CFP2: Mainstream 100 G Deployment Drives New Test and Measurement Requirements

Since 2009, we have seen early 100 G client deployment based on the CFP form factor. The module host interface was based on a more conservative 10 x 10 G interface based on established 10 G technology. While the CFP form factor is flexible (supporting a wide range of technologies, from 40 G serial through 40 G parallel and 100 G in both 4 x 25 G and 10 x 10 G), its large form factor and relative complex technology keeps it from meeting the density and price expectations required for mainstream 100 G.

Figure 1 shows the 100 G client-side form-factor evolution. The CFP2 form factor realizes the density at a price for more mainstream applications. It shares the same photonic (LR4) interface as the 100 G CFP at twice the faceplate density. The host interface moves from 10 x 10 G to 4 x 25 G. It also continues to use the MDIO control interface that the CFP uses.

When the first 100 G CFPs appeared in 2008, 10 G data rate electronics were relatively well established, making the host interface (10 x 10 G) relatively risk-free compared to the 25 G required for the photonics. The CFP required a so-called "gearbox" IC to multiplex/demultiplex the 10 x 10 G (CAUI) host interface into the 4 x 25 G lanes that the photonics required. The 25 G data path was limited to the short distances within the CFP module eliminating the need to traverse pluggable connectors or longer traces on the host/module boundary. Using 10 G technology allowed for connector and PCB designs to be based on existing practices. Challenges on issues such as dynamic skew and signal integrity have been overcome allowing early adopters to use CFP.

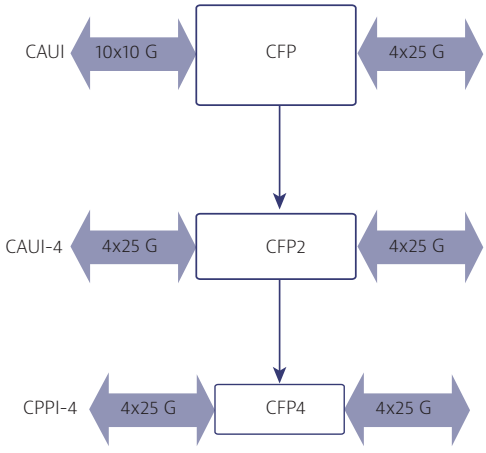

Figure 1. Evolution of the 100 G form factor

The transition to CFP2 requires 25 G data-rate signals to traverse the pluggable host/module interface, likely including several inches of PCB trace, presenting a considerable challenge as most 25 G ICs are still first generation; and we lack the extensive knowledge and experience we have with 10 G. We expect to face the challenges we experienced with the CFP at 10 G (crosstalk, signal integrity, and dynamic skew), but they will be compounded with challenges of the 25 G data rate.

CFP2, and the closely related CFP4, will likely be the dominant form factor for 100 G pluggables from 2013 on. Therefore, it is important to fully understand the challenges of developing equipment that exploits the density and price of CFP2. Figure 2 provides a CFP2 block diagram, and Figure 3 shows a commercial CFP2.

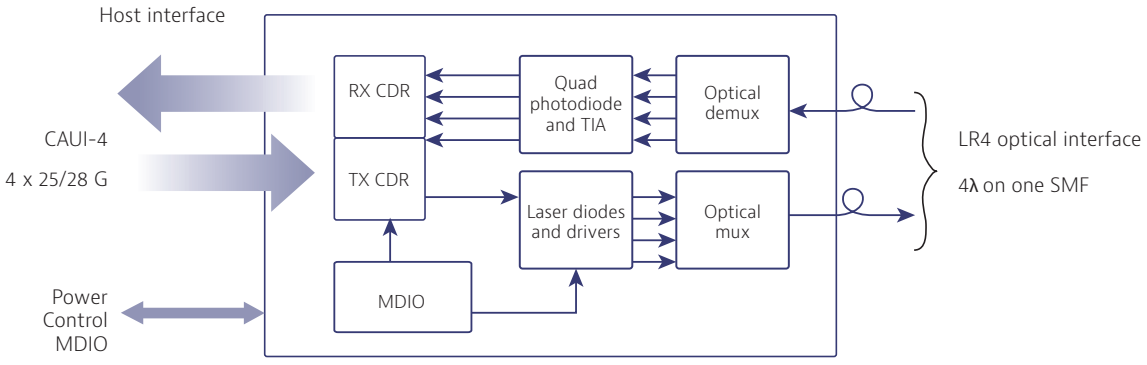

Figure 2. CFP2 Block Diagram

Figure 3. Viavi CFP2

25 G I/O

High-speed I/O is extremely challenging and even the third-generation 'mature' 10 G technology used in CFP gave designers many headaches maintaining signal integrity and reliable performance, especially when products moved into the production phase.

Although many vendors have shown 25/28 G ICs and connectors on test boards, it is much harder to design and build a product that a contract manufacturer must then mass produce.

Traditional high-speed, unframed BERTs were the only suitable tools for board-design validation and troubleshooting applications. However, shortcomings were quickly exposed with the need for greater insight into errors, lack of native form-factor support, and the inability to validate performance with real-world Ethernet and OTN signals.

Viavi Solutions was the first company to announce a 100 G CFP2-based test set with native 25/28 G I/O. During the course of its R&D and validation phase, Viavi gained extensive experience turning up these novel 25 G interfaces.

The use of a CFP2-based test set with native 25/28 G interface for design verification, eliminates the traditional challenges that come with module and system validation, such as signal integrity on the test fixture, control and support for modules (such as MDIO), and real-world traffic. The very act of connecting a traditional high-speed BERT to a CFP2 evaluation board required 16-phase, matched 40 GHz 2.4 mm cables, power supplies. and laptops with MDIO control. These are expensive, take a great deal of time, and still leaves the end user with many signal-integrity issues. Even with this test setup, the user is completely unable to conduct many of the tests and processes required to successfully validate CFP2 optics. Figure 4 shows the complexity of a 25/28 G test environment employing traditional BER analyzers.

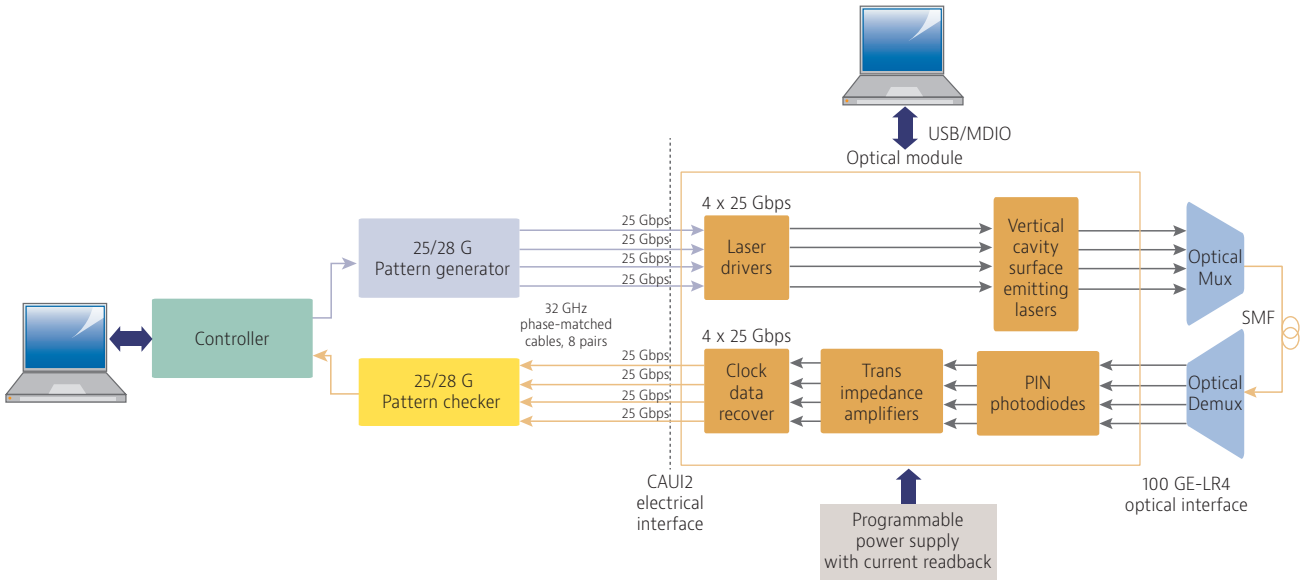

Figure 4. 25/28 G test environment

With native CFP2 support and a full range of applications including physical layer and Ethernet and OTN traffic, CFP2 debug, MDIO, dynamic skew and advanced error analysis, Viavi ONT CFP2-based test module provides complete coverage to validate CFP2.

Critical Tests

A CFP2 module is a closely-coupled system that requires a systems approach to validation and test. Interaction between performance, temperature, supply voltage, and control status need close examination, as shown in Figure 5.

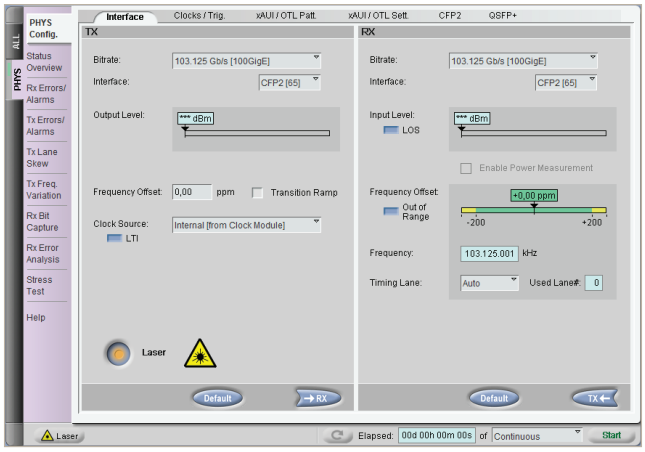

Figure 5. CFP2 overview

The most basic test is a simple loopback traffic test. Typically, this involves inserting the DUT into the ONT CFP2 and running a range of unframed (PRBS31) and framed (Ethernet traffic pattern) over a range of conditions. Many of the dynamic components such as a clock and data recovery (CDR) devices in the data chain can be impacted by the signals' various spectral nature. It is important to stress them with real signals (for example, Ethernet) under dynamic clock situations , as Figure 6 shows. The ONT CFP2 lets the user set a range of clock offsets which the test set can cycle through with variable transition and dwell times. This can simulate the dynamic nature of clocking in the real world, and shows bit errors occurring if dynamic elements can track the clock changes. The ONT CFP2 lets the user conduct these tests over a range of supply voltages (including ranges well outside the normal MSA range). Over long traces and the full load in real systems supply voltages, voltages below 2.7 V can be encountered, making validation important at these extremes.

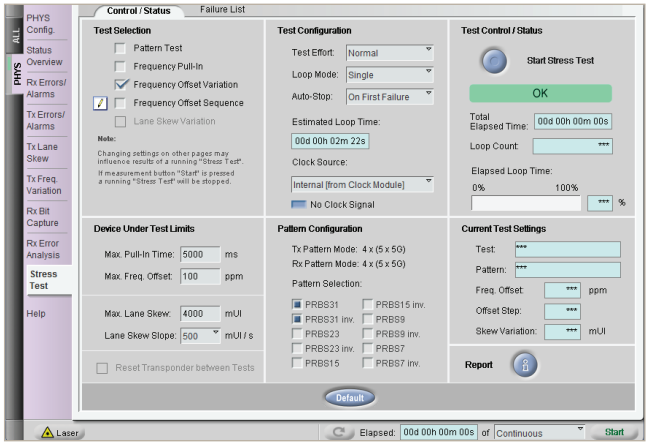

Figure 6. CFP2 stress test

Dynamic skew caused by subtle but dynamic changes in inter-channel delay is specifically called out in IEEE 802.3ba. This is a very real-world problem and understanding how elements perform at or beyond the limits is critical in delivering a reliable system. Tests originally developed with the ONT CFP have been expanded to allow dynamic skew testing out to 512 UI on 25/28 G interfaces to allow comprehensive testing and validation to well above the 92 UI called out in the standard, as Figure 7 illustrates.

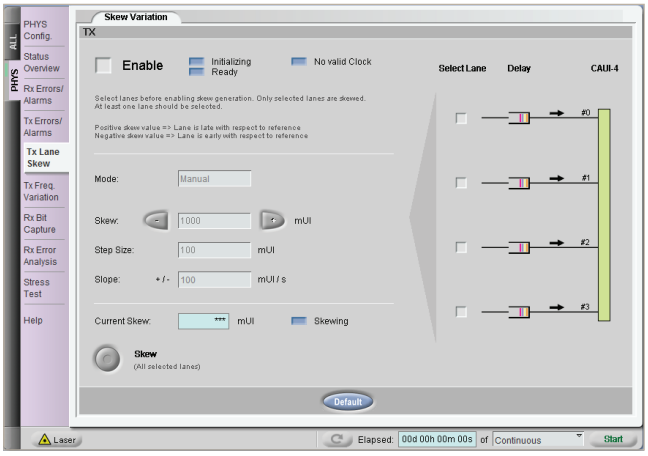

Figure 7. CFP2 dynamic skew

Table 1 gives examples of the critical tests required to ensure reliable CFP2 validation.

Required Test	Why	How
Traffic loopback	Connectivity and basic module functionality (MDIO, power and control)	Insert the CFP2 under test into the ONT CFP2 to check the MDIO and control validation status. Check the module's power draw. Run the test with unframed (PRBS) and then Ethernet traffic. Errors are recorded and checked against requirements. Manipulate and monitor the module's control pins using the MDIO application to ensure correct functional behavior.
Dynamic skew	Validate skew variation tolerance against IEEE 802.3ba	Use the dynamic skew application to set appropriate limits (92 µl). Run the test using different skew rates (from 1 UI/s down to 10 mUI/s) to check performance under likely operating conditions. Note: This is a very important test that should be conducted to validate full interoperability with CFP modules on the LR4 (photonic) side.
Module power supply margin	Ensure module operates correctly over potential power supply range	Run a traffic loopback test but adjust module voltage (2.5 V – 3.7 V possible using the ONT CFP2). Look for issues, especially with clock dynamics.
MDIO validation	Ensure MDIO registers can be accessed correctly	Use the integrated MDIO application to read blocks of MDIO address space to ensure you can write to the registers using MDIO peek/poke.
Stress test	Full stress of module to validate performance in real-world situations including: · Traffic signal integrity (crosstalk) · Clock dynamics (clock offset and pull-in) · Skew variation tolerance · Traffic sensitivity	Insert the CFP2 into the ONT CFP2 and run the CFP2 stress test. The user can set parameters such as clock dynamics, skew range, and traffic types. The test will validate the module's performance across the full dynamic range.

Table 1. Critical tests to reliably validate CFP2

Physical-Layer Challenges

Signal Integrity

Bit Errors

Deterministic (pattern-dependent) — Frequency-dependent effects, such as coupling capacitors and driver and receiver bandwidth limitations, can lead to pattern dependency. Long-sequence pseudorandom bit sequences (PRBSs) (typically 231) have long runs of zeros and ones with broad frequency content that stresses bandwidth dependencies. BER testing and monitoring with an oscilloscope (from a stable harmonically-related trigger signal) can detect and show the impact that they cause.

Crosstalk — The four parallel (differential) 25 G lanes require close routing across the host module PCB, across the connector, and inside the CFP2 module. Crosstalk can couple energy from adjacent lanes causing bit errors that are highly dependent on pattern and interlane timing. Clearly, if all lanes carry the same pattern (usually a PRBS) with the same offset, the crosstalk will likely re-enforce the desired pattern which is far from ideal during stress testing. Ideally a "victim" lane is assigned a basic pattern and other lanes run with aggressive patterns (often with slight sub-UI skew) to maximize the potential for destructive crosstalk-inducing patterns.

Random — Random bit errors are always possible during transmission, especially on photonic links at lower signal levels. External optical (per lambda) attenuation enables plotting of the BER against optical power to gain important information about the photonic receiver/laser combination. This capability is especially important with second-generation 100 G optics as they move from an externally modulated laser (EML)- to directly modulated laser (DML)-based technology (DML may not have a crisp, eye-like EML). A BERT system that tracks errored zeros and ones also helps to validate optical automatic gain control (AGC) among other things; typically, the errors should have a balanced zero and one distribution.

Bit Slips

CDR (short length slips) — Clock and data recovery (CDR) circuits are a critical part of high-speed links. A wide range of effects influence the performance of such circuit blocks, but the typical outcome is a bit slip as the CDR retimes. In classic test sets, a bit slip typically drives a loss of sync requiring hundreds of microseconds for the test set to re-sync. The end user must determine the cause of the error burst and potentially may need to re-sync. Advanced test equipment can indicate that the bit error is a short slip (typically 1 to 2 bits), indicating a CDR-related issue.

FIFO re-centering (commonly seen as sync loss) — First-in/first-out (FIFO) buffers are vital when moving data between different clocking domains. In rare cases, FIFO issues can corrupt large blocks of data and then re-synchronization occurs as the FIFO refills and centers. The bit slippage is far longer than the typical 1 to 2 bits that occur with CDR-based issues. This slippage can be a design issue triggered by physical-layer conditions.

Dynamic Skew

IEEE Standard Defines Skew Variation Tolerance at Specific Interface Points

Dynamic skew generation (the ability to continuously move individual lanes relative to each other over a large range in sub-UI steps) is important for validating timing circuit correction operation at various interface points. The 4 x 25 G interface that the CFP2 uses must tolerate a dynamic skew of 92 UI (per IEEE 802.3ba, table 80.5). Dynamic skew helps highlight crosstalk issues and also helps adjust the relative lane timing to get the worst-case coupling between lanes. This test is also very useful for checking the 25 G side of CFP modules to ensure that the internal gearbox is functioning correctly.

Influence of Real-Data Signals on CDRs, Dispersion Control, and Other Issues

Pattern Dependency (PRBS vs. Real Signals)

Traditional testing based around PRBS patterns and the wide choice of patterns (typically 27...231) cannot identify many hardware issues. These two issues occur when using PRBS to validate data links:

- The PRBS contains different spectral content than real signals (such as framed OTN) and, therefore, frequency-dependent systems such as CDRs behave differently, even if the data signal is less aggressive than the PRBS.
- As the integration level increases, components such as a MAC IC cannot work with unframed PRBS and need a genuine data signal to validate and pass through internal data paths. The ability to test the true physical layer (skew, jitter, and eye control) with real signals is critical when validating component performance.

Therefore it is necessary to migrate from the traditional board design verification process with BER analyzers to a more sophisticated test procedure that employs a test set that can verify physical layer performance and compliance while generating framed (Ethernet/OTN) signals used in real-world telecommunication networks.

Clocking Range and Dynamics

The various CDRs and clocking systems used within the 100 G system often use a large number of phase-locked loops (PLLs). The ability to alter the base clock rate over large ranges (typical x 10-fold normal parts per million [PPM] offset) and to allow clock jumps and ramps can stress the clocking chain, therefore, the design requires significant margin. The ability to jitter modulate the data signal is also important for validating the performance and design margin of the receiver, especially the CDR.

Of course, test equipment does not exist in isolation, so core test equipment requires appropriate triggers and clock I/Os for it to synchronize with other equipment, especially fast oscilloscopes.

Clock and trigger I/O should be low noise and low jitter, which is especially critical for trigger signals as even a little uncertainty in the trigger can cause excessive eye closure, see Figure 8.

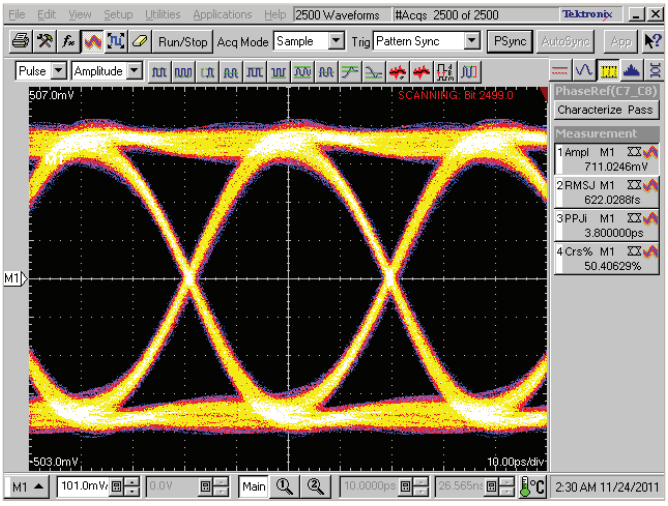

Figure 8. Timing margins at 25 G are extremely tight and require low jitter triggers and low phase-noise clocking.

Conclusion

For 100 G technology to transition to mainstream, it must meet price and form-factor expectations. To achieve this, it must be based on reliable, stable hardware. During the transition to this phase, it is critical that the test equipment used during development and validation meet the challenges the physical layer demands of 25/28 G.

Contact Us **+1 844 GO VIAVI**
(+1 844 468 4284)

To reach the Viavi office nearest you,
visit viavisolutions.com/contacts.

cfp2-wp-opt-tm-ae
30173262 901 0313