WHITE PAPER

infinera®
what THE NETWORK will be

Super-Channels DWDM Transmission Beyond 100Gb/s

Will Bandwidth Growth Ever Stop? No.

A number of different industry surveys indicate that total internet demand is growing at about 40% per year. This growth is driven mainly by increasing video traffic in the network—Netflix now takes up to 30% of the internet's bandwidth at peak hours, and new competitors like Amazon, Hulu, Youku, and the BBC iPlayer are growing rapidly. This growth is now further accelerated by mobile access, with video clients shipping on an all smart phones and tablets, enabling video to be consumed more conveniently via network connections anywhere, anytime.

Furthermore, the video and rich media files are migrating into the network, stored in cloud architecture such as Microsoft Azure or Apple's iCloud, which tend to produce significant replication of data for resilience and performance purposes, and in which every user experience, whether consumer or enterprise, accesses the network. This provides an opportunity and a threat for service providers. Those that can provide the best user experience in this dynamic and media-rich mobile environment can capture market share. But they must be able to scale their networks dramatically, and do so while lowering capital and operational costs per gigabit per second. The place to start is with the Transport Network, which forms the foundation of long distance internet communication. It is clear that in addition to a move toward larger, more powerful transport switches, the mechanisms of DWDM optical transmission have to change too. A new approach to DWDM capacity—the super-channel—promises an effective solution to the challenges posed by internet growth.

Figure 1: Video, Mobility and Cloud are helping to drive around 40% annual growth in internet demand

What is a Super-Channel?

DWDM is an innovation that enables multiple optical carriers to travel in parallel in a fiber to more efficiently use the expensive fiber assets that have been installed over thousands of kilometers in the ground and on aerial poles. Many in the industry view the "state of the art" in DWDM in 2012 and 2013 to be 100 Gb/s. However, the growth in the internet is demanding additional bandwidth scale without increasing operational complexity, and asking the question, "What is beyond 100 Gb/s?" One answer is a super-channel, an evolution in DWDM in which several optical carriers are combined to create a composite line side signal of the desired capacity, and which is provisioned in one operational cycle. Later, this paper will explain why it becomes more practical to combine multiple carriers into a super-channel to move beyond 100 Gb/s than it is to simply increase the data rate of an individual carrier. However, from the point of view of the client-side services that use them, such as a 10G, 100G or 1Tb Ethernet connection from a service provider to an enterprise customer, super-channels are indistinguishable from a single carrier channel of the same data rate.

In the past, transponder-based DWDM line systems operated with a fixed relationship between client data rates and line side data rates (e.g. a 10 GbE service would be carried by a 10G DWDM carrier). However, another innovation that is offering new levels of flexibility is software selectable modulation formats. This new capability is critical to allowing service providers to trade off DWDM reach (how far the signal can travel) with underlying optical transport data rate (raw bandwidth capacity), which in turn requires that this fixed relationship be broken. It is now essential to be able to map client data rates services (<10GbE, 10GbE, 100GbE and beyond) into any rate of super-channel through a dynamic abstraction or virtualization mechanism.

What Problems Will Super-Channels Address?

Super-channels address three fundamental issues:

- Scaling bandwidth without scaling operational procedures
- Optimizing DWDM capacity and reach
- Supporting the next generation of high speed services

Moreover, it must be possible to implement super-channel technology using existing DWDM engineering techniques, without requiring longer term technology advances.

First, as mentioned above, service providers are faced with the need to turn up ever more capacity without any real increase in the number of engineers to do the work. With the introduction of Infinera's Digital Optical Network in 2004, service providers have been able to turn up 100 Gb/s in a single operational procedure. In 2010 other vendors in the industry introduced 100 Gb/s optical capacity. However, in the meantime internet demand has grown at 40% year over year and is driving more and more service providers to say, "I need more than 100 Gb/s, and I need it now." By offering 300G, 400G, 500G or even a terabit of operational capacity, super-channels allow network engineers to keep pace with demand without increasing OpEx.

> "Bringing up multiple wavelengths [in the form of a super-channel] on the same card means less patching of fibers so there are fewer opportunities for something to go wrong.
>
> Over the past few years we've increased our lit kilometers across Europe by a factor of four—but our OpEx costs have only risen by about one point five. Infinera has had a part to play in this success story and it creates a cost advantage we can pass on to our customers to help us win business."
>
> **Matthew Finnie**
> **CTO, Interoute**

Another problem is the need to achieve this scale without compromising reach and, in fact, with the ability to flexibly trade off capacity for reach based on the particular application and optical route. The most recent innovation that enables a practical increase in bandwidth with long haul reach is coherent detection —allowing 40G and 100G carriers to operate over the same or even greater reach compared to 10G Intensity Modulation Direct Detection (IMDD, also known as Intensity Non-Return to Zero (NRZ) or On/Off Keying. See " A Primer on Modulation" for further explanation). Super-channels will help move beyond 100 Gb/s but must do so while maintaining long haul reach, and the multi-carrier approach appears to be the obvious choice. A flexible coherent (FlexCoherent) modulation capability will be essential in achieving these goals.

The first generation "split spectrum" super-channels (See: "Implementing Super-Channels") will offer around 8 Tb/s in the C-band—between five and 10 times the spectral efficiency offered by 10G IMDD systems that use a traditional 50 GHz grid. When service providers are ready to move outside the current ITU DWDM fixed grid spacing (defined in ITU G.694.1), the next generation of "gridless" super-channels will offer even higher spectral efficiency, enabling a gain of up to 25% in net fiber capacity through more efficient spectrum use.

Finally, super-channels will allow service providers to support the next generations of Ethernet, video and storage area network client-side services that operate at data rates in excess of 100 Gb/s (which is the highest data rate of Ethernet defined in IEEE802.3ba). At this stage it's not clear what that data rate will be (the two main candidates being 400 Gb/s and 1 Tb/s), so super-channels must be flexible enough to deal with this challenge when it arrives.

A Primer on Modulation

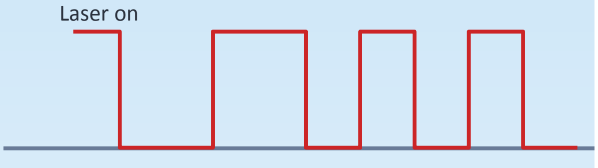

An IMDD Signal

Modulation is the process of imposing a digital signal onto a carrier in the analog domain—in this case a beam of light.

A remarkably simple and effective form of modulation has been used in the entire optical industry for decades. This is Intensity Modulation Direction Detection (IMDD)—sometimes also referred to as On/Off Keying (OOK). An IMDD signal encodes a single bit (a 1 or 0) in each symbol, with each symbol representing one cycle of a clock. It is simple to implement and uses very few optical components.

While optical fiber is an amazingly good material, in which light travels for hundreds or thousands of kilometres on the line side, the fiber impairments (such as attenuation, chromatic dispersion (CD), polarization mode dispersion (PMD), polarization dependent loss (PDL) and non-linearity) must be factored into network design.

Although IMDD is simple to implement, there are two problems. First, half of the transmitted power is in the carrier, which does not contain data information but results in 3 dB loss in signal to noise ratio (SNR). Second, an IMDD receiver is nonlinear because of the square-law detection, which makes it very inefficient to do electronic compensation of linear impairments such as CD and PMD. Note that CD is an all-pass transfer function with a parabolic phase against frequency and as such is proportional to the square of the symbol rate. The impairment due to PMD is proportional to the symbol rate.

So a 100G IMDD signal (which is ten times faster than 10G IMDD) actually experiences a significantly higher impairment level because it is very hard to compensate for PMD and residual CD. Broadly speaking, this is why 10G IMDD remained such a sweet spot for the optical transmission industry for over a decade.

The industry has responded by moving to modulation techniques that carry more bits in each symbol. Since each symbol is one cycle of a clock, the more bits there are per clock cycle, the more data can be throughput.

PM-QPSK

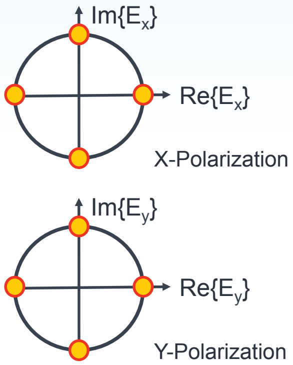

Polarization Multiplexed Quadrature Phase Shift Keying (PM-QPSK, also known as Dual-Pol QPSK, DP-QPSK) actually carries four bits per symbol (two in each polarization)—and so will experience significantly less impact from fiber impairments compared to an IMDD signal at the same data rate, because IMDD would have to run the clock 4 times faster to get the same throughput.

At the same time, coherent detection, in which a local reference laser is used to "tune" the receiver to the precise signal it is supposed to be receiving, results in far greater sensitivity, and the ability to more effectively reject noise from adjacent carriers. Note also that coherent detection is linear and all the linear impairments can be corrected using DSP algorithms/techniques.

Super-channels are critical in delivering a practical and efficient solution so that large service providers are able to move beyond 100 Gb/s optical transport rates as soon as possible, without requiring radical advances in engineering to be made. This will be covered in the next section.

Implementing Super-Channels

At the time of this writing, there are no finalized standards for super-channels, and therefore issues such as the number of carriers, their data rates, whether they consist of contiguous or non-contiguous carriers, and the level of component integration are all open to implementation decisions. In order to include wider industry efforts on super-channels, descriptions in this whitepaper will focus on the advantages of combining multiple carriers in order to solve all three fundamental problems described earlier, and delivering that solution in a timely manner.

There are two obvious implementation options for developing single-carrier transponders that operate at data rates above 100 Gb/s. One is to transmit more modulation symbols per second and the other is to encode more bits into a modulation symbol (or some combination of the two). Super-channel technology adds a third option—the ability to treat multiple carriers as a single operational unit.

For simplicity, this description will assume that a 1Tb/s unit of capacity is required.

Transmit More Symbols per Second

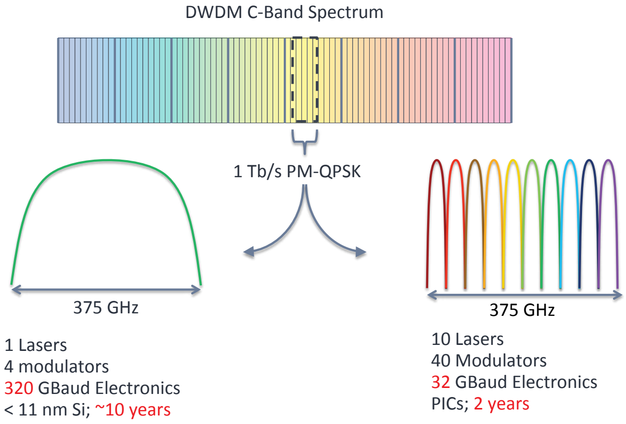

Figure 2: A single, 1Tb/s carrier is impractical because the electronics will not exist for another decade. In addition, the single carrier solution would experience significantly greater fiber impairments compared to a super-channel solution

The Primer on Modulation describes how we can encode data bits into modulation symbols for transmission. A 100G PM-QPSK transponder will send symbols at the rate of about 32 GBaud. Thus a single carrier, 1 Tb/s transponder could simply multiply the symbol rate by 10 for a total of 320 GBaud. This is shown on the left hand side of Figure 2.

There are two drawbacks in doing this. First is that the electronics that drive the interface would also need to operate at 320 GBaud—and this level of integrated electronics performance will probably not be available for another decade. Second is that the high symbol rate will encounter significantly higher implementation penalty with today's electro-optic technology for the same modulation type. Implementing a 1 Tb/s super-channel using 10 carriers (shown to the right of Figure 2) divides both the required electronics performance and the Baud rate on the fiber by a factor of 10, and so 32 GBaud electronics is all that is required. Moreover, because the symbol rate of each carrier is (in this example) the same as a 100G PM-QPSK transponder, the optical performance is more than adequate for most terrestrial long haul and ultra-long haul links. Where transpacific optical reach is needed, FlexCoherent modulation allows the super-channel to use PM-BPSK by simply changing a configuration parameter.

A key point is that both single carrier and super-channel implementations have the same spectral efficiency, but the super-channel has far better optical performance and is possible to build using technology that will be available in the near future.

> **"Generally, doubling the spectral efficiency results in one quarter the distance.**
>
> **Distance between cities is something we can't change…so we have to trade capacity for reach."**
>
> **Director of Network Design, Tier 1 Carrier**

Encode More Bits per Symbol

PM-QPSK encodes four bits per symbol—which is four times more than conventional IMDD modulation. The combination of this encoding efficiency, coherent detection and high-gain FEC technology allows a 100G signal to have the same or even better reach compared to 10G IMDD. So why not move to higher order modulation such as 8QAM, 16QAM or 32QAM? These techniques are widely used (at much lower data rates) in WiFi, cable modem and xDSL technology today. In addition, moving from PM-QPSK to PM-16QAM will halve the Baud rate required for the electronics, bringing a practical implementation forward in time.

Higher order modulation will certainly be a useful tool to allow service providers to optimize the total spectral efficiency for certain routes. But the penalty for this is reduced reach. Figure 3 illustrates part of the problem. It shows four common modulation types: BPSK, QPSK,

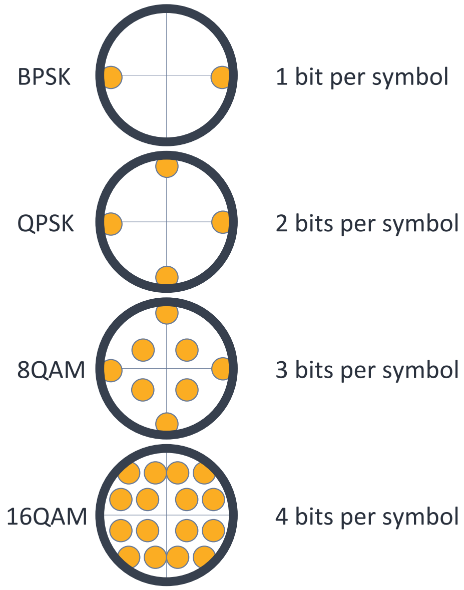

Figure 3: Adding more bits to a symbol increases spectral efficiency, but the total power per symbol (before non-linear threshold is reached) is indicated by the thick black circle.

8QAM and 16QAM. Notice the thick black circles. To a first approximation the area of these circles represents the maximum optical power carried by this symbol. If the power exceeds the non-linear threshold for this fiber path, then optical penalty will rise rapidly.

A pair of yellow circles represents an encoded bit—and the more of these that are put into the black circle, the less optical power is available per bit. Table 1 shows this numerically. In simple terms, modulation techniques like 16QAM may be limited to regional network use.

Modulation	Normalized Reach	C-Band Capacity (Split-Spectrum)	C-Band Capacity (Gridless)
PM-BPSK	5000 km	4Tb/s	5Tb/s
PM-QPSK	3000 km	8Tb/s	10Tb/s
PM-8QAM	1500 km	12Tb/s	15Tb/s
PM-16QAM	700 km	16Tb/s	20Tb/s
PM-32QAM	350 km	24Tb/s	30Tb/s
PM-64QAM	175 km	32Tb/s	40Tb/s

Table 1: Reach vs Total Capacity for a selection of phase modulation types. Note this table is illustrative only—not all modulation types shown are practical, and other modulation types may be available in final products.

Unfortunately, higher order modulation on its own is not a solution in moving to Terabit capacities—and there is no single correct modulation technique for any given route. Once again, this highlights the need for a multi-carrier super-channel implementation with FlexCoherent modulation so that service providers can optimize the combination of reach and spectral efficiency without having to order multiple part numbers from their system supplier.

The Importance of Photonic Integration

Super-channels allow a Terabit of DWDM capacity to be turned up in a single operational cycle, without any penalty in terms of spectral efficiency and with the same optical reach as today's generation of 100G coherent transponders.

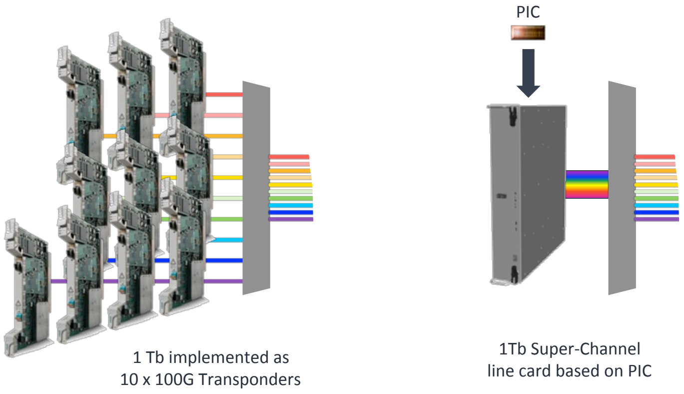

Figure 4: A super-channel built with a PIC. The PIC enables hundreds of optical functions to be collapsed into two small chips the size of a fingernail enabling, in this example, 10 x 100 Gb/s in a single line card.

It is clear that a 10 carrier super-channel requires 10 sets of optical components per line card. Implementing such an interface using discrete optical components would seem totally unrealistic, and Figure 4 shows the scale of the problem. On the left we see 10 individual 100G transponders. These will contain in total around 600 optical functions that are probably implemented in discrete optical components.

> **"Photonic integration is the optical industry's best hope for scaling to meet future bandwidth requirements while similarly reducing cost per bit."**
>
> **Sterling Perrin**
> **Senior Analyst–Heavy Reading**

On the right of Figure 4 is a Terabit super-channel line card. All of the major optical functions on all 10 100G line cards have been integrated into a single pair of Photonic Integrated Circuits (PICs)—one to transmit and one to receive. All 10 carriers can now be implemented in a compact line card where the super-channel is brought into service in one operational cycle, consuming far less power than 10 discrete transponders and resulting in far greater service reliability.

PICs bring the same kind of engineering practicality to super-channels that electronic integration brings to multi-core CPUs or graphics engines. Within reason, the more carriers are in the super-channel, the simpler the electronics and the better the optical performance. PICs remove the limitation of optical component complexity, and allow the right engineering balance to be chosen.

Flexibility is the Key to Super-Channel Success

We have seen that super-channels offer a number of benefits over single carrier implementations for line side DWDM transmission. In the example above a 1 Tb/s PM-QPSK super-channel was compared to a 1 Tb/s single carrier implementation. In the real world super-channels will have to be extremely flexible in a number of properties:

- What type of modulation should be used?
- What is the best way to optimize spectral efficiency and reach?
- What spacing will be used between the carriers?
- What is the total width of the super-channel?

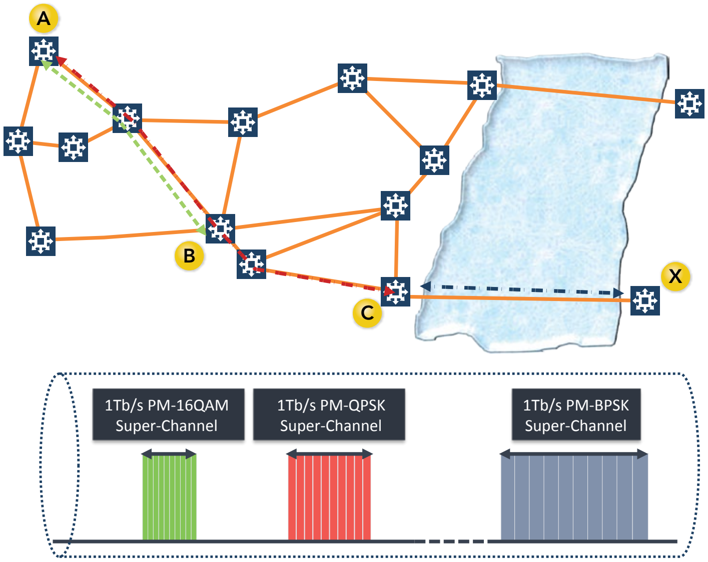

Figure 5: The ideal super-channel line card would offer flexible modulation to allow the service provider to trade reach for capacity without making spares management more complex.

There is no simultaneous answer to all of these questions. In fact an ideal super-channel implementation should allow, for example, a service provider to select the modulation type, the resulting spacing and the total super-channel capacity in software. Infinera's FlexCoherent modulation technology allows just such a capability—supporting up to six different coherent modulation types in a single line card to enable the optimization of link capacity versus optical reach based on the requirements of the application. Figure 5 shows the operational flexibility offered by FlexCoherent. In this case there are two super-channels that originate at Node A, and terminate at Nodes B and C respectively. The first path from A to B is relatively short (less than 700 kilometers) and can use PM-16QAM modulation and thereby trade the shorter reach to maximize spectral efficiency on that span. The second path is longer (less than 3,000 kilometers), and may require PM-QPSK, which reduces spectral efficiency in favor of optical reach to close it successfully. The Path from Node C to Node X is a long submarine link (more than 5,000 kilometers), and the service provider takes advantage of FlexCoherent modulation to use PM-BPSK, thereby trading spectral efficiency for reach. (Please note that in this example, the unlabeled intermediate nodes are all-optical from the point of view of the super-channels and offer no digital regeneration.)

The line cards used on all paths are identical—and the FlexCoherent configuration is chosen by the network operator in software.

Super-Channel Implementation Progress

In 2011 Infinera announced a series of super-channel trials, each one designed to showcase a specific aspect of PIC-based super-channel technology.

AUGUST 2011: THE SEACOM 500G SUPER-CHANNEL TRIAL **Infinera's trial with SEACOM, a leading pan-African telecommunications provider, involved lighting a 500G super-channel (five 100G PM-QPSK sub-channels) in South Africa over a 1,732 kilometer link between Johannesburg and Mtunzini in KwaZulu Natal. This was the first demonstration of a PIC-based super-channel operating as a single unit. Note that the trial also showcased FlexCoherent capability by switching between PM-QPSK and PM-BPSK modulation.**

OCTOBER 2011: CLOSING THE PACIFIC–THE PC-1 TRIAL **The most challenging fiber routes are those in subsea networks, and the Pacific Ocean is the most challenging of all. This 9,500 kilometer trial on the PC-1 cable from Japan to California not only showcased Infinera's ability to close Pacific routes with 100G PM-BPSK FlexCoherent modulation (with SD-FEC), it also showed the power of Bandwidth Virtualization by allowing 100 GbE client services to be sent over 40G PM-QPSK channels.**

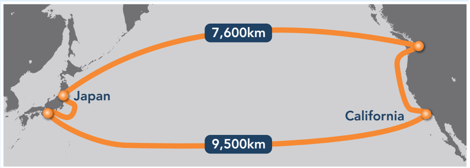

NOVEMBER 2011: THE TELIASONERA TERABIT SUPER-CHANNEL TRIAL
In November 2011 TeliaSonera International Carrier successfully completed the world's first Terabit optical transmission based on two 500 Gb/s super-channels. The trial was over a 1,105 kilometer route from Los Angeles to San Jose, California. This was a production fiber, and in addition to the Terabit of super-channel capacity, there was 300 Gb/s of 10G IMDD traffic on the route. This shows that a split-spectrum super-channel can operate very effectively with grid-based IMDD traffic.

Abstraction of Line Side and Client Side

As mentioned earlier, the move to FlexCoherent is expected to be an industry trend, giving service providers the ability to trade capacity for reach. The implication of this is that another major change in the industry is needed. Historically vendors have tied the client services data rate that connects to the enterprise or wholesale customer, such as 10 GbE or 100 GbE, to the line side (the long distance DWDM transmission side). In other words, a 100 GbE client service would be hard wired to a 100 Gb/s line side DWDM channel. However as we saw earlier in Table 1, changing modulation will not only impact reach but also data rate. Therefore the client side and the line side must be abstracted or virtualized. Currently there is a trend in the industry to standardize on 400 Gb/s as the next line side data rate because many believe that 400 Gigabit Ethernet will be the next Ethernet standard. However it is imperative that the line side standardization be defined to be flexible because if one has a 400 Gb/s line side solution at 16 QAM which is changed via software to be QPSK, it suddenly becomes 200 Gb/s. In fact, there are scenarios in which some of the super-channel carriers could be using QPSK and some could be using BPSK to deal with non-linear impairments.

Standards are needed that allow, and systems that deliver, an abstraction between the line side and client side. For example, Infinera has implemented this in what we call Bandwidth Virtualization. Bandwidth Virtualization will aggregate the available capacity on any link, regardless of how many carriers or what modulation is used, and allow all of that capacity to be used for any service. For example, let's say there are two 500 Gb/s super-channels on a particular link: super-channel one has 100 G available and super-channel two has 300 Gb/s available. Bandwidth Virtualization would allow the mapping of a 400 GbE service across these two super-channels. This functionality was proven with the Infinera DTN platform, and is available on the DTN-X platform. Bandwidth Virtualization is essential to ensuring the efficient use of network resources if operators wish to use software selectable modulation.

Migrating To Super-Channels: When to Go Gridless?

The ITU frequency grid defined in G.694.1 has been used in most of the world's DWDM networks for many years. However, some phase modulated carriers do not require such wide grid spacing as, for example, 50 GHz. In addition, multi-carrier super-channels do not require such rigid 50 GHz guard bands between carriers, and this fiber spectrum could be reclaimed by going "gridless".

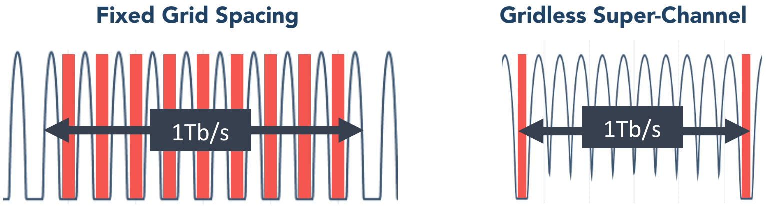

Figure 6: Split-Spectrum vs Contiguous Super-Channels. By removing the inter-channel " guard bands" a contiguous spectrum super-channel could occupy about 25% less fiber spectrum. A split-spectrum super-channel would be more compatible with existing grid-based transmission systems.

This is shown in Figure 6, in which about 25% of the fiber spectrum could potentially be reclaimed by a contiguous carrier super-channel that ignores the current ITU grid. This type of super-channel may well be incompatible with existing WSS ROADMs since these must be designed on a fixed grid (usually 50 GHz). The ITU is now working on a new "flex-grid" spacing (based on multiples of 12.5GHz) that would enable support for contiguous, or "gridless", multi-carrier super-channels. Since this work is still in progress, and service providers must be able to operate super-channel traffic along with existing 10G IMDD carriers in the same fiber, the first generation of super-channels will be "grid-based", or "split-spectrum". A split-spectrum super-channel solves all three key problems (scaling operational processes, maximizing fiber capacity and supporting next generation services), but it will not be able to "reclaim" the inter-carrier spectrum in such a flexible way—resulting in about 25% less total fiber capacity compared to a gridless approach.

A split-spectrum super-channel trial was recently completed by TeliaSonera International Carrier (see sidebar), in which interoperability with existing 10G IMDD, grid-based carriers was demonstrated.

Services Today and Tomorrow

The most common service type for most carriers today has a data rate of 10G—and this is forecasted to continue for the near future. So it's essential that Gigabit Ethernet, FiberChannel, 10 GbE, 40 GbE and 100 GbE services can all be supported efficiently in a super-channel architecture. This tends to imply the need for OTN switching in order to maintain efficient use of capacity across the network. Services will often originate at the network edge using lower data rate, single carriers (e.g. metro 10G), and this traffic pattern is likely to remain common in metro networks for several years. The services would then be aggregated using OTN multiplexing into super-channels for long haul transport. But the need for switching may not be limited to these aggregation points because as network meshing increases there will be frequent requirements to groom the super-channel en route. In this way super-channel capacity can be kept efficiently filled; and electronic protection and restoration techniques can be more effectively implemented.

Infinera already offers a capability like this in the form of Bandwidth Virtualization, and this will be enhanced as super-channels evolve to become the primary transmission mechanism in long haul networks.

Conclusion

Service providers are faced with compounded growth in demand for network capacity. Technological advances to date have allowed an increase in DWDM data rates to 100 Gb/s per optical carrier, but these technologies will soon be close their practical or theoretical limits. Super-channels offer a practical and timely solution to the problems of scaling operational processes, optimizing fiber capacity and reach, and supporting the next generation of client services beyond 100G.

Integrated OTN switching everywhere in the network allows highly efficient use of capacity, and the abstraction of client data rates from the super-channel data rate. These features lower both CapEx and OpEx, and allow new services to be turned up rapidly, and without complex design or operational intervention.

By implementing a FlexCoherent modulation technique, super-channels can offer a dramatic increase in spectral efficiency, or they can be configured to have ultra-long haul optical reach.

Large scale photonic integration is a critical technology in producing super-channel line cards that are compact, power-efficient and reliable.

Infinera Corporation
169 Java Drive
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 via the email addresses below.

Americas:	sales-am@infinera.com
Asia & Pacific Rim:	sales-apac@infinera.com
Europe, Middle East, and Africa:	sales-emea@infinera.com
General E-Mail:	info@infinera.com

www.infinera.com

Document Number: WP-SC-12-2011