Consortium For On-Board Optics
The Use of On-Board Optic Compliant Modules in Coherent Applications

COHERENT WORKING GROUP PARTICIPANTS

3M Company
ADVA Optical Networking SE
Applied Optoelectronics Inc.
Arista
Celestica
Ciena
Cisco
Credo Semiconductor
DustPhotonics Ltd.
East Point Communication Technology Co., Ltd.
Fujikura Ltd.
Fujitsu Optical Components
Furukawa Electric
GlobalFoundries
Hirose Electric (USA)
HUBER+SUHNER AG
IBM Corporation
Inphi
Jess-Link Products Co., Ltd.
Lumentum
Luxshare Precision Industry Co., Ltd.
Mellanox
Molex Electronic Technologies
Microsoft
Oclaro
Optec Technology
Optomind
Panduit
PETRA
Quanta
RANOVUS
Reichle & De-Massari AG
Rosenberger
Samtec
Semtech
SENKO
Sicoya
Sumitomo
TE Connectivity
US Conec
Yamaichi ElectronicsInnovium

CONTENTS

1.0 Introduction...1

2.0 The Relevance of COBO to Coherent Applications Compliant Modules in a Coherent Application.......3

2.1 Definition of the 400ZR Use Case and Its Requirements.........................3

2.2 Other Applications...4

3.0 Host System Design Considerations...5

3.1 Requirements and Features..5

3.1.1 Environmental, Power and Size Targets for a Reference Application..........5

3.2 Module Placement on Host Board...5

3.2.1 Host Board Placement Considerations.......................................6

3.2.2 Host Board Placement Examples...8

3.3 Additional Considerations..8

3.3.1 Electrical Connectors on Host Board.......................................9

3.3.2 Retention of Modules onto Host Board......................................9

3.3.3 Assembly and Latching of OBOs on Host Board..............................10

3.3.4 Attachment of Heat Sinks...11

3.4 Modelled Thermal Capabilities for 1RU DCO Application.....................12

3.4.1 Thermal Model of COBO-Compliant Coherent OBO..............................12

3.4.2 Heatsink and Airflow Assumptions...13

3.4.3 Thermal Modelling Results...15

4.0 Coherent DCO COBO Module Design Considerations...............................19

4.1 Class C COBO Specification is Intended for Coherent Applications..........19

4.2 Example Coherent Module Implementations...................................19

4.2.1 Fully Integrated Coherent Optical Subassembly.............................19

4.2.2 Integrated High-speed Coherent Optics with External Laser.................21

4.2.3 Fiber Management Inside a Coherent COBO-compliant OBO.....................22

5.0 COBO Support for Additional Functions..23

5.1 Sync for Coherent Using COBO-Compliant Modules............................23

5.1.1 SyncE 400ZR Example...23

5.1.2 SyncO FlexO Example...24

5.1.3 SyncE over SGMII..24

5.2 GCC for OTN...25

5.3 Precision Timing Protocol..25

5.4 Fast Communications with Host..26

CONTENTS

6.0 Management Interface for Coherent COBO Modules....27
  6.1 Management Interface Specification Status....27
  6.2 Relationship to Management of Optical Subassemblies and ICs....27
  6.3 Features to Support Coherent Performance Management....28
7.0 Appendix: Fiber Reliability Background....29
  7.1 Overview....29
  7.2 Reliability of Thirty-Two Class C Modules on a Host Board....30
    7.2.1 Methodology....30
    3.2.1 Host Board Placement Considerations....30
    7.2.2 Reliability Estimate for Example System Placement....30
8.0 References....33

FIGURES

Figure 2.1: Data Center Network Interconnect Example.......... 3
Figure 3.1: Host Board Application Concept.......... 7
Figure 3.3: Class C On-Board Optic Dimensions Including Exit Fiber.......... 8
Figure 3.4: Example Host Board Placement Options.......... 9
Figure 3.5: Module Install and Latching with High Speed Connector Latch.......... 10
Figure 3.6: Module Install and Latching with Low Speed Connector Latch.......... 11
Figure 3.7: Example Heat Sink Attachment Options for C x8 Modules.......... 12
Figure 3.8: COBO Compliant Coherent DCO Module Thermal Model.......... 13
Figure 3.9: Trade Offs Between Heat Sink and Module Height.......... 14
Figure 3.10: Chassis Front Panel Showing LC Connectors Below Air Intake.......... 15
Figure 3.11: Thermal Simulation Results for 2 Rows of 15 W COBO Modules.......... 16
Figure 3.12: Thermal Simulation Results for 2 Rows of 17.5 W COBO Modules.......... 16
Figure 3.13: Thermal Simulation Results for 3 Rows of 15 W COBO Modules.......... 17
Figure 3.14: Thermal Simulation Results for 3 Rows of 17.5 W COBO Modules.......... 17
Figure 3.15: Airflow for 3 Row Configuration.......... 18
Figure 4.1: Example Module Design Using Fully Integrated Optical Subassembly.......... 20
Figure 4.2: Example Module Design with High Speed Optical Subassembly and External Laser.......... 21
Figure 4.3: Detail of Fiber Management for External Laser.......... 22
Figure 5.1 Synchronization to an External Clock Using REFCLK Pin.......... 24
Figure 5.2: Synchronization Using SGMII.......... 24

TABLES

Table 3-1: Example Host System Parameters for Reference Application........5
Table 5-1: Timing Performance for PTP........26
Table 7-1: Estimated Failure and Fit Rate as a Function of a Lifetime........27
Table 7-2: Estimated Failure and Fit Rate as a Function of Bend Radius........31
Table 7-3: Estimated Failure and Fit Rate as a Function of Proof Test........31
Table 7-4: Estimated Failure and Fit Rate as a Function of Fiber Cladding Diameter........32

1. INTRODUCTION

The COBO (Consortium for On-board Optics) specification [1] was written to cover both intra-datacenter and inter-datacenter applications. The latter, also commonly referred as Data Center Interconnect (DCI), is a rapidly growing application and the one discussed in this document. Ever increasing content within datacenters is driving exploding traffic volumes between datacenters. In the next generation, DCI growth will be accelerated by the new trend to edge computing. New applications require time-sensitive performance served out of distributed edge data centers connected with low latency links. Service robustness requires workloads and data to be moved swiftly between datacenters.

Starting about five years ago, specialized network equipment for DCI application started to become mainstream. This type of equipment continues to evolve rapidly to meet the bandwidth, cost and density requirements of the application, particularly to meet the needs of the biggest content providers. As with servers, switches and storage inside the datacenter, the dramatic technology progress in DCI system capabilities driven by these aggressive early adopters also benefits the broader set of datacenter operators.

Networking equipment continues to become denser, cheaper, faster and consume lower power. But, that's not all. The incredible pace of growth and the unpredictability of new applications mean that datacenters-- and by extension, datacenter equipment-- must be more flexible as well. New solutions are needed to ease the ever-increasing system implementation challenges. With coherent optical engines scaled down in size, cost and power for reaches under 100 km, coherent solutions can support a wider variety of network applications. The use of on-board optics (OBOs) is another improvement which allows more efficient implementation of high-density optical communications systems. The COBO specification provides a common and standards-based guide for component manufacturers, coherent module manufacturers, system integrators and end users.

This document describes some examples of how Digital Coherent Optics (DCO) and systems can be implemented with COBO-compliant OBOs. It includes examples of placement and fiber routing for host board and module. It presents results of thermal modelling based on several host board placements use case. While other types of OBOs do exist in the market, all mention of OBOs in this document are assumed to be COBO-compliant OBOs, also sometimes called COBO modules.

This application note is one in a series of documents published by COBO to supplement the specification. Further work on COBO thermal experimental measurements and electrical signal integrity using 56G-VSR-PAM4 and considerations for future 112G-VSR-PAM4 interfaces developed by the OIF (Optical Internetworking Forum) will be published as a white paper [2]. More details on considerations for fiber and connector choices can be found in another white paper [3].

It should be noted that the descriptions in this Application Note are intended as examples of applications and possible implementations. In no case are they intended to be prescriptive. While the COBO specifications includes features were crafted using the DCI use case as a guide and therefore may be particularly useful for coherent DCO functionality, they are also relevant to telecom transport and other applications. This document also calls out additional flexibility in the COBO specification that could be useful for future use cases. This document is not intended as an exhaustive description.

For compliance to the standard, please refer to the specification document [1].

2. THE RELEVANCE OF COBO IN COHERENT APPLICATIONS

2.1 - Definition of the 400ZR Use Case and Its Requirements

One target application for the COBO specification is DCI between datacenters. In this application, there are low-latency point-to-point links between multiple sites. As shown in Figure 2-1, regional network gateways (RNG) serve as hubs interconnecting the datacenters (DC) to each other and to the longer-distance backbone network. Depending on the specific network design, the regional gateways may also serve as datacenters, or the datacenters may directly connect to each other instead of through the hubs. Typical reaches are about 70 km from each datacenter to the hubs.

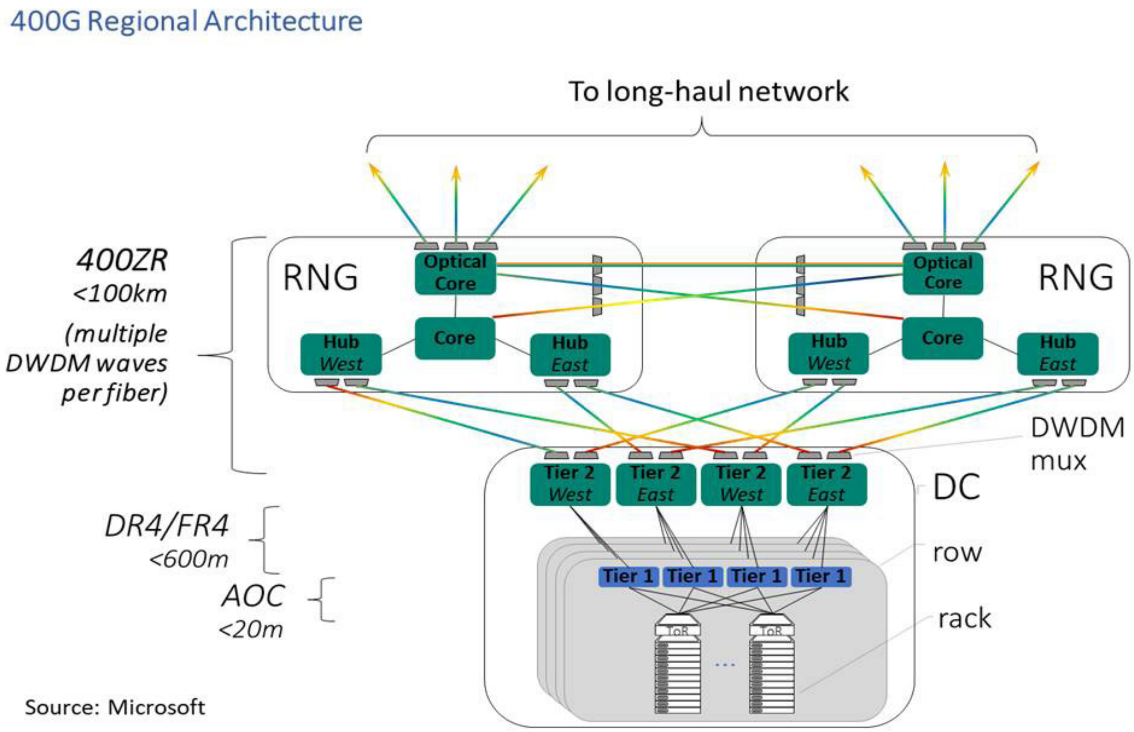

Source: Microsoft

FIGURE 2.1: DATACENTER NETWORK INTERCONNECT EXAMPLE

The key applications that helped inform development of the COBO specification included 400ZR, an interoperable optical standard being developed by the OIF (www.oiforum.com). The IEEE P802.3ct Task Force is developing a 400GBASE-ZR specification with anticipated alignment to the OIF 400ZR specification. Both 400ZR and 400GBASE-ZR propose to support approximately 400 Gbps per wavelength using coherent detection. The terminal optics specification will apply to DWDM applications over an unamplified link or amplified DWDM networks with amplifiers for reaches up to approximatively 120 km. This application is a natural fit for the COBO specification, which is defined for applications of 400 Gbps and beyond.

The target application is the links between the switches within the RNG and the DC shown in Figure 2-1. Each switch has grey optics links facing downwards into the datacenter and DCI links facing out from the building. Multiple 400ZR links coming from each switch (not shown explicitly in the figure) are multiplexed onto the DWDM links shown as colored lines. 400ZR links can also interconnect the RNG directly to another RNG as shown in the figure. Links within the RNG are not expected to use 400ZR. DWDM uplinks labelled in the figure as heading into the long-haul network are not in the current scope of the ZR example described in this document.

Individual network designs will dictate their own mix of uplinks and downlinks. A given system design can efficiently support the required variety of customer configurations through a flexible design in which DCI optics fit in the same sockets as grey client optics. This implies a very challenging requirement that the coherent module have the same density despite greater complexity of the internal optical components and higher power consumption for signal processing. The wavelength-tunable coherent module requires more management information to be exchanged with the host than grey client optics.

2.2 - Other Applications

Telecom applications could also benefit from COBO-compliant implementations. While COBO, OSFP and QSFP-DD are all form-factors that intend to satisfy the requirements for both grey and coherent versions, COBO has distinct advantages for thermal capacities and optional features support as discussed in Sections 3.4 and 5.0.

3. HOST SYSTEM DESIGN CONSIDERATIONS

3.1 Requirements and Features

3.1.1 Environmental, Power and Size Targets for a Reference Application

The use case described in this document is some considerations for implementation of a 1 RU design 12.8 Tbps data center switch with a design that supports a flexible mix of COBO modules. The application and the example implementations here are not intended to be prescriptive or exclusive. The configuration detailed in this document was chosen as the worst-case from a module placement and thermal perspective where the system is fully populated with COBO-compliant modules for 400ZR application and constraints.

This Application Note describes examples that meet the application requirements in terms of physical density, thermal density and power consumption. A set of typical system parameters is shown in Table 3-1. This is only an example: other system designs will have a different set of constraints and solutions. Substituting a different mix of media types is straightforward given that the COBO specified electrical connectors are common across media types. The smaller Class A and Class B COBO form-factors can be substituted without changing the board layout or connectors.

Parameter	Value
Environmental Ambient Max Temp	45˚C
Airflow Direction	Front to Back
Power Consumption for 400G OBO	15W
Number of COBO Modules	32
System Capacity	12.8 Tbps
System Size	1RU in 19" Rack
Airflow	~165 CFM

TABLE 3-1: EXAMPLE HOST SYSTEM PARAMETERS FOR REFERENCE APPLICATION

3.2 Module Placement on Host Board

The COBO specification defines two types of OBOs: a narrower single-wide size (also called 400G OBO) that has eight duplex electrical lanes and supports approximately 400 Gbps and a double-wide size (also called 800G OBO) with sixteen duplex electrical lanes and that supports approximately 800 Gbps. As depicted conceptually In Figure 3-1, the host board connects the switching ASIC to the OBOs. In the specific example, a 12.8 Tbps switching ASIC is connected to thirty-two 400G OBOs each via high-speed electrical interface compliant to 400GAUI-8 C2M specification. Alternatively, a total of sixteen double-wide OBOs can be used to support the same 12.8 Tbps ASIC.

OBOs. In the specific example, a 12.8 Tbps switching ASIC is connected to thirty-two 400G OBOs each via high-speed electrical interface compliant to 400GAUI-8 C2M specification. Alternatively, a total of sixteen double-wide OBOs can be used to support the same 12.8 Tbps ASIC. As detailed in Figure 3-2, the board also supplies the modules with power and a two-wire low-speed management interface. Since this is a DCO application, the coherent DSP as well as all the opto-electronics are internal to the OBO.

3.2.1 Host Board Placement Considerations

The design example here assumes a single 200 W Ethernet switch ASIC. (Note that some systems may have multiple switching ASICs.) A reasonable design approach is to locate the switch at towards the back of the chassis to allow temperature-sensitive optics to have first access to the front-to-back cooler air.

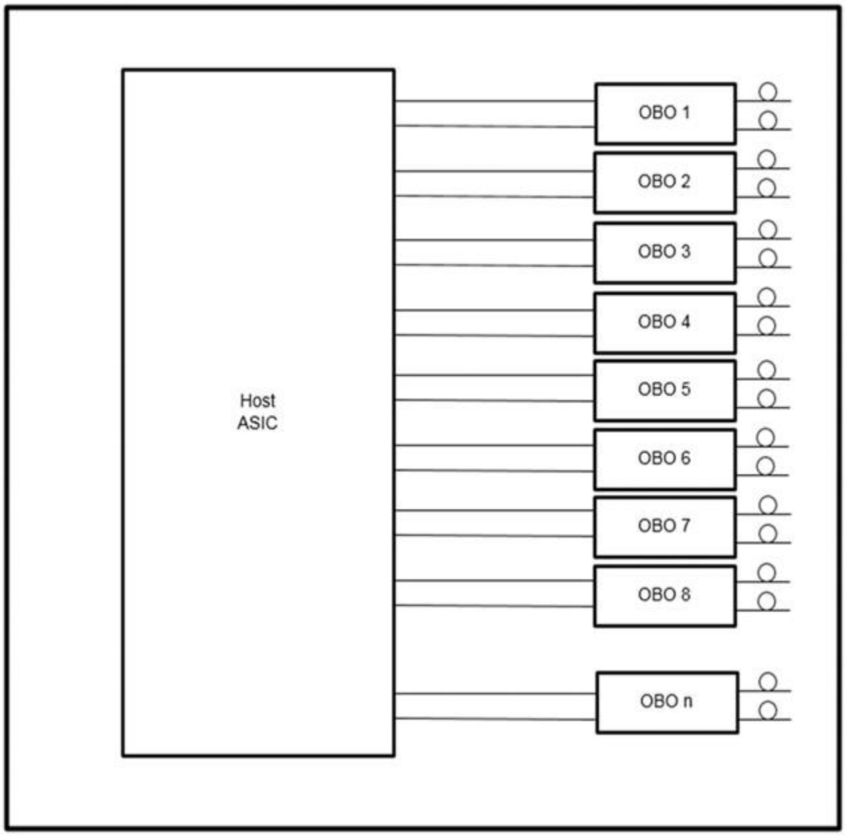

FIGURE 3-1: HOST BOARD APPLICATION CONCEPT

DCO coherent optics implementations require more components (e.g. DSP, polarization splitting and tunable laser control) than either direct-detect optics or analog coherent optics hence requires more module area. Class C, the largest of the three modules sizes allowed under the COBO specification was developed with support of DCO optics in mind. The Class C size of 20 mm x 60 mm was chosen to balance the need to provide as much balance the need to provide as much board space internal to the module yet allow at least thirty-two modules to fit on a 1RU system board.

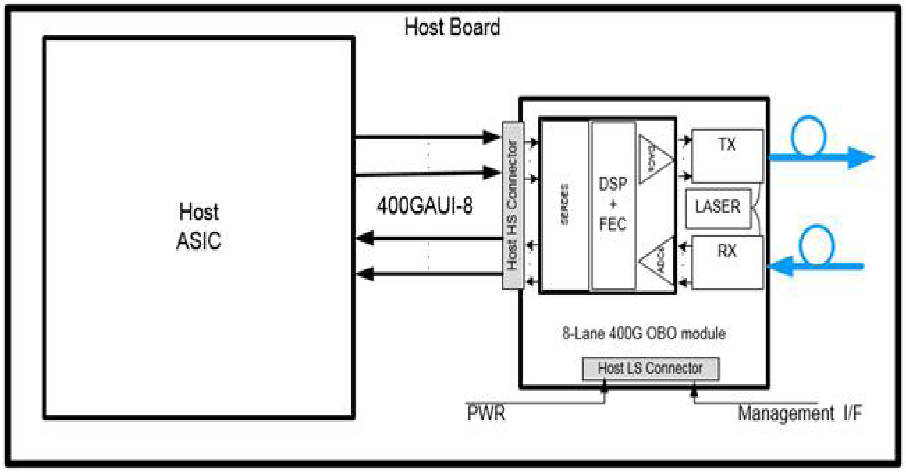

8-LANE COHERENT 400G ON-BOARD OPTICS

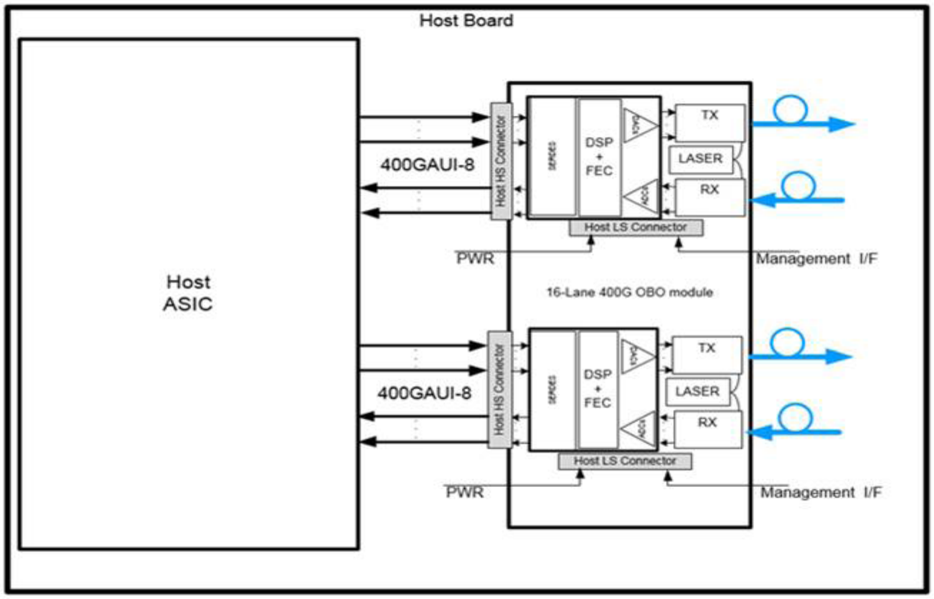

16-LANE COHERENT 800G ON-BOARD OPTICS

In the case of 400ZR or likely any longer-reach application, there will be two fibers (one for each direction) exiting an eight-lane 400G OBO or four fibers exiting a sixteen-lane 800G OBO. These would typically terminate in dual LC (or two dual LC or one CS for the 800G OBO) connectors on the front panel. While the COBO specification also allows for modules to be connectorized at the module either using a connector or pigtail, the design example uses the pigtailed option.

A Class C OBO including keep-out for the high-speed connector takes up a length of 68.3 mm (Figure 4-3 of the COBO specification). When considering how densely to pack the modules on the board, an additional length allocation is needed for the fiber exit with reasonable bend radius. For placement purposes, each module was allocated 77 mm as shown in Figure 3 3.

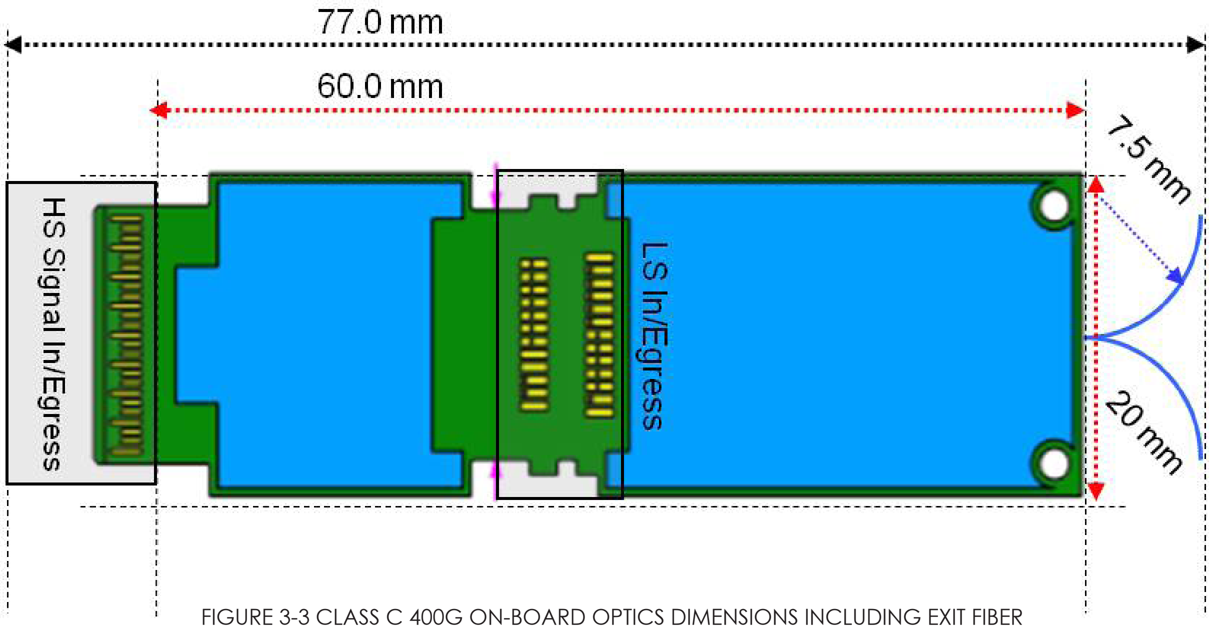

FIGURE 3-3 CLASS C 400G ON-BOARD OPTICS DIMENSIONS INCLUDING EXIT FIBER

This is consistent with the use of 7.5 mm bend radius for the fibers, a tighter bend than commonly seen in host applications but extensively used inside high-density modules. A system reliability measured in failures in time (FIT) for all 64 fibers at this bend radius is estimated to be less than 0.6 and can be further lowered to 0.02. Further details on this FIT estimate is given in Section 7.0.

3.2.2 Host Board Placement Considerations

Figure 3-4 shows a few host board placement options. At the top of the figure, corresponding to the back of the chassis, the fan assembly is centrally located with the power supplies and their separate airflow ducts on either side. The switch ASIC is a square seen within the wide light gray rectangle that depicts its overhanging heat sink. Two rows of Class C OBOs are oriented such that the fiber pigtails (not shown) face the other row with a total of 2 x 7.5 mm allowed between rows. In practice, it would be possible to have these fibers with 7.5 mm bend radius overlap and thus reduce the row spacing. Finally, the front panel connector housings can be seen at the bottom.

The module placement design on the left has three rows totaling thirty-two single-width 400G OBOs. The design on the right has two rows totaling sixteen double-width 800G OBOs. There is not much difference in the module placement and electrical layout for the host board and none for the front panel between double and single-width modules. Note that there may be more advantage for OBO implementation since a double-wide module allows more design flexibility for fiber routing internal to the module. Other placement options are discussed in Section 3.4 where their thermal performances are compared.

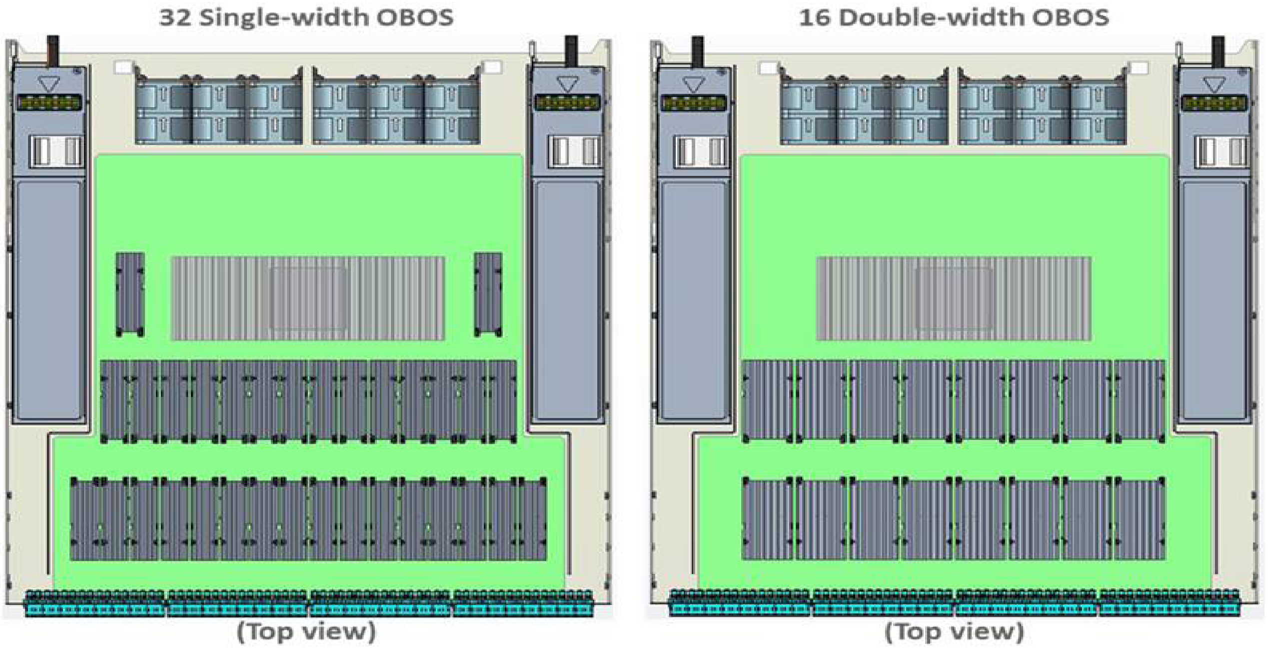

FIGURE 3-3 CLASS C 400G ON-BOARD OPTICS DIMENSIONS INCLUDING EXIT FIBER

3.3 Additional Considerations

3.3.1 Electrical Connectors on Host Board

There is no special accommodation for the high-speed or low-speed connectors required for COBO-compliant modules for 400ZR application with power draw up to 15 W. Higher power modules can also be supported. The low speed COBO connector is designed to supply up to 20 W with a method to increase to 29 W through use of the "for future use" pins.

3.3.2 Retention of Modules onto Host Board

The Class C OBO is longer than Class A and B OBOs. The designer should take care to support the longer module adequately. In addition to latching the module to either the high-speed connector at the rear or the low-speed connector located in the middle of the module, the Class C specification shows an extra mounting feature on either side at the front (fiber exit face) to provide additional retention. This feature is shown as a screw labelled "Class C supplemental screw" on the left side of Figure 3 5 and Figure 3 6. Note that the additional two screws are external to the module body.

The COBO module board is elevated from the host PCB due to the design of the connectors allowing surface layer routing and low-profile component placement. Hence, for Class C it is recommended to use a stand-off as shown under the left side of the module in Figure 3-5 and Figure 3-6. Not visible in each diagram is the screw and standoff for the other side of the module.

3.3.3 Assembly and Latching of OBOs on Host Board

The specification provides implementers the choice of latching to either the high-speed connector (at the right of module as shown in Figure 3-5) by means of a latch-arm or the low-speed electrical connector by means of a clip (in the middle of the module in Figure 3-6).

The recommended sequence to install modules and latch them is shown for the case using a high-speed connector latch in Figure 3-5 and for the case using a low-speed connector latch in Figure 3-6. Do not use both the high-speed connector latch and the low speed connector latch simultaneously, only use one or the other.

The COBO connector is designed to facilitate placement of modules on densely populated boards. Each module should be inserted straight down onto the host board. The low speed connector has features that allow the module to be mechanically engaged before electrical contacts are made. This allows the correct seating to be verified. Then the module is slid horizontally towards the high-speed connector. This action makes electrical contact to both the high-speed edge connector and the low-speed connector. Then the latch should be engaged, followed by installation of the supplemental retention screws.

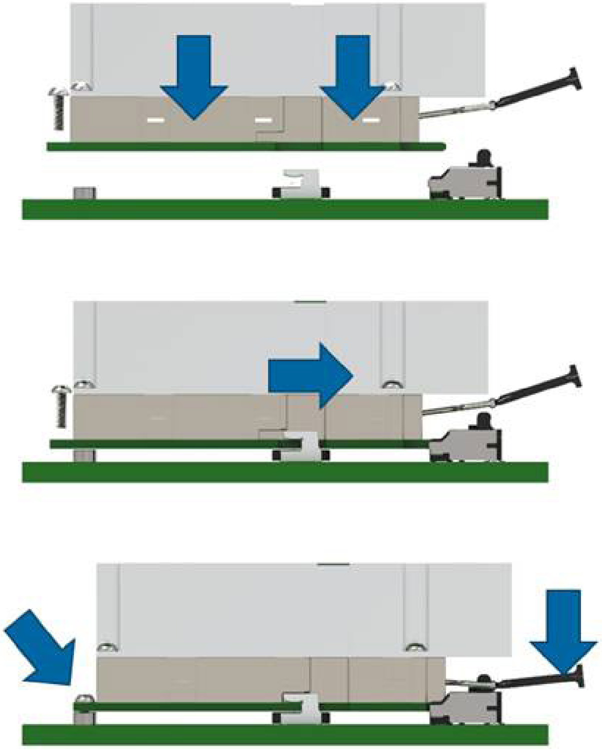

Vertically lower the COBO module onto high speed and low speed connectors

Horizontally slide the COBO module into the high speed connector

Rotate the high speed connector spring latch arm onto the high speed connector hooks. Install Class C supplemental screws

FIGURE 3-5 MODULE INSTALL AND LATCHING WITH HIGH SPEED CONNECTOR LATCH

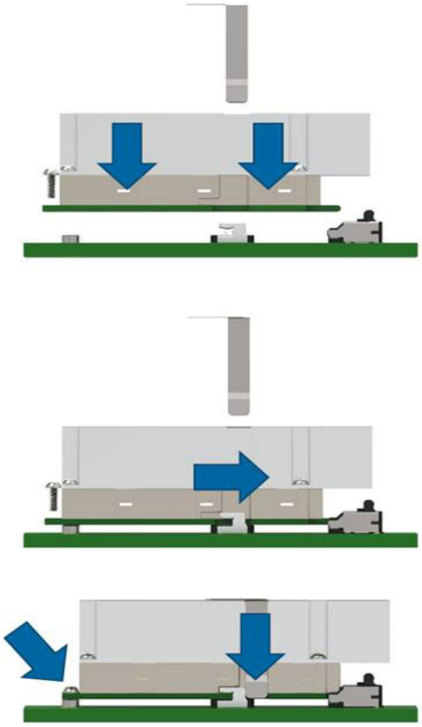

Vertically lower the COBO module onto high speed and low speed connectors

Horizontally slide the COBO module into the high speed connector

Insert the low speed connector U-clip latch piece into the module and low speed connector channel. Install Class C supplemental screws

FIGURE 3-6 MODULE INSTALL AND LATCHING WITH LOW SPEED CONNECTOR LATCH

3.3.4 Attachment of Heat Sinks

The specification does not describe methods for heat sinking. Common methods such as screws or clips are shown for eight-lane examples in Figure 3-7 below. Similar methods would apply to sixteen-lane modules as well.

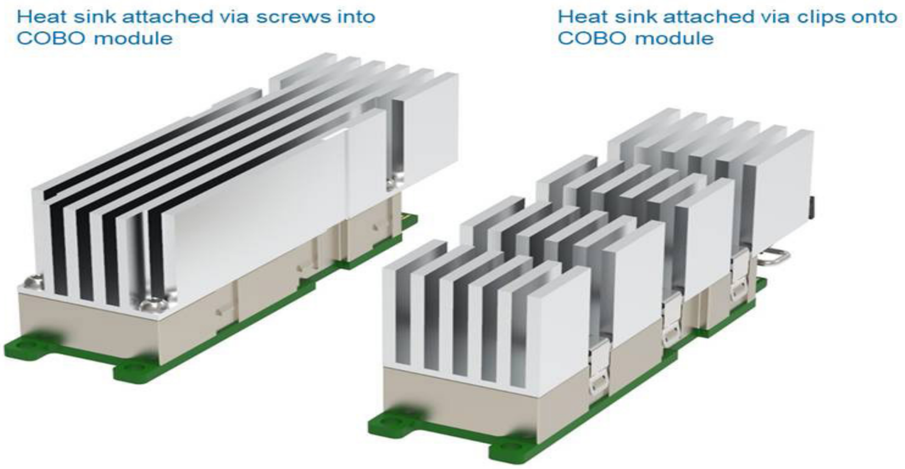

FIGURE 3-7 EXAMPLE HEAT SINK ATTACHMENT OPTIONS FOR CLASS C x8 MODULES

3.4 Modelled Thermal Capabilities for 1RU DCO Application

This section presents some modelling results for the use case of a 1 RU chassis cooled using front-to-back airflow, populated with thirty-two eight-lane OBOs. The use case targets 15 W per module. The results show that a simple two row design for 15 W modules keeps the worst-case temperature located on top of DSP to slightly above 70 °C. Moving the two worst-case modules to a third row brought all modules under the target 70 °C. Further optimizations are possible but were not explored.

The thermal study also did a sensitivity analysis of the impact of raising the power consumption per module to 17.5 W and a study of the impact of decreasing the heat sink height.

The study used these assumptions about the thermal load, module sensitivity and system sensitivity to demonstrate feasibility based on typical values. The model was purposely kept simple, as details will differ in actual designs of both modules and systems. The thermal design shown here was not optimized. System designs are expected to have different component placement, cooling mechanism and airflow constraints as determined by the equipment manufacturer.

3.4.1 Thermal Model of COBO-Compliant Coherent OBO

The simplified reference module is an eight-lane 20 mm x 60 mm Class C OBO with two critical thermal blocks as shown in Figure 3-8. These blocks represent the key heat generating, and temperature sensitive, electronic and electro-optical components. For a DCO module rated for a maximum 15 W total power consumption, the coherent DSP is assumed to consume 7 W while the optoelectronics consume 8 W.

The optoelectronic components assumptions include the tunable laser, modulator driver, modulator, photo-detectors and TIA amplifiers. In practice, the tunable laser may be located further away from the DSP. For a module with 17.5 W total power, the two blocks increase to 8.5 W and 9 W respectively.

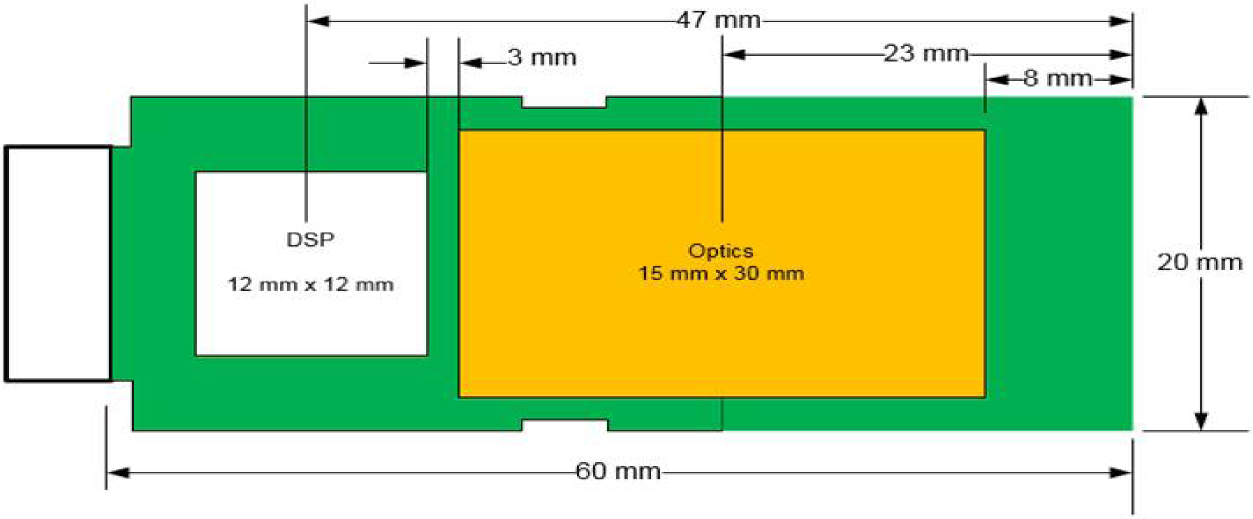

.

FIGURE 3-8 COBO COMPLIANT COHERENT DCO MODULE THERMAL MODEL

The case temperature location is not defined in the specification; in practice the location of the thermal sensor that defines case temperature varies by module design.

Optical components are typically developed for use in modules with case temperatures up to 70 °C. While the closely associated electronics such as driver may be able to tolerate a higher temperature, this modelling sought to keep the entire opto-electronic block to a case temperature of 70 °C. The DSP is typically able to withstand a higher operating temperature than optoelectronics.

3.4.2 Heatsink and Airflow Assumptions

The model assumes individual aluminum heat sinks with vertical fins, the same lateral size for all modules and all cases. The model assumes all heat generated in the coherent OBOs are dissipated via heat sinks; no heat is assumed to exit downwards to the host printed circuit board (PCB). Modelling results used a module height of 4.5 mm and a heat sink height of 22 mm. The thermal study also looked at the impact of decreasing the heat sink height by 1.5 mm to 20.5 mm. The shorter heat sink resulted in a rise of 1 °C.

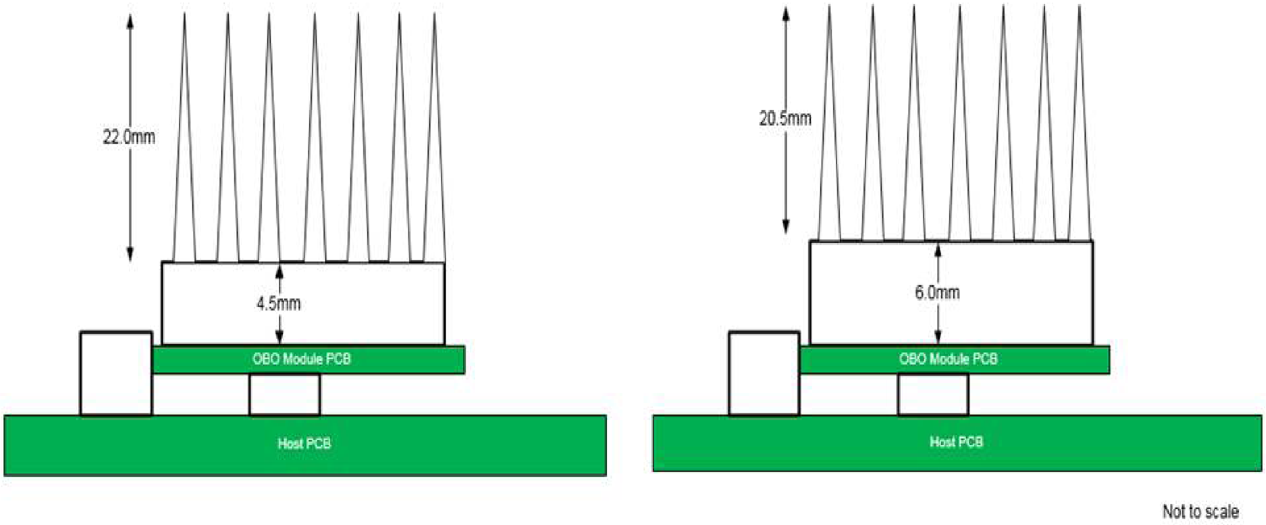

FIGURE 3-9 A TALLER MODULE CONSTRAINS THE HEIGHT OF THE HEAT SINK
Note that a board-mounted optic may sit higher in the chassis than a front-panel pluggable

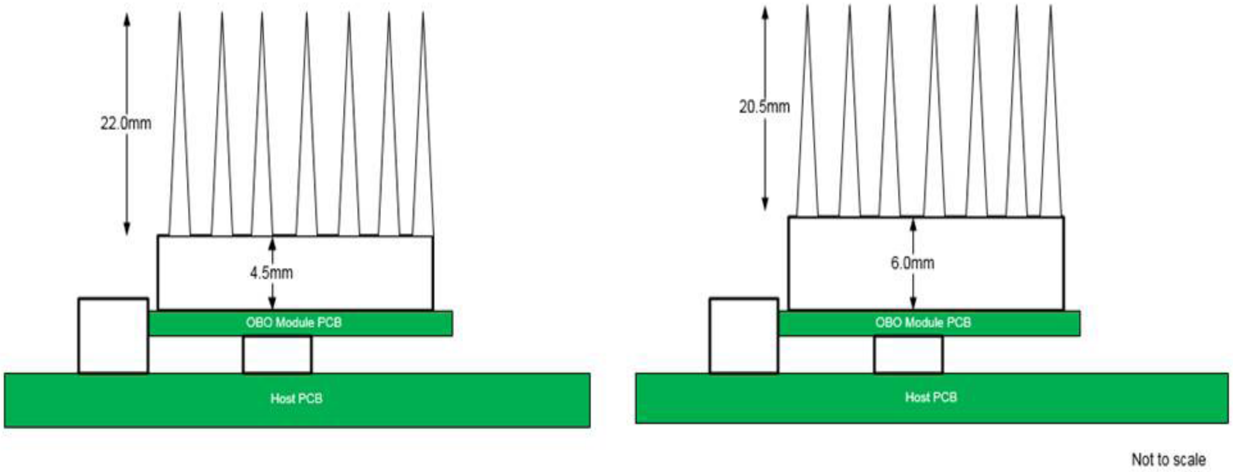

FIGURE 3-9-2. TRADE-OFF BETWEEN HEAT SINK AND MODULE HEIGHT

A further optimization available for use in a system design is to use shorter heat sinks for the front row of modules, improving the cooling for the second row in a front-to-back airflow to offset the effect of pre-heated air reaching the second row. The shorter front row fins provide less cooling capacity but the inlet air temperature here is lower. This option was not modelled. Given that the heat sinks are constrained in height, a solution which seeks to optimize based on multiple fin heights likely results in an equal module temperature for both rows.

The system is cooled with front-to-back airflow with a flow rate of 165 CFM. The fibers from the OBO modules to the front panel were not modelled. It is assumed that these are low enough and routed so as not to obstruct airflow significantly.

Figure 3-10 shows the front panel detail with LC connectors. Note that the air intake is above the connectors such that airflow is not blocked by the connectors and leads directly to the fins of the heat sinks.

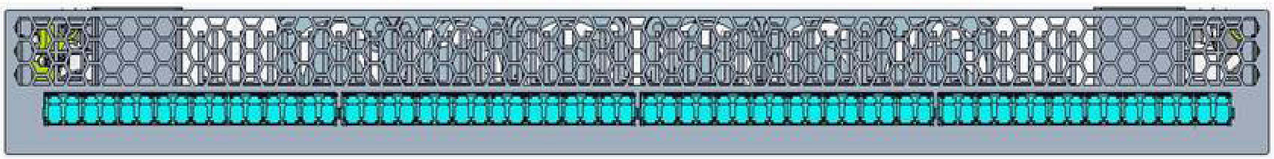

FIGURE 3-10 CHASSIS FRONT PANEL SHOWING C CONNECTORS BELOW AIR INTAKE

3.4.3 Thermal Modelling Results

In this model, the switch ASIC is assumed to dissipate 200 W. The modelling shows the ASIC case temperature only reached the mid-60 °C which leaves headroom for it to operate hotter. The ambient environment is modelled assuming sea level 45 °C air intake.

As seen previously in Figure 3-4, the AC-DC or DC-DC converters have a separate airflow system independent of the main system and were excluded from the simulations. In the thermal modelling results below, the power supplies air intakes are in the blue areas on either side of the chassis; the wider area in the back corresponds to the actual power supplies.

The analysis first looked at a simple two row placement design of thirty-two single-width modules. Results for 15 W and 17.5 W modules are shown Figure 3-11 and Figure 3-12 respectively.

As expected, the front row of modules is cooler than the back row. The worst locations are on the edges of the back row where the corners resulting from the power supply placement creating turbulence. For a 17.5 W system, these worst-case placements exceeded 70 °C right above the DSP.

An alternative three-row placement moves these modules back next to the ASIC. This is the preferred placement described previously in Section 3.2.2. Despite the location next to the ASIC, the results in Figure 3-13 and Figure 3 14 show a 2 °C decrease in the maximum module case temperature compared to the two row configurations in Figure 3-11 and Figure 3-12.

Further design optimization is possible but was not pursued for this study. Figure 3-15 shows the airflow for the three-row configuration. The high airflow shown in orange on either side of the third-row OBOs could be used more effectively if directed to module heatsinks. In Figure 3-14, the ASIC is well below its temperature limit. One possibility is to narrow the ASIC heatsink and move two OBOs from the second row to the third row to better utilize this airflow.

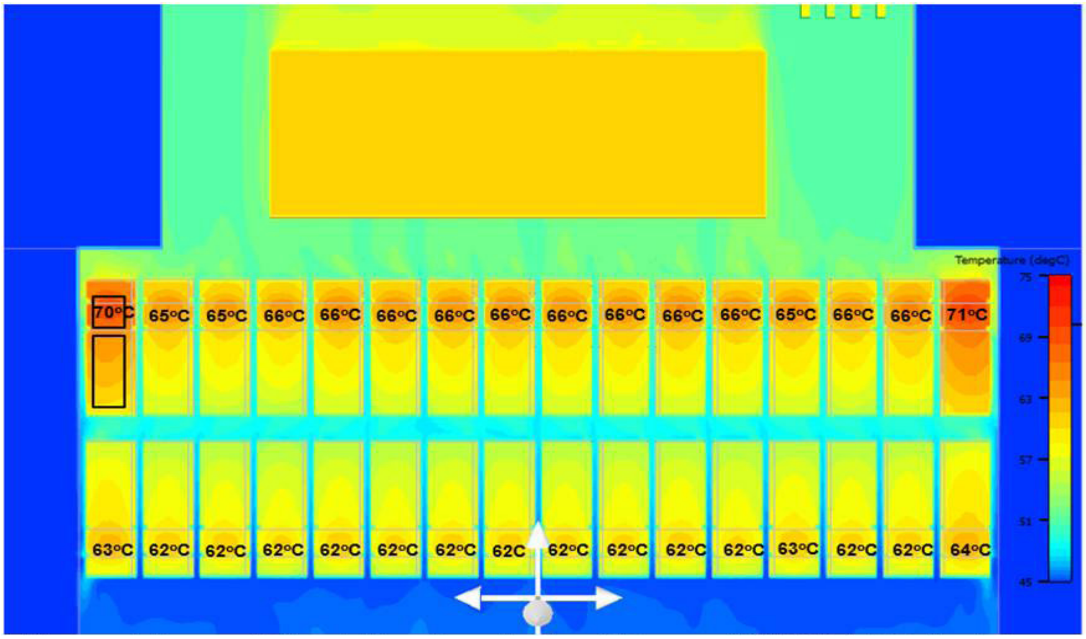

FIGURE 3-11 THERMAL SIMULATION RESULTS FOR 2 ROWS OF 15 W COBO MODULES

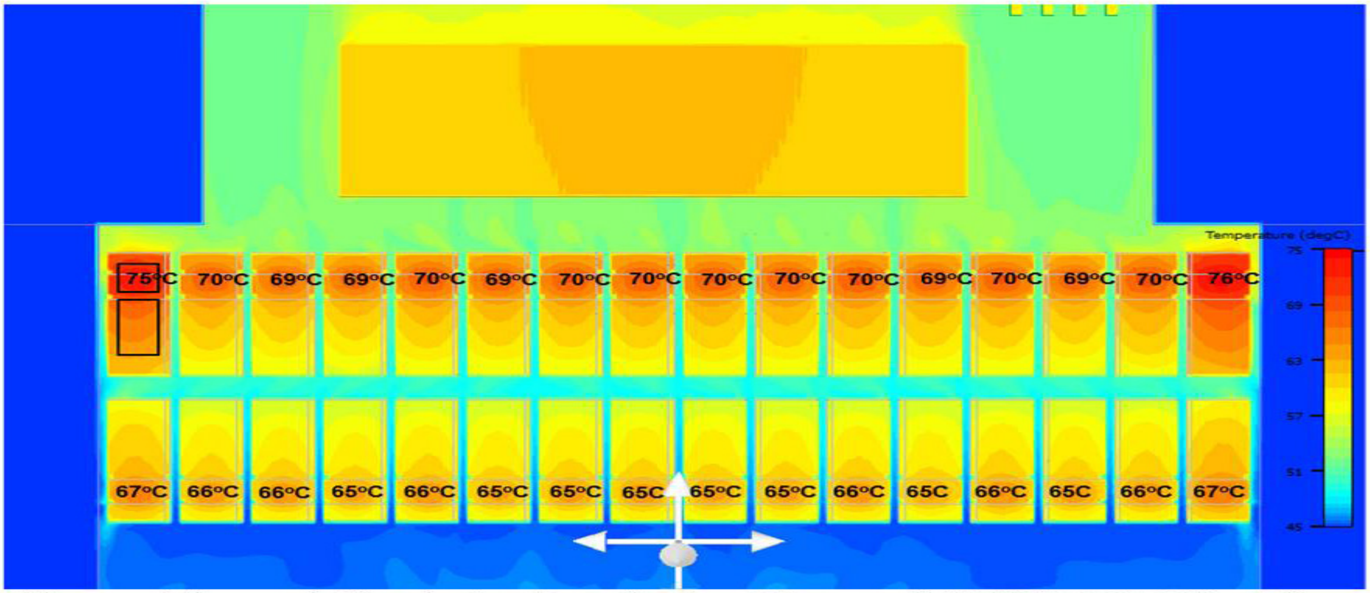

FIGURE 3-12 THERMAL SIMULATION RESULTS FOR 2 ROWS OF 17.5 W COBO MODULES

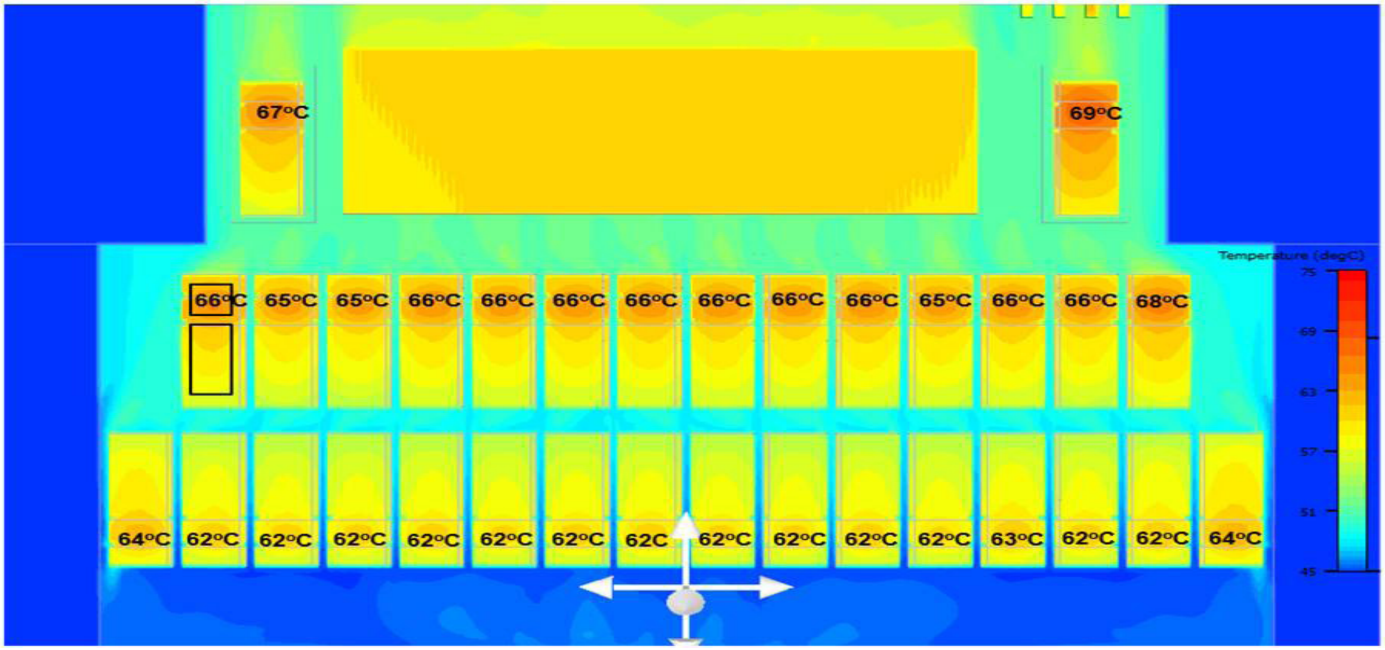

FIGURE 3-13 THERMAL SIMULATION RESULTS FOR 3 ROWS OF 15 W COBO MODULES

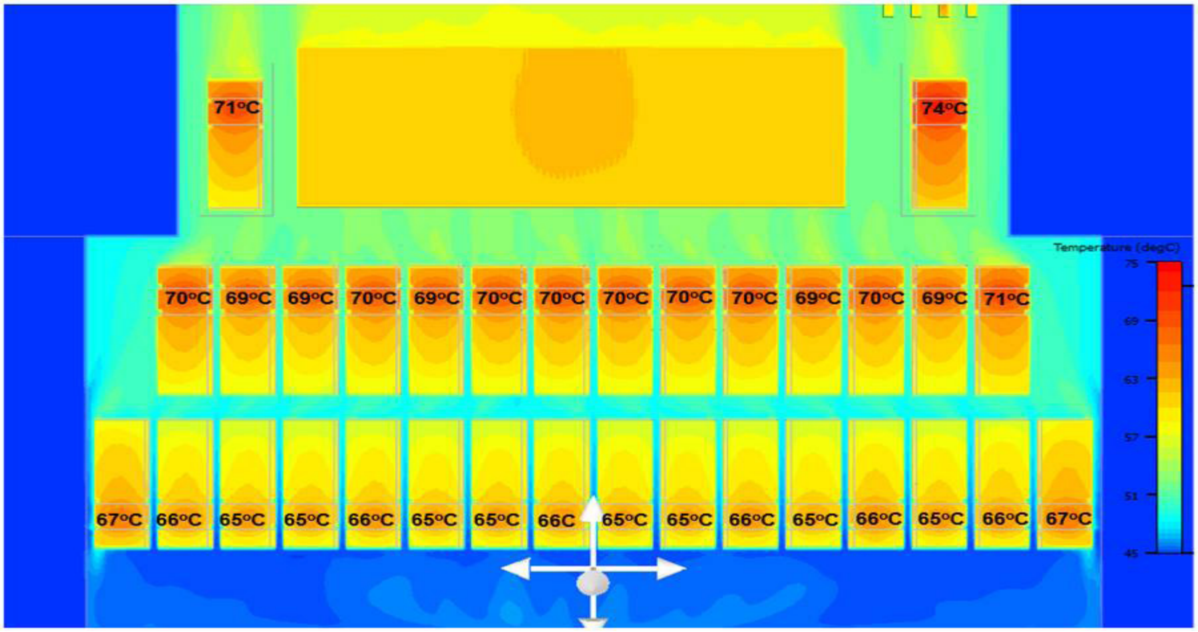

FIGURE 3-14 THERMAL SIMULATION RESULTS FOR 2 ROWS OF 17.5 W COBO MODULES

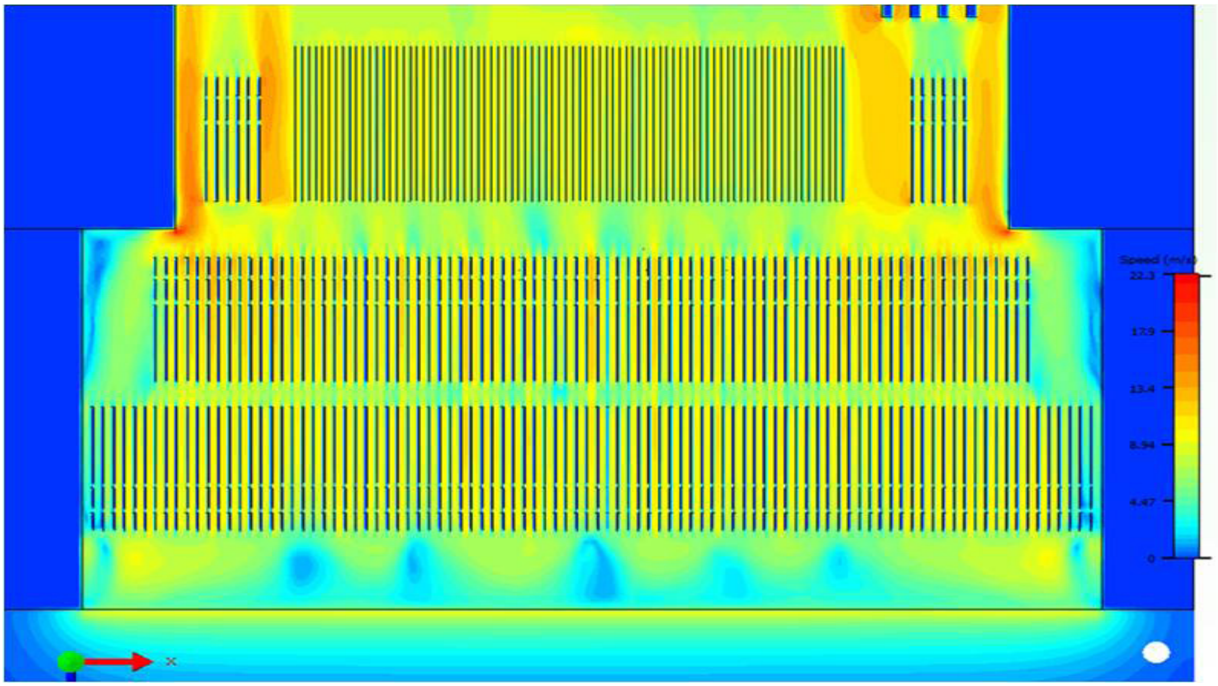

FIGURE 3-15 AIRFLOW FOR A 3-ROW CONFIGURATION

4. COHERENT DCO COBO MODULE DESIGN

4.1 Class C COBO Specification is Intended for Coherent Applications

The 400G DCO module example described below contains a laser, a PM-Q modulator with drivers, and an intradyne coherent receiver as well as associated control circuitry, power converters necessary to support the optical components and a microcontroller for management. This section describes how these can be made to fit into the eight-lane COBO module form-factor with external dimensions of 60 mm x 20 mm. The sixteen-lane double-wide module is not discussed: it has twice the component content, twice the width and is expected to have slightly greater placement flexibility.

There exist different levels of component integration into subassemblies that go into the module. The following section will describe examples of how a variety of integrated components could be used in a coherent COBO implementation.

The COBO specification was developed in the expectation that multiple implementation technologies will co-exist and continue to advance. Class C, the largest COBO module size specified, was chosen to support current silicon photonics and InP implementations. If denser components become available over time, the specification does not prohibit implementing DCO modules in COBO Class B or Class A form-factors as well. Due to the identical electrical connectors, these denser modules could be used with system boards initially designed for Class C modules.

Implementation examples below assume a coherent DSP intended for DCO modules that target DCI performance as implemented in 7nm CMOS. While not strictly required, it is probable that many of the coherent DCO implementations will be wavelength tunable. The inclusion of a tunable laser and its associated control circuitry adds to the design challenge. This section describes some placement options that demonstrate feasibility for the required density.

4.2 Example Coherent Module Implementations|

4.2.1 Fully Integrated Coherent Optical Subassembly

This approach starts with a fully integrated optical subassembly, making it easier for the module designer to achieve very high density in the module. The fully integrated subassembly includes all the optics and their associated high-speed electronics as well as circuitry to control the operating conditions of the optics.

Essentially the density challenge is pushed down to the subassembly maker. The only other major blocks that the module designer needs to add are the coherent DSP chip and the module microcontroller. In the most integrated version, these two digital chips might be further combined into one if the DSP ASIC has enough extra processing headroom to handle the management firmware for the DCO application.

Figure 4-1 shows an exploded view of how a module using such an integrated subassembly example. Not shown are ancillary electronics such as power supplies. For signal integrity performance, the DSP is immediately next to the high-speed electrical connector on the left end side and attaches to the integrated optical subassembly.

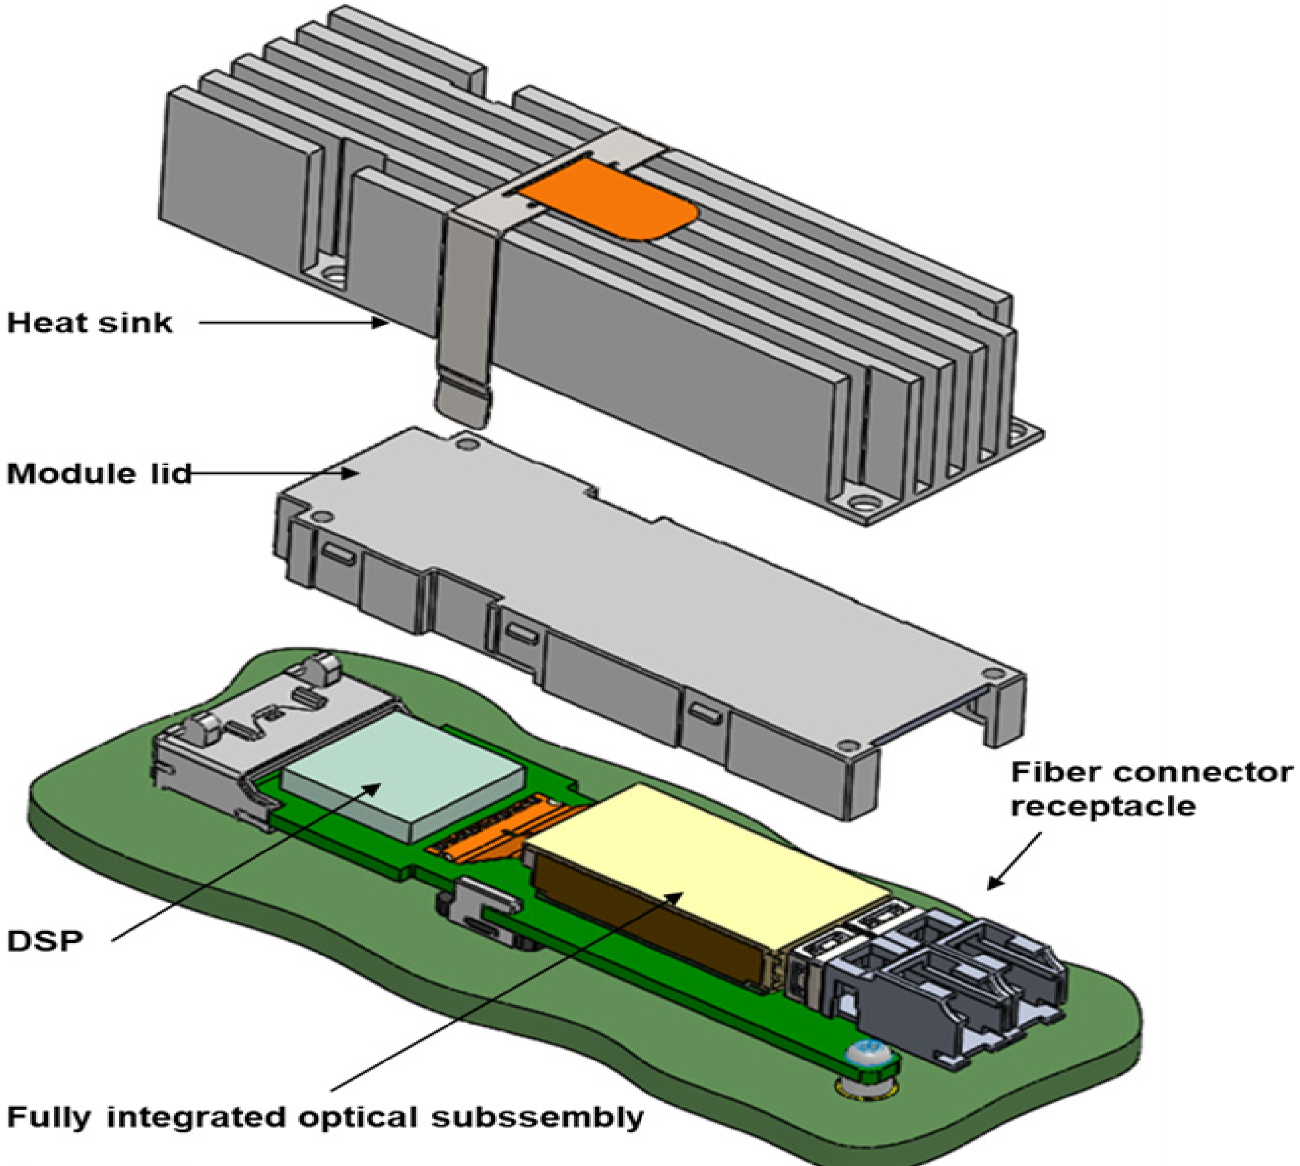

Source: Finisar

FIGURE 4-1 EXAMPLE MODULE DESIGN USING FULLY INTEGRATED OPTICAL SUBASSEMBLY

4.2.2 Integrated High-speed Coherent Optics with External Laser

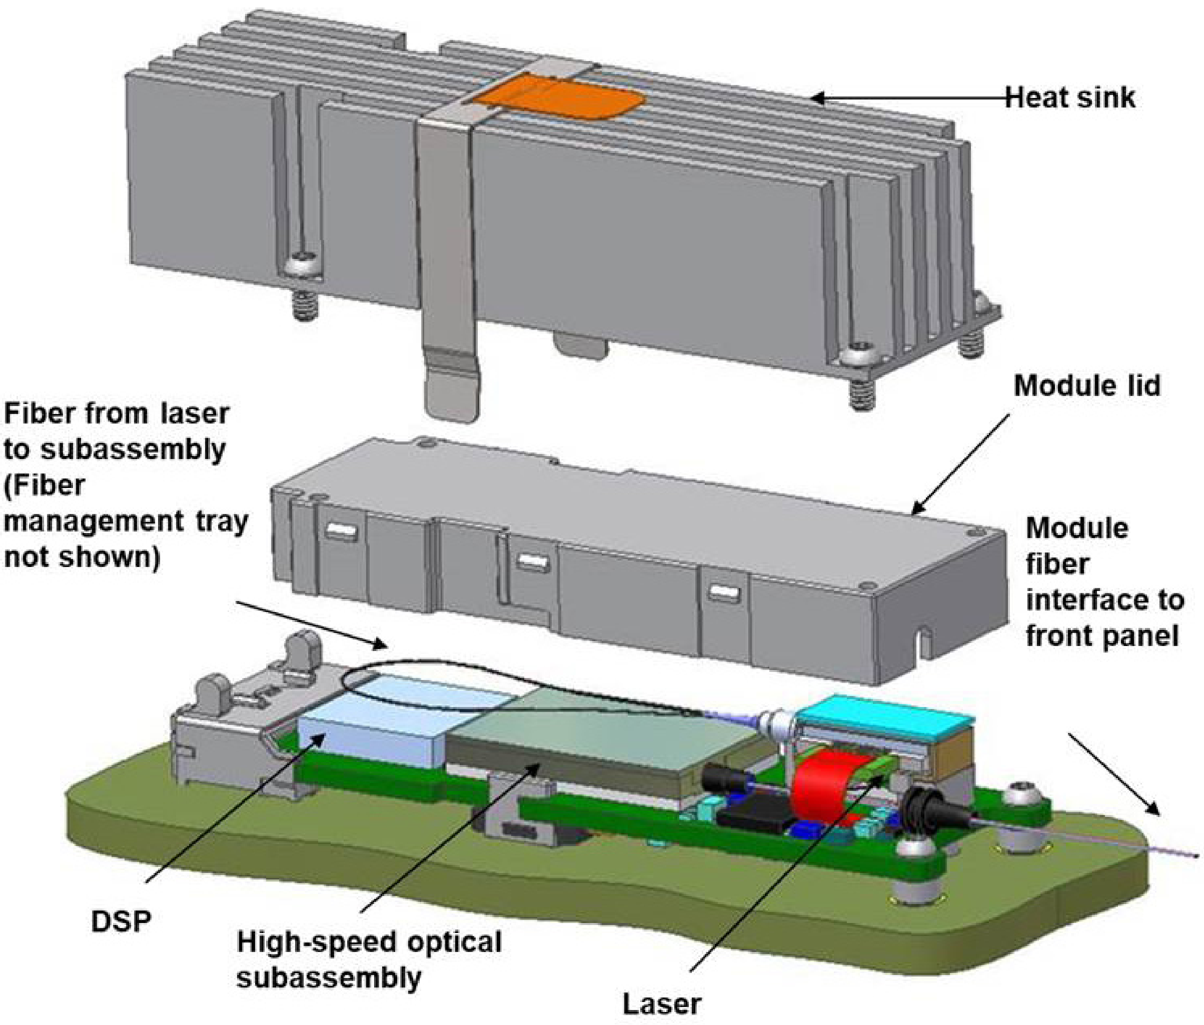

Source: ADVA

FIGURE 4-2 EXAMPLE MODULE DESIGN WITH HIGH SPEED OPTICAL SUBASSEMBLY AND EXTERNAL LASER

Another possible packaging style for the opto-electronics is to integrate all the high-speed components into one sub-assembly while the tunable laser and its control electronics are a separate subassembly. Note that the laser is continuous-wave and it therefore its placement constraints are very different. This approach may make sense if the module designer already has a preferred tunable laser or if the high-speed optoelectronics are implemented in silicon photonics chips without lasers.

This approach has the benefit of giving the module designer some flexibility to place the laser far away from the hot DSP but at the cost of having to find space for the laser package and route an additional fiber between the laser and the high-speed optoelectronic subassembly. Figure 4-2 shows an exploded view of how such a design might be done. The top of the laser package contacts the module lid for heatsinking. Note that in practice there would be a fiber tray over the DSP and high-speed optical subassembly to provide thermal contact to the module lid. This has not been shown in the figure.

4.2.3 *Fiber Management Inside a Coherent COBO-compliant OBO*

Care must be taken with the fiber routing as well as the component placement. While the module width accommodates 7.5 mm radius bend, fiber that supports bends down to 5 mm can be used. (See Section 7.0 for reliability calculations). Strain relief elements are shown both at the high-speed subassembly exit and the COBO module exit.

In this example of fiber routing, the fiber between the external laser and the high-speed subassembly exits the laser to the left, is looped back and under the laser before entering the optical subassembly. Figure 4-3 shows the two ends of this loop. Because the laser requires no high-speed signals, it can be raised from the PCB. Electrical contact is made via the red flex circuit. Note that care needs to be taken for the fiber routing design. The loop is controlled by a fiber tray, which is omitted from the drawing for clarity.

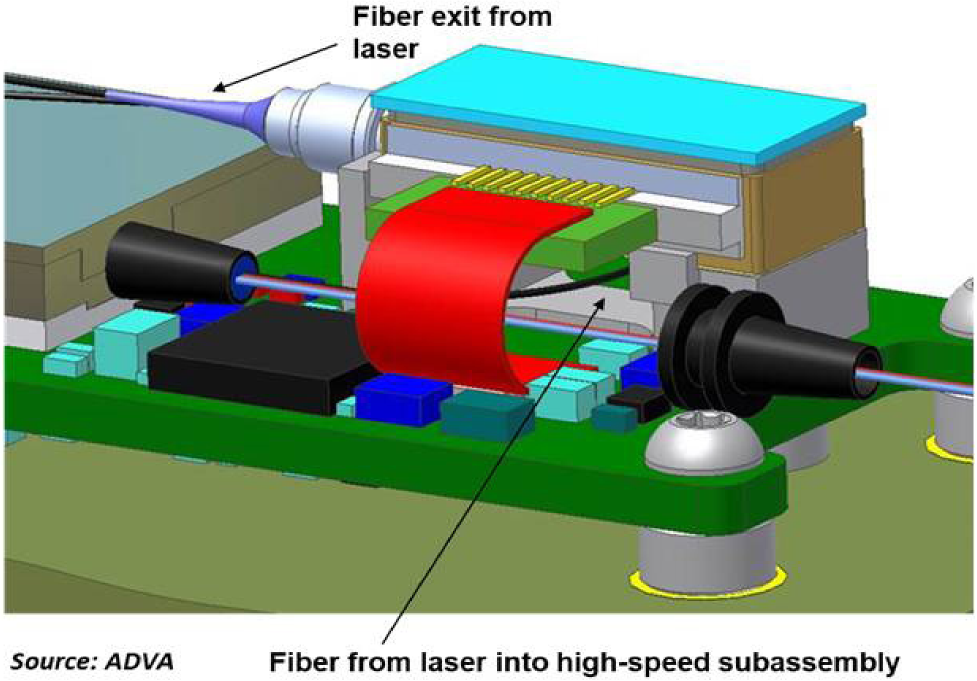

Source: ADVA

FIGURE 4-3 DETAIL OF FIBER MANAGEMENT FOR EXTERNAL LASER

5. COBO SUPPORT FOR ADDITIONAL FUNCTIONS

Although the specification was developed with Ethernet datacenter networking and inter-datacenter interconnect applications as guidance, the COBO form-factor is intended to be a general platform. The mechanical specification includes hardware pins that are not found in other high-density modules. The functions of these extra pins are not defined in the specification but intended to provide additional flexibility and future-proofing. Below are some examples of how these pins can support additional hardware functions as needed by some applications. Note that the examples below cannot be all done simultaneously; they represent choices available to the module designer.

5.1 Sync for Coherent Using COBO-Compliant Modules

Ethernet extension applications including 400ZR and FlexO/OTN can be used in networks that require timing distribution through the implementation of SyncE or SyncO (e.g. using the FlexO/OTN server layer). A network element such as a switch will receive at least two master clocks from higher in the network and likely distribute the clock to additional slave network elements. In addition, in the case of failure of the incoming clocks for some reason, it should have the ability to generate its own local reference clock and continue to operate in holdover mode. The host typically distributes a clock signal to all its modules. Below are some methods by which a COBO-compliant module can receive the clock signal, pass on the clock signal to other modules and communicate faults in timing communications.

Form-factors such as CFP2 have hardware pins that support REFCLK input and TX and RX monitor clock outputs. Smaller form-factors such as QSFP-DD and OSFP do not. The COBO specification shows optional designations of pins to match these well-known module functions: reference input clock REFCLK (B4/B5), recovered clock-Out/RxMCLK (B7/B8). RXLOSAlt (A6) can be used to signal loss of the clock signal.

5.1.1 SyncE 400ZR Example

One example application could be the support of SyncE over Coherent application (G.8275.1) by using REFCLK as a source reference input and RXMCLK as the recovered clock output as shown in Figure 5-1. SyncE allows the use of either a set of dedicated hardware pins (as in CFP) or via the 400GAUI data interface (as in QSFP-DD or OSFP). Either method is possible for a coherent DCO OBO provided the relevant hardware pins have not already been defined for another purpose such as OUT for a higher speed communications channel.

5.1.2 SyncO FlexO Example

Another use of COBO-defined pins is by using REFCLK (B4/B5) to provide a timing reference to synchronize the FlexO adaptation layer rather than the Ethernet payload. The host provides the timing information and the module DSP incorporates it appropriately into the OTN optical data stream. The diagram of Figure 5-1 still applies: the difference is how the timing information informs the optical signal.

In this example, Out/RxMCLK (B7/B8) could be used to monitor either the transmitter or receiver for optical input measurements.

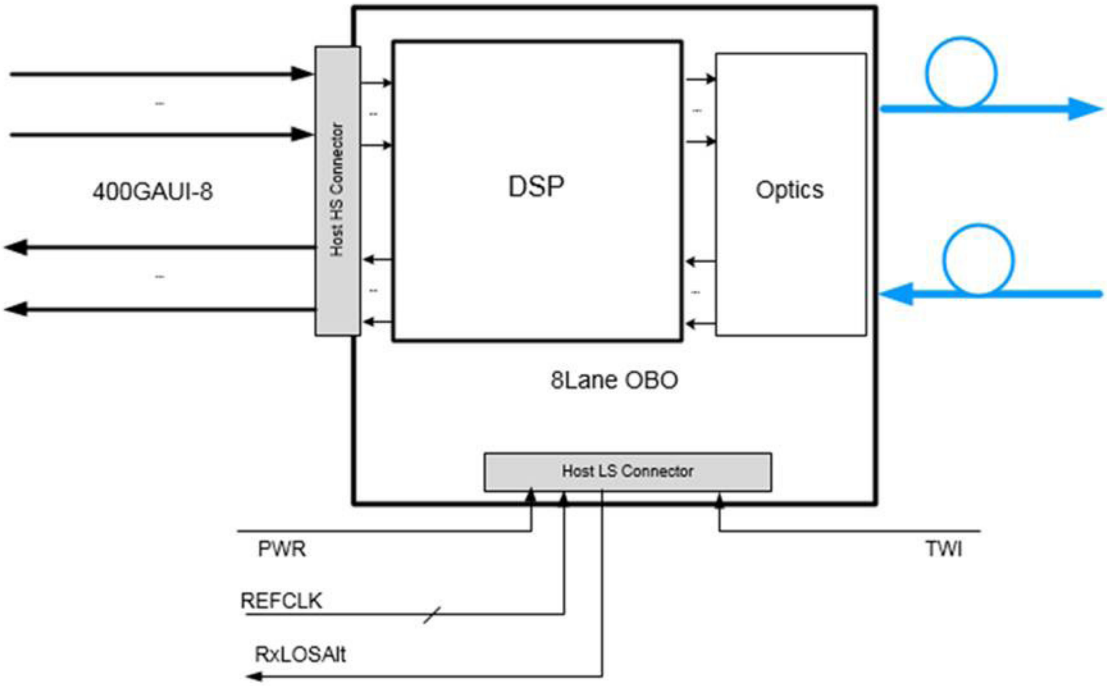

FIGURE 5-1 SYNCHRONIZATION TO AN EXTERNAL CLOCK USING REFCLK PIN

5.1.3 SyncE over SGMII

Serial Gigabit Media Independent Interface, SGMII, defined in IEEE 802.3 could also be used for SyncE/O as shown in Figure 5-2. More than a just a clock signal, SGMII is an Ethernet-based communications interface, but it can be synchronous. It can provide either system clock or a separate clock source. RxLOSAlt could be used to declare a Loss of Signal for SGMII.

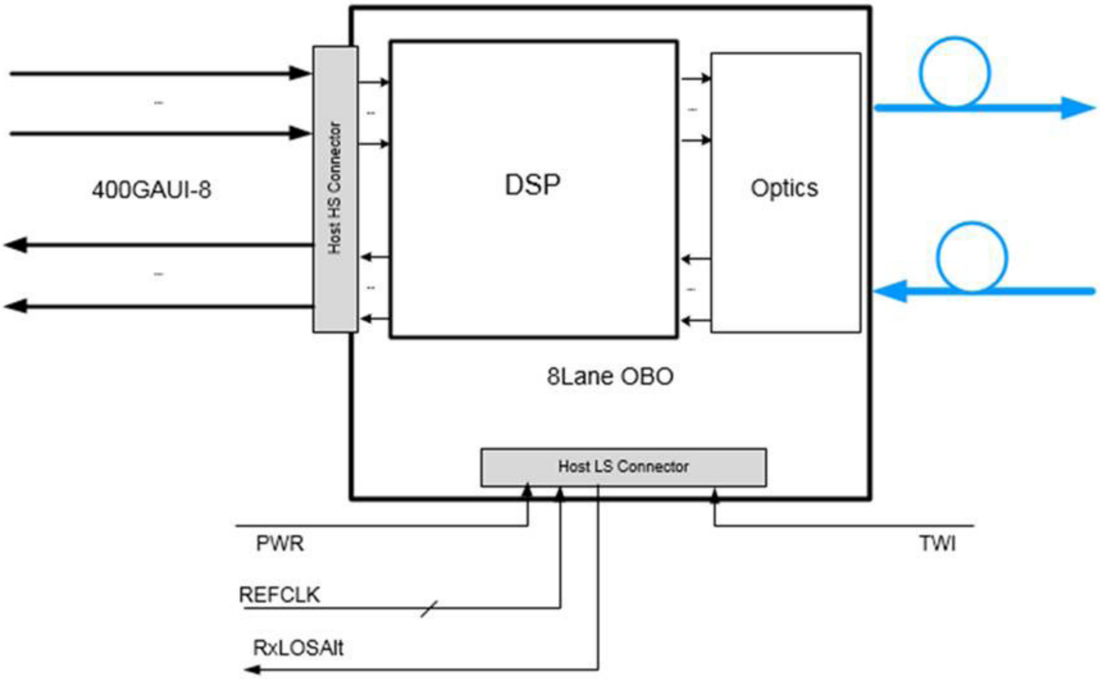

FIGURE 5-2 SYNCHRONIZATION USING SGMII

5.2 GCC For OTN

A different use for the same pins is for directly routing the GCC overhead information directly to the host board. This could be done using an Ethernet based SGMII communication interface.

5.3 Precision Timing Protocol

The precision time protocol (PTP) provides full time synchronization and relies on an absolute time-stamp. To compensate for delays, the protocol measures the round-trip time between adjacent nodes and assumes symmetric latency in forward and reverse direction.

Coherent modems/optics typically have asynchronous adaptation to a FEC layer, which requires elastic buffering that can introduce latency variations. When these optics are used to transport PTP packets embedded in the payload, this latency variation shows up as time error in the clock that is being distributed. The effect of this time error in a clock adversely affects the efficacy and operation of the application(s) for which the clock is distributed. The effects described impact the measurement in two ways:

1. PTP is not aware of a sudden change in latency hence would not attempt to re-calibrate latency expectation such as from reset of the module and its buffers

2. The change in latency is different in each direction, hence it creates a variation in the symmetry which cannot be calibrated by PTP

The contribution SG15-C.1131 to ITU-T SG15 provides a measurement based on existing implementations. The following observations are made.

1. It is observed that PDV is very minimal – on the order of 2 to 3 ns.
2. One-directional Latency variation is as high as 162 ns over two modules.
3. cTE varies as high as 78 ns – based on 20 iterations performed on 200G-16QAM mode of operation.

Performance targets for “devices” are outlined in G.8273.2, whereby an optical module would be considered part of a “device” and contribute to the generated Time Error.

T-BC Class	Permissable Range of Constant Time Error cTE (ns)
A	±50
B	±20

TABLE 5-2 SYNCHRONIZATION USING SGMII

5.4 Fast Communications with Host

Some future applications may require very fast communications. For example, this need may arise if the host is responsible for feedback loops within the module. Speeds up to 3 Gbps are possible using the extra pins.

6. MANAGEMENT INTERFACE FOR COHERENT COBO MODULES

6.1 Management Interface Specification Status

Management information is exchanged between the COBO module and the host system over a two-wire interface (TWI) which uses dedicated pins on the low-speed connector. The information model for COBO--sometimes still referred to as the "EEPROM registers" though modern implementations use a variety of memory hardware in practice--conforms to the Common Management Interface Specification (CMIS) developed by the CMIS Advisory Group [4]. Some management functions that were handled with dedicated hardware pins in CFP and CFP2 are now done over the TWI.

Basic functions that apply to all copper and optics modules have been defined and published. At time of publication of this Application Note, the CMIS functions for coherent have not yet been published. The information model has allocated space for the extra DWDM and coherent fields, but detailed assignment of the fields remains to be defined in detail.

The CMIS also defines state machines for module and datapath initialization. The module advertises its capabilities and then turns control over to the host which selects an application to instantiate, whereupon the module will turn up the transmitter. However, the coherent module could require extra intermediate states due to the turn-up process for the coherent receiver.

For a DCO coherent module the DSP is internal. Therefore, some coherent registers in CFP2 that were introduced to support other types of modules such as analog coherent optics (ACO) are not required.

6.2 Relationship to Management of Optical Subassemblies and ICs

In future, as optics and associated electronics become ever more complex, optical subassemblies are increasingly likely to also be managed via digital interfaces. For a COBO coherent module implemented with a high functionality optical subassembly inside, the host will likely communicate over one TWI with the COBO module's microcontroller which will in turn communicate over another TWI with the subassembly's microcontroller. The module microprocessor may need to re-map values between the subassembly and CMIS register addresses or accommodate different types of TWI for host-module or module-subassembly communication. Fast, high priority signals (e.g. alarms) need to be handled without incurring delay. This functionality should be kept in mind when choosing a microcontroller for the COBO module.

Another trend is for COBO, optical subassembly and DSP to have in-field (though not necessarily hitless) firmware update capability. The COBO module's TWI needs to be able to handle the substantial data transfers for not only its own updates but those of the subassembly and DSP. One way that was used for CFP modules that may be useful for coherent COBO or even client COBO modules is the transfer of large blocks of data or complex commands as command data blocks rather than simple writing of individual registers.

6.3 Features to Support Coherent Performance Management

The OIF has defined management of coherent modules at 100 Gbps [5]. It is anticipated that similar functions apply at 400 Gbps and 800 Gbps.

Coherent modules need highly accurate transmitter and local oscillator frequency control. It may be desirable to monitor the laser temperature as well as the module temperature.

Unlike for a fixed wavelength client module, the host is also responsible for setting a wavelength-tunable module's transmitter wavelength. 400ZR and 400GBASE-ZR specifications prescribe the module output power but there may be proprietary interfaces that would also require host control of output power through the available control register. These host requested values and module response as to whether they have successfully been satisfied need to be communicated via the management interface.

7. APPENDIX: FIBER RELIABILITY BACKGROUND

7.1 Overview

The density of components inside COBO modules and the density of COBO modules on the host board implies that bend radii of fiber routing needs to be small. There are two separate impacts of tight bending: loss and reliability.

Loss of the fiber in tight bends is not addressed in detail here. It is typically simple to calculate from the manufacturer's guidelines. High-density implementations inside a module and on a host board can take advantage of various types of bending-loss insensitive (BIF) single-mode optical fiber such as those compliant to ITU-T Recommendation G.657 [6]. There can be a small trade-off between reducing macrobending loss and increasing coupling loss to standard transmission fiber, depending on the mode field diameter of the fiber selected. The designer should consider the possibilities and trade-offs involved in using various bend-insensitive fiber types. Differences in fiber types with regards to bend loss is related to the optical index profile. Variations in index design typically play a minor effect in the mechanical strength of the fiber, which must be considered separately.

For reliability, the fiber choice, the fiber routing and strain relief to maintain proper routing need to be considered. More detail on design options and reliability considerations can be found in the COBO Optical Connectivity White Paper [3].

This section describes the considerations behind the use of 7.5 mm bend radius in the host and module examples described in Sections 3 and 4 of this App Note. This bend radius is now commonly used in optical modules including those with many turns of fiber. This section also provides an introduction on how to make the assessment for fiber reliability should implementers want to choose a different bend radius or target failure rate. Fibers with tolerable loss down to 5 mm or even less are available.

The takeaways from fiber reliability studies undertaken by COBO are:

- Failure probability goes up sharply as bend radius decreases, so total failure probability is often dominated by the tightest bends even when milder bends involve much (orders of magnitude) longer fiber lengths.

- Often, the most cost-effective way to improve the reliability of the fiber in a tight bend on-board optics application is to use fiber screened with a 2% proof test. The cost premium of the fiber is often negligible considering the limited length of fiber within the equipment. This criterion does not restrict any other design choice as proof testing is a standard process for commercially available fiber types.

- Another way to increase reliability in tight bend applications is by reducing stress by reducing the fiber diameter (e.g. to 80 µm cladding diameter from 125 µm cladding diameter). This option is typically costlier because fiber economies of scale are smaller, but also because more expensive specialty connectors and splice equipment are needed. In addition, yields can be reduced due to increased problems handling the smaller fiber.

- Often, the most cost-effective way to improve the reliability of the fiber in a tight bend on-board optics application is to use fiber screened with a 2% proof test. The cost premium of the fiber is often negligible considering the limited length of fiber within the equipment. This criterion does not restrict any other design choice as proof testing is a standard process for commercially available fiber types.

Further detail can be found in IEC TR 62048 [7]. A Corning white paper [8] gives some rules of thumb which may be helpful. This white paper considers a 7.5 mm bend radius in the context of an FTTH installation.

7.2 Reliability of Thirty-Two Class C Modules on a Host Board

7.2.1 Methodology

Bending creates tensile and compressive stresses in fiber. Stress in fibers due to tension cause slow growth in very small surface flaws over the life of the fiber eventually resulting in sudden fracture of the fiber. Failure probability is a function of lifetime, proof stress, bend radius (allowable stress), continuous fiber length, and empirically determined fiber strength characteristics.

Due to the empirical nature, estimates of failure rates will vary but agree within an order of magnitude. The calculations below are based on the IEC TR 62048 [7] report which uses a power law equation to demonstrate dependence on varying bend radii, proof stress and fiber diameter. Implementers who wish to use more stringent conditions should use this technical report as a starting point to their own calculations. The original purpose of the IEC document was to assess reliability of long fiber spans. Leading fiber vendors have reported that for short lengths below one meter the empirical failure rates are below the conservative bound provided by the power law.

Note that the fiber length used in the power law should not be understood to be the total fiber length used in the system. The length parameter used in the power law is defined as the fiber effective length under uniform stress such as the length of continuous fiber that is under test in the proof test equipment. For the system level reliability, the failure probability per individual fiber should be determined and the FIT rates summed over the number of fibers just as FIT is summed for all system components.

7.2.2 Methodology

Based on the host board placement shown in Figure 3-4, thirty-two 400G OBO modules can be fit onto the host board of the target 1 RU switch application with the most stringent bends which dominate the reliability calculation being the quarter-turn exiting the modules. The system-level failure rate multiplies the per pigtail failure rate by the sixty-four fibers and a quarter-turn. System FIT (failures in time) is the number of failures per billion system-hours of operation. A calculation for a different placement choice or number of modules per system would scale the calculation accordingly. The design condition in the example (2% proof test fiber with 7.5 mm bend radius and 5 years lifetime) corresponds to a system FIT rate contribution of 0.02 from the fiber bends.

Bend Radius (mm)	Estimated Failure Rate per m (ppm)	Length of 1 Turn (mm)	Failure Rate Per Turn (ppm)	System Failure Rate (ppm)	System FIT Rate
3	32	47	1.5	24	0.9
5	35	47	1.6	26	0.6
25	43	47	2.0	32	0.1

TABLE 7-1 ESTIMATED FAILURE AND FIT RATE AS A FUNCTION OF A LIFETIME

Table 7-1 shows the results of calculated failure rates based on the power law equation given in the IEC paper over different lifetimes. It shows there is little dependence on system lifetime in this range of values. All conditions in Table 7-1 are for 7.5 mm bend radius, 1% proof test and 125 µm fiber cladding diameter.

Bend Radius (mm)	Estimated Failure Rate per m (ppm)	Length of 1 Turn (mm)	Failure Rate Per Turn (ppm)	System Failure Rate (ppm)	System FIT Rate
5	90	31	2.8	45	1.0
7.5	35	47	1.6	26	0.6
15	3	94	0.2	4	0.1

TABLE 7-2 ESTIMATED FAILURE AND FIT RATE AS A FUNCTION OF BEND RADIUS

Table 7-2 shows a stronger dependence of failure rates as a function of bend radius. All conditions in Table 7-2 are for 5 years lifetime, 1% proof test and 125 µm fiber cladding diameter.

Proof Test	Estimated Failure Rate per m (ppm)	Length of 1 Turn (mm)	Failure Rate Per Turn (ppm)	System Failure Rate (ppm)	System FIT Rate
1%	35	47	1.6	26	0.6
2%	0.9	47	0.04	0.7	0.02

TABLE 7-3 ESTIMATED FAILURE AND FIT RATE AS A FUNCTION OF PROOF TEST

Table 7-3 shows that shows that applying a 2% proof test substantially improves the reliability. Both conditions are for a 7.5 mm bend, 5 years lifetime and 125 µm fiber cladding diameter.

Fiber Cladding Diameter (µm)	Estimated Failure Rate per m (ppm)	Length of 1 Turn (mm)	Failure Rate Per Turn (ppm)	System Failure Rate (ppm)	System FIT Rate
1.25	35	47	1.6	26	0.6
80	6	47	0.3	4	0.1

TABLE 7-4 ESTIMATED FAILURE AND FIT RATE AS A FUNCTION OF FIBER CLADDING DIAMETER

Table 7-4 shows that a fiber with a smaller cladding diameter that reduces the bend stress is also helpful but not as much as increasing the proof test criterion as in Table 7-3. Both conditions are for 5 years lifetime, 1% proof test and 7.5 mm bend radius.

8. REFERENCES

[1] "COBO 8-Lane & 16-Lane On-Board Optics Specification Release 1.1," Consortium for On-Board Optics (COBO), 2018.

[2] "COBO PCB Laminate Study and Thermal Assessment White Paper," Consortium for On-Board Optics (COBO), 2019.

[3] "Optical Connectivity Options for 400Gbps and Higher On-Board Optics," COBO Group, 2019.

[4] "Common Management Interface Specification for 8x/16x Pluggable Transceivers Rev 3.0," QSFP-DD MSA, 2018.

[5] "Implementation Agreement for Generation 2.0 100G Long-Haul DWDM Transmission Module -- Electromechanical," Optical Internetworking Forum, 2015.

[6] "ITU-T Recommendation G.657 Characteristics of a bending-loss insensitive single-mode optical fibre and cable," International Telecommunication Union, 2016.

[7] "IEC Technical report 62048 Ed. 3.0, Optical fibres – Reliability – Power law theory," International Electrotechnical Commission, 2014.

[8] P. D. M. E. &. J. D. G. Scott Glaesemann, ""The Mechanical Reliability of Corning Optical Fiber in Small Bend Scenario""," Corning Inc., 2007.

[9] "Class C COBO Thermal Study," Consortium for On-Board Optics (COBO), 2019.