WHITE PAPER

THE ULTIMATE GUIDE TO HIGHER BAUD RATES

Since its emergence in the late 2000s, coherent technology has undergone several transitions. First it went from 40G to 100G, then from hard-decision forward error correction (FEC) to soft-decision FEC, and finally from single-rate interfaces to flexible-rate interfaces with additional modulation schemes such as 8QAM (150 Gb/s) and 16QAM (200 Gb/s). With no end in sight to the traffic growth from video, cloud, and DCI, and with the full impact of 5G yet to be felt, network operators need to scale capacity cost-effectively while minimizing power consumption and footprint. To address this need, optical vendors are evolving coherent optics technology with ever-higher baud rates, first to ~60-70 Gbaud with 600G generation coherent, and now to ~90-100 Gbaud with 800G generation coherent, with even higher baud rates expected to follow.

DRIVERS FOR COHERENT TRANSPORT EVOLUTION

The need to contain CapEx and OpEx is driving the evolution of coherent technology, as IP traffic is growing at 26% year-over-year and peak data rates are growing even faster at around 40%, according to the Cisco VNI (2/2019). Transceiver cost per bit, for both the endpoints and any regeneration points, is the most direct driver of ongoing optical network CapEx, while power consumption and footprint are the primary drivers of network operational costs. The key to reducing these costs is achieving more bits per second for a given reach requirement, with the same hardware investment in the transceiver's ASIC/DSP, analog electronics, photonic components, and packaging. An additional key driver for evolving coherent optics is the need to transport higher-speed client interfaces, including 400 GbE, cost-effectively with a single wavelength over a wide range of network scenarios, including terrestrial long-haul and trans-oceanic submarine.

HIGHER BAUD RATES VS. HIGHER-ORDER MODULATION

These market drivers are leading optical equipment and component vendors to evolve coherent transport to support higher-capacity wavelengths. They have two primary levers for doing this. One lever is the baud rate, otherwise known as the symbol rate, which is the number of symbols per second. The other is the modulation scheme, which determines the number of bits per symbol. Each has its advantages, as shown in Table 1.

		Higher Baud Rates	Higher-order Modulation
Reach at 2x Data Rate		~80%-100%	~25% (QPSK->16QAM)
Increased Spectral Efficiency		x	✓
Lower Cost per Bit		✓	*At Shorter Distances*
Lower Power and Footprint		✓	*At Shorter Distances*
Fewer Wavelengths to Manage		✓	*At Shorter Distances*
Applications	Flexible-grid Networks	✓	✓
	75 GHz/100 GHz Point-to-Point	✓	✓
	50 GHz Point-to-Point	*Limited*	✓
	100 GHz Mesh ROADM	✓	✓
	50 GHz Mesh ROADM	x	✓

TABLE 1: *Higher baud rates vs. higher-order modulation*

Increasing the wavelength capacity with the baud rate has far less impact on reach than increasing it with higher-order modulation. Higher baud rates therefore offer the best potential for reducing cost per bit, power, and footprint. By enabling higher-capacity wavelengths, higher baud rates also simplify operations, thus reducing OpEx, as network operators have fewer wavelengths to provision and manage for a given amount of optical bandwidth.

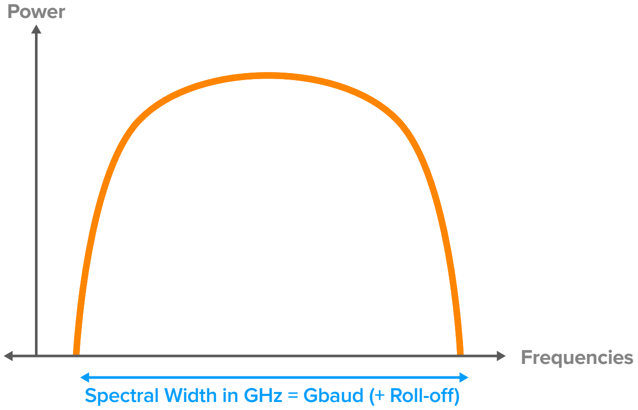

***FIGURE 1**: Spectral width is proportional to baud rate*

As shown in Figure 1, the spectral width of a wavelength is directly proportional to the baud rate, with the exact spectrum equal to the baud rate plus a percentage of "roll-off," which is a function of its spectral shape. For example, a 32 Gbaud wavelength with 10% roll-off occupies 35.2 GHz. Higher baud rates are therefore typically unable to increase spectral efficiency and can even decrease it to the extent that they decrease reach. This also limits the role of higher baud rates in 50 GHz fixed-grid networks, as will be discussed in more detail later.

Higher-order modulation, on the other hand, does not change the spectral width of the wavelength, so it is able to increase spectral efficiency and is equally applicable to fixed-grid and flexible-grid DWDM networks. However, it comes at a significant cost in terms of reach, and therefore its benefits are largely limited to shorter distances.

Doubling Wavelength Capacity: Modulation vs. Baud Rate

To understand the relative advantages of higher baud rates and higher-order modulation, let us look at what happens if we double wavelength capacity, starting with a 100G wavelength based on QPSK modulation. The exact baud rate is also impacted by the FEC overhead and the exact spectrum is impacted by the roll-off, but for the purposes of this example, we will use 32 Gbaud and 35 GHz as the base baud rate and spectral width for 100G wavelengths based on QPSK.

	QPSK	16QAM
Bits per Symbol (per Polarization)	2 ---- x 2 ---->	4
Constellation Points	4 ---- x 4 ---->	16

TABLE 2*: QPSK vs. 16QAM: bits per symbol and constellation points*

One option for 200G is to double the bits per symbol on each of the two polarizations by going from QPSK (two bits per symbol per polarization) to 16QAM (four bits per symbol per polarization). However, as shown in Table 2, this increases the number of constellation points by a factor of four, making the wavelength approximately four times (~6 dB) more sensitive to noise and nonlinearities, such as four-wave mixing, self-phase modulation, and cross-phase modulation. The reach is therefore decreased to approximately 25% of the 100G QPSK wavelength's, as shown in Figure 2. However, ignoring any grid constraints, the spectral width of the channel stays the same, so we have doubled the spectral efficiency.

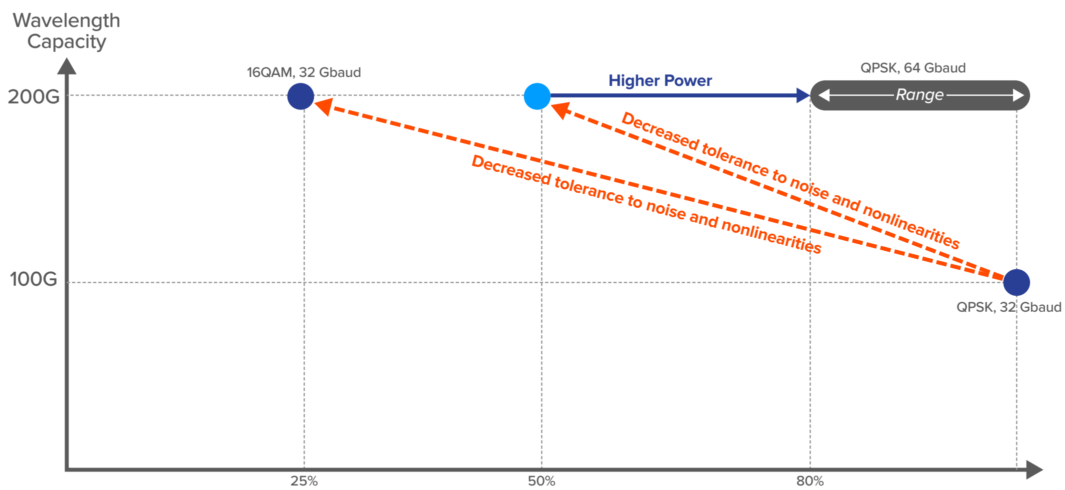

FIGURE 2*: Impact of doubling wavelength capacity*

On the other hand, if we go to 200G by doubling the baud rate to 64 Gbaud and keeping the modulation at QPSK, we double the sensitivity to noise, as well as to nonlinearities, which by itself would reduce the reach to 50%. However, because the spectrum of the wavelength has doubled from 35 GHz to 70 GHz,

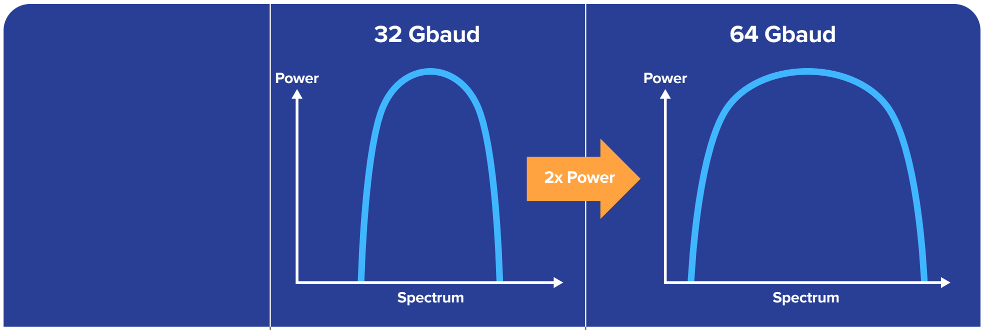

	32 Gbaud	64 Gbaud
Total Power per λ (dBm)	1x	2x (+3 dB)
Spectrum (Hz)	1x (35 GHz)	2x (70 GHz)
Power Spectral Density (dBm/Hz)	Same	Same
Nonlinearities Generated	Same	Same

__TABLE 3__: Twice the baud rate enables twice the power

we can double the power (+3 dB) for the same power spectral density, and therefore the same quantity of nonlinearities generated, as shown in Table 3.

Doubling the power halves our sensitivity to OSNR, which typically puts the reach of the 64 Gbaud 200G wavelength in the range of 80-100% of the original 32 Gbaud 100G wavelength, depending on the specifics of the route/link, including the type of add/drop and fiber. The reach of 200G based on 64 Gbaud QPSK is therefore between 3.2 and four times the reach of the 200G based on 32 Gbaud and 16QAM. The price to be paid for this increase in reach is a loss of spectral efficiency relative to the higher-order modulation option, as the 64 Gbaud QPSK wavelength has half the spectral efficiency of the 32 Gbaud 16QAM wavelength.

HIGHER BAUD RATES AND 800G

As the industry moves to the 800G generation of coherent technology, the question becomes, what baud rate do we need for 800G wavelengths? Well, 800G with 20% FEC can be achieved with less than 84 Gbaud. However, for the sake of simplicity, we will use 84 Gbaud and 64QAM as our starting points in the examples shown in Table 4.

So, if we can get 800G with full 64QAM and <84 Gbaud, the next question becomes, is there any value to

Modulation	64QAM	PCS-64QAM	PCS-64QAM	32QAM
Baud Rate (Gbaud)	84	90	96	100.8
Bits per Symbol per Polarization	6	5.6	5.25	5
Spectrum (GHz, 10% Roll-off)	92.4	99	105.6	110.9

__TABLE 4__: 800G examples: baud rate and modulation

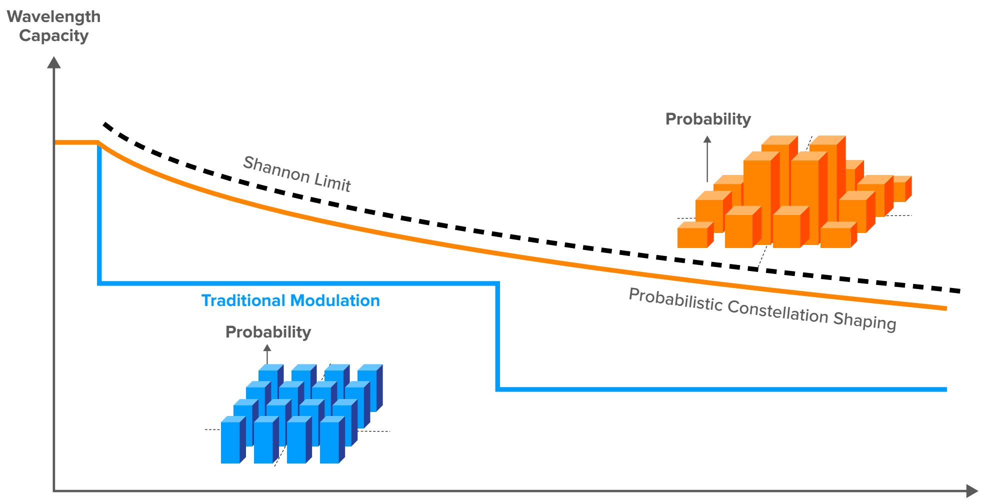

***FIGURE 3**: Probabilistic constellation shaping*

baud rates that exceed this figure? The answer to that question is yes. Increasing the baud rate beyond 84 Gbaud will let us use lower-order modulation, which increases reach. For example, 32QAM has 5 bits per symbol per polarization and would require 100.8 Gbaud for 800G, with reach of at least double that of 64QAM. But if we can't get to 100.8 Gbaud, do higher baud rates still have value? The answer to that is still yes, with a little help from probabilistic constellation shaping (PCS).

PCS reduces the probability of higher-power outer constellation points, while increasing the probability of lower-power inner constellation points. It provides extremely granular control of the number of bits per symbol and can improve OSNR tolerance, lifting the capacity-reach curve closer to the Shannon limit, as shown in Figure 3.

As we increase the baud rate beyond 84 Gbaud, PCS lets us reduce the probability of the outer constellation points, which increases OSNR tolerance (at constant power) and extends reach, with the exact improvement also dependent on factors such as the PCS codeword length. 96 Gbaud will therefore be able to achieve significantly better 800G reach than 84 Gbaud. However, because the higher baud rate increases the spectrum of the wavelength, a 96 Gbaud wavelength will require just over 14% more spectrum than an 84 Gbaud one.

HIGHER BAUD RATES AND NYQUIST SUBCARRIERS

Though the increased spectrum of high-baud-rate wavelengths enables higher power for the same number of nonlinearities generated, higher baud rates are more sensitive to nonlinearities at the receive end. Nyquist subcarriers can mitigate this increased sensitivity by leveraging advanced digital signal processing to divide a single high-baud-rate carrier into multiple lower-baud-rate subcarriers, as shown in Figure 4. This can improve reach by approximately 10%.

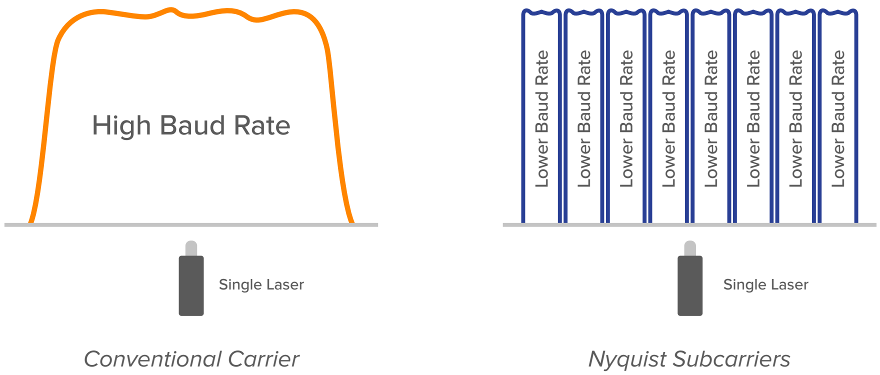

***FIGURE 4**: Nyquist subcarriers*

Another benefit of Nyquist subcarriers is reduced chromatic dispersion. As shown in Figure 5, 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. As the spectral width of the signal is proportional to the baud rate, a high-baud-rate signal has a bigger gap between its lowest and highest frequencies and therefore experiences greater variation in the speed of its frequencies

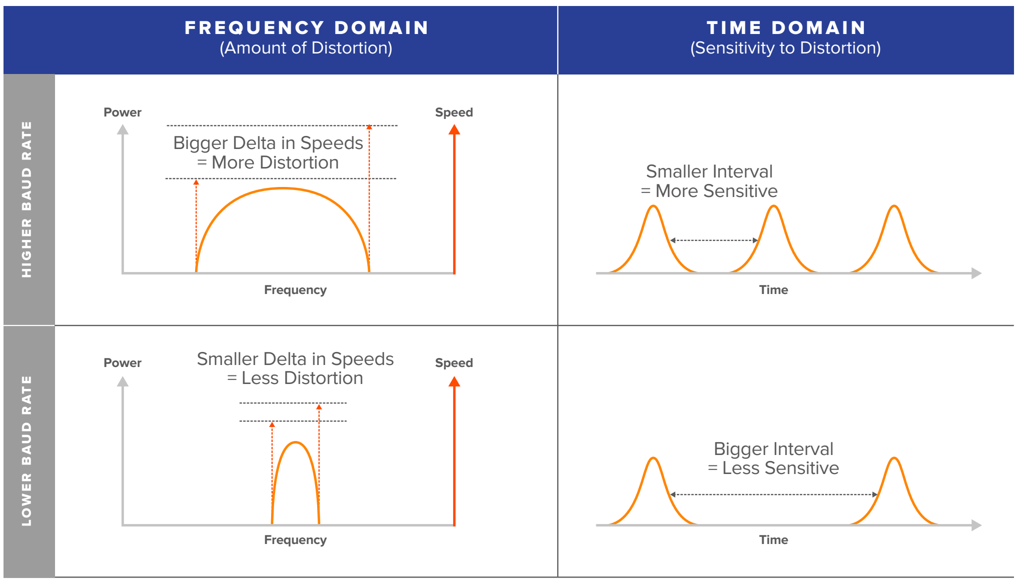

***FIGURE 5**: Higher baud rates and chromatic dispersion*

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 16 with four subcarriers and by a massive factor of 64 with eight subcarriers. Even if the single-carrier chromatic dispersion is within the capabilities of the DSP, compensating chromatic dispersion has a cost in terms of additional noise, more specifically phase noise. Reducing chromatic dispersion has a significant benefit in terms of reducing this noise and therefore improving performance.

LOWER BAUD RATE USE CASES

While the highest possible baud rate is normally the best strategy for minimizing cost per bit, power, and footprint while simplifying operations with fewer wavelengths, there are still a number of scenarios in which lower baud rates are either required or have advantages. This creates a need for coherent transceivers with the flexibility to support a wide range of baud rates.

Use Case 1: 50 GHz Fixed-grid ROADM Networks

Higher baud rates have limited applicability to 50 GHz fixed-grid ROADM networks for a couple of reasons. First, as discussed previously, increasing the baud rate results in a proportional increase in the spectral width of the channel. Second, many wavelength selective switches (WSSs) have a limited passband for each channel, and this passband reduces significantly as the number of WSSs in the path increases, as shown in Figure 6.

A typical 50 GHz WSS has a passband of around 46 GHz per channel. However, any differences in width, center wavelength, or shape of the passband will cause the effective width to reduce in an effect called filter narrowing. After only 10 WSSs (e.g., a colored/directional add, four route-and-select ROADMs, and a colored/directional drop), the effective width would be reduced by around 10 GHz to around 36 GHz, limiting the baud rate to around 33 Gbaud. In 50 GHz mesh ROADM networks, the only real role for higher baud rates is increasing FEC.

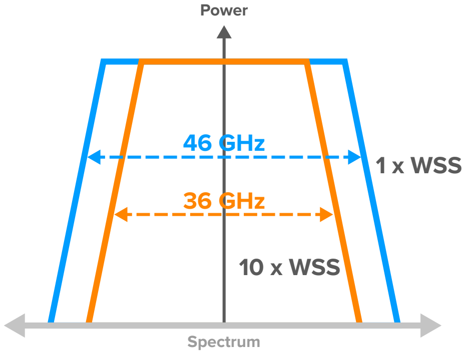

FIGURE 6: *Filter narrowing and WSS cascades*

Use Case 2: Fixed-grid Point-to-Point Networks

50 GHz mux/demux filters based on a 50 GHz interleaver typically have a passband of around 46 GHz per channel, similar to many 50 GHz WSSs. 400G wavelengths with 64QAM and around 43 Gbaud can therefore be used for point-to-point 50 GHz fixed grid applications. Here, the higher baud rate also increases spectral efficiency as less spectrum is wasted, as shown in Figure 7.

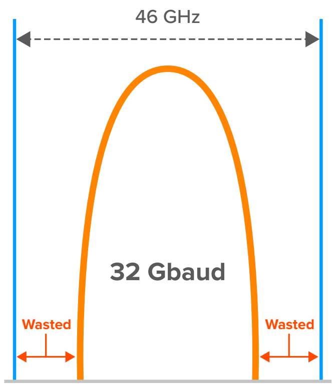

***FIGURE 7**: 50 GHz Fixed-grid point-to-point is limited to ~43 Gbaud*

In the same way, 75 GHz mux/demux filters will limit the baud rate of the wavelength to approximately 70 Gbaud, enabling 600G wavelengths, while 100 GHz mux/demux filters will limit the baud rate to approximately 90 Gbaud, which is sufficient for 800G wavelengths.

Use Case 3: Better Alignment with Available Flexible-grid Spectrum

There may be cases where a lower baud rate can better align with the available spectrum. For example, if a 150 GHz block of spectrum is available between the two endpoints, a super-channel with two 68 Gbaud wavelengths would deliver better spectral efficiency than a single 96 Gbaud wavelength, as shown in Figure 8.

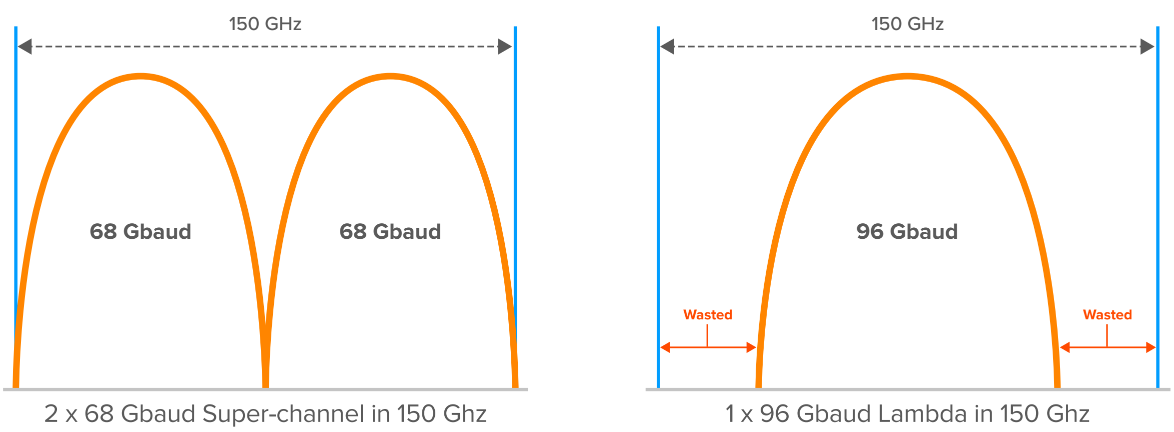

***FIGURE 8**: Lower baud rates can align better with available spectrum*

Use Case 4: Maximum 800G Spectral Efficiency

As discussed previously, 800G at baud rates beyond the minimum required based on the full 64QAM can dramatically increase reach. However, when the reach requirement can still be met by the full modulation (i.e., 64QAM), then dialing down the baud rate to the minimum will deliver the best spectral efficiency, as shown previously in Table 4.

Use Case 5: Simplified Flexible-grid Operations

A lower baud rate may align better with operationally simplified flexible-grid granularity. For example, a 90 Gbaud wavelength would likely fit into 100 GHz, while a 96 Gbaud wavelength might require 112.5 GHz, based on 12.5 GHz flexible-grid granularity, as shown in Figure 9. For some operators, managing spectrum in 100 GHz increments may be preferable to the alternative of 112.5 GHz increments.

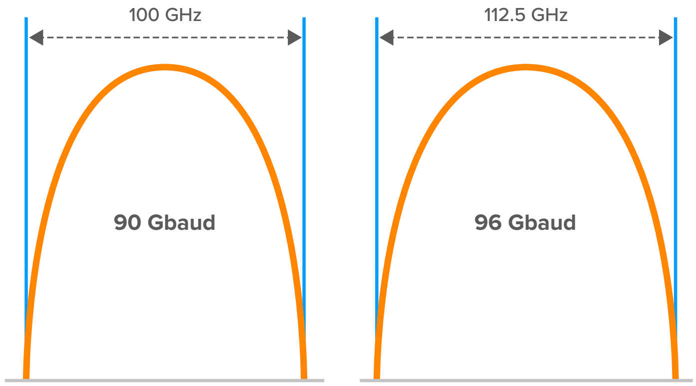

***FIGURE 9**: Simplified Flexible-grid Operations with 90 Gbaud and 100 GHz*

Use Case 6: Trans-oceanic Submarine and Chromatic Dispersion

As discussed previously, baud rate and chromatic dispersion have a squared relationship. With trans-Pacific fibers of up to 14,000 km, and 96 Gbaud having nine times the chromatic dispersion of 32 Gbaud, trans-oceanic chromatic dispersion may require lower than maximum baud rates.

SUMMARY

Increasing the baud rate enables a proportional increase in the data rate of a wavelength with minimal impact on its reach. Higher baud rates therefore provide a key lever for reducing cost per bit, power consumption, and footprint over a wide range of distances, while also reducing OpEx with fewer individual wavelengths to provision and manage. Nyquist subcarriers can address two key challenges of high baud rates: increased sensitivity to nonlinearities and the squared relationship between baud rate and chromatic dispersion. Meanwhile, PCS provides the option to “overclock” the baud rate while dialing down the number of bits per symbol in order to maximize reach. However, there are several use cases for lower than maximum baud rates, thus requiring next-generation coherent transceivers with baud rate tuneability.