

Maximizing the Capacity-Reach of 800G Generation Coherent: Baud Rates, Features, and Modem SNR

High-end coherent optical technology is evolving to an 800 Gb/s generation characterized by 7 nm process node ASICs/DSPs, 90+ Gbaud symbol rates, advanced modulation techniques including probabilistic constellation shaping (PCS), and wavelength data rates of up to 800 Gb/s. However, not all 800G generation coherent optical engines are able to deliver the same levels of performance. This white paper explains how an ultra-high baud rate, innovative features, and a high modem signal-to-noise ratio (SNR) enable Infinera's 800G generation Infinite Capacity Engine (ICE6) optical engine to maximize wavelength capacity-reach – the maximum wavelength data rate that can be achieved for a given reach requirement.

Wavelength Capacity-Reach Impact on Optical Total Cost of Ownership

Wavelength capacity-reach has a significant impact on the total cost of ownership of an optical transport network. It directly impacts network CapEx in terms of the number of coherent interfaces that are required and therefore the cost per bit. If, for example, for the same given reach requirement optical engine A can deliver 600 Gb/s (150%) while optical engine B can deliver 400 Gb/s (100%), then optical engine A will reduce the cost per bit by 33% based on the assumption that the cost per interface is the same, which typically holds true over high-performance coherent generations. If optical engine A can deliver 800 Gb/s (200%) compared to 400 Gb/s (100%) for optical engine B, then the cost per bit is reduced to 50%, as shown in Figure 1. Using fewer wavelengths/interfaces also provides an OpEx benefit, not just in terms of lower power consumption and smaller footprint, but also in terms of the operational costs of installation, provisioning, and management.

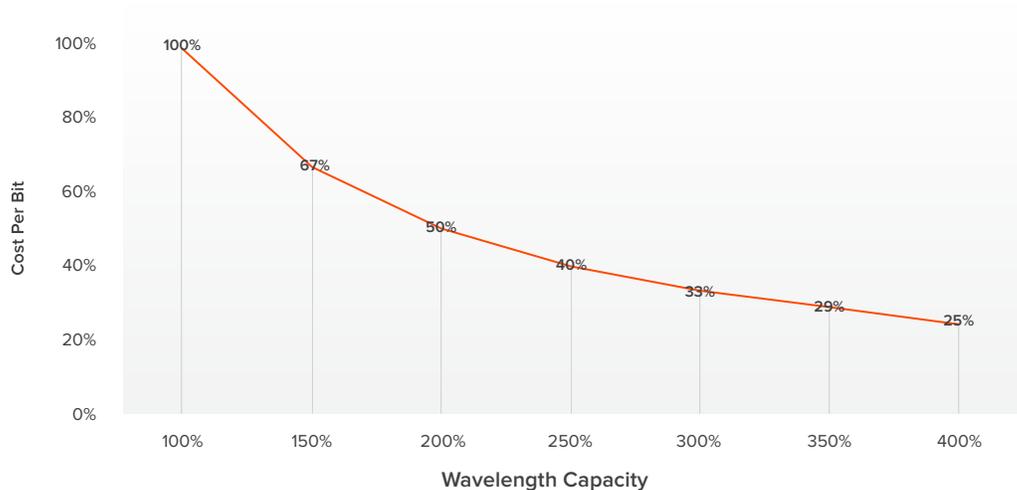


Figure 1: Increased wavelength capacity reduces the cost per bit

What Determines Optical Performance?

At ultra-high baud rates, optical performance in terms of wavelength capacity-reach is determined by both external and internal factors. Key external factors include optical penalties, such as amplified spontaneous emission (ASE) noise from optical amplifiers that accumulates as the wavelength traverses the network, thus reducing the optical signal-to-noise ratio (OSNR) at the receiver. Lots of factors influence the amount of accumulated noise, including amplifier types, span lengths/losses, and fiber quality/age factors such as splice and repair losses. A second key external factor is fiber nonlinearities such as cross-phase modulation (XPM) and self-phase modulation (SPM). Factors that influence nonlinearities include the power spectral density of each wavelength and the chromatic dispersion profile of the fiber type, with more chromatic dispersion helping to reduce nonlinearities.

Key External Factors	Key Internal Factors
<ul style="list-style-type: none"> • Optical noise (i.e., OSNR) • Fiber nonlinearities (i.e., XPM, SPM) • Filter penalties 	<ul style="list-style-type: none"> • Modem SNR (noise/distortions inside the optical engine) <ul style="list-style-type: none"> • Noise/distortions at high baud rates • Noise created when digitally compensating for impairments such as chromatic dispersion • Frequency-dependent performance across the wavelength

Table 1: Key external and internal performance factors

Performance depends on how big these impairments are and how tolerant the optical engine is to them. An additional external factor is filter penalties, the loss of energy at the edge of the wavelength due to the limited passband of any individual filter and the filter narrowing effect that occurs due to filter cascade, as shown in Figure 2.

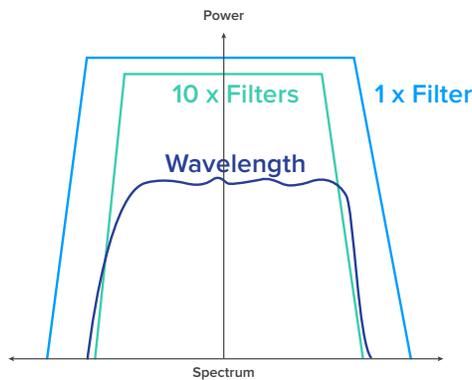


Figure 2: Filter penalties and filter narrowing

The other factor that can critically impact optical performance is the modem SNR, the amount of noise/distortions created inside the optical engine. This includes noise/distortions when operating at very high baud rates. It also includes the noise that is created when the DSP compensates for impairments such as chromatic dispersion. Plus, it includes the frequency-dependent performance, across the same wavelength, of the digital and analog electronics in the optical engine.

Three Key Enablers of Maximized Wavelength Capacity-Reach

ICE6 leverages three enablers to maximize wavelength capacity-reach: ultra-high baud rate, innovative features, and high modem SNR. Figure 3 indicates the relationships between these enablers and the external/internal performance factors described in the previous section.

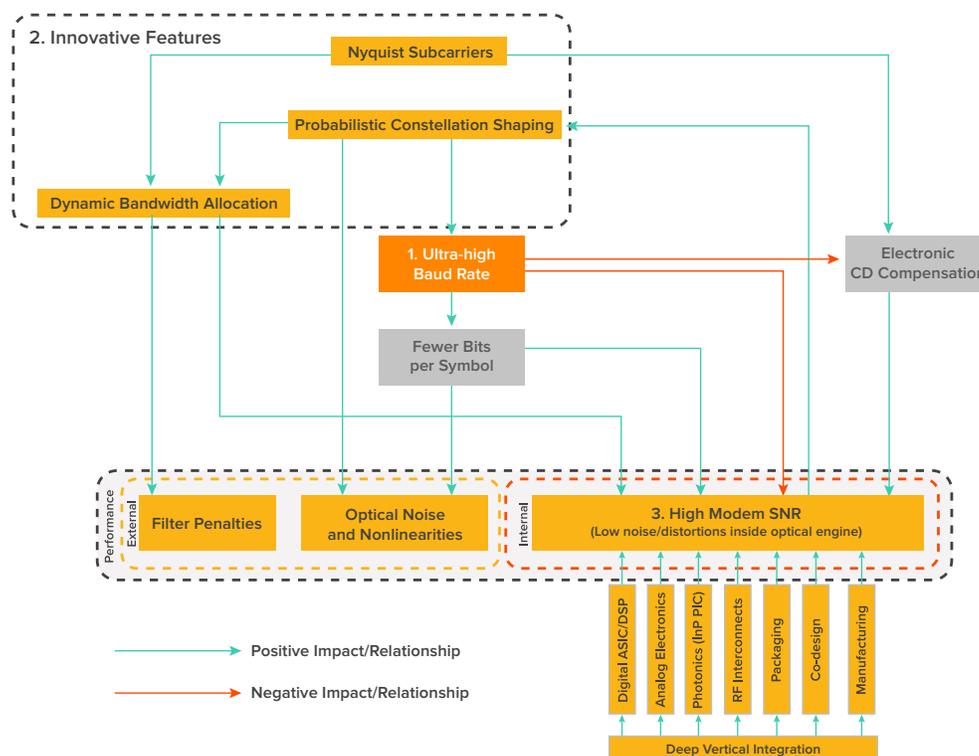


Figure 3: Key enablers of ICE6 wavelength capacity-reach

Enabler 1: Ultra-high Baud Rate (96 Gbaud)

As explained in the Infinera white paper *“The Ultimate Guide to Higher Baud Rates,”* higher baud rates provide the key lever for increasing wavelength capacity-reach. Higher baud rates enable the use of lower-order modulation to achieve the same data rate. Lower-order modulations benefit from greater Euclidean distance between constellation points, making them easier to distinguish in the presence of noise. At the same time, as the spectrum of the wavelength is proportional to the baud rate, a higher-baud-rate wavelength can leverage higher power for the same power spectral density and therefore the same level of nonlinearities. Together, lower-order modulation and higher power more than offset the increased sensitivity to noise and nonlinearities of the higher baud rate itself, resulting in significant capacity-reach improvements.

A well-known rule of thumb with higher-order modulation and coherent transmission is the “1 bit = 3 dB rule” – doubling the number of constellation points (QPSK → 8QAM, 8QAM → 16QAM) adds one bit per symbol per polarization and approximately halves reach with a 3 dB increase in the required OSNR. In fact, the Shannon limit itself follows this rule at higher SNR values, where we can ignore the 1 in its famous equation, $C/B = \text{Log}_2(1+\text{SNR})$. Adding 1 bit to the spectral efficiency requires us to double (+3 dB) the required SNR, thus halving the reach.

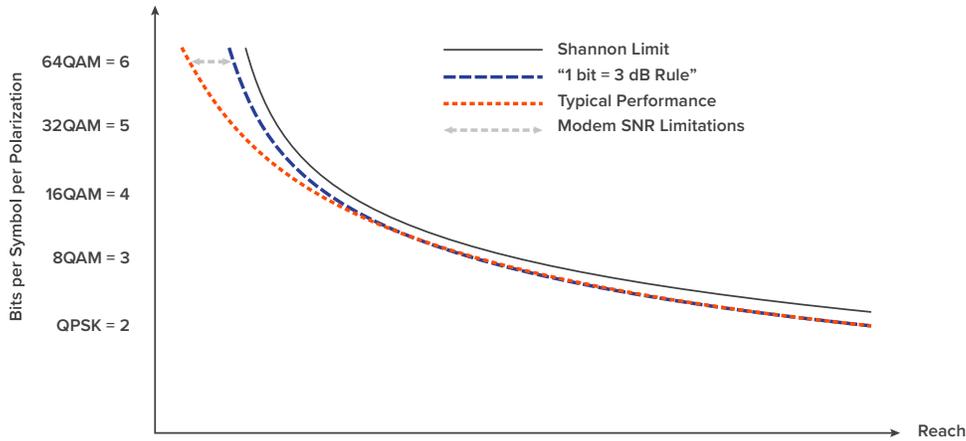


Figure 4: Modulation (bits per symbol per polarization) vs. reach

However, as shown in Figure 4, with higher-order modulations (32QAM, 64QAM, etc.), modem SNR, the amount of noise and distortions inside the optical engine, becomes a key limiting factor that further reduces reach. With low-order modulation, the SNR limit is relatively low and optical noise and nonlinearities are the primary limitation on SNR and therefore reach. With the higher SNR limit of high-order modulation, modem SNR takes up a larger portion of the available noise limit, allowing for a much smaller amount of external noise (i.e., OSNR and nonlinearities).

Figure 5 shows, for Infinera’s ICE6 optical engine, the relative reach of an 800 Gb/s wavelength with the base baud rate required for 800 Gb/s and the full 64QAM. Increasing the baud rate by around 8% increases the reach by around a factor of three, while increasing the baud rate by around 15% increases the reach by around a factor of four. These dramatic increases in reach can be explained as follows. As we move away from the full 64QAM, modem SNR becomes less of a limiting factor. Also, with a high number of bits per symbol, increasing the baud rate by 8% reduces the number of bits per symbol per polarization by almost half a bit, while increasing it by 15% reduces the number of bits per symbol by almost three-quarters of a bit, enabling us to benefit from a large portion of the 1 bit = 3 dB rule. Finally, as we move away from the full 64QAM, we also start to benefit from PCS gain.

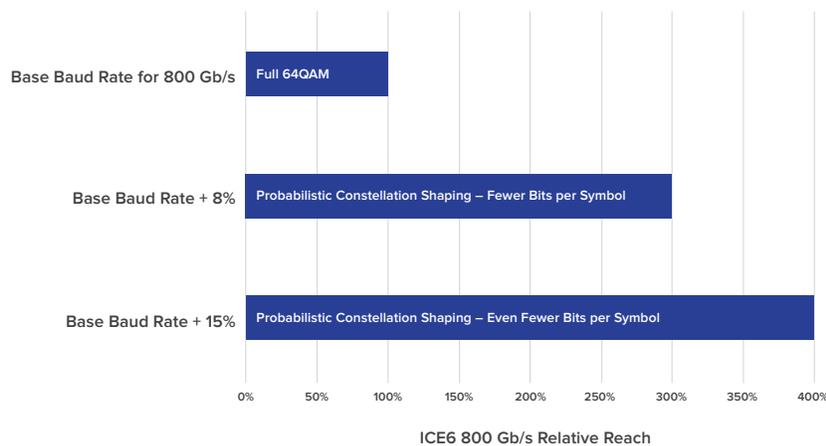


Figure 5: 800 Gb/s wavelengths: ICE6 reach vs. baud rate

At lower data rates, a marginal increase in the baud rate has a less dramatic but still significant impact on the reach: increasing the baud rate by 15% (from the same base as the 800 Gb/s example) increases the 600 Gb/s reach by up to 40% and the 400 Gb/s reach by up to 20%. One reason these increases are less dramatic is because modem SNR is no longer such a key limitation. A second reason is that the absolute reduction in bits per symbol is proportionally lower, giving us less gain from the 1 bit = 3 dB rule. For example, at 800 Gb/s with the full 64QAM (6 bits per symbol per polarization), a 20% increase in the baud rate reduces the number of bits per symbol per polarization by 1, while at 400 Gb/s (3 bits per symbol per polarization), the same 20% increase in baud rate would reduce the number of bits per symbol per polarization by only half a bit. A third reason is that at 600 Gb/s and 400 Gb/s, even at the base baud rate we already benefit from PCS-64QAM gain.

Enabler 2: Innovative Features

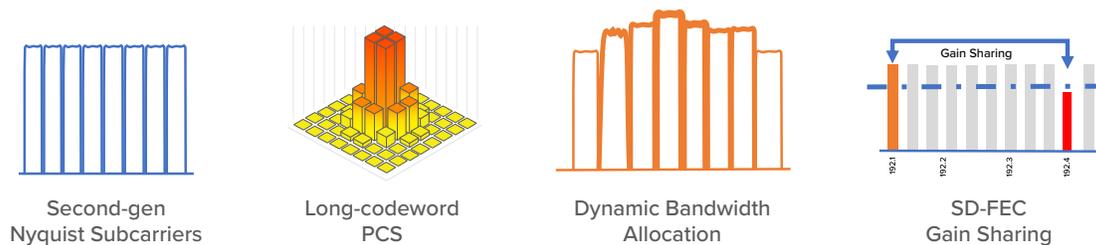


Figure 6: ICE6 innovative features

The second factor that has a key influence on wavelength capacity-reach is innovative features. For ICE6, these include Nyquist subcarriers, long-codeword probabilistic constellation shaping (LC-PCS), dynamic bandwidth allocation (DBA), and SD-FEC gain sharing.

2A – Nyquist Subcarriers

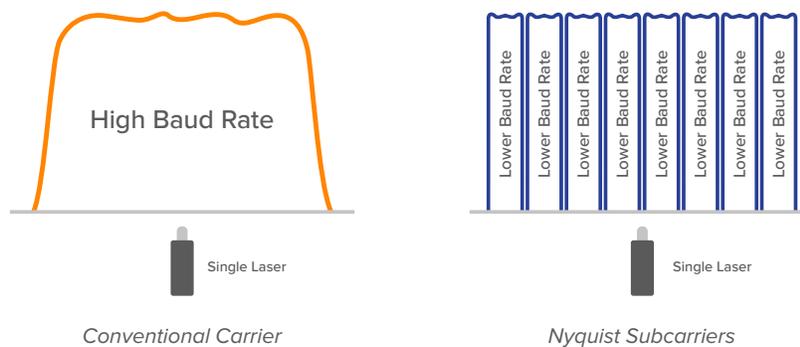


Figure 7: Nyquist subcarriers

As explained in the Infinera white paper “*The Ultimate Guide to Nyquist Subcarriers*,” Nyquist subcarriers take a single high-baud-rate carrier and digitally divide it into multiple lower-baud-rate subcarriers. In terms of ICE6 wavelength capacity-reach, the primary benefit of Nyquist subcarriers is reduced chromatic dispersion. As shown in Figure 8, chromatic dispersion occurs because different frequencies travel at different speeds through the fiber – even different frequencies of the same wavelength travel at slightly different speeds and eventually distort the signal.

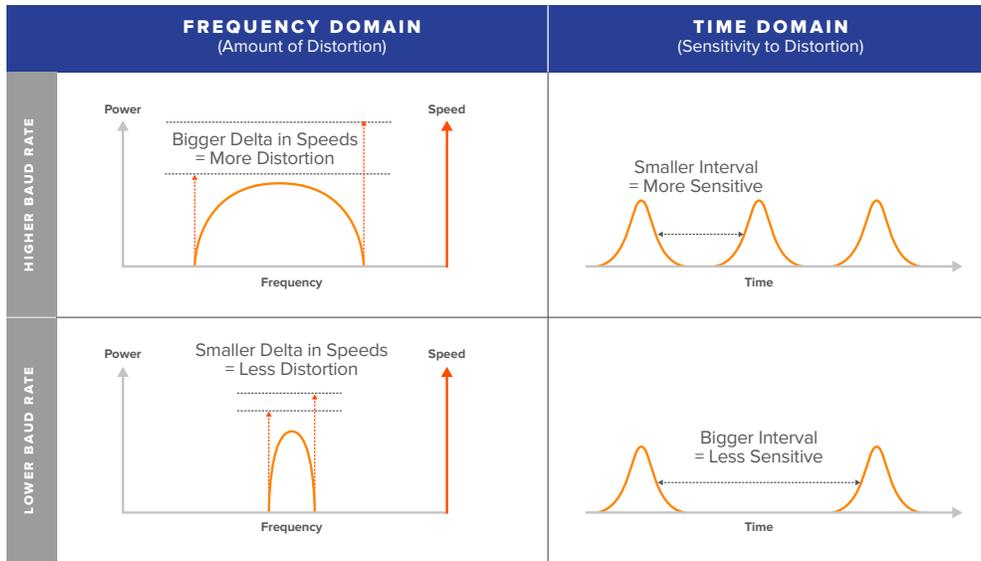


Figure 8: Baud rate and chromatic dispersion

As the spectral width of the signal is proportional to the baud rate, a high-baud-rate signal has a bigger delta between its lowest and highest frequencies and therefore experiences greater variation in the speed of its frequencies through the fiber and more spreading or distortion in the time domain. In addition, more symbols per second means a shorter time interval between symbols, so distorted symbols can more easily overlap. These two factors combine to create a squared relationship between baud rate and chromatic dispersion. Nyquist subcarriers can therefore dramatically decrease the effect of chromatic dispersion – by a factor of 64 with eight subcarriers, as is the case with ICE6. Even if the single-carrier chromatic dispersion is within the capabilities of the digital ASIC/DSP, compensating chromatic dispersion has a cost in terms of additional noise inside the optical engine, reducing the modem SNR. Reducing chromatic dispersion has a significant benefit in terms of reducing this noise and therefore improving performance.

2B – Long-codeword Probabilistic Constellation Shaping

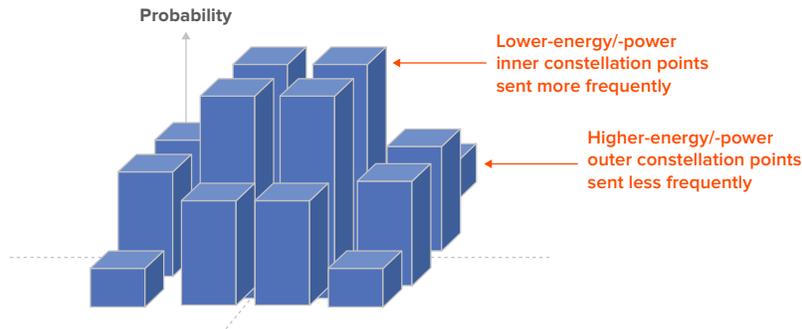


Figure 9: Probabilistic constellation shaping

As explained in the Infinera white paper “*Faster, Further, Smoother: The Case for Probabilistic Constellation Shaping*,” probabilistic constellation shaping uses the lower-energy/-power inner constellation points more frequently and the higher-energy/-power outer constellation points less frequently, as shown in Figure 9, as opposed to conventional modulation where each constellation point has the same probability of being used. For the same average power and spectral efficiency, there is greater Euclidean distance between the constellation points relative to conventional QAM, which increases tolerance to noise. Figure 10 illustrates this showing 16QAM and PCS-64QAM, both with 8 bits per symbol and the same total power, with the probability of each constellation point represented by its area. In addition, the receiver can use the probability distribution of the PCS constellation points to further enhance noise tolerance. This increase in noise tolerance comes without having to increase the wavelength’s power (i.e., power spectral density) and therefore increasing the nonlinear penalties. Alternatively, for the same noise tolerance, less power is required, and therefore nonlinearities are reduced. For high-power scenarios, ICE6 also supports a super-Gaussian probability distribution that results in less variation in the power levels of the symbols and therefore lower nonlinear impairments.

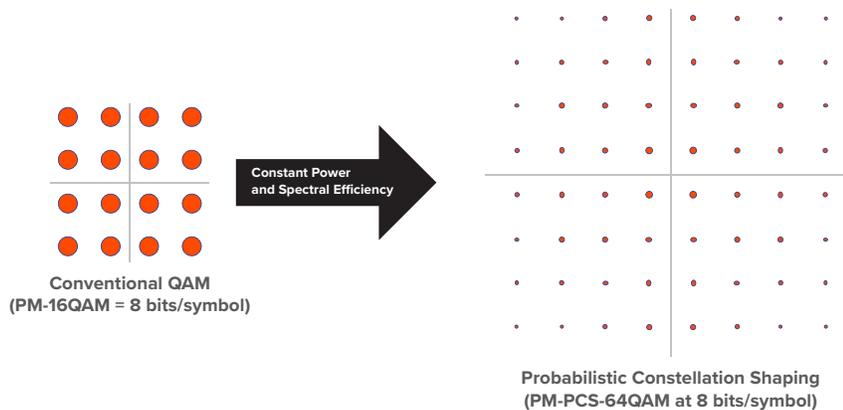


Figure 10: PCS improves noise tolerance

As also explained in the PCS white paper, Infinera’s PCS implementation uses a long codeword (>1,000 symbols), which delivers almost all the theoretical gains of PCS, close to double that of a short codeword (~100 symbols), though modem SNR, as discussed later in the Modem SNR section, also has an impact on PCS gain. Furthermore, PCS provides the bits per symbol granularity that enables the use of the highest baud rate, even where that baud rate would not have aligned with conventional modulation to deliver the required wavelength data rate.

2C – Dynamic Bandwidth Allocation

Dynamic bandwidth allocation combines Nyquist subcarriers and probabilistic constellation shaping, enabling the data rate of each subcarrier to be set independently. This improves performance in two ways. It helps mitigate filter penalties, which tend to be higher on the outer subcarriers and lower on the inner subcarriers, by using PCS to set a lower data rate on the outer subcarriers and a higher data rate on the inner subcarriers. DBA can also address frequency-dependent performance across the same wavelength of the digital and analog electronics inside the optical engine, which has a similar pattern of better performance for the inner subcarriers and worse performance on the outer subcarriers.

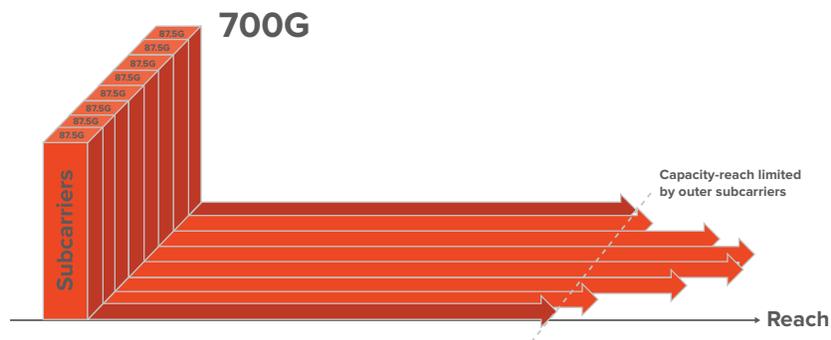


Figure 11: Without DBA, capacity-reach is limited by outer subcarriers

For example, if all the subcarriers have to be at the same data rate, then for the reach requirement shown in Figure 11, we are limited by the outer subcarriers to 87.5 Gb/s per subcarrier, giving a total of 700 Gb/s for the wavelength. With DBA, we can increase the inner subcarriers to 95 Gb/s, 105 Gb/s, and 112.5 Gb/s, increasing the total wavelength capacity to 800 Gb/s, as shown in Figure 12.

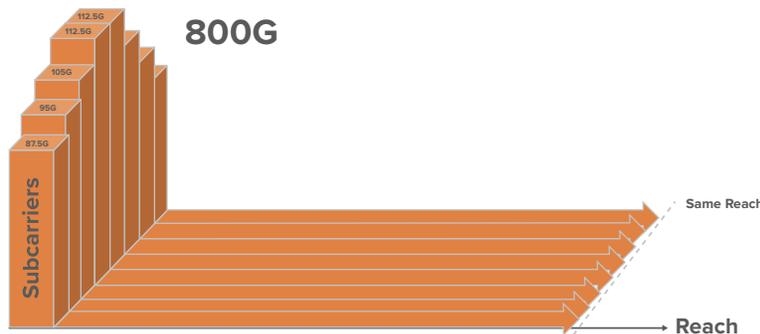


Figure 12: Dynamic bandwidth allocation enables increased capacity

2D – SD-FEC Gain Sharing

Another innovative feature in the ICE6 optical engine is SD-FEC gain sharing. While this feature cannot improve the capacity-reach of a single wavelength, it can maximize the combined capacity-reach of a pair of wavelengths. The need for SD-FEC gain sharing arises because even two wavelengths traveling along the exact same A-Z path will experience different impairments related to amplifier tilt, chromatic dispersion, polarization mode dispersion (PMD), and polarization-dependent loss (PDL).

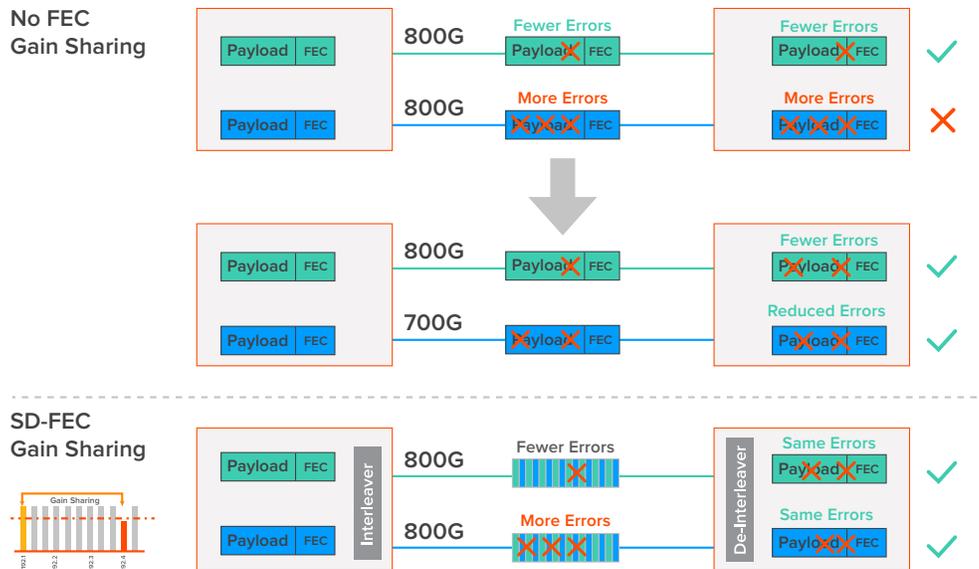


Figure 13: SD-FEC gain sharing

Rather than perform forward error correction on the frames sent over each wavelength independently, as shown in Figure 13, SD-FEC gain sharing interleaves the two frames, payload and FEC overhead, so that half of each frame goes over each wavelength, and therefore they each experience a statistically identical number of pre-FEC errors. Now, when the original frames are reassembled, the FEC decoders have the same amount of work to do, and the gain is the same. For example, as shown in Figure 13, while the first wavelength can achieve 800 Gb/s with margin to spare without SD-FEC gain sharing, the second wavelength is limited to 700 Gb/s. With SD-FEC gain sharing, both wavelengths can operate at 800 Gb/s.

Enabler 3: High Modem SNR

A third factor that has a big impact on optical performance, especially at very high data rates, is the amount of noise/distortions created inside the optical engine, distortions that reduce the modem SNR. As the baud rate increases, distortions and noise inside the optical engine become an increasing challenge.

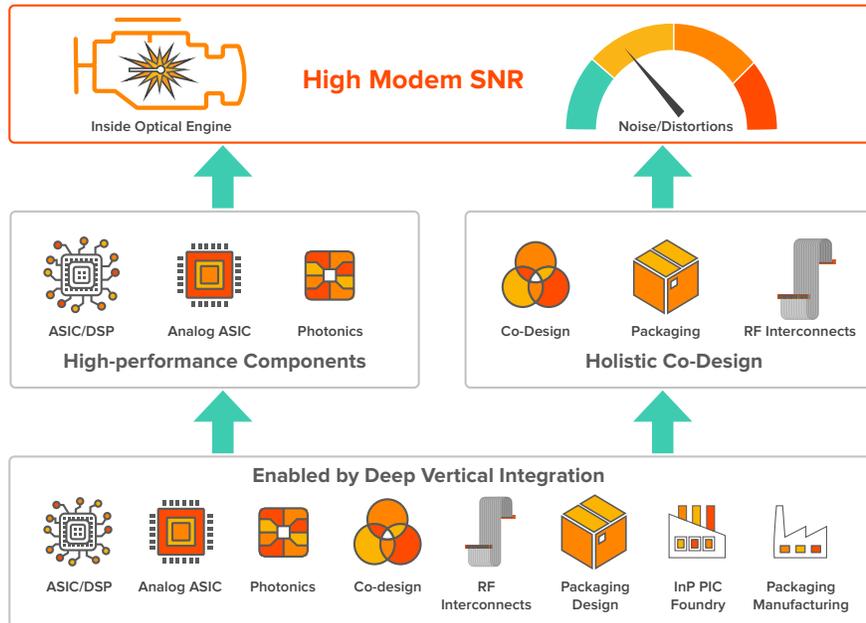


Figure 14: High modem SNR enabled by deep vertical integration

The key factors, as shown in Figure 14, that can impact modem SNR include:

3A – The Performance of Individual Components

The quality of the individual components of the digital ASIC/DSP, including the DAC and ADC, analog electronics, and photonics, have a large influence on modem SNR.

For example, as described in the Infinera white paper *“The Advantages of Indium Phosphide Photonic Integration in High-performance Coherent Optics,”* ICE6 leverages an indium phosphide single photonic integrated circuit (PIC) that includes the critical modulator function where indium phosphide enables the highly efficient electro-optic effect. Furthermore, ICE6 integrates multiple photonic functions including lasers, modulators, semiconductor optical amplifiers, photodetectors, and various passive functions into a single PIC. This improves modem SNR by minimizing coupling losses by connecting optical functions with waveguides inside the PIC, as opposed to coupling optics between discrete components.

Analog electronics also have a critical impact on modem SNR. Analog electronics include the drivers that convert lower voltages from the DSP/DAC to the higher voltage required by the modulator at the transmit end, and transimpedance amplifiers (TIAs) that convert current from the photodetectors to the voltages required by the ADC/DSP at the receive end. ICE6 leverages a single analog ASIC for two wavelengths, transmit and receive, with a total of eight drivers and eight TIAs. It is made with high-performance silicon germanium (SiGe), fabricated with a 180 nm Bipolar CMOS process. The drivers leverage a two-stage amplifier design featuring built-in equalization, while the TIAs include automated gain control (AGC) amplifiers and built-in equalization.

3B – Holistic Co-Design Including the RF Interconnect

Modem SNR is also determined by holistic co-design, packaging, and the electrical/radio frequency (RF) interconnects between the ASIC/DSP and analog electronics, and between the analog electronics and the photonics/PIC. Holistic co-design enables the design of each individual component, the RF interconnect, and packaging to be done with consideration for the impact on other components and overall optical engine performance, optimizing any trade-offs to maximize performance.

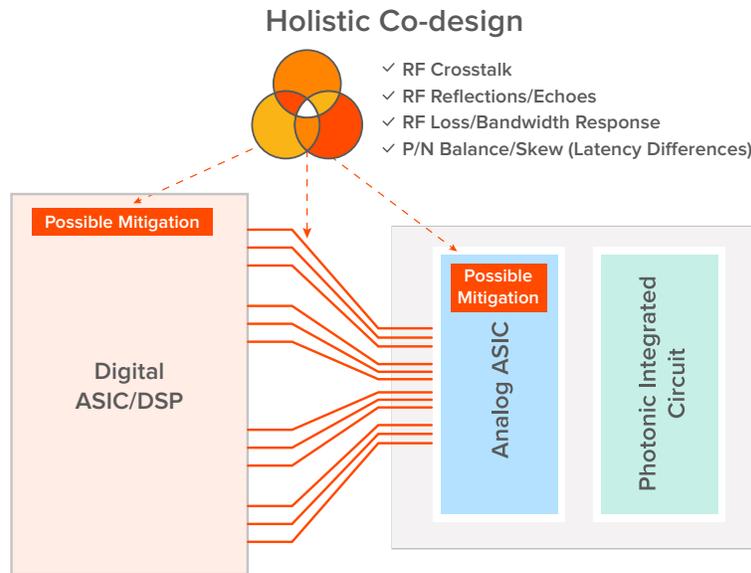


Figure 15: Holistic co-design of the critical RF interconnect

For example, the RF interconnects are critical to the performance of the optical engine, especially at ultra-high baud rates. As shown in Figure 15, optimal design of the RF interconnect between the digital ASIC/DSP and the analog electronics needs to consider multiple factors including RF crosstalk, reflections/echoes, P/N balance/skew (latency differences between positive and negative voltages), and RF loss/bandwidth response. Electronic mitigation can be located in the digital ASIC/DSP or the analog ASIC, so an optimized design considers the digital ASIC/DSP, analog ASIC, and RF interconnect, sharing the load of impairment mitigation.

3C – Enabled by Deep Vertical Integration

The individual high-performance components, holistic co-design, and innovative features are all themselves enabled by the deep vertical integration provided by the Infinera Optical Innovation Center (OIC). OIC disciplines include coherent ASIC/DSP design, photonic integrated circuit design and manufacturing, analog ASIC design, advanced packaging design and manufacturing, and holistic co-design, including the RF interconnect.

In addition, as discussed previously, while higher baud rates can exponentially increase the effect of chromatic dispersion, ICE6 leverages Nyquist subcarriers to dramatically reduce the amount of noise created inside the optical engine when compensating for chromatic dispersion. Modem SNR is also one of the key factors that determine the PCS gain of a practical coherent transceiver, with both a long codeword and high modem SNR required for good performance.

ICE6 Capacity-Reach

ICE6 has demonstrated 950 km at 800 Gb/s and is capable of 2,500+ km at 600 Gb/s and 6,500+ km at 400 Gb/s. These figures are at the top of the range of industry claims for 800G generation coherent, as shown in Figure 16, which also shows the industry ranges for 600G generation (i.e., 16 nm DSP, 60-70 Gbaud, max 600 Gb/s) at 600 Gb/s and 400 Gb/s, and 400G generation (i.e., 28 nm DSP, 45 or 56 Gbaud) at 400 Gb/s.



Figure 16: ICE6 capacity-reach comparison

However, wavelength capacity-reach claims need to be backed up with demonstrations and trials. Factors that can impact demonstrable wavelength capacity-reach include the amount of required margin, the number of wavelengths and channel spacing, the available fiber path length, span distances, amplifier types, fiber type, fiber quality/age, the number and types of ROADMs, and the add/drop structure. For this reason, ICE6 has demonstrated its ability to maximize wavelength capacity-reach across a range of scenarios and network environments, as shown in Table 2.

Date	Demo/Trial	Data Rate	Reach
March 2020	Lab Demo with Corning Corning G.654.E TXF fiber	800 Gb/s	800 km
March 2020	North American Network Third-party optical line system with live traffic	800 Gb/s	950 km
June 2020	Windstream, U.S. G.652 fiber	800 Gb/s	730 km
		700 Gb/s	1,460 km
July 2020	Verizon, U.S. G.655 LEAF fiber	800 Gb/s	667 km
		600 Gb/s	2,283 km
		400 Gb/s	4,091 km

Table 2: ICE6 public demos and trials as of July 2020

Summary

Increased wavelength capacity-reach can have a dramatic impact on optical network total cost of ownership. Infinera’s 800G generation optical engine, ICE6, maximizes wavelength capacity-reach with an ultra-high baud rate (up to 96 Gbaud); innovative features including Nyquist subcarriers, long-codeword PCS, and dynamic bandwidth allocation; and a high modem SNR leveraging high-performance individual components and holistic co-design, both enabled by deep vertical integration.