [1]. Figure 1 shows the connectivity between spine, fabric and rack layers that constitute the interconnect.

Each connection seen in figure 1 represented by multiple colors, on a physical level, are made up of optical links supporting ever increasing demand of the bandwidth and data rate. The number of optical modules deployed in such a hyper scale data center is large and the reliable detection of link failure and timely resolution of these failures gets challenging.

In this paper, we present a novel methodology, BER Illusion methodology (BIM), aimed at faster, reliable and scalable link failure detection and resolution leveraging the underlying Open Compute Platform (OCP) hardware and software stack of the data center. BIM, a Pseudo Random

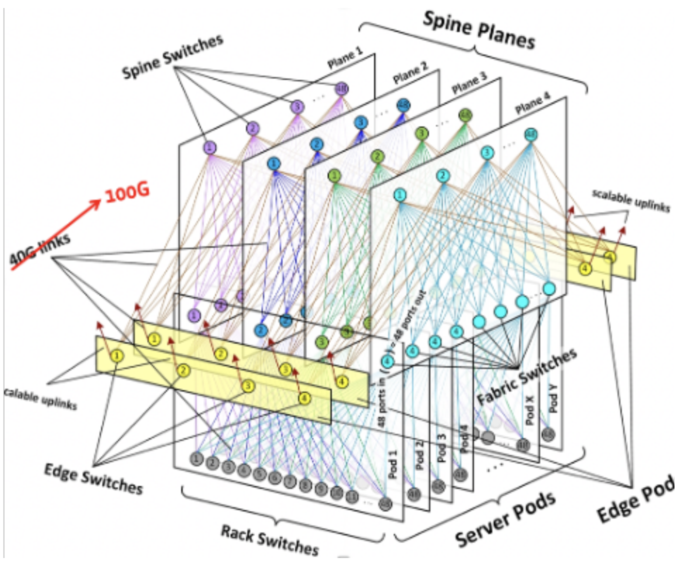

FIGURE 1: FACEBOOK DATA CENTER NETWORK TOPOLOGY

Binary Sequence (PRBS) based Bit Error Rate (BER) approach, is also a resource efficient solution that significantly reduces the time taken to troubleshoot pluggable optical interconnects as well as next-gen On-Board optical modules (OBO). Typically aimed at lowering the No-trouble found (NTF) rate and optimally utilizing the available resources, BIM also provides diagnostics data of the optical module for faster failure analysis. This data includes Transmitter Power, Transmitter Bias Current, Receiver Power, Case Temperature, Bit Error Rate result per channel, Vendor information and Manufacturing part number.

Detection and troubleshooting of failed optical module becomes highly automated by simply plugging in the Unit Under Test (UUTs) on OCP hardware platform and connecting loopback cables. Failure Analysis shows typical failures relate to side mode suppression ratio and laser optical power failures. The hardware runs a PRBS generator as well as a PRBS checker to calculate the BER over a user defined test time. The extensive report generated along with true failure indication builds trust with the suppliers and helps them gain confidence with the analysis. The rest of the paper is divided into 4 sections. In section 2 of the paper, we discuss some of the challenges associated with current debugging and link failure detection techniques. Section 3 gives an introduction to the BIM development and underlying hardware stack. The process flow for a failed optic is discussed in detail in section 4. The results and advantages seen as compared to conventional methodology is discussed in section 5.

2. CHALLENGES

The scale at which optical interconnects are deployed in the Facebook data centers is enormous [2]. Even a low failure rate accounts to a significantly large number of interconnects that need to undergo troubleshooting and be replaced. The tools currently available in the market to troubleshoot are expensive and the process is tedious allowing to test a single part at a time. This deployment scale at Facebook cannot be supported using these equipments.

As this tedious process is completed, the units are then shipped to the suppliers for Failure analysis. With little information available, the suppliers need longer time as well as resources to identify the root cause. More often than not, these units are shipped back to the data center spares as No trouble found. Although there are testing tools available that allow us to identify NTF prior to shipping the units to the supplier, the process of deployment is largely slowed down and causes operational overhead.

As the NTF rates keep getting higher and the mean time to resolve (MTTR) the failure takes longer, another challenge is in identifying and understanding product reliability. The failure with NTF can skew the data, not providing enough insights on the product performance in the data center. There is a strong need to address these challenges with key aspects being true failure identification, lesser manual efforts and simplified process for lower resolution time. BIM is a solution that targets these three pressure points in addition to providing flexibility and a large range of advantages.

3. BIM DEVELOPMENT AND ARCHITECTURE

The development of the BIM solution started on our Wedge 100S platform. Wedge 100S is open-sourced and disaggregated top of the rack switch deployed at a large scale in Facebook data centers. The software stack of this switch uses FBOSS and OpenBMC which is Facebook's networking stack and baseboard management controller, providing flexibility to introduce new features and innovate. This flexibility allows us to use some of the key features to troubleshoot optics.

Leveraging these open-source features along with the software development kit (SDK) provided by our PHY (gearbox) vendors, we developed the solution that can enable a PRBS generator as well as a PRBS checker to measure the BER. The gearbox generates a PRBS sequence and acts as the sender. This sequence traverses through the optical interconnect and the receiver then uses this bit stream to calculate the BER of the transmission. The receiver could be a different hardware device on the remote side, such as a different gearbox on the same switch or a gearbox in a different switch. Or the receiver can be the sender itself by looping the traffic back to the sender using loopback cables on the optics.

For the purpose of troubleshooting optics in our fleet, we use the PRBS based approach to triage and diagnose link flaps and failures. Since we are testing individual optical transceivers and to better isolate the issue, we use the loopback cables and connect the Transmitter of the optical module to the Receiver of the same module. Thus, the gearbox on the switch acts as a PRBS sender as well as receiver. Although PRBS is supported on forwarding ASIC as well, theoretically lowering the segment length between sender and receiver helps pinpoint the failure. The sender generates a binary sequence using a PRBS31 polynomial with a user defined test time defaulting to 120s for a confidence level of over 90%.

Since the Wedge 100S architecture makes up of 3.2Tb/s aggregate switching bandwidth and 32 port availability on a single switch, the maximum data rate per port is 100G. The test station is stacked with 4 such devices allowing 128 QSFP28 units to be tested simultaneously. This significantly improves the performance in 2 vectors: mean time to troubleshoot a single transceiver and number of transceivers than can be tested at any given point of time. The BER is observed for each channel of a QSFP28 module where there are 4 channels per module [3, 4].

The increasing demand of data center bandwidth saw a shift from 100G to 200G links. As the initial development was limited to 100G, we moved the hardware stack from Wedge 100S to Minipack to support QSFP56 along with QSFP28 testing. Minipack is a successor to Facebook's modular switch Backpack with significant power saving. Primarily aimed at the Fabric layer of the datacenter, Minipack hardware as well as software stack is also open sourced with 128 ports capable of supporting 100G and 64 ports supporting 200G data rate supported by a 12.8Tb/s aggregate switching ASIC [5]. The next generation linecard developed by Facebook consisting of 4 On-Board optical (OBO) modules per linecard can be plugged into the Minipack chassis (8 per chassis) thus troubleshooting 32 OBO modules and providing a single hardware platform for testing QSFP28, QSFP56 and OBO modules.

These OCP switches include digital optics monitoring (DOM) acceleration function. The DOM FPGA, present on each linecard, periodically polls the optical modules for transceiver information through low-speed I2C buses. The switch main board consists of Input/output block (IOB) FPGA which is connected to the CPU by a peripheral component interconnect express (PCIe) link. This IOB FPGA communicates with the DOM FPGA to fetch Transmitter Power, Transmitter Bias Current, Receiver Power, Case Temperature and Manufacturing part number for each module plugged into the linecard.

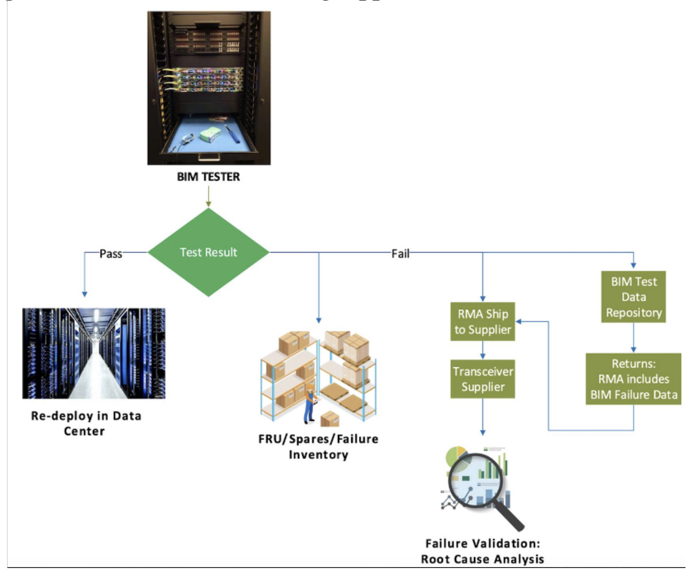

FIGURE 2: BIM IDEOLOGY

In figure 2, we show the ideology where BIM fits in the troubleshoot process at the datacenter. All the failed optics from the data center are tested in a test station. Based on the test result, if the result is a pass, we re-deploy these units in the data center as and when needed or keep them as spares. If the test result shows a failure, we have a true failure indication and the unit along with the test report is sent to the supplier for Failure Analysis (FA). The information we collected pertaining to the Transmitter and Receiver (Current and power) when the module failed provides suppliers a better insight if the issue is related to driver or lasers within the module. This helps suppliers provide a faster root cause analysis. Since we are sending the units to the supplier for FA only when a true failure is detected, the No trouble found rate reduces significantly.

4. TROUBLESHOOT PROCESS

The BIM solution is developed as a CLI based tool where the optic speed to be tested, test time, username, password and IP address is passed as arguments. The help command is user-friendly and allows the user to see the possible optional arguments that can be passed along with required arguments. Figure 3 shows a screenshot of the possible arguments that can be passed along with the tool. Most of the options are self-explanatory. The speed option is used to configure the port speed at which the user desires to test the optic and the mode option allows interoperability tests between different types of optics such as QSFP28 and QSFP56 operating at 100G.

Table 1 summarizes the different speed and mode combinations that can be used for troubleshooting.

The process flow for BIM is summarized in figure 4. Once the switch chassis is loaded with optics to be tested, the modules or LC connectors are connected with loopback cables. Running the CLI command with desired arguments initiates the test. The highly automated approach then connects to the test station, with IP address passed as required field, via SSH and copies the configuration file that initializes the optics with desired speed and mode. The tool verifies if ports are up and sets the pre-emphasis values.

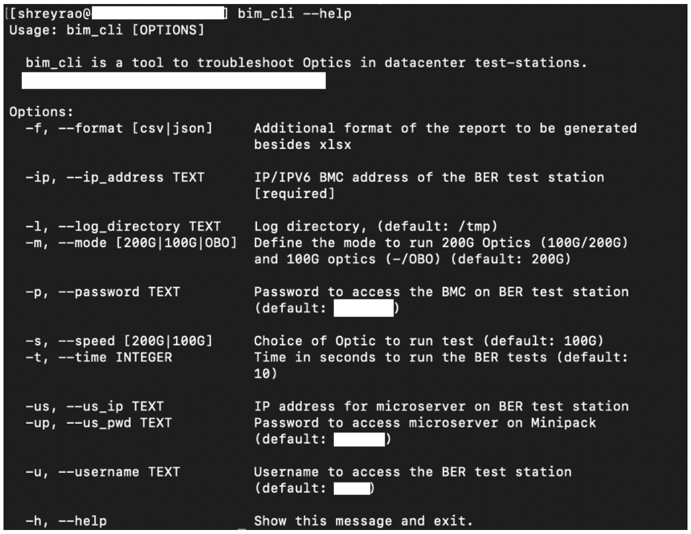

```
[shreyrao@[redacted]] bim_cli --help
Usage: bim_cli [OPTIONS]

  bim_cli is a tool to troubleshoot Optics in datacenter test-stations.
  [redacted]

Options:
  -f, --format [csv|json]     Additional format of the report to be generated
                              besides xlsx

  -ip, --ip_address TEXT      IP/IPV6 BMC address of the BER test station
                              [required]

  -l, --log_directory TEXT    Log directory, (default: /tmp)
  -m, --mode [200G|100G|OBO]  Define the mode to run 200G Optics (100G/200G)
                              and 100G optics (-/OBO) (default: 200G)

  -p, --password TEXT         Password to access the BMC on BER test station
                              (default: [redacted])

  -s, --speed [200G|100G]     Choice of Optic to run test (default: 100G)
  -t, --time INTEGER          Time in seconds to run the BER tests (default:
                              10)

  -us, --us_ip TEXT           IP address for microserver on BER test station
  -up, --us_pwd TEXT          Password to access microserver on Minipack
                              (default: [redacted])

  -u, --username TEXT         Username to access the BER test station
                              (default: [redacted])

  -h, --help                  Show this message and exit.
```

FIGURE 3: BIM CLI USER-INTERFACE

Mode | Speed	100G	200G	OBO
100G	QSFP28 at 100G	Physically not possible	OBO module testing
200G	QSFP56 at 100G	QSFP56 at 200G	Not supported currently

TABLE 1: OPERATING MODE AND SPEED FOR BIM

Post initialization, the tool enables the PRBS generator on each gearbox and runs the test for user defined time. During this interval we do not perform any action. After the test time, the tool probes for BER value from each channel of each port. As a next step, we query the DOM FPGA for optic health and data which is parsed into excel report. This report is generated on the local machine as well as stored to a database/central repository. The status during the entire duration of the test can be tracked by real-time updates on the terminal and detailed logs are also stored at /tmp folder unless otherwise specified by the user.

As seen in figure 4, some additional tests are run for Miniphoton. Miniphoton is the next generation On board optic linecard developed by Facebook. Unlike pluggable optical transceivers, it is not possible to replace a single port or LC connector in case of failure. This leads to entire linecard replacement. More stringent tests are run on these OBO modules. Apart from the line side BER test, we run a system side BER test to ensure the link between gearbox and the switching ASIC is functioning correctly. The OBO modules provide an option to test the electrical loopback as well as optical loopback within the module providing an additional test

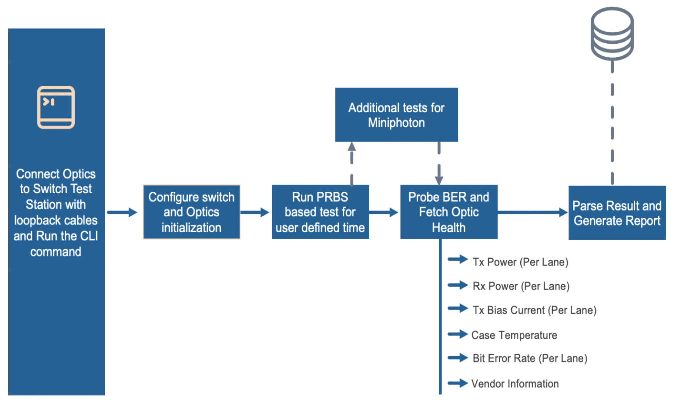

FIGURE 4: PROCESS FLOW FOR BIM

to ascertain the electrical path or optical path prior to the front panel is not faulty. One of the three paths that a signal can take depends on the configuration at the start of the test. This helps isolate if the issue is with drivers, modulator, laser or the front panel.

5. RESULTS AND ADVANTAGES

A proof-of-concept test station was set up in one of the Facebook labs. Figure 5 shows a stack of 4 Wedge 100S with 128 QSFP28 parts loaded and connected with loopback connectors.

FIGURE 5: PROOF OF CONCEPT SET UP AT FB LAB WITH WEDGE100S

BER test was run for 180 seconds, and a sample report was generated as shown in figure 6. The report contains vendor information, technology, optical health data, BER result and a Pass or Fail status. The optics that show failure are true failures and sent to the supplier for Failure analysis along with this report. The optical module that shows Pass, which would have eventually returned by the supplier as a NTF, goes into spares saving time and resources.

A pilot run was introduced with one of the suppliers and failure analysis data shows 50-60% reduction in NTF rate along with faster resolution. Moreover, compared to the conventional troubleshooting method the time saved by BIM to troubleshoot 128 optics goes down from ~900 minutes to about 15 minutes. The data from the sample report as shown in figure 6 is synced to a central repository for data collection that can be used for future trend analysis and product reliability. The feedback from suppliers showed higher confidence as well as faster MTTR.

Parameters	Port1	Port2	Port3
Vendor Name	NA	Supplier A	Supplier B
Vendor SN	NA	AAAAAAAAAAAA	BBBBBBBBBBBB
Technology	NA	100G_CWDM_LITE	CWDM_MSA_100G
Vendor PN	NA	XXXXXXXXXXXX	YYYYYYYYYYYY
Vendor rev	NA	1.0	2.3
TX1 Power(mW)	NA	0.7800	0.6219
TX2 Power(mW)	NA	0.7542	0.6487
TX3 Power(mW)	NA	0.5819	0.6559
TX4 Power(mW)	NA	0.7731	0.6253
RX1 Power(mW)	NA	0.8325	0.0313
RX2 Power(mW)	NA	0.8684	0.0312
RX3 Power(mW)	NA	0.6840	0.0298
RX4 Power(mW)	NA	0.8507	0.0320
TX1 Bias(mA)	NA	60.950	60.950
TX2 Bias(mA)	NA	60.950	60.950
Tx3 Bias(mA)	NA	60.950	60.420
Tx4 Bias(mA)	NA	60.950	60.420
Temperature(degC)	NA	35.33984375	41.11718750
Lane1_BER	NA	better than5E-5 (PASS)	>5E-5(FAIL)
Lane2_BER	NA	better than5E-5 (PASS)	>5E-5(FAIL)
Lane3_BER	NA	better than5E-5 (PASS)	>5E-5(FAIL)
Lane4_BER	NA	better than5E-5 (PASS)	>5E-5(FAIL)
Test Results	NA	PASS	FAIL

FIGURE 6: SAMPLE REPORT

This methodology has numerous advantages over the conventional method. BIM is automated and highly scalable. One example of scalability can be seen in the discussion, how by just changing the hardware platform we were able to support multiple optical technologies as well as support a higher number of UUTs on a single platform. The tool is extremely efficient as the Bill of material (BOM) consists of in-house switch, optics and loopback cables. No external device is needed for troubleshooting. The test time reduction and NTF rate numbers mentioned in the results are the key highlights of BIM. Storing reports and optical data in a central repository improves the inventory management and tracking. As no human intervention is needed after plugging in the optics, BIM can also be used for long term reliability testing of optical modules by exposing the module to a longer test time and PRBS sequence [6].

6. CONCLUSION

In this paper, a novel methodology is introduced and developed using the PRBS based BER approach for troubleshooting and identification of link failure with an optical interconnect. This method shows significant improvement in the effectiveness and efficiency compared to the conventional methodology. With the scalability offered by the tool, BIM can be extended to 400G as well as 800G solutions seamlessly.

REFERENCES

[1] Alexey Andreyev, “Introducing data center fabric, the next-generation Facebook data center network:, Facebook Engineering, 2014

[2] Venkataraman and Thayer, Tang, "Facebook’s Large Scale Monitoring System Built on HBase", Strata + Hadoop World Conference, 2012

[3] Abhijit Chakravarty, Katharine Schmidtke, Srinivasan Giridharan, John Huang, and Vincent Zeng, “100G CWDM4 SMF Optical Interconnects for Facebook Data Centers”, Conference on Lasers and Electro-Optics, OSA Technical Digest, paper STu4G.1, 2016

[4] Katharine Schmidtke, Vincent Zeng, Abhijit Chakravarty “Facebook: CWDM4-OCP, 100G Optical Transceiver Specification”, Open Compute Project © Facebook, 2017'

[5] Facebook, Inc. , “Facebook Minipack” , OCP Global Summit, 2019

[6] Abhijit Chakravarty, Srinivasan Giridharan, Matt Kelly, Ashwin Poojary, Vincent Zeng, “Characterizing Large-Scale Production Reliability for 100G Optical Interconnect in Facebook Data Centers”, Frontiers in Optics / Laser Science (FiO/LS), 2017