WHITE PAPER

Infinera

The Three Myths of Optical Capacity Scaling

How to Maximize and Monetize Subsea Cable Capacity

MYTH 1

Scaling **BAUD RATE** increases spectral efficiency, thereby increasing total subsea fiber capacity

BUSTED

MYTH 2

Scaling **MODULATION ORDER** increases spectral efficiency, thereby increasing total subsea fiber capacity

PLAUSIBLE

MYTH 3

Scaling **INTEGRATION** has no effect on spectral efficiency or subsea fiber capacity

BUSTED

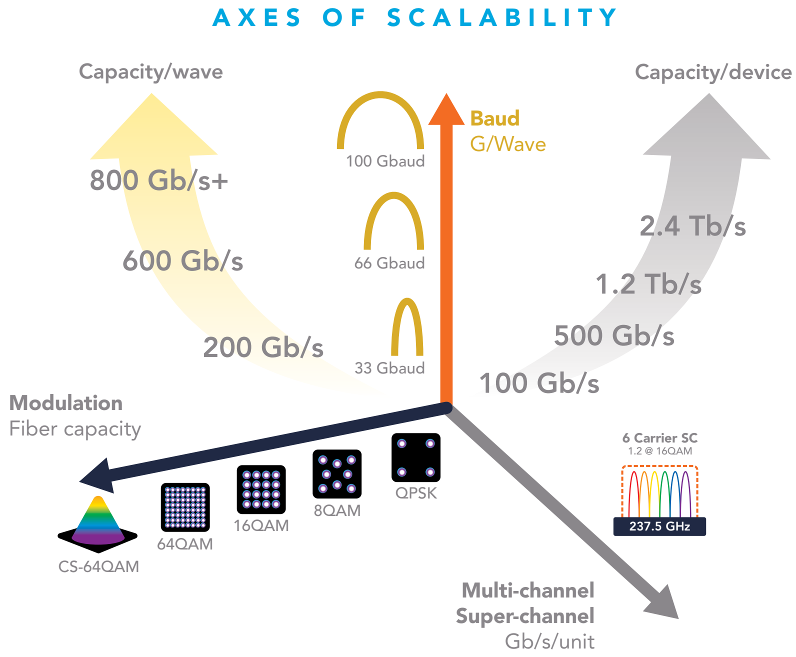

Figure 1: The Three Axes of Optical Scaling

The Three Axes of Optical Scaling

Figure 1 shows the three axes that are needed for optical scaling—baud rate, modulation order and integration. This paper will discuss scaling along each of these axes, but it is important to clarify exactly what is being scaled. There are at least three interpretations:

- Increasing the data rate per wavelength
- Increasing the capacity per line card or appliance
- Increasing the total fiber capacity, or spectral efficiency

These can potentially be scaled independently, or in concert as transponder technology evolves. In subsea deployments total fiber capacity, or spectral efficiency, is usually the dominant factor in determining overall network economics, so it becomes extremely important to understand the most effective way to achieve total capacity, especially when scaling baud rate has become a focal point for many dense wavelength-division multiplexing (DWDM) vendors today.

Scaling Baud Rate

The baud rate for an optical transmitter is the rate at which symbols are sent. For example, polarization-multiplexed (PM) 16 quadrature amplitude modulation (16QAM) carries 8 bits in each symbol, and at a baud rate of 33 gigabaud (GBaud) the wavelength data rate is 200 gigabits per second (Gb/s)—note that the additional bits are required for Optical Transport Network (OTN) framing and forward error correction (FEC).

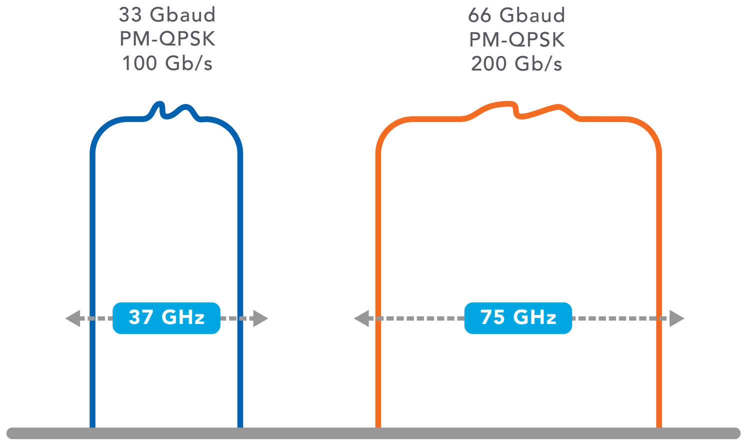

Figure 2: Double the Baud Rate for a Given Modulation Order Doubles the Data Rate per Wavelength, but Also Doubles the Spectral Width, Resulting in No Change in Total Fiber Capacity

Doubling the baud rate to 66 GBaud doubles the wavelength data rate to 400 Gb/s. This axis of scaling is useful for a number of reasons. First is that doubling the transponder data rate effectively doubles the "productivity" for an engineer installing it—in other words, they can bring twice the capacity into service in a given period of time. Factors like cost and power consumption will also trend down over time, and density will improve.

These are all excellent advantages of scaling the baud rate, but one thing that does not change is the spectral efficiency of this wavelength. This is explained in Figure 2. On the left is a pulse-shaped 33 GBaud signal, which has a spectral width around 35 to 37 gigahertz (GHz).

On the right of Figure 2 we see that the 66 GBaud signal occupies exactly twice the spectral width of the 33 GBaud signal. This is purely a matter of optical physics, and it means there is no improvement in spectral efficiency, and thus total fiber capacity, from increasing baud rate alone. From an optical impairment point of view, increasing the baud rate can have a negative impact, as described later.

The Value of Subcarriers

It's tempting to assume that Moore's Law will help drive higher baud rates, but for the past 10 years the serial processing rate for high-performance electronics has remained essentially constant. It is the same in the computing industry, where central processing unit (CPU) speeds plateaued some years ago. In applications like CPU or graphics processing unit applications, increases in processing power have been delivered by adding cores—scaling in parallel, rather than using a serial approach.

A single-carrier transponder would require higher serial processing rates for the optoelectronics that drive it, and the only practical way to achieve this is by moving to smaller feature sizes for the application-specific integrated circuits—from 40 nanometers (nm) to 28 nm, 14 nm and soon to 7 nm node sizes—for practical power dissipation. However, increasing rates to deliver higher baud still doesn't improve total system capacity.

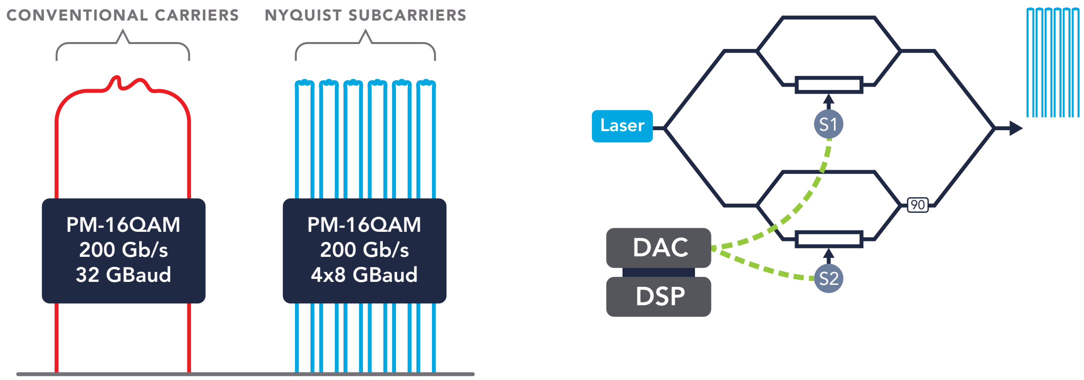

Figure 3: Nyquist-shaped Subcarriers Occupy the Same Spectral Width and Deliver the Same Aggregate Baud Rate as a Conventional Carrier, but Each Subcarrier Can Be Modulated Independently. The Transmitter Circuit, Including the DSP and DAC, Is Shown on the Right.

A parallel carrier approach has been pioneered by Infinera, creating multiple subcarriers from a single laser output using the capabilities of the coherent transmission technology in the fourth-generation Infinite Capacity Engine (ICE4) chipset. In this case, the subcarriers are created by the transmitting digital signal processor (DSP) and digital-to-analog converter (DAC). Thanks to the sophisticated pulse shaping algorithms included in ICE4, these subcarriers can be very tightly spaced, and we refer to them as Nyquist subcarriers.

Each subcarrier can be independently modulated by a separate parallel circuit within the transmitter, thus allowing the effective baud rate to scale using parallel processing. Subcarriers result in better optical performance because, for a given effective baud rate, more processing power (and thus more optical impairment compensation) can be applied to the signal compared to a single-carrier approach.

Figure 4 shows another reason that subcarriers can give a huge boost to subsea transmission. The horizontal axis shows the baud rate, while the vertical axis shows the "distortion penalty"—in this case, the magnitude of the two most common nonlinear penalties, self-phase modulation

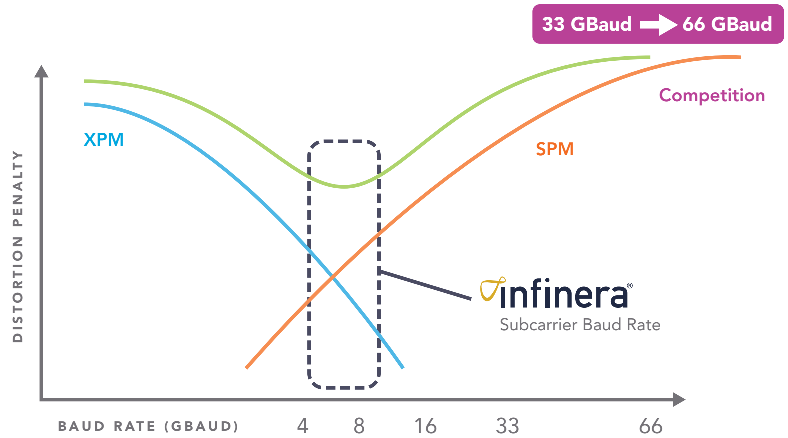

Figure 4: Baud Rates Between 4 GBaud and 12 GBaud Experience Lower Nonlinear Penalties

(SPM) and cross-phase modulation (XPM). The blue plot shows that XPM decreases as baud rate increases, while the red plot showing SPM does the opposite. The green plot is the compound distortion penalty, which shows a minimum distortion penalty around 8 Gbaud—note that the exact value of the minimum varies with fiber type but has been observed to be between 4 GBaud and 12 GBaud.

Infinera's competitors tend to focus on increasing baud rate on a single wavelength—which at best delivers no improvement on fiber capacity—and then invest power and complexity in trying to correct the resulting nonlinear impairments. Infinera's pioneering use of subcarriers takes the opposite approach: don't create the nonlinear penalty in the first place, and spend less complexity and power correcting the smaller impairments that result. This approach increases reach and total system capacity. Note that the number of subcarriers can be increased as the aggregate baud rate rises, maintaining this optimum performance to ever higher wavelength capacities.

In Summary for Myth #1: Scaling Baud Rate…BUSTED

Increasing the baud rate for a line card or appliance may increase the capacity per wavelength and eventually reduce the power per bit. But the spectral efficiency, and therefore total fiber capacity, at best stays the same. Nyquist subcarriers are a vital capability to optimize optical performance as baud rates increase, and can actually increase total system capacity.

Scaling Modulation Order

Phase modulation techniques are well established in optical transmission, and Infinera's was the first commercial system to allow the modulation type to be selected in software to optimize the capacity and reach for a given fiber span with FlexCoherent® modulation.

Scaling modulation order will certainly increase fiber capacity, but in subsea networks it may not be that simple. To understand why this is, consider Table 1.

We see that PM-binary phase-shift keying (BPSK) carries 2 bits per symbol, PM-3QAM carries 3, PM-quadrature phase-shift keying (QPSK) carries 4 and so on. If these signals are modulated at the same rate, let's say 33 GBaud, the signals will occupy the same spectral width, but a higher-order modulation symbol will carry more bits. Thus, the spectral efficiency, and therefore the total fiber capacity, will increase with modulation order.

Modulation	Bits per Symbol	Optical Power per Bit	Fiber Capacity Factor	Reach Factor
PM-BPSK	2	1	1	>150
PM-3QAM	3	0.67	1.5	100
PM-QPSK	4	0.5	2	65
PM-8QAM	6	0.33	3	24
PM-16QAM	8	0.25	4	14
PM-32QAM	10	0.2	5	2
PM-64QAM	12	0.17	6	1

Table 1: A Comparison of Modulation Order Characteristics

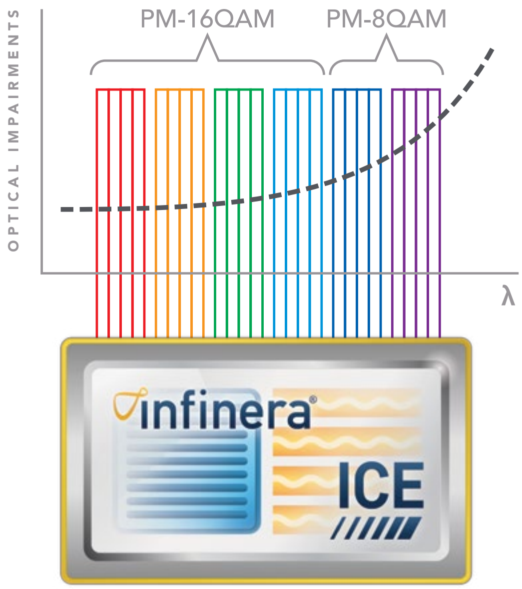

Figure 5: Different ICE4 Carriers Can Operate with Different Modulation Types per Carrier in Order to Optimize Fiber Capacity

However, optical fiber is a nonlinear medium, with effects such as SPM, XPM and four-wave mixing resulting in serious nonlinear transmission penalties. This means that there is a maximum threshold for the power of an optical symbol before it incurs a nonlinear penalty. If we assume that each of the signals in Table 1 are transmitted at a power level that is just below the nonlinear threshold, we can see that increasing the modulation order will decrease the optical power factor per bit, and exponentially reduce the optical reach. As a comparison, PM-64QAM can deliver three times the fiber capacity compared to PM-QPSK, but with about 65 times shorter reach!

Infinera's FlexCoherent technology allows the individual wavelengths generated by the photonic integrated circuit (PIC) to be configured to use the optimum modulation. Figure 5 shows the six wavelengths generated by an ICE4 module. In this example, the first four wavelengths are operating in a fiber region of relatively low nonlinear penalties and are thus able to use PM-16QAM. But the last two wavelengths are in a region of higher non-linear penalties and need to be modulated using PM-8QAM. The ability to select the modulation per wavelength allows the maximum fiber capacity to be extracted.

The challenge in subsea deployments is that cable length tends to be extremely long. This means that while increasing the modulation order could certainly increase fiber capacity, these higher-order modulations may not have the necessary optical reach to close the span.

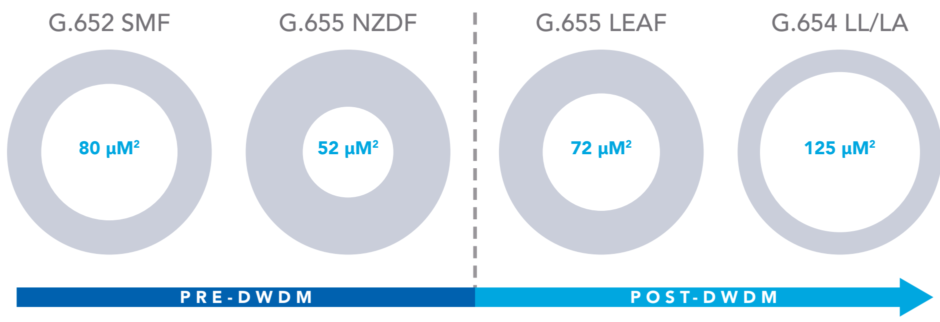

Figure 6: Before the Advent of DWDM, Fiber Evolution Resulted in Reduced Effective Area in Fiber Cores. In the DWDM Era, the Pressure Was to Reduce Nonlinear Penalties by Increasing the Effective Area of the Fiber Core.

Evolution of Optical Fiber

Nonlinear penalty is proportional to the optical power density in a fiber or the optical power per unit area of fiber, so the higher optical powers needed for longer reach drive up nonlinear impairments for a given diameter of fiber core. Optical fiber manufacturers have responded to this problem by evolving the design of optical fiber. In Figure 6, we see that the original G.652 fiber had an effective core area of about 80 square micrometers (µm2). The move to G.655 non-zero dispersion-shifted fiber (NZSDF) resulted in a much smaller effective area, which means that early G.655 fibers tend to have higher nonlinear penalties than G.652. Large effective area fiber (LEAF) was developed to mitigate this problem, but LEAF still has a smaller effective area than G.652. Now that chromatic dispersion can be fully compensated for by digital signal processing, fiber manufacturers do not need to limit the total dispersion in the fiber, and can now focus on reducing power density by increasing core area independent of dispersion. Contemporary systems employ pure silica core fibers with very large effective areas exceeding 125 µm2 to reduce nonlinear penalty.

In Summary for Myth #2: Scaling Modulation Order…PLAUSIBLE

Increasing the modulation order will certainly increase spectral efficiency and thereby fiber capacity, but it will also reduce the reach exponentially. Modern large area fiber types can help support longer reach for higher-order modulation, but scaling modulation order alone may not be particularly effective at increasing capacity in older subsea cables.

Scaling Photonic Integration

The first two axes of scaling are available to all DWDM vendors, but only Infinera has access to a third dimension—the ability to integrate the optical circuits for multiple wavelengths on a single chip, a PIC. This axis can be scaled in addition to the first two axes of increased baud rate and higher-order modulation, and scaling the number of wavelengths on the same line card or appliance will clearly increase operational scalability and make installation engineers more productive. It also unlocks the potential for software defined capacity, as described in the next section of this paper. But why would it have any effect on fiber capacity?

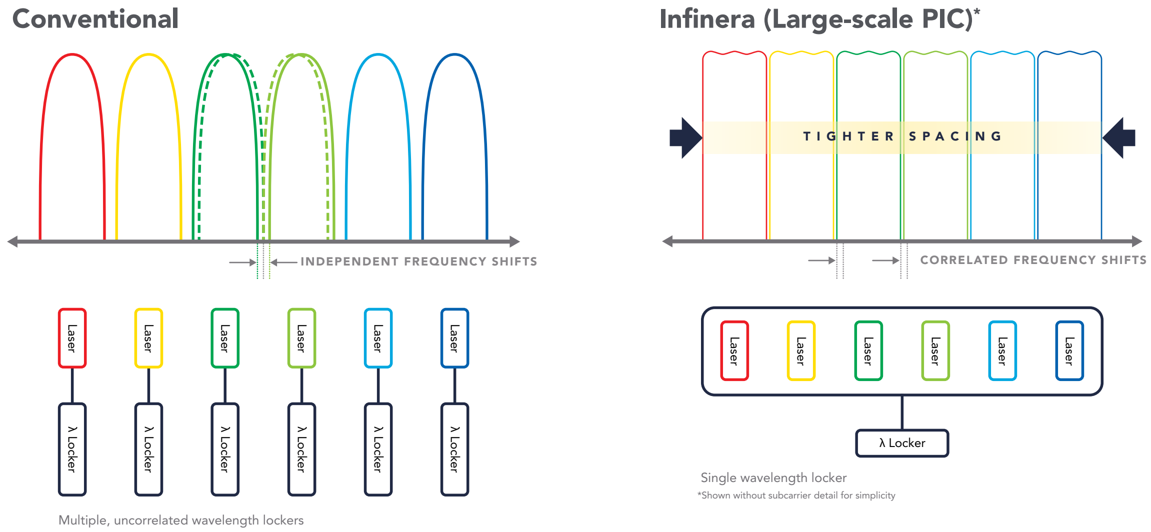

Figure 7: Tighter Channel Spacing Thanks to Correlated Wavelength Locker in PIC Implementation

Tighter Channel Spacing

The left side of Figure 7 shows a set of six conventional single-wavelength transponders that are multiplexed into the subsea fiber. In this configuration, there is a wavelength locker circuit on each transponder card, which is used to keep the transponder locked to its operational wavelength. Each wavelength locker operates independently, which means that each individual carrier may drift to the "left" or "right" on the wavelength axis independently. For this reason, network designers must include a specific tolerance, in the form of wider channel spacing, so that each wavelength will have a permitted drift within the fiber spectrum. These drifts are caused by ambient temperature variations in points of presence and landing stations, which can be a significant challenge in both terrestrial and subsea deployments.

The right side of Figure 7 shows the configuration of a PIC-based design, in which a single wavelength locker is used to monitor and control all six wavelengths on the line card or appliance. Drift may still occur in this wavelength locker, but in this case all six carriers will drift together, either right or left. The result is that there is no need to operate with large wavelength tolerances, enabling the delivery of significantly higher fiber capacity.

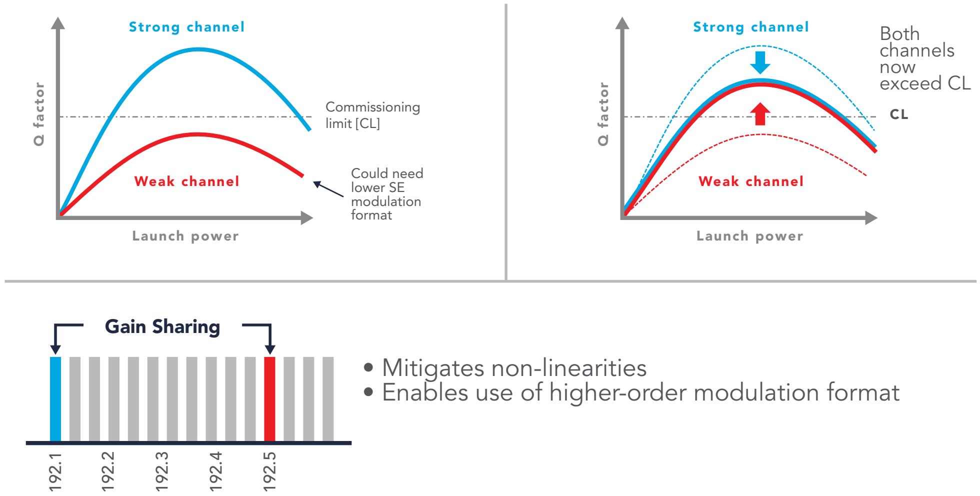

Figure 8: On the upper left we see a traditional, multi-transponder approach in which each wavelength is processed by a separate FEC chip. Excess FEC margin from the blue wavelength cannot be "donated" to the red wavelength.

On the right is a dual-channel approach in which the excess FEC margin of the blue wavelength can be used to pull up the red wavelength so that both are now above the commissioning limit.

SD-FEC Gain Sharing

FEC has been widely used in DWDM transmission for almost two decades and has proven to be extremely effective in recovering errors in received data streams. However, in conventional implementations, the data from a single optical carrier is processed by a given FEC chip. A further advantage of implementing multiple carriers on the same line module or appliance is that multiple data streams can be passed through the same FEC processor. Figure 8 shows the advantage of this approach.

On the left side of Figure 8 we see two carriers, red and blue. The service provider would like to deploy both carriers using PM-8QAM modulation, but the red carrier is operating in a region of high impairments and the FEC cannot recover sufficient errors to make this possible. The blue carrier is operating in a more benign region, and there is actually an excess of FEC gain available to the PM-8QAM modulation. However, in single-carrier implementations this excess gain simply goes to waste—it cannot be used to assist the red wavelength.

Fortunately, a PIC-based multi-wavelength implementation can leverage this excess gain since the two signals are processed by the same FEC engine. This allows the FEC margin to be shared, which effectively "pulls up" the red wavelength using the excess FEC gain from the blue

wavelength, allowing both to be above the PM-8QAM commissioning limit. Thus, SD-FEC gain sharing can achieve higher capacity over a longer distance, or avoid compromising capacity for a given cable distance. Note that these shared wavelengths do not have to be in close proximity—they can be spaced anywhere across the extended C-band.

In Summary for Myth #3: Scaling Photonic Integration...BUSTED

The conventional thinking is that large-scale PICs have no impact on subsea fiber capacity. However, by taking advantage of tighter channel spacing enabled by a common wavelength locker, and by allowing excess gain from favorable carriers to be shared by less favorable carriers, fiber capacity can be increased significantly.

Real-world Proof

In a series of subsea deployments and trials since late 2017, Infinera's ICE4-based subsea solution has proven itself to deliver between 30 and 50 percent more capacity than our competitors over the same cable systems. In the case of the Seabras-1 cable, over a distance of 10,500 kilometers, ICE4 achieved 18.2 terabits per second (Tb/s) for a record-breaking spectral efficiency of 4.5 bits per second per hertz with commercially deployable margins.

The fiber capacity capabilities of Infinera's ICE4 are being extended with almost every field trial and deployment we make. We publish our latest results in press releases and on our blog, so please refer to those to learn about new trials and deployments as they are announced.

Summary

Total capacity is the primary factor driving the economics and competitiveness of a submarine cable business model. Increasing baud rate has no effect on fiber capacity, and while moving to higher-order modulation will drive more capacity, this comes at the expense of optical reach. Scaling integration level has been shown to enable higher spectral efficiency through tighter channel spacing, even at trans-oceanic distances. Trials and deployments since late 2017 and underway today continue to demonstrate these advantages, delivering world-record capacities.

Figure 9: Summary of ICE4 Performance over the Seabras-1 Cable System

Infinera Corporation
140 Caspian Court
Sunnyvale, CA 94089 USA
Telephone: +1 408 572 5200
Fax: +1 408 572 5454
www.infinera.com

Have a question about Infinera's products or services?
Please **contact us**.

Specifications subject to change without notice.

Document Number: 0182-WP-RevA-0519