Everything You Always Wanted to Know About Optical Networking – But Were Afraid to Ask

Richard A Steenbergen <ras@turkbergen.com>

Updated: May 1, 2017

Purpose of This Tutorial

Why give a talk about optical networking?

- The Internet as an industry is largely based around fiber.
- Yet many router jockeys don't get enough exposure to it.
- This leads to a wide variety of confusion, misconceptions, and errors when working with fiber optic networks.

Will this presentation make me an optical engineer?

- Maybe, but just remember, I omitted almost all the math. ☺
- The purpose of this tutorial is to touch on a little bit of every topic, from the mundane to the advanced and unusual.
- But it helps to have a basic understanding of how and why things work, even if you aren't designing fiber networks.

The Basics of Fiber Optic Transmission

What is Fiber, and Why Do We Use It?

Fiber is ultimately just a “waveguide for light”.

- Basically: light that goes in one end, comes out the other end.
- Most commonly made of glass/silica, but can also be plastic.

So why do we use fiber in the first place?

- Very low-cost to produce (silica is cheap).
- Extremely light (relative to copper), flexible material.
- Carries tremendous amounts of information (20 Tbps+ today).
- Can easily carry large numbers of completely independent signals over the same fiber strand, without interference.
- Can carry signals thousands of kilometers without regeneration
- Technology continues to radically improve what we can do with our existing fiber infrastructure, without digging or disruption.

Hold It Down Like I'm Giving Lessons in Physics

A quick flashback to High School physics class:

- Light propagating through a vacuum is (theoretically) the maximum speed at which anything in the universe can travel.
  - That speed is 299,792,458 meters per second, otherwise written as "*c*".
  - For doing shorthand math, you can round this up to 300,000 km/s.
- But when light passes through materials that **aren't** a perfect vacuum, it actually propagates much slower than this.
  - The speed of light in any particular material is expressed as a ratio relative to "*c*", known as that material's "refractive index".
  - Example: Water has a refractive index of "1.33", or 1.33x slower than "*c*".
- And when light tries to pass from one medium to another with a different index of refraction, a reflection can occur instead.
  - This is why you will see a reflection when you look up from under water.

Fiber Works by "Total Internal Reflection"

Fiber optic cables are internally composed of two layers.

- A "core", surrounded by a different material known as the "cladding".
- The cladding always has a higher "index of refraction" than the core.

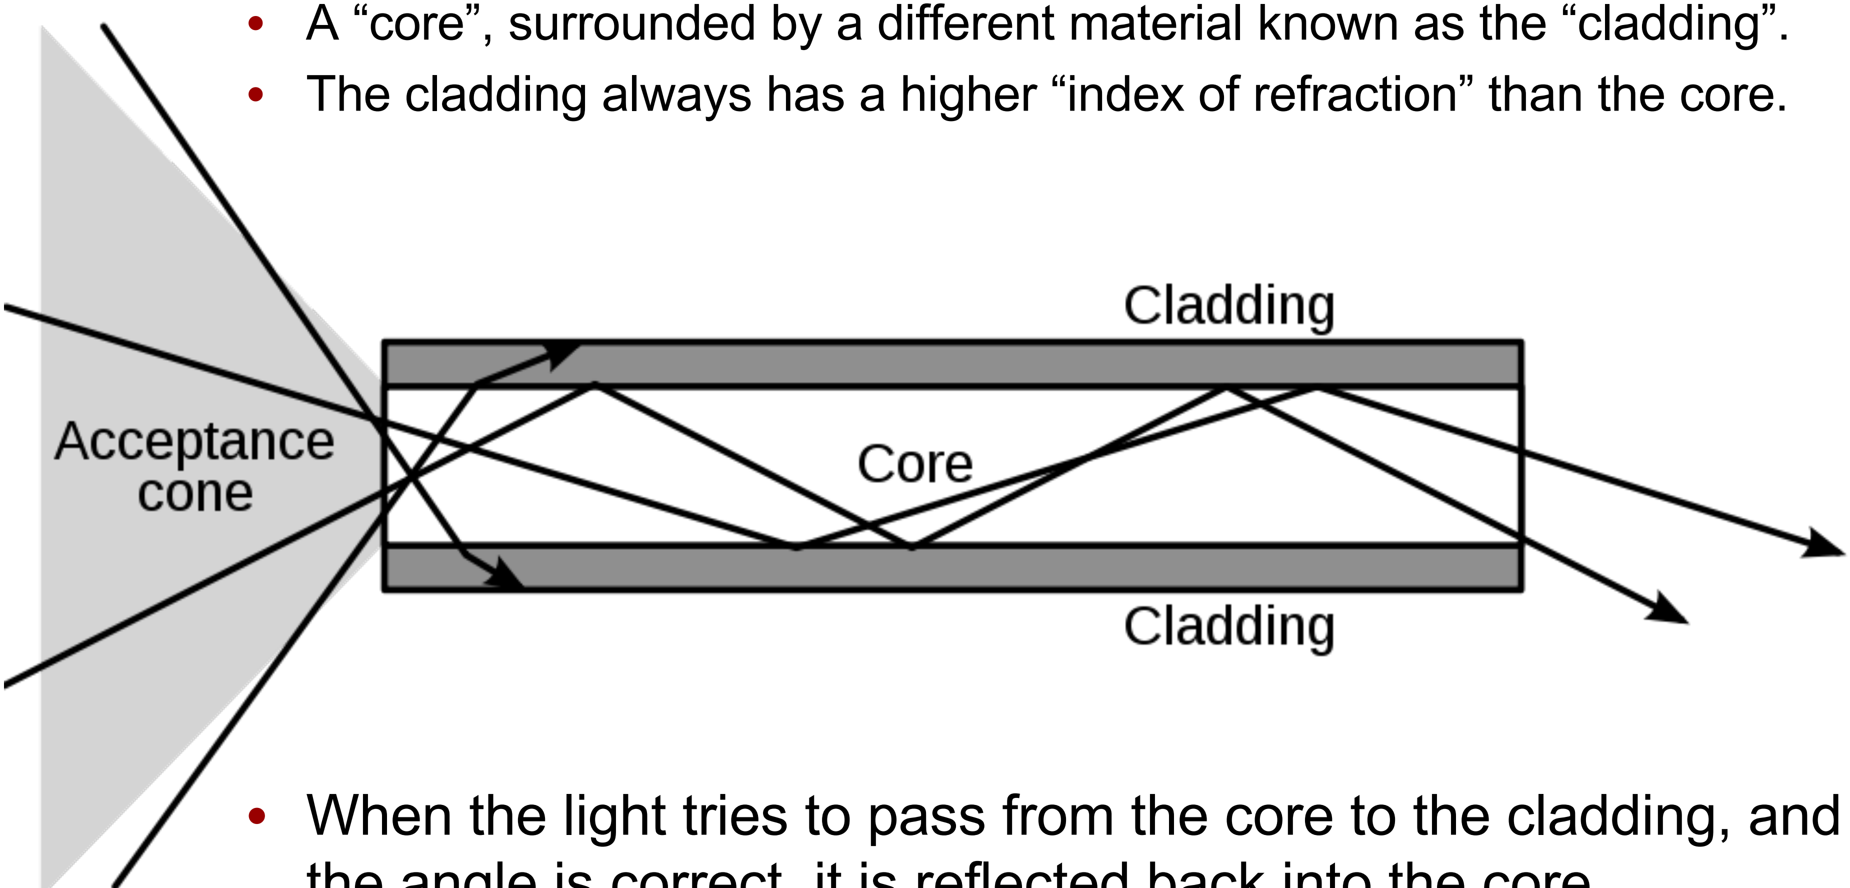

- When the light tries to pass from the core to the cladding, and the angle is correct, it is reflected back into the core.

Demonstration Using a Laser Pointer

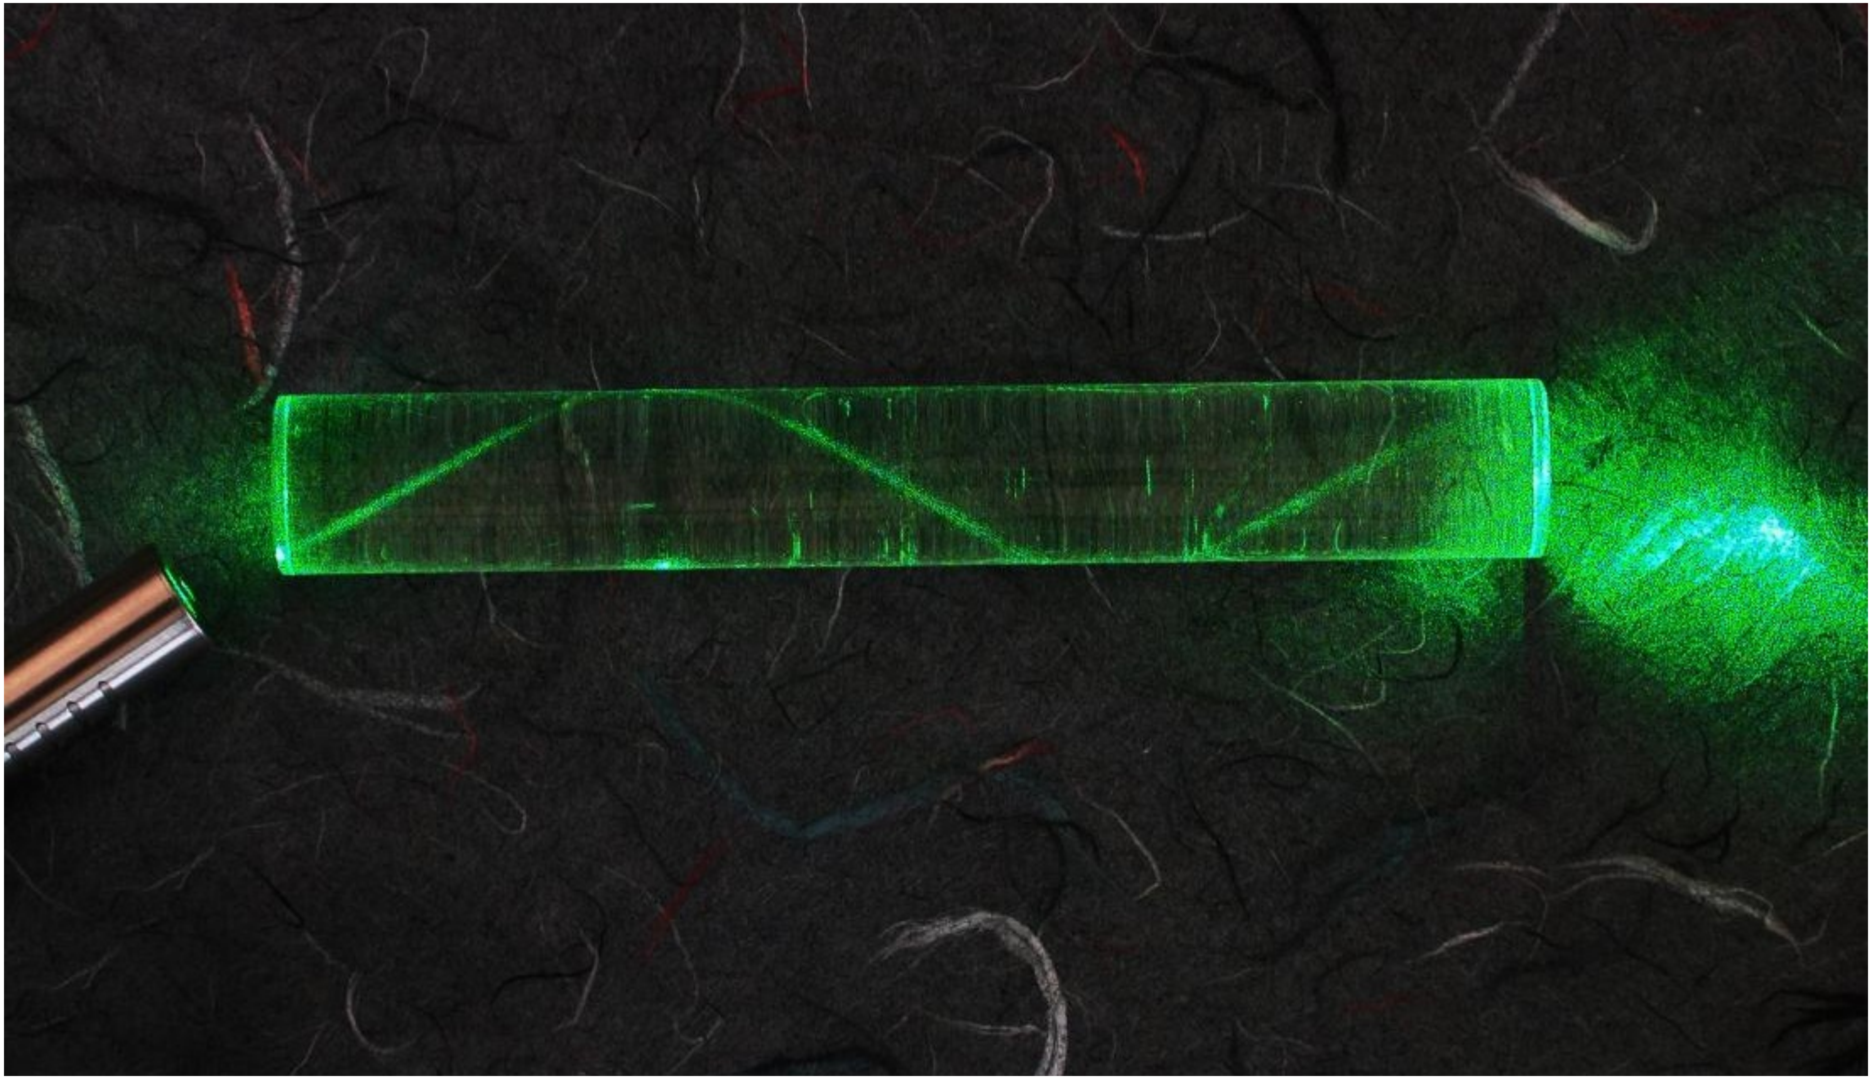

The Inside of a Common Fiber Cable

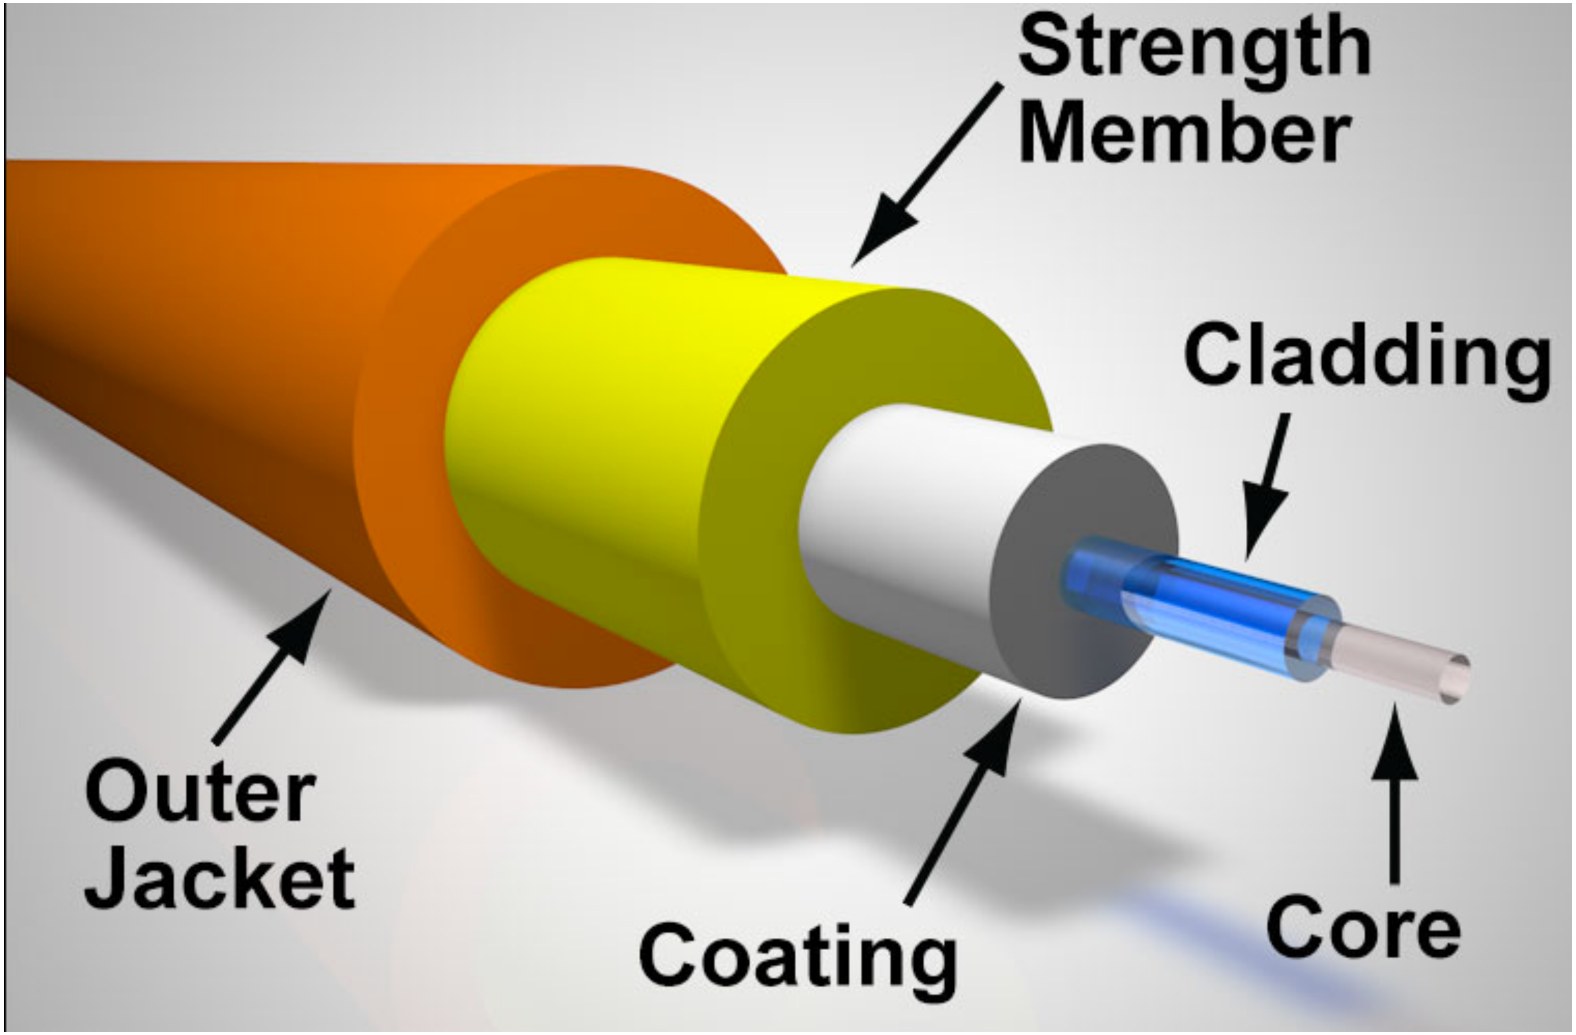

How Do We Actually Use The Fiber?

- The vast majority of deployed fiber optic systems operate as “duplex”, or as a fiber pair.
  - One strand is used to transmit a signal, the other to receive one.
  - This results in the simplest and cheapest optical components.
  - And usually holds true whenever the fiber is relatively cheap.
- But fiber is perfectly capable of carrying many signals, in both directions, over a single strand.
  - It just requires more expensive optical components to do so.
  - Which is generally reserved for when the fiber is expensive.
  - As with most things in business, cost is a deciding factor behind the vast majority of the technology choices we make.

What Do We Actually Send Over Fiber?

- Our digital signals must be encoded into analog pulses of light
  - The simplest (and cheapest) method is known as “IM-DD”.
    - Which stands for ”Intensity Modulation with Direct Detection”.
    - The most common version of which is “NRZ”, or “Non-Return to Zero”.
    - Which is really just a fancy way of saying “bright for a 1, dim for a 0”.
  - This modulation (called “baud”) can happen billions of times/sec.
    - The receiving end “sees” these flashes, and turns it back into 1s and 0s.
    - This technique was used for essentially all optical signals up to 10Gbps.
  - Beyond 10GBaud, this technique gets increasing hard to scale.
    - Today we barely squeak out 25GBaud, with FEC, for short reach links.
- But “better than NRZ” systems are becoming more pervasive.
  - QPSK – 16QAM is the basis for most long-haul links today.
  - PAM4 (Pulse Amplitude Modulation 4) is becoming viable for QSFP28 DWDM optics, and will be used in 400GE client optics.

What We Transmit Over Fiber

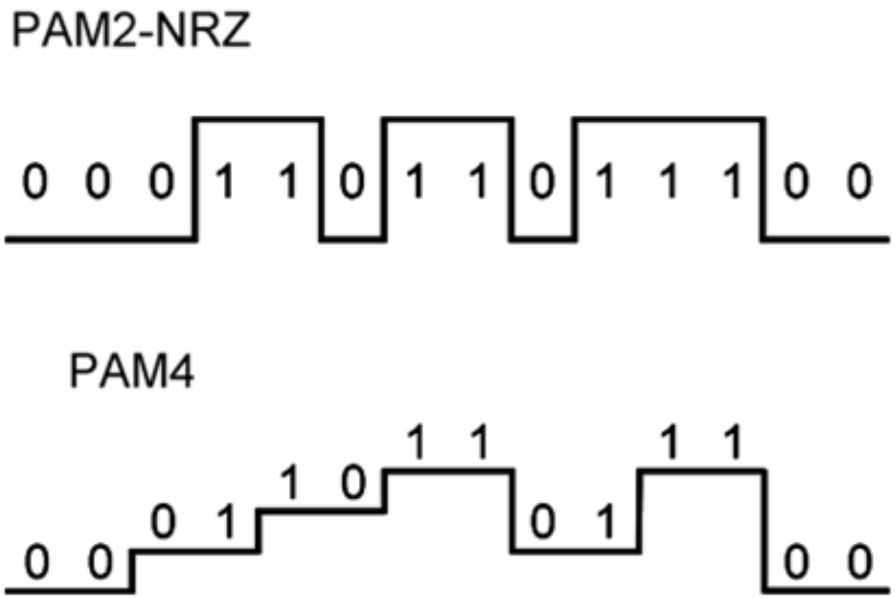

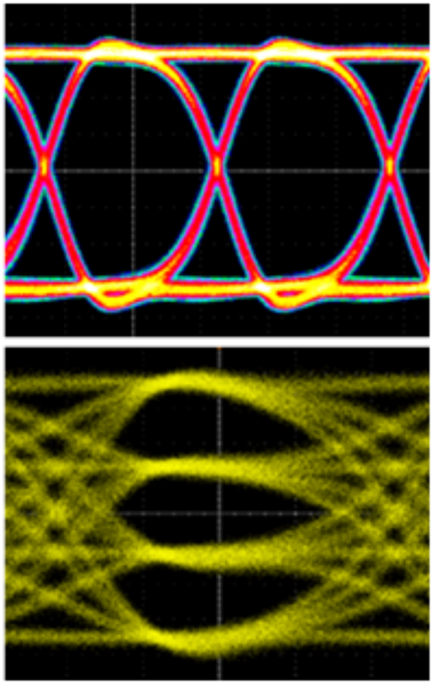

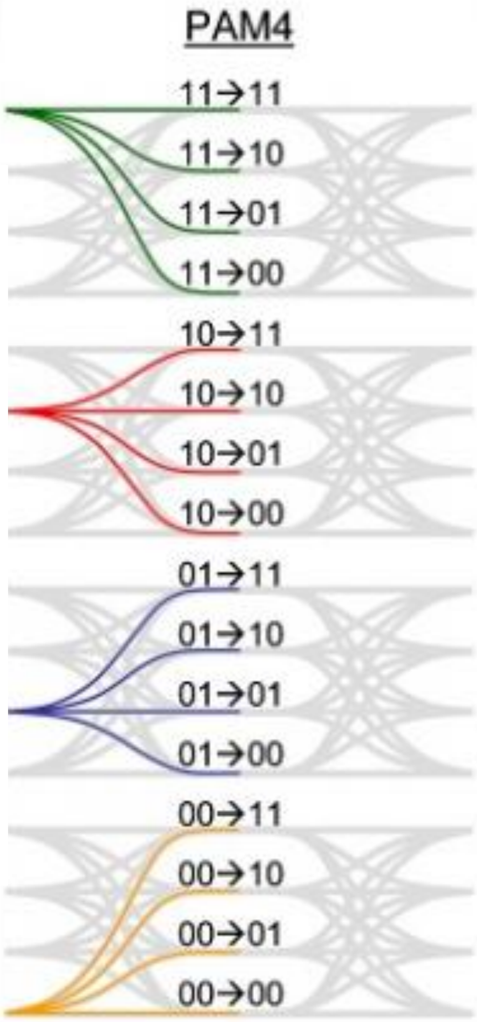

The Most Basic Distinction in Fiber: Multi-Mode vs Single Mode

Multi-Mode Fiber (MMF)

- Specifically designed for use with “cheaper” light sources.
  - Has an extremely wide core, allowing the use of less precisely focused, aimed, and calibrated light sources.
  - But this comes at the expensive of long-distance reach.
    - Fiber is so named because it allows multiple ”modes” of light to propagate.
    - “Modal dispersion” typically limit distances to “tens to hundreds” of meters.
- Types of Multi-Mode
  - OM1/OM2 aka “FDDI grade”: found with orange fiber jackets.
    - OM1 has a 62.5 micron (μm) core, OM2 has a 50μm core.
    - Originally designed for 100M/1310nm signals, starts to fail at 10G speeds.
  - OM3/OM4 aka “laser optimized”: found with “aqua” fiber jackets.
    - Specifically designed for modern 850nm short reach laser sources.
    - Supports 10G signals at much longer distances (300-550m, vs 26m on OM1).
    - Required for 40G/100G signals (which are really 10G/25G signals), 100-150m.

Single Mode Fiber (SMF)

- The fiber used for high bandwidths, and long distances.
  - Has a much smaller core size, between 8-10 microns (µm).
  - No inherent distance limitations caused by modal dispersion
    - Typically supports distances of ~80km without amplification.
    - With amplification, can transmit a signal several thousand kilometers.
- SMF has an even broader array of types than MMF.
  - Also has “OS1” and “OS2”, but they’re packaging, not fiber type.
    - OS1 ”tight buffered” for indoor use, OS2 “loose” to be blown into ducts.
  - ”Classic” SMF can be called “SMF-28” (a Corning product name)
  - But it comes in many different formulations of Low Water Peak Fiber (LWPF), Dispersion Shifted Fiber (DSF), Non-Zero Dispersion Shifted Fiber (NZDSF), Bend Insensitive Fiber, etc.

Single Mode vs Multi-Mode Fiber

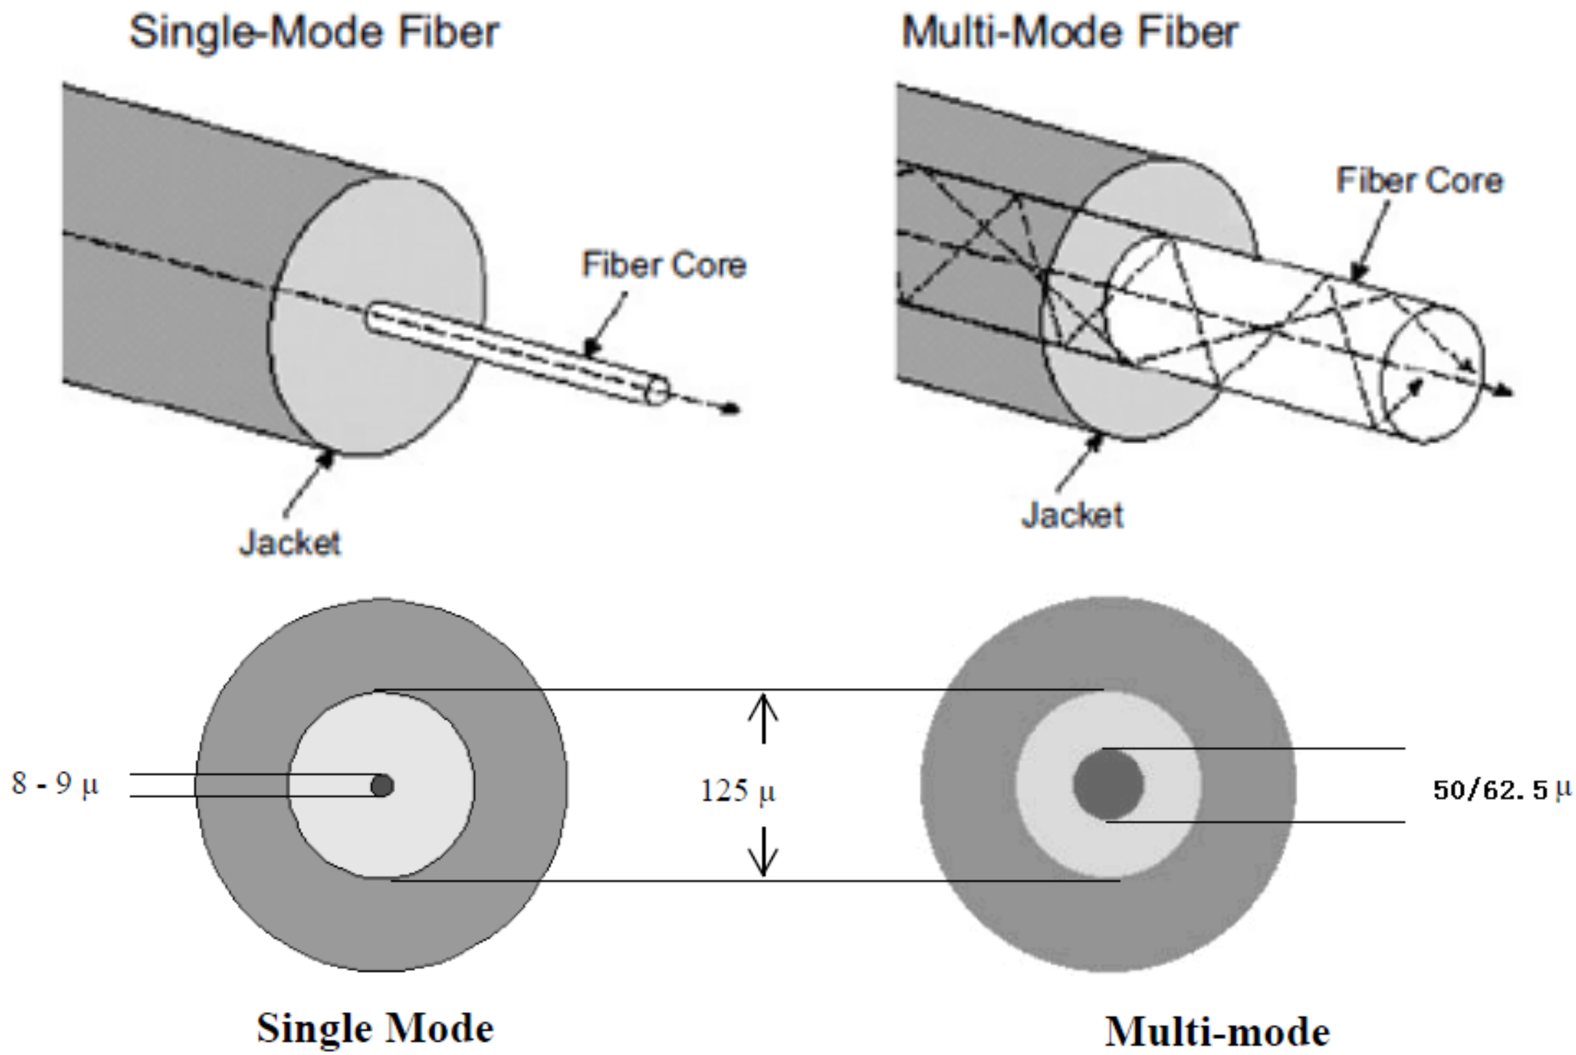

Basic Optical Networking Terms and Concepts

Optical Power

- What is optical power?
  - Quite simply, the brightness (or "intensity") of light.
  - As light travels through fiber, some energy is lost.
    - It can be absorbed by glass particles, and converted into heat;
    - Or scattered by microscopic imperfections in the fiber.
  - This loss of intensity is called "attenuation".
- We typically measure optical power in "Decibels"
  - A decibel (dB, $1/10^{th}$ of a Bel) is a logarithmic-scale unit expressing the relationship between two values.
  - The decibel is a "dimensionless-unit", meaning it does not express an actual physical measurement on its own.

Optical Power and the Decibel

- A decibel is a logarithmic ratio between two values
  - -10dB is 1/10th the signal, -20dB is 1/100th the signal, etc.
  - Another easy one: +3dB is double -3dB is half.
  - But remember, this doesn't tell you "double of what?"
- To express an absolute value, we need a reference.
  - In optical networking, this is known as a "dBm".
    - That is, a decibel relative to 1 milliwatt (mW) of power.
  - 0 dBm is 1 mW, 3 dBm is 2 mW, -3 dBm is 0.5mW, etc.
  - So what does this make 0mW? Negative Infinity dBm.
- Confusion between dB and dBm is probably the single biggest mistake made in optical networking!

Optical Power and the Decibel

- Why do we measure light with the Decibel?
  - Light, like sound, follows the inverse square law.
    - The signal is inversely proportional to the distance squared.
      - A signal travels distance X and loses half of its intensity.
      - The signal travels another distance X and loses another half.
      - After 2X only 25% remains, after 3X only 12.5% remains.
  - Using a logarithmic scale simplifies the calculations.
    - A 3dB change is approximately half/double the original signal.
    - In the example above, there is a 3dB loss per distance X.
    - At distance 2X there is 6dB of loss, at distance 3X it is 9dB.
    - This allows us to use elementary school addition/subtraction when measuring gains/losses, which is easier on the humans.

Decibel to Power Conversion Table

Table 1 - Decibel to Power Conversion			
dB (loss)	**Power Out as a % of Power In**	**% of Power Lost**	**Remarks**
1	79%	21%	---
2	63%	37%	---
3	50%	50%	1/2 the power
4	40%	60%	---
5	32%	68%	---
6	25%	75%	1/4 the power
7	20%	80%	1/5 the power
8	16%	84%	1/6 the power
9	12%	88%	1/8 the power
10	10%	90%	1/10 the power
11	8%	92%	1/12 the power
12	6.3%	93.7%	1/16 the power
13	5%	95%	1/20 the power
14	4%	96%	1/25 the power
15	3.2%	96.8%	1/30 the power
16	2.5%	97.5%	1/40 the power
17	2%	98%	1/50 the power
18	1.6%	98.4%	1/60 the power
19	1.3%	98.7%	1/80 the power
20	1%	99%	1/100 the power
25	0.3%	99.7%	1/300 the power
30	0.1%	99.9%	1/1000 the power
40	0.01%	99.99%	1/10,000 the power
50	0.001%	99.999%	1/100,000 the power

Dispersion

- Dispersion simply means "to spread out".
  - In optical networking, this results in signal degradation.
  - As the signal is dispersed, it is no longer distinguishable as individual pulses at the receiver.

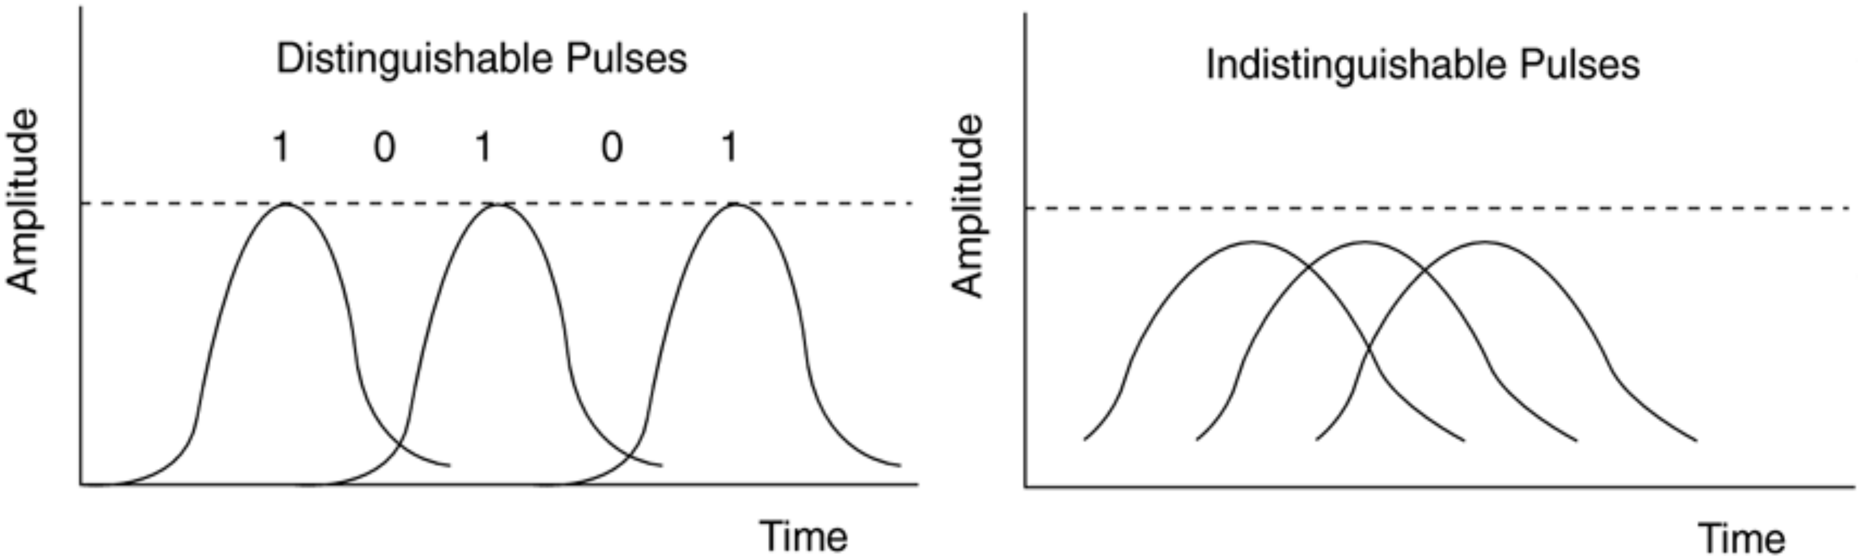

Chromatic Dispersion (CMD)

- Different frequencies propagate through a non-vacuum at different speeds. This is how optical prisms work.
- The wider your signal linewidth, the more CMD affects it.
- The faster your baud rate, the more CMD affects it.
- So for every doubling of baud rate, CMD increases by 4.

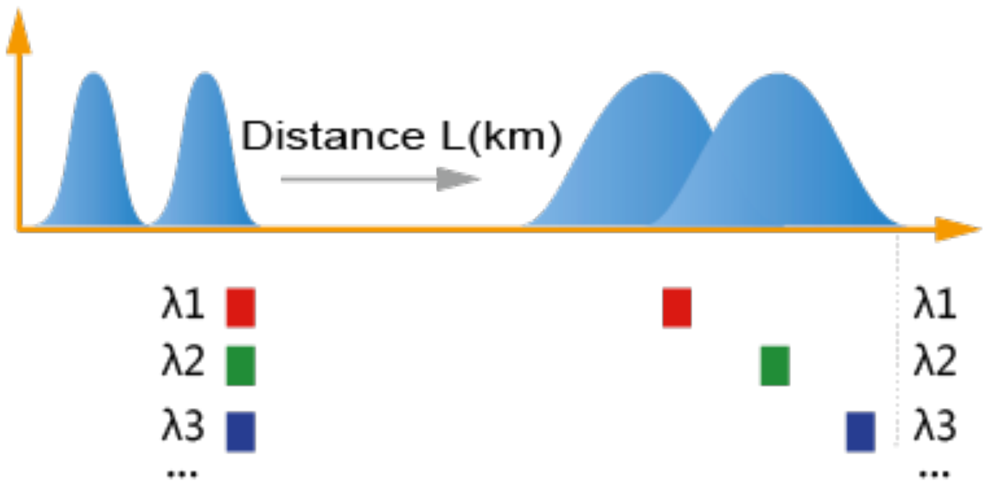

Polarization Mode Dispersion (PMD)

- Caused by imperfections in the shape of the fiber.
- Not perfectly cylindrical fiber causes one polarization of light to propagate faster than the other.
- The difference in arrival time between the polarizations is called “Differential Group Delay” (DGD).

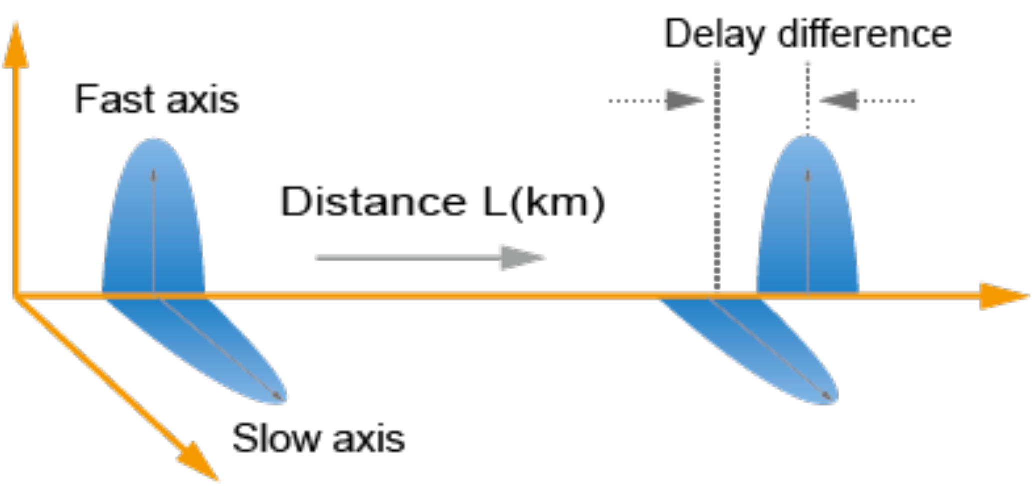

Fiber Optic Transmission Bands

- There are several frequency “windows” available
  - 850nm – The First Window
    - Highest attenuation, only used for short reach applications today.
  - 1310nm – The Second Window (O-band)
    - The point of zero dispersion on classic SMF, but high attenuation.
    - Primarily used for medium-reach applications (~10km) today.
  - 1550nm – Third Window (C-band)
    - Stands for “conventional band”, covers 1525nm – 1565nm.
    - Has the lowest rate of attenuation over SMF.
    - Used for almost all long-reach and DWDM applications today.
  - Fourth Window (L-band)
    - Stands for “long band”, covers 1570nm – 1610nm.

Fiber Optic Transmission Bands

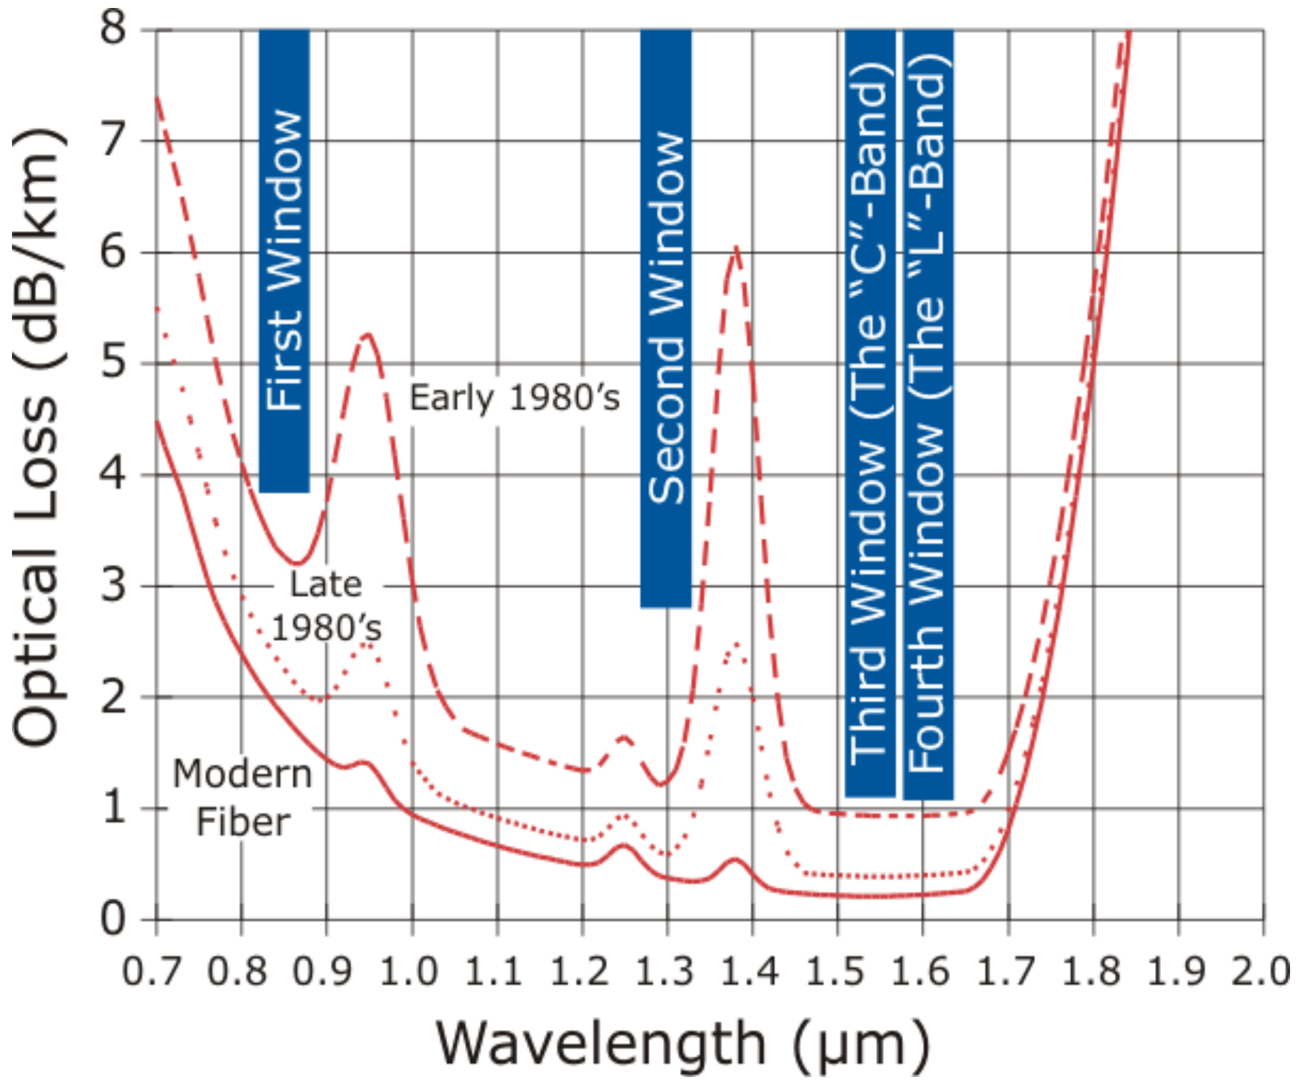

Optical Signal to Noise Ratio (OSNR)

- Modern high-bandwidth systems are often more limited by OSNR than any other optical parameter.

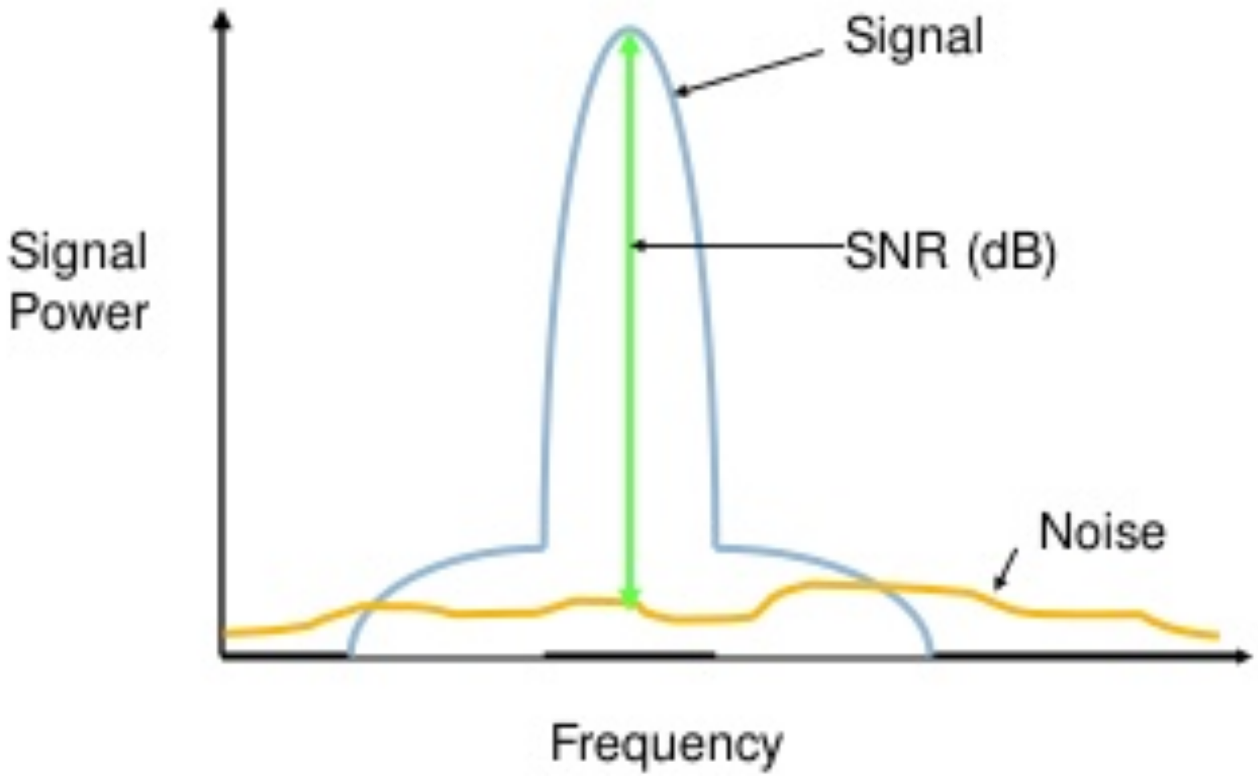

Wave Division Multiplexing

Wave Division Multiplexing (WDM)

- What is Wave Division Multiplexing (WDM)?
  - We know that light comes in many different "colors".
    - What we perceive as "white" is actually just a mix of many wavelengths.
  - These different colors can be combined on the same fiber.
  - The goal is to put multiple signals on the same fiber without interference ("ships in the night"), thus increasing capacity.

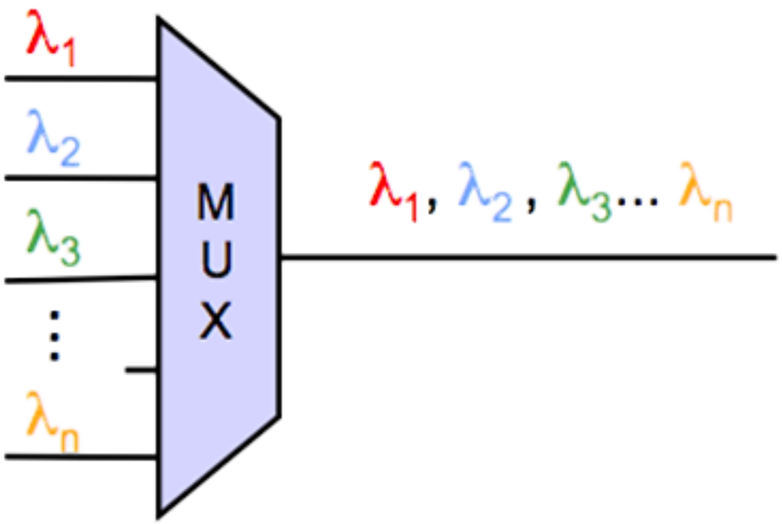

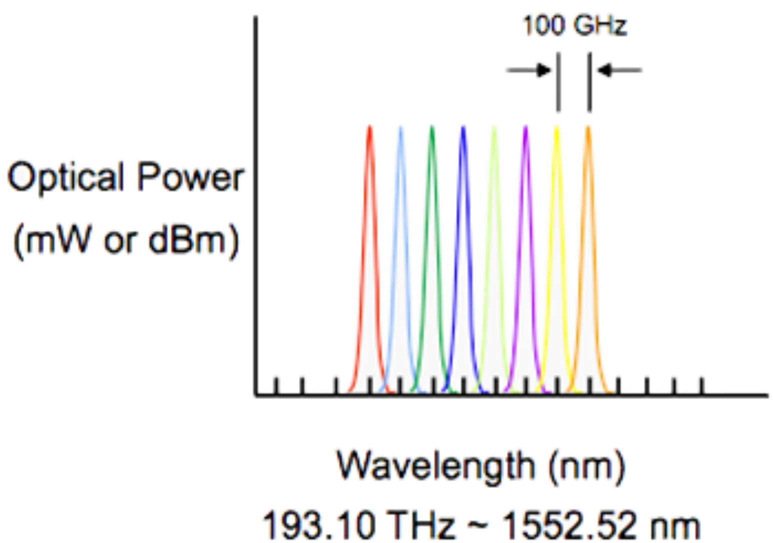

Wave Division Multiplexing Channels

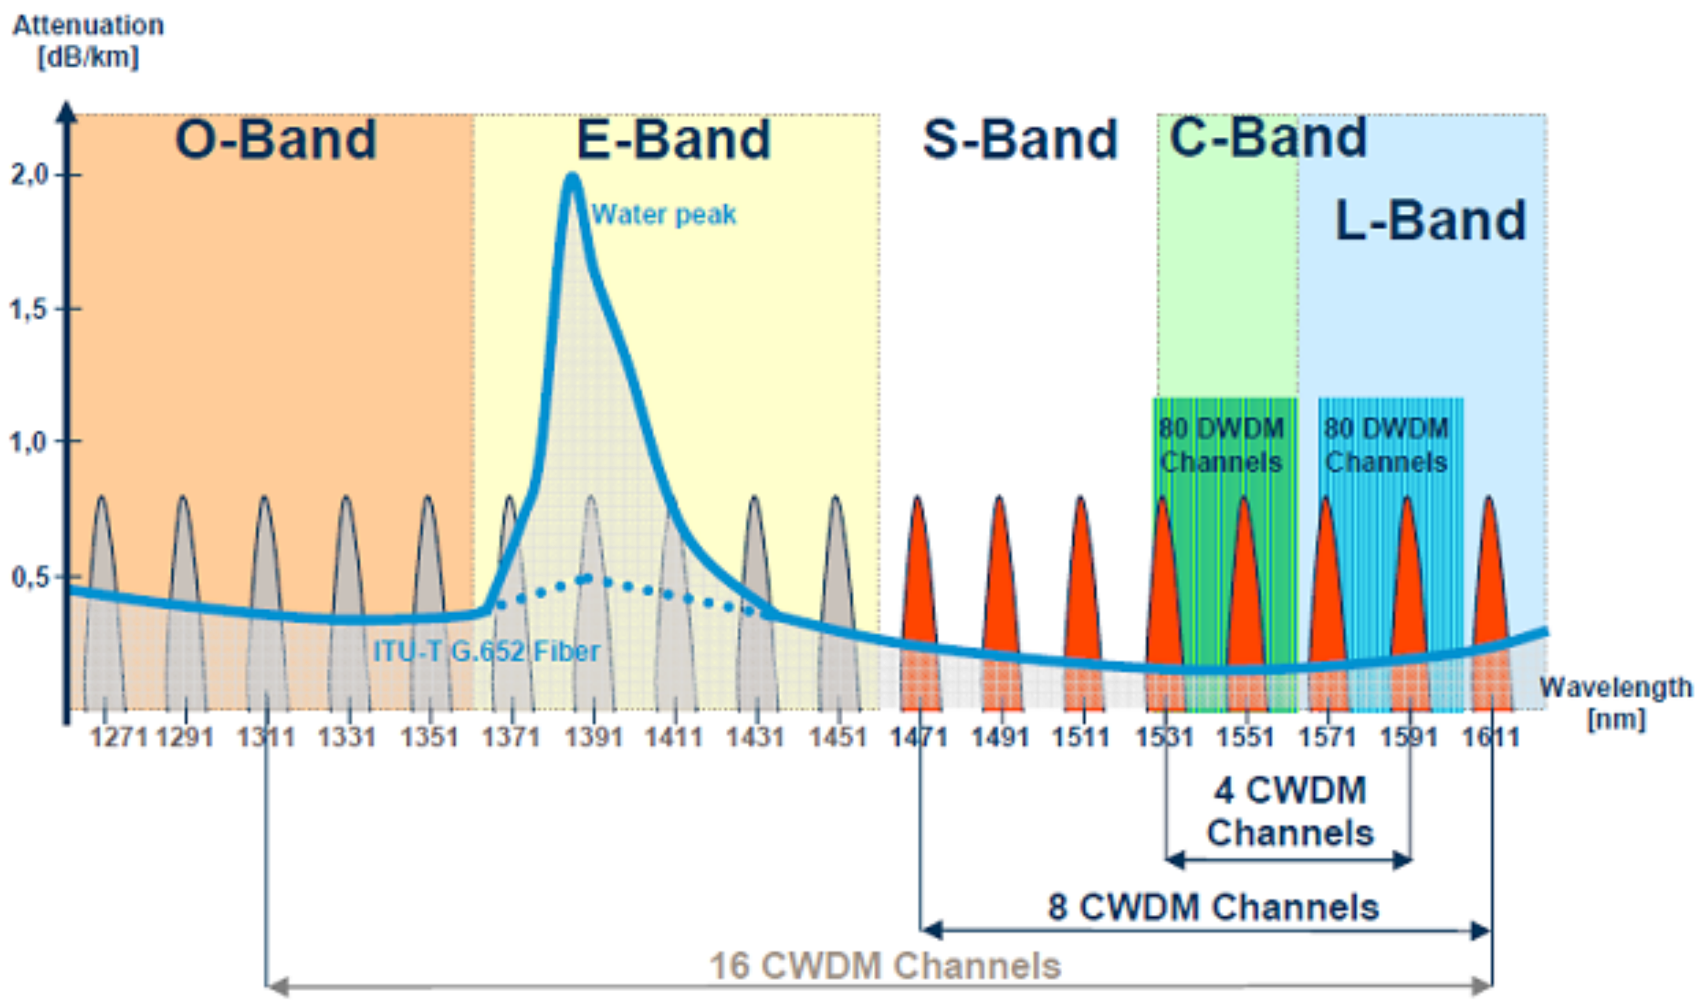

Coarse Wave Division Multiplexing

- CWDM is loosely used to mean "anything not DWDM"
  - One "popular" meaning is 8 channels with 20nm spacing.
    - Centered on 1470 / 1490 / 1510 / 1530 / 1550 / 1570 / 1590 / 1610

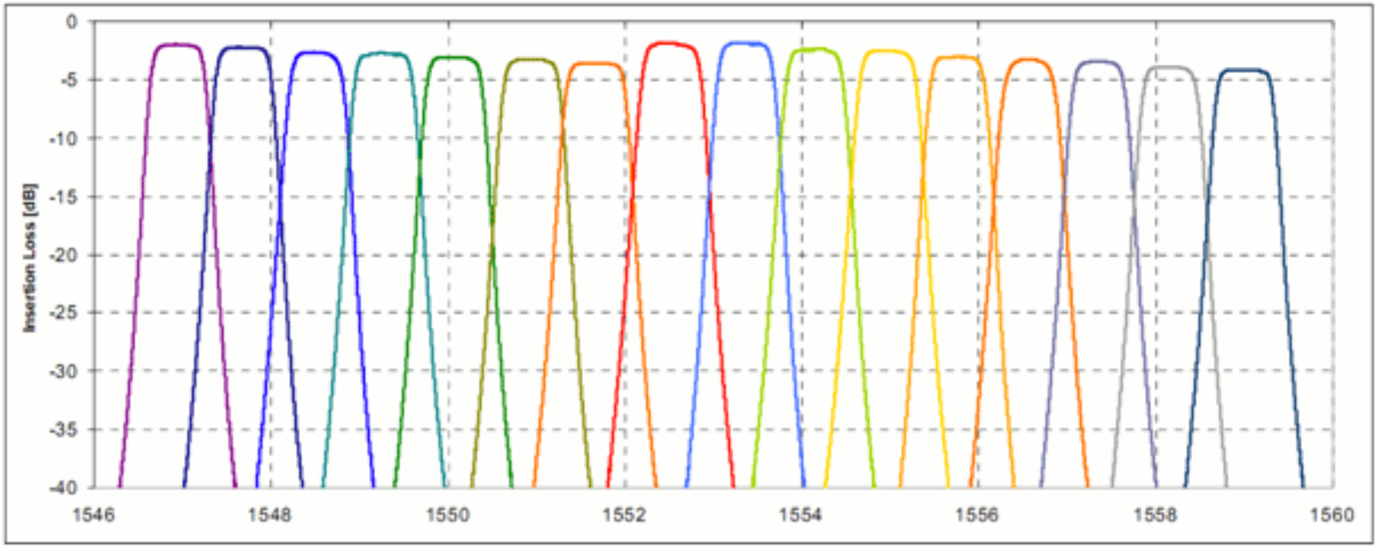

  - With Low Water Peak fiber, another 10 channels are possible
    - Centered on 1270/1290/1310/1330/1350/1370/1390/1410/1430/1450.
  - Can also be used to refer to a simple 1310/1550nm mux.

Dense Wave Division Multiplexing

- So what does that make Dense WDM (DWDM)?
  - Defined by the ITU-T G.694.1 as a "grid" of specific channels.
  - Within C-band, the follow channel sizes are common:
    - 200GHz – 1.6nm spacing, 20-24 channels
    - 100GHz – 0.8nm spacing, 40-48 channels
    - 50GHz – 0.4nm spacing, 80-96 channels
    - 25GHz – 0.2nm spacing, 160-192 channels
  - A rough guideline to the technology:
    - 200GHz is "2000-era" old tech, rarely seen in production any more.
    - 100GHz is still quite common for metro DWDM tuned pluggables.
    - 50GHz is common for commercial, long-haul, and 100G systems.
    - 25GHz was used briefly for high-density 10G systems, before the move coherent 100G systems shifted back to 50GHz.
    - Modern systems are flexible, in 12.5GHz increments or smaller.

What Are The Advantages to Each?

- CWDM
  - Cheaper, less precise lasers can be used.
    - The actual signal in a CWDM system isn’t really any wider.
    - But the wide channel allows for large temperature variations.
    - Cheaper, uncooled lasers can more easily stay within the window.
- DWDM
  - Far more channels are possible within the same fiber.
    - 160 channels (at 25GHz) in 32nm of spectrum, vs. 8ch in 160nm.
  - Can stay completely within the C-band
    - Where attenuation and dispersion are far lower that other bands.
    - Where simple Erbium Doped Amplifiers (EDFAs) work.
  - But can also be duplicated within the L-band.

CWDM vs. DWDM Relative Channel Sizes

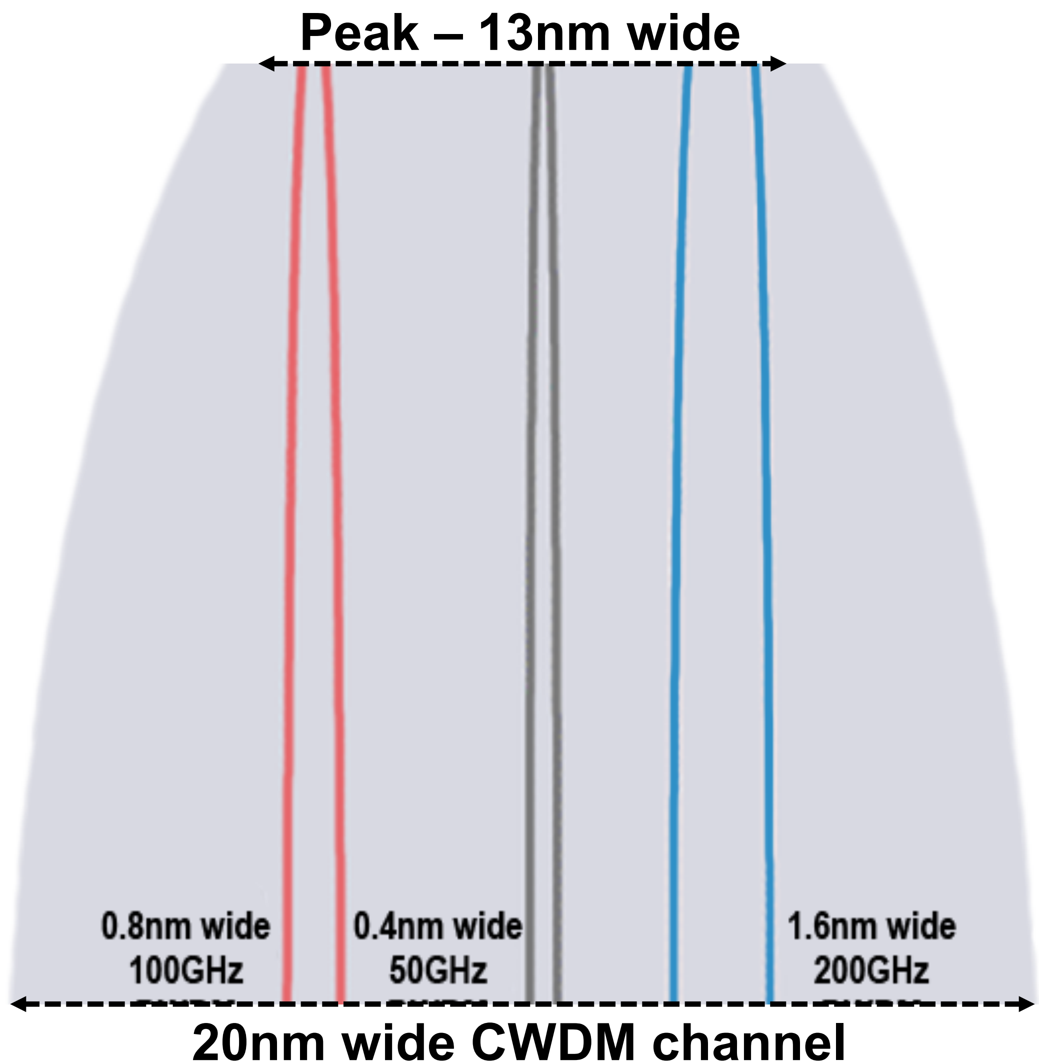

Other Uses of Wave Division Multiplexing

- But other forms of WDM can be used as well
  - The classic 1310/1550 muxes.
  - 4-lane “Grey” Optics
    - New high speed interfaces often start using multiple WDM lanes.
      - Cheaper to implement, or supports older fiber technology.
    - 10GE had 10GBASE-LX4 (4x 2.5G channels, rather than 1x 10G)
    - 40GE has LR4 (4x 10G, 1270nm / 1290nm / 1310nm / 1330nm)
    - 100GE has LR4 (4x 25G, 1295nm / 1300nm / 1305nm / 1310nm)
  - Single Strand Optics (BX “bidirectional” standards)
    - E.g. 1310 / 1490nm mux integrated into a pluggable transceiver.

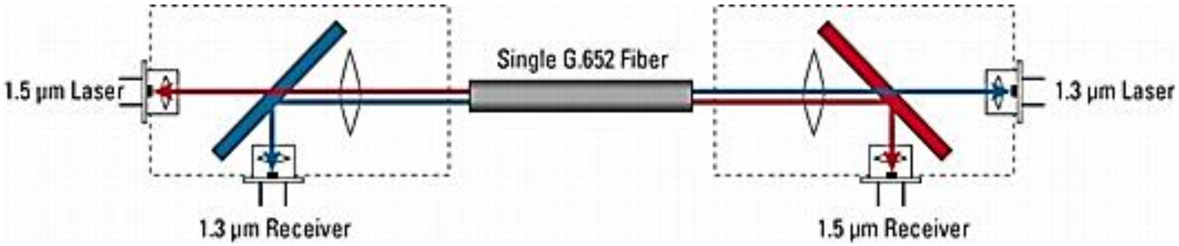

DWDM Channel Terminology

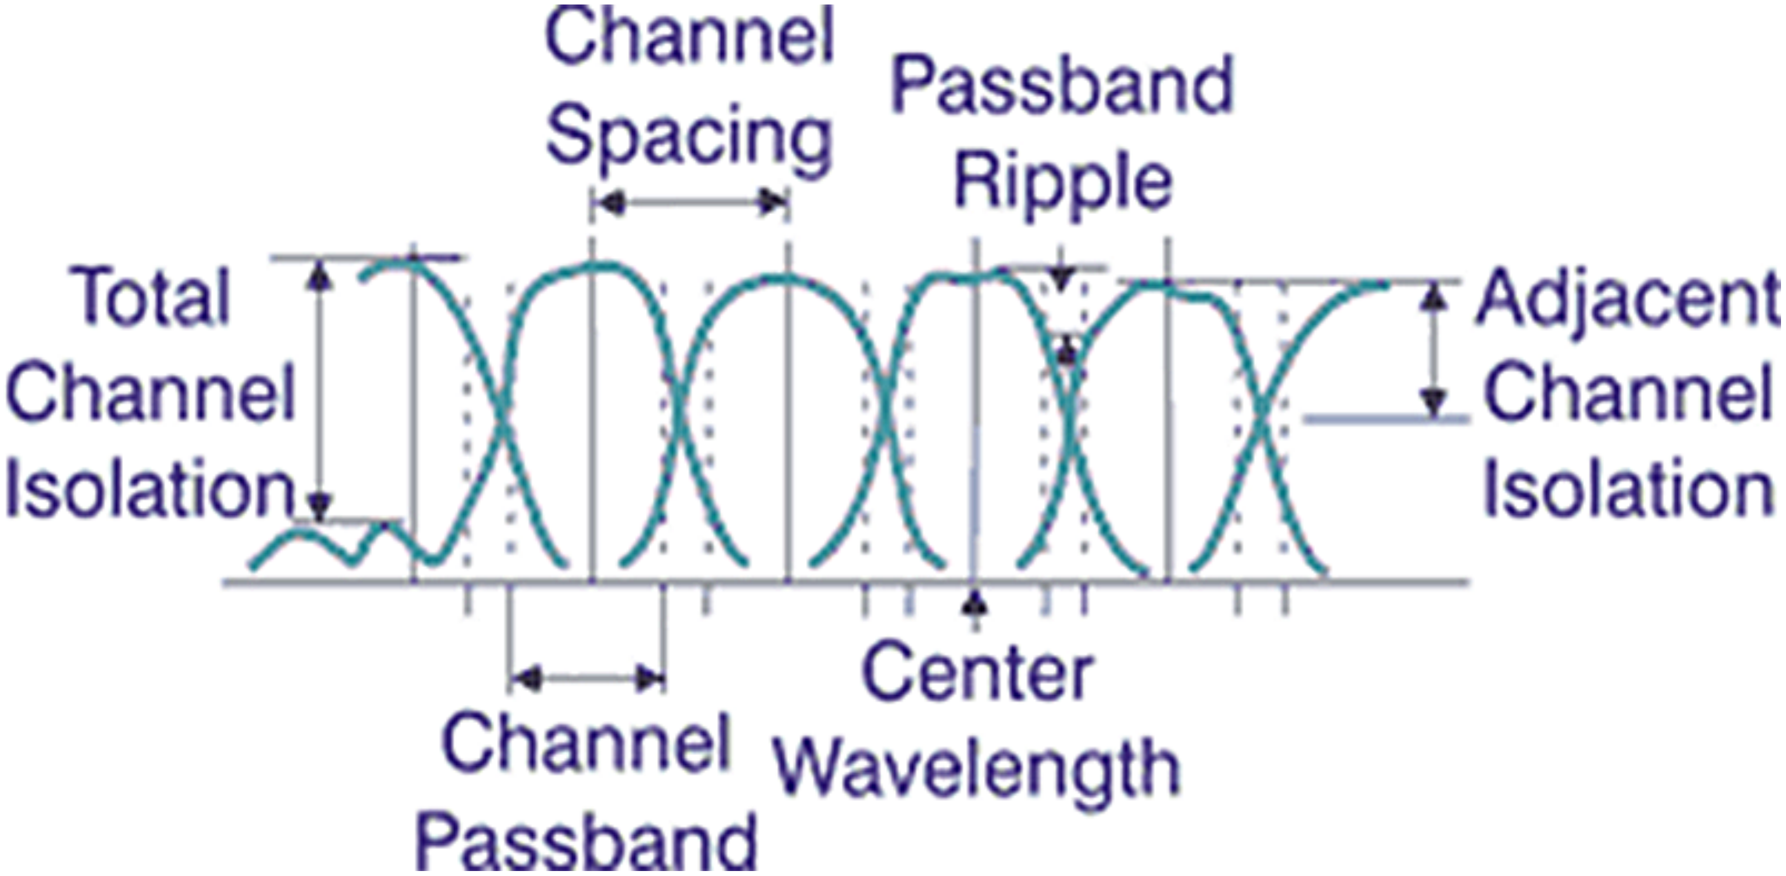

WDM Networking Components

WDM Mux/Demux

- A simple, passive (unpowered) device, which combines/splits multiple colors of light to/from a single “common” fiber.
- Short for “multiplexer”, sometimes called a “filter”, or “prism”.
  - A “filter” is how it actually works, by filtering specific colors.
  - But people conceptually understand that a prism splits light into its various component frequencies.
- A complete system requires both a mux and a demux, for the TX and RX operation.
- Often sold as a single package containing both units, for simplicity (and use on duplex fiber).

The Optical Add/Drop Multiplexer (OADM)

- Selectively Adds and Drops certain WDM channels, while passing other channels through without disruption.
- While muxes often used at major end-points to break out all channels, OADMs are often used at mid-points within rings.
- A well-constructed OADM ring can reach any other node in the ring, potentially reusing some wavelengths multiple times on different portions of the ring.
- Also an entirely passive and unpowered device.

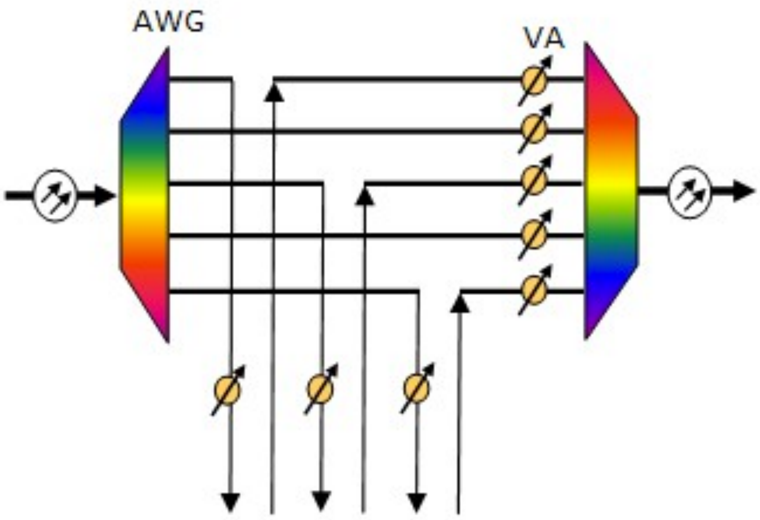

Passive Optical Filter Technology

Passive filters (Mux/OADM) can be build in several ways

- Thin Film Filter (TFF)
  - Typically used for low channel counts.
- FBG (Fiber Bragg Grating)
  - An “etched” fiber core, which causes certain frequencies to be reflected.

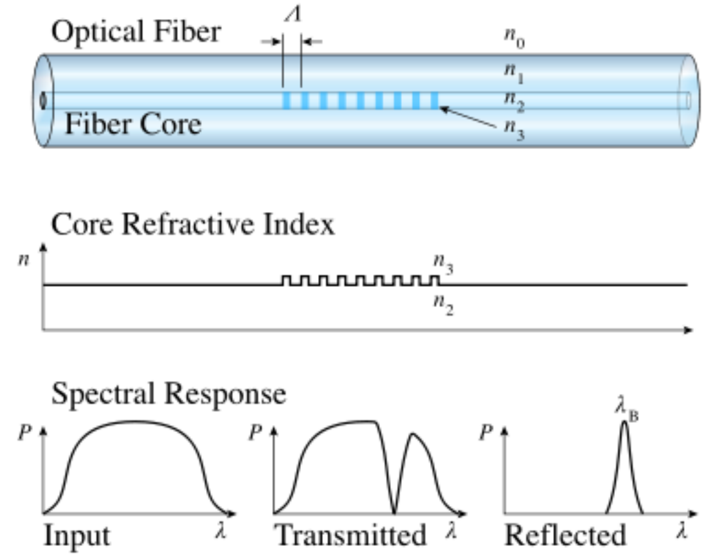

- AWG (Arrayed Waveguide Grating)
  - Typically used for high channel counts.
  - Essentially a very fancy interferometer.
  - Cheapest and lowest IL, but not flat-top.
  - Lowest loss versions have specific temperature constraints.
  - Most common versions are AAWG (Athermal AWG) today.

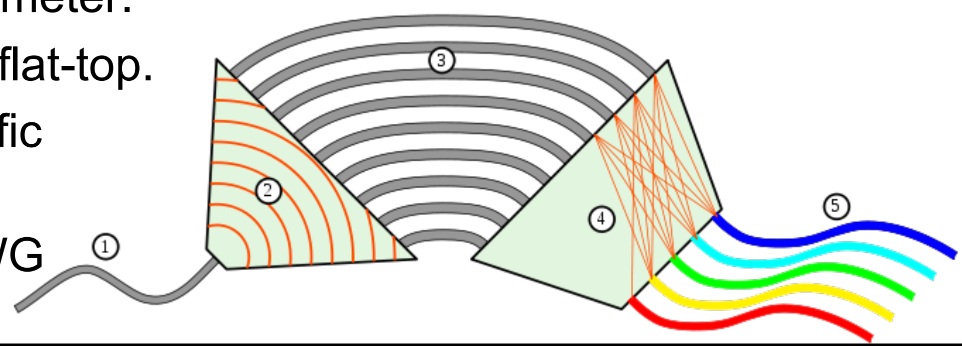

Reconfigurable OADM (ROADM)

A ROADM is a software reconfigurable OADM

- Often capable of building many different “degrees”.
  - 2-degree is the simplest OADM style, east/west and add/drop
  - 4-degree, 8-degree, and 20-degree are also common designs.

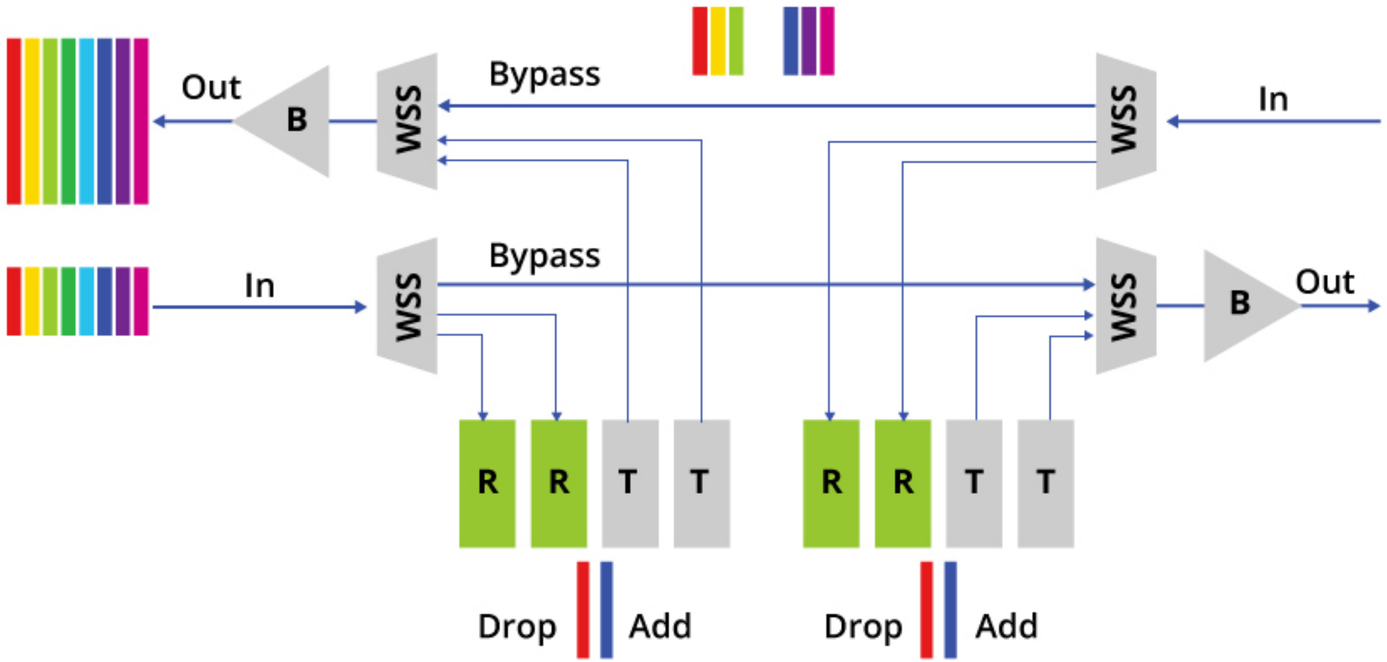

The Evolution of the ROADM

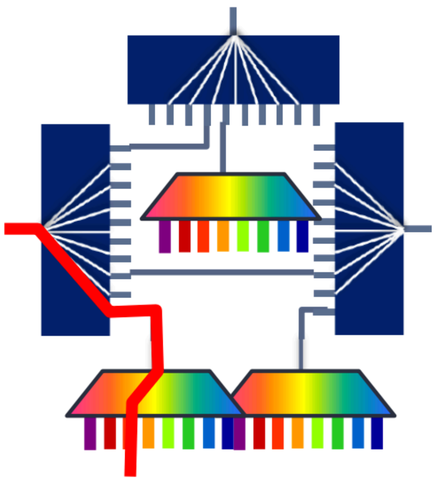

Basic ROADM

- Reconfigurable, but add/drop still goes to a standard fixed mux.
- Specific frequencies must be connected to specific ports.
- The network must be recabled in order to change or move frequencies.

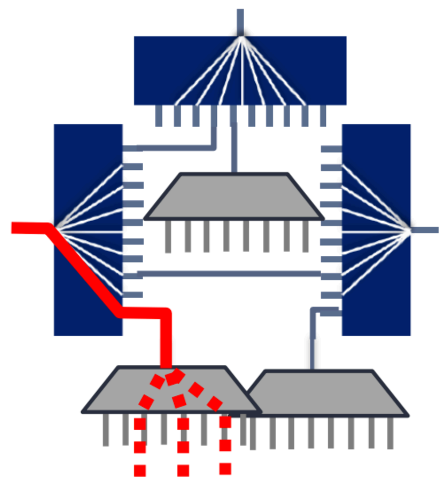

Colorless ROADM

- Eliminates the need to map specific frequencies to specific ports.
- But still limited to muxing in one direction at a time.

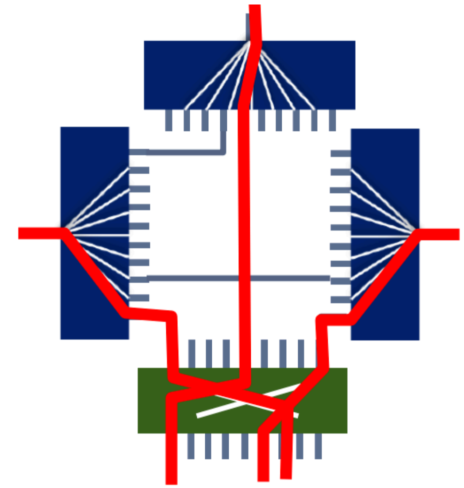

CDC ROADM

- Colorless – Any channel can be add/dropped on any port.
- Directionless – Any channel can be sent to any direction.
- Contentionless – The same channel can be reused on different directions without causing internal contention

Modern Networking and the CDC ROADM

- The goal is to move optical channels entirely in software.
  - Transponders can be reallocated onto different physical paths as traffic patterns change (due to time-of-day changes, or during fiber cuts), potentially increasing efficiency and reducing costs.
  - Eliminates the need to physically move cables to reconfigure.
  - Allows dynamic bandwidth allocation at an optical level.
- Potentially the entire process could be automated.
  - IETF pushing for vendor interoperability, and signaling via mechanisms like PCEP (Path Computation Element Protocol).
  - Routers could “request” additional bandwidth from a pool of underlying transponders on demand, based on real-time traffic requirements.

What Goes Into a Modern CDC ROADM

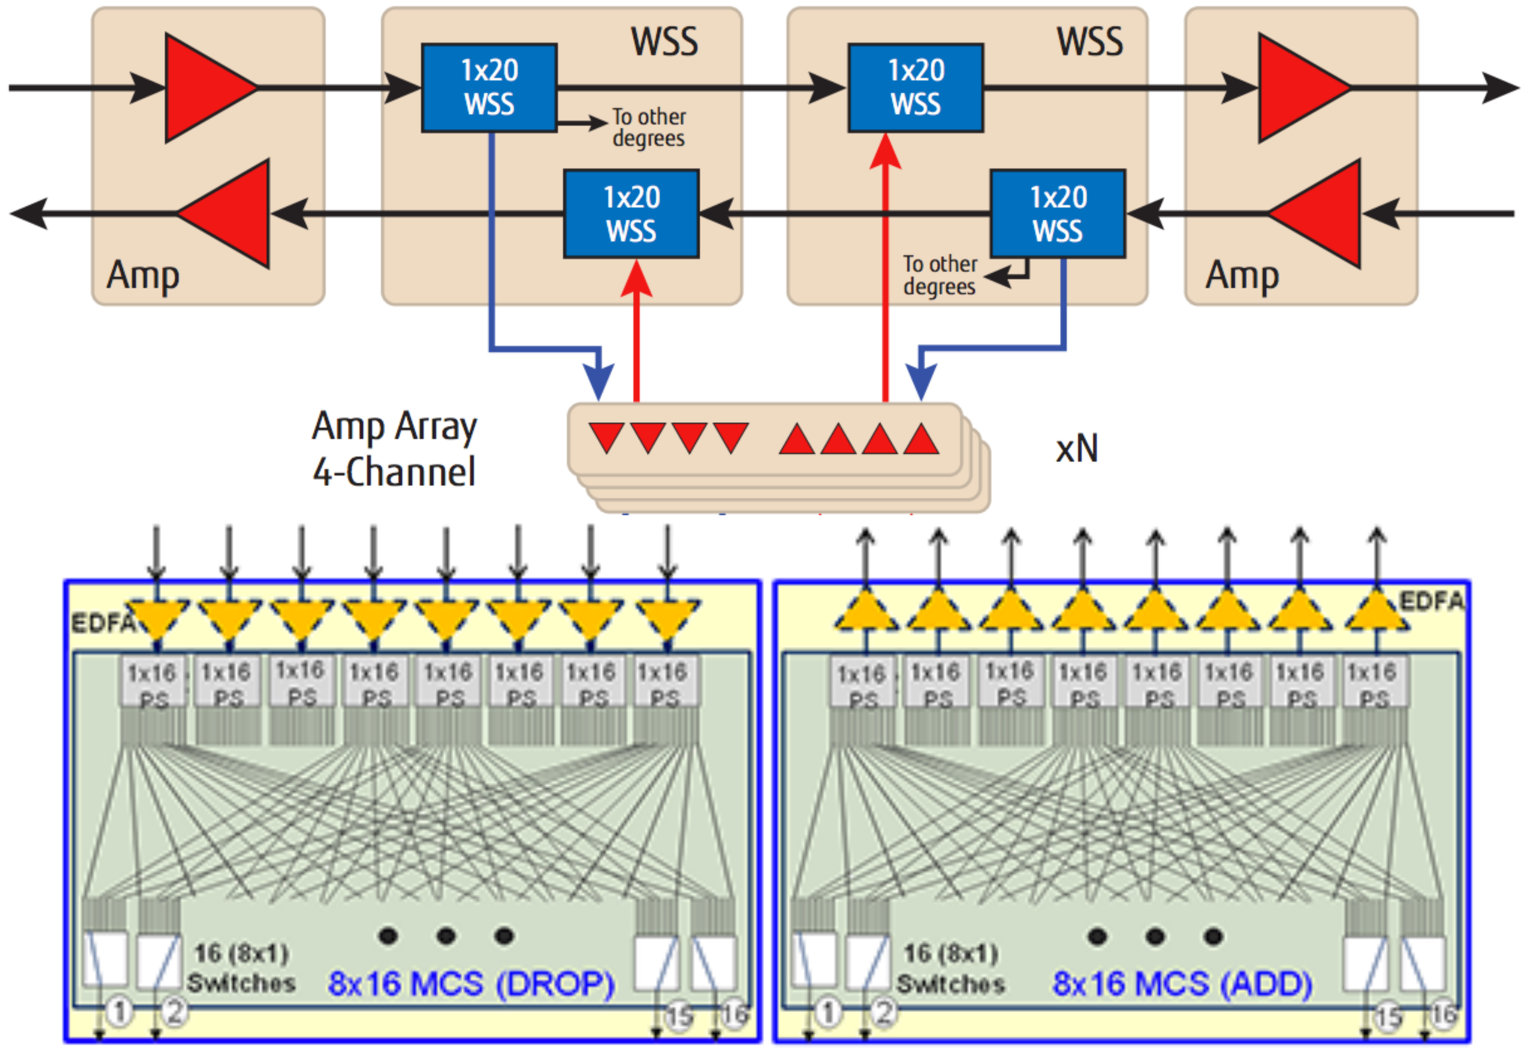

DWDM Superchannels

What if we want performance a single carrier can't deliver?

- Superchannels pack multiple carriers together in a single channel, enabling more bandwidth, longer reach, and cheaper components.
- Often you can pack the carriers together tighter than if you were using standard channels too, adding spectral efficiency.
- In this example, we deliver 1 Tbps using existing technology.

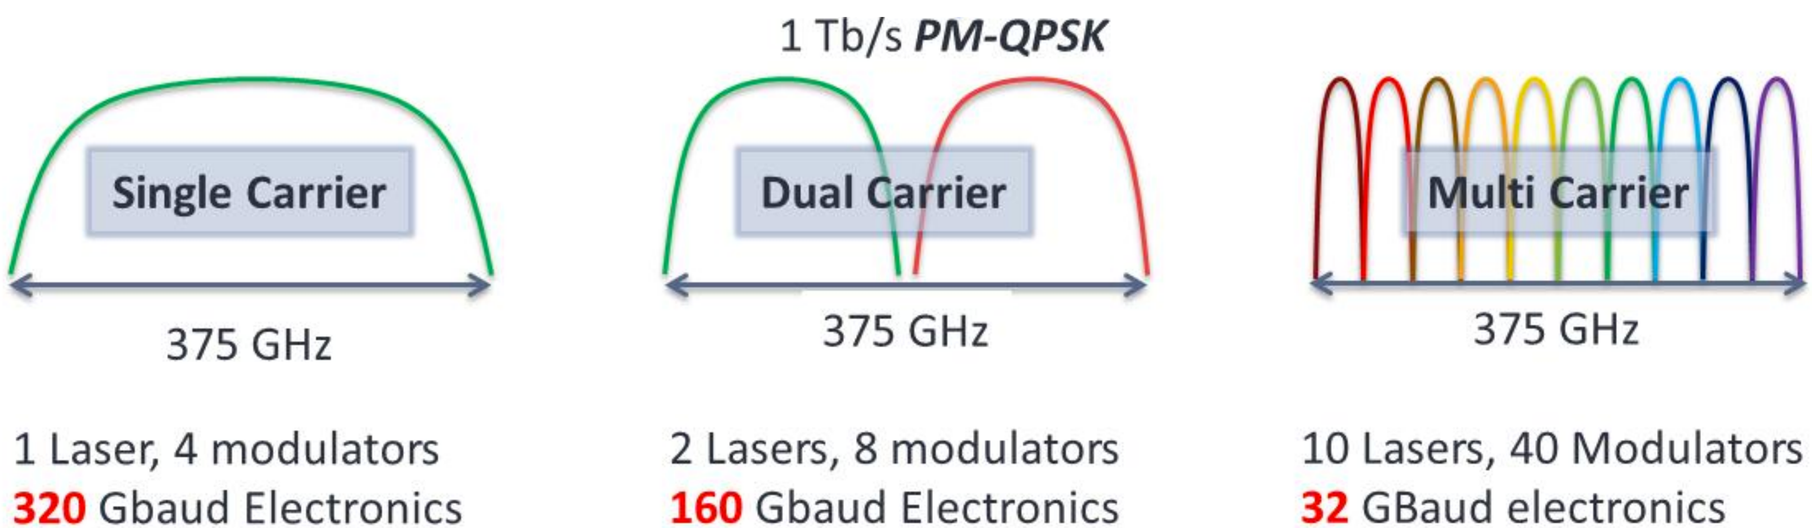

The Evolution of DWDM Channels

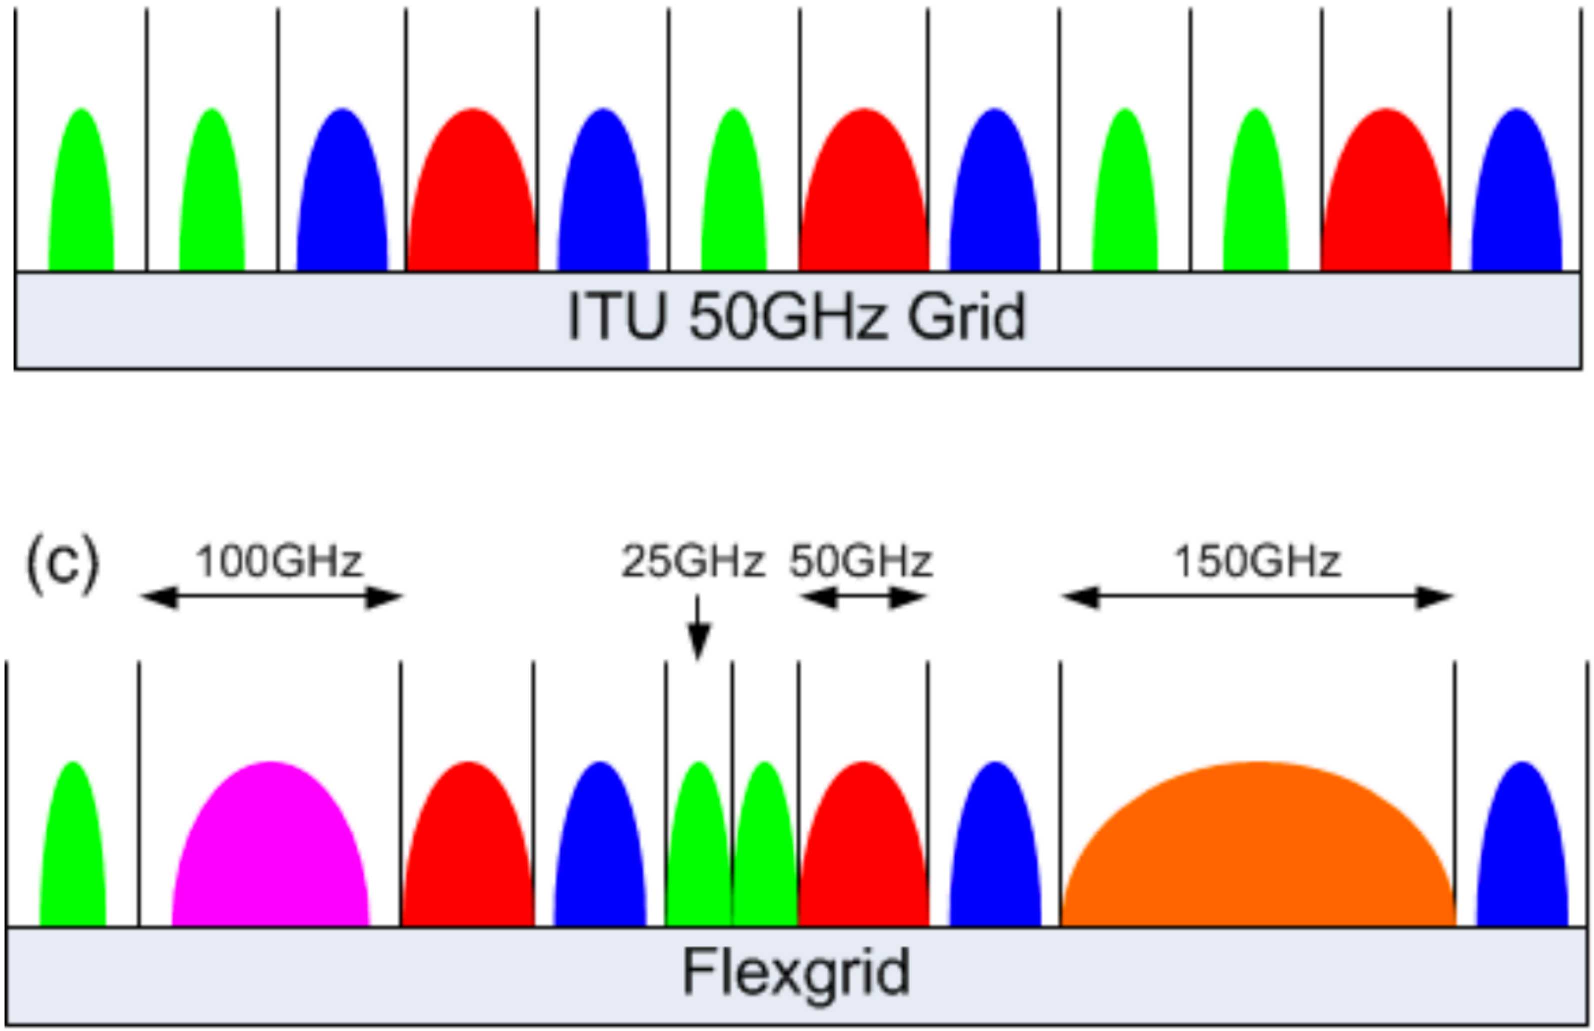

Optical Amplification

Optical Amplifiers

Optical amplifiers increase the intensity of a signal

- Purely optical way to extend signal reach, no regeneration.
- There are different types, for different frequencies of light.
- Different designs, for different positions within the span.
  - Booster Amplifiers are designed for high total output powers.
  - Pre-Amplifiers are designed for low input powers with minimal noise.
  - Inline Amplifiers strike a balance between the two.

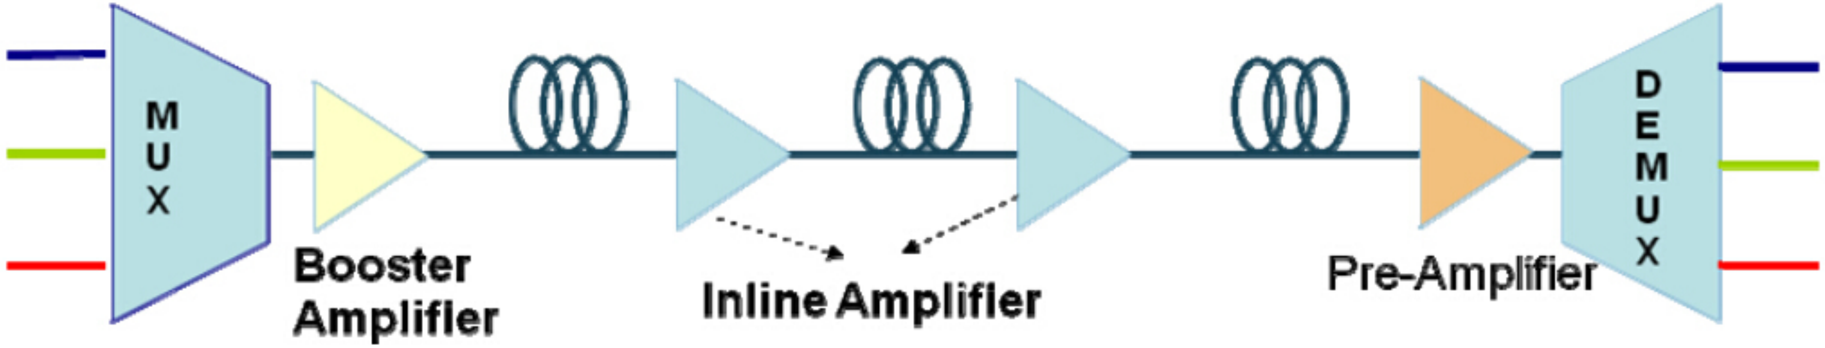

Erbium Doped Fiber Amplifier (EDFA)

- The most basic/common fiber amplification system.
  - In an EDFA, a piece of fiber is "doped" with Erbium ions.
  - Another laser (980/1480nm) is pumped in via a coupler.
  - Pump laser puts Erbium electrons into higher energy state.
  - 1550nm photons cause the Erbium electrons to decay to their ground state, and emit a clone of the original photon.

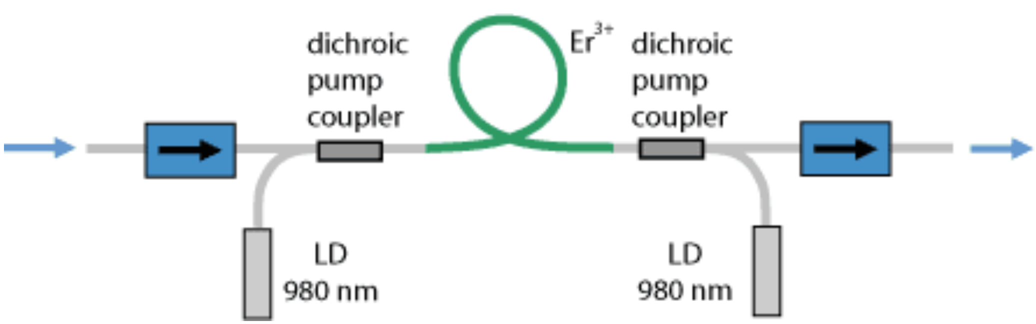

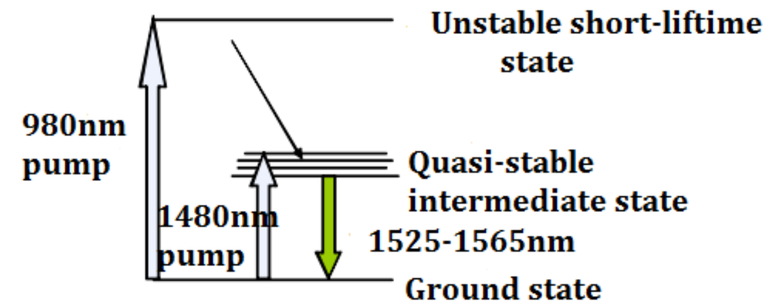

EDFA Noise

- So why can't we use this to boost a signal forever?
  - In addition to the intended boosting of signal, EDFAs also generate noise ("Amplified Spontaneous Emission", ASE).
  - Whenever an excited Erbium electron fails to encounter a "good" photon within ~10ms, it falls back to its ground state spontaneously, emitting "noise" photons.
  - Once generated, this noise is indistinguishable from the original signal.
  - After enough hops, the noise ruins the original signal.

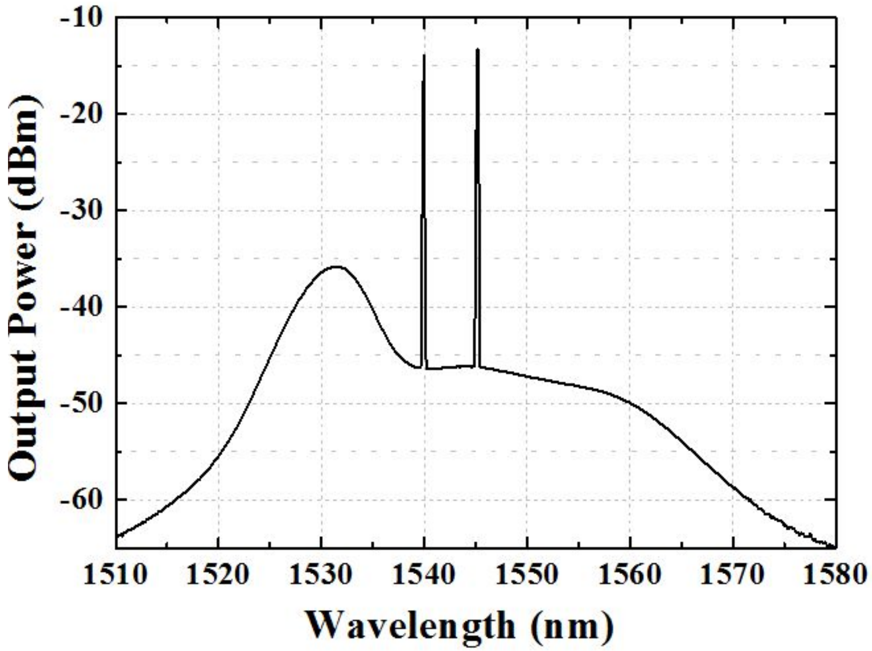

EDFA Noise – Why Input Power Matters

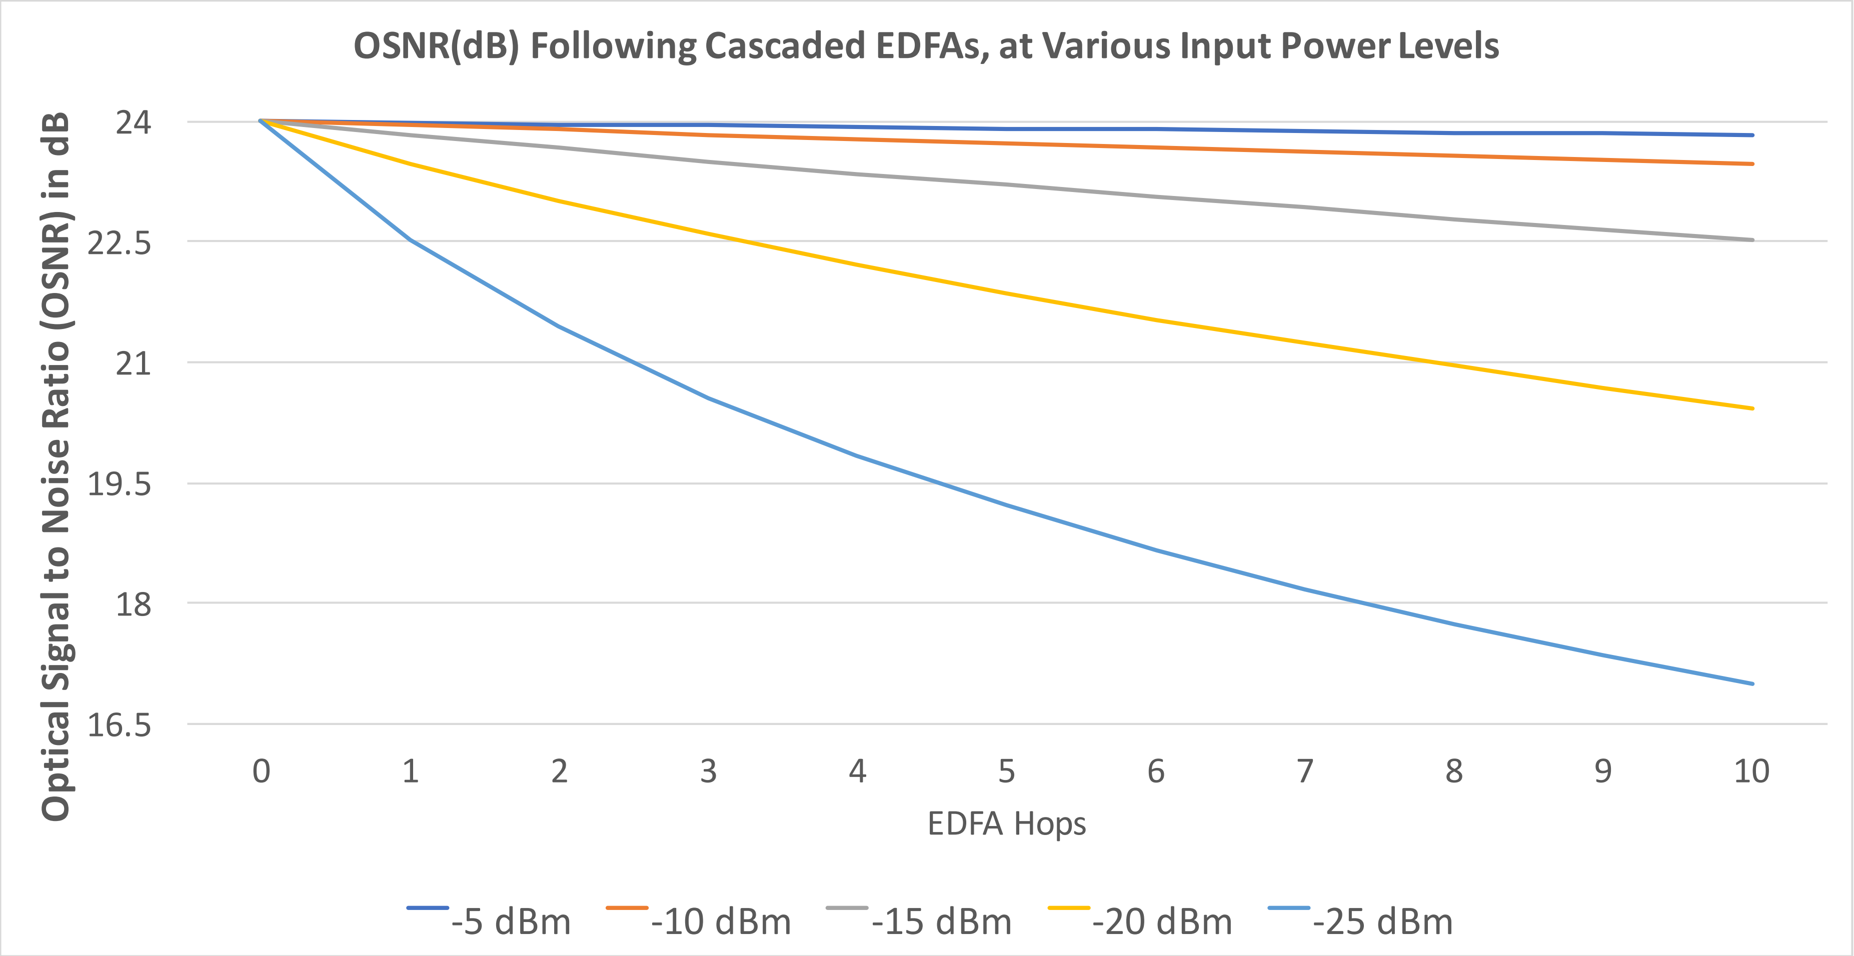

Raman Amplification

- The other major optical amplifier type is Raman.
  - Works on a principal of “Stimulated Raman Scattering”.
  - Requires very high power pump lasers, long gain mediums.
    - EDFAs used “lumped” design, gain media of 1-20 meters.
    - 0+ km.

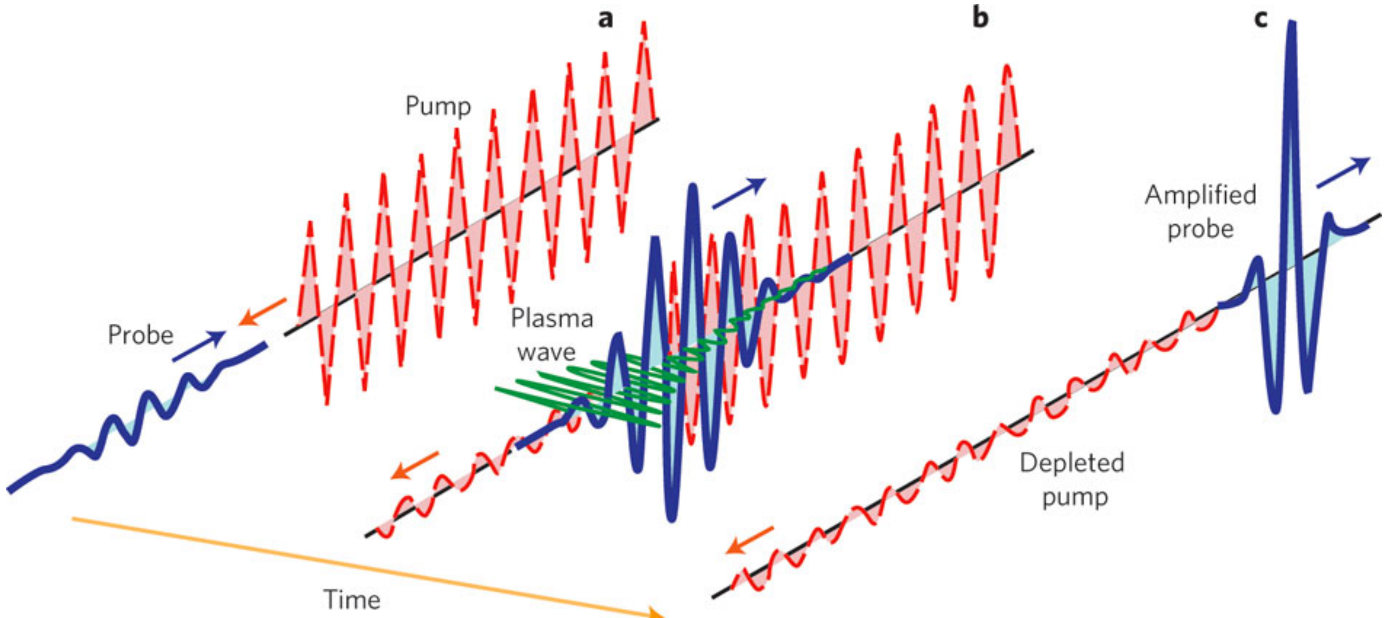

EDFA Only Amplification

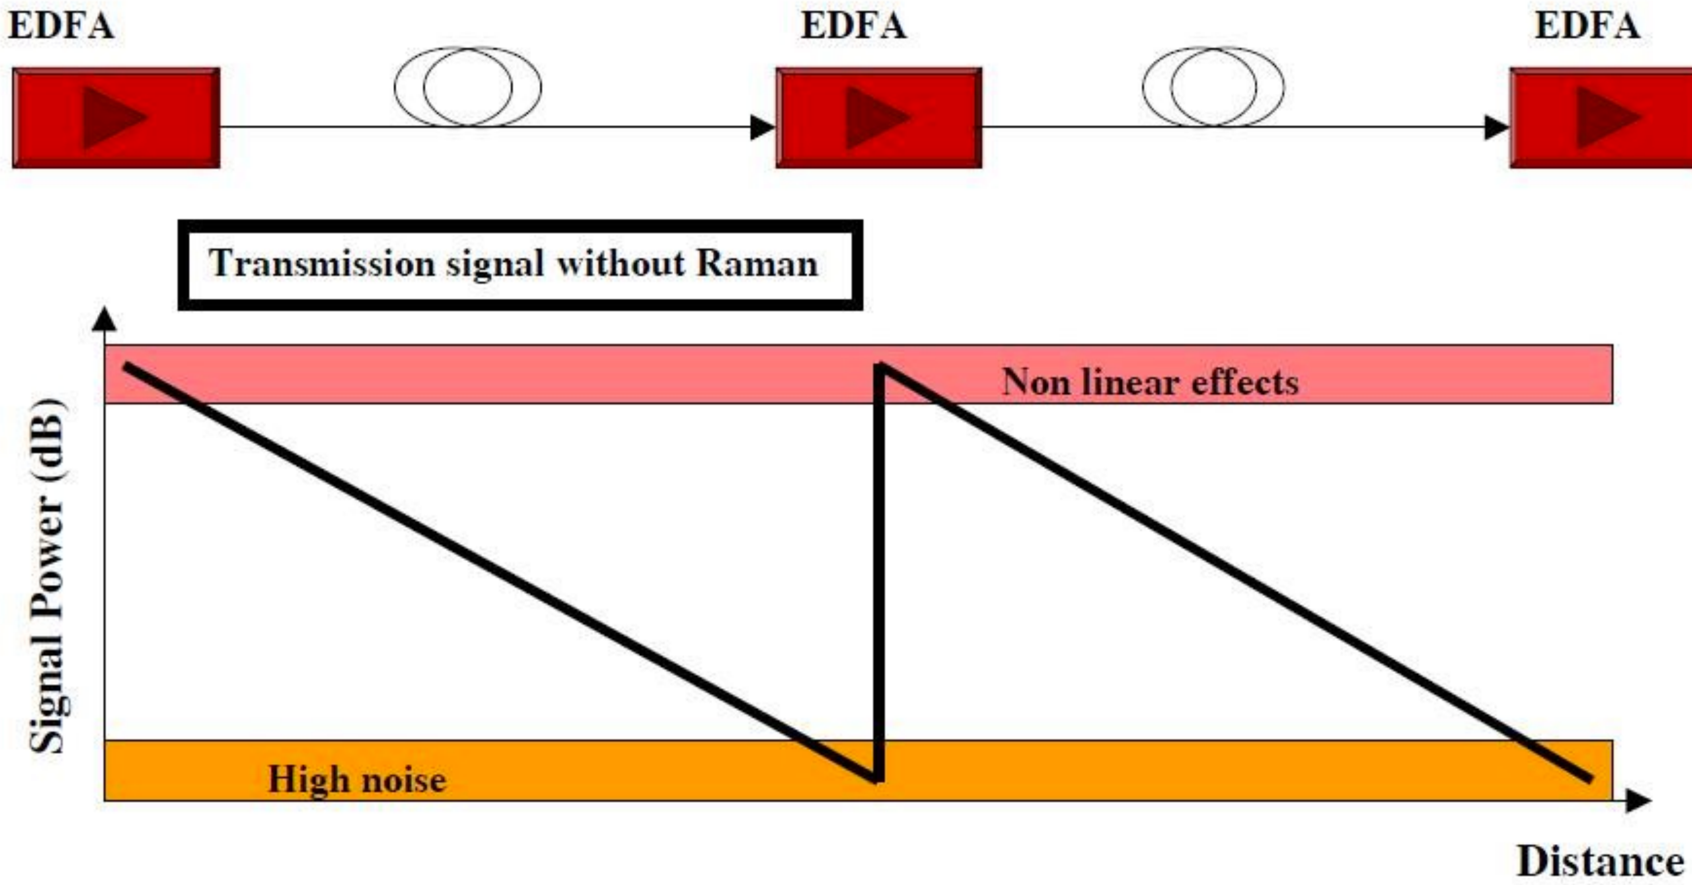

Hybrid EDFA + Raman Amplification

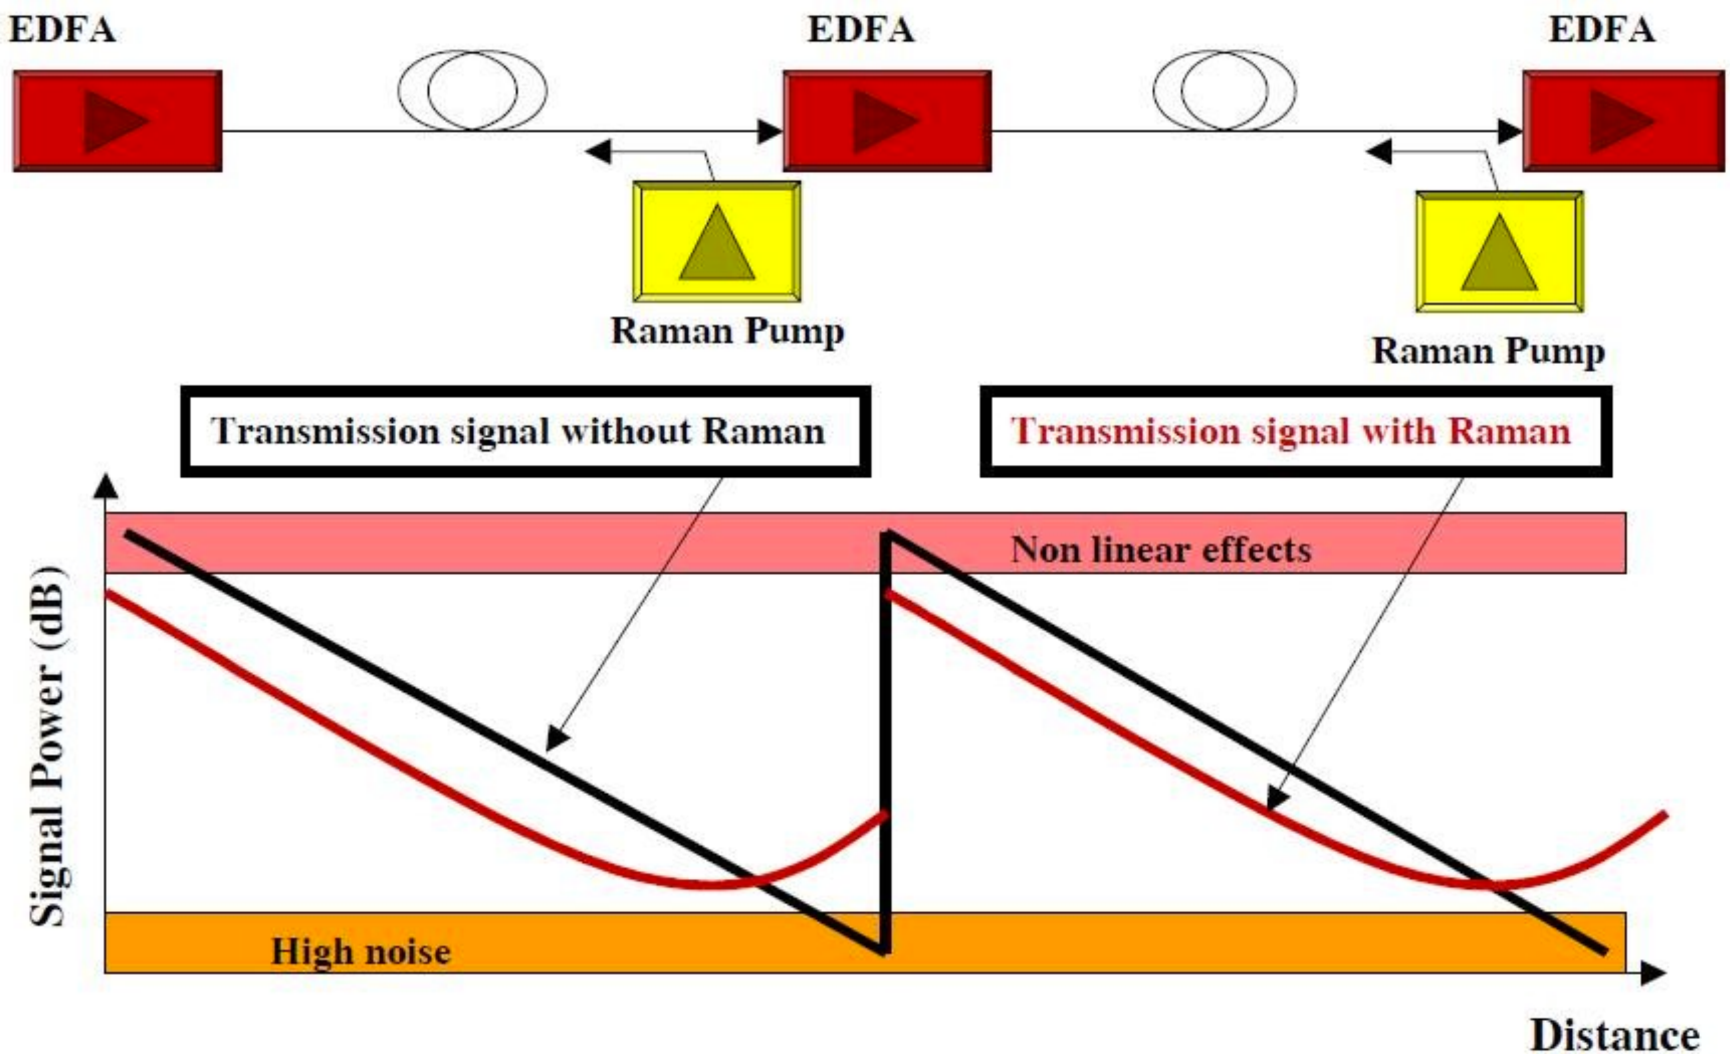

Hybrid EDFA + Raman Performance

Adding Raman extends EDFA reach significantly!

- In this example (21dB OSNR): from 7 hops to 20 hops
- At 100km/each, we go from only 700km to doing 2000km.

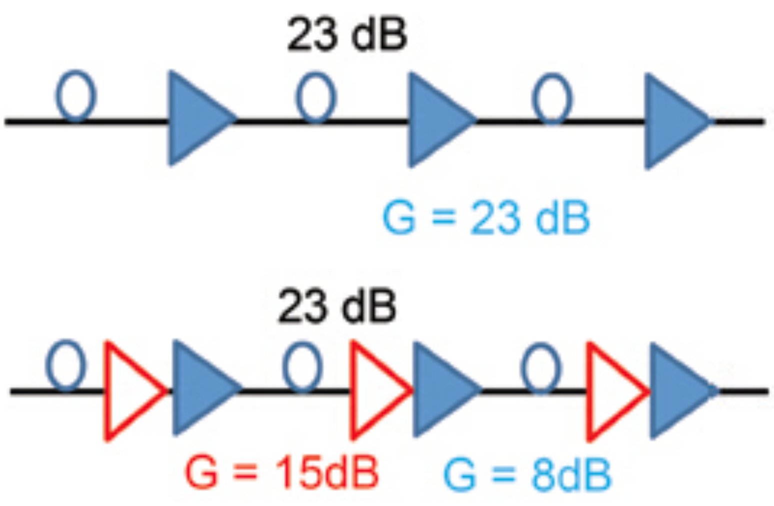

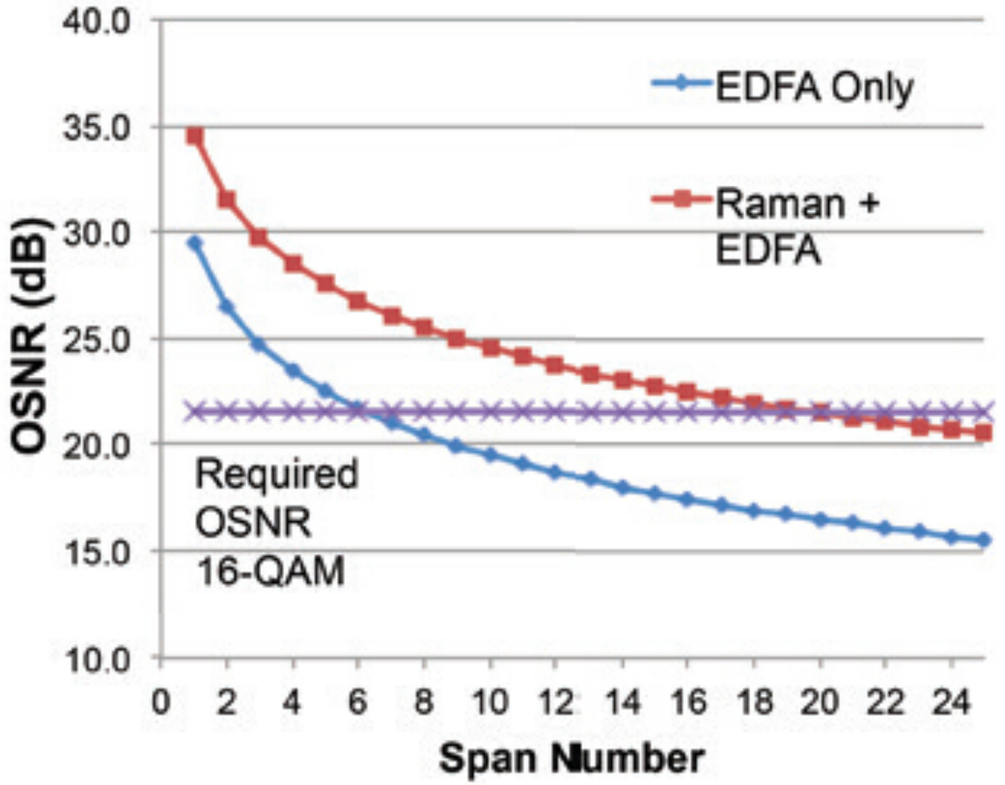

Amplifiers and Power Balance

- Amplifiers introduce some of their own unique issues
  - Amplifier gain is often varies significantly across frequencies
    - Gain Flattening Filters try to compensate for this property.
    - Typical gain variations between channels (“ripple”) are still < 1dB.

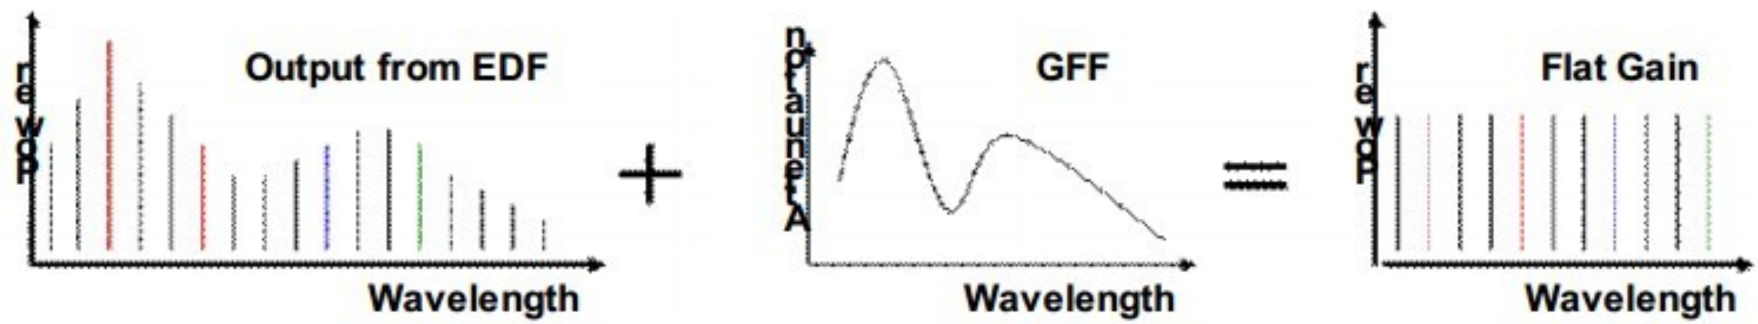

- Unbalanced channel power causes OSNR penalties
  - Even small power variations can add up after several hops.
  - Dynamic Gain Equalization (“DGE”) is required periodically.
  - ROADMs are often used in this role, to balance every channel.

Amplifiers and Total System Power

- Amplifiers also have limits on their total system power
  - Both what they can output, and what they can take as input.
  - And the total input power changes as you add channels.
    - A single channel at +0 dBm is only 1mW of input power.
    - 40 DWDM channels at +0 dBm/each is 40mW, or +16dBm of power.
    - If your amplifier's maximum input power is -6dBm, and you run 40 DWDM channels through it, each channel must be below -22dBm.
    - Failing to plan for this can cause problems as you add channels.
  - The total input power also changes as you lose channels.
    - Imagine power fails to a POP, and many channels are knocked offline.
    - Suddenly the total system power has changed significantly.
    - A good EDFA needs to constantly monitor and adjust power levels.
    - The best EDFAs will communicate with others on the line system.

Other Optical Networking Concepts

Optical Switches

- Optical Switches
  - Let you direct light between ports, without doing O-E-O conversion.
  - Built with an array of tiny mirrors, which can be moved electrically.
  - Allows you to connect two fibers together optically in software.
  - Becoming popular in optical cross-connect and fiber protection roles.

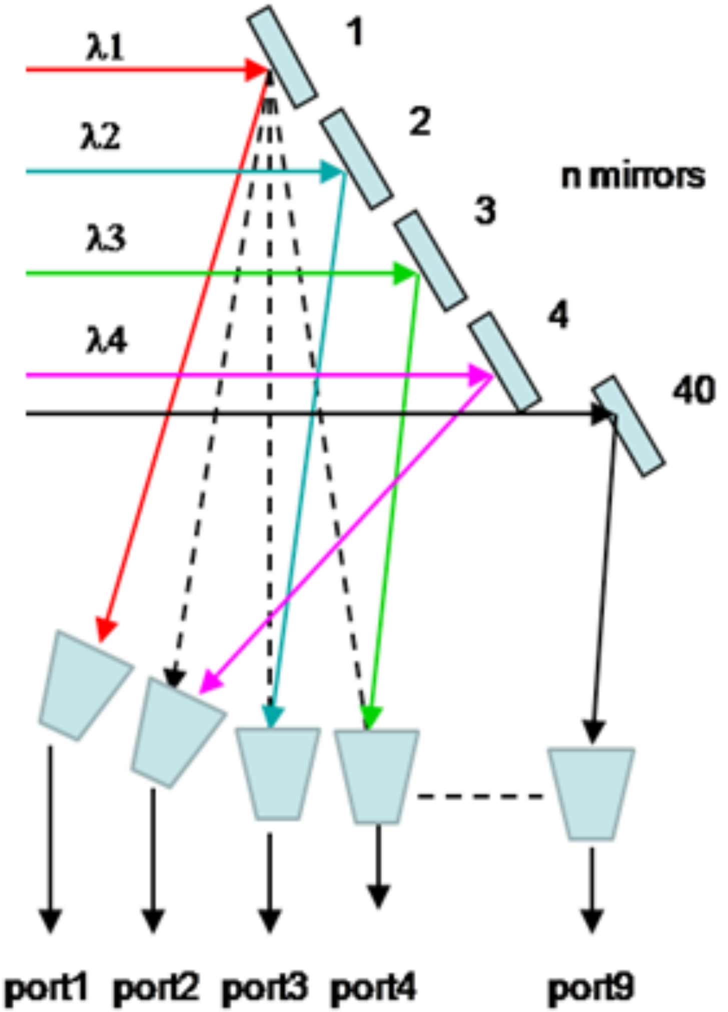

Wavelength Selective Switch (WSS)

- Lets you "route" an individual wavelength across ports
  - The WSS is a key component inside of a ROADM.
  - First generation WSS' used 3D MEMS optical switches.
  - Modern WSS' use Liquid Crystal on Silicon (LCoS).

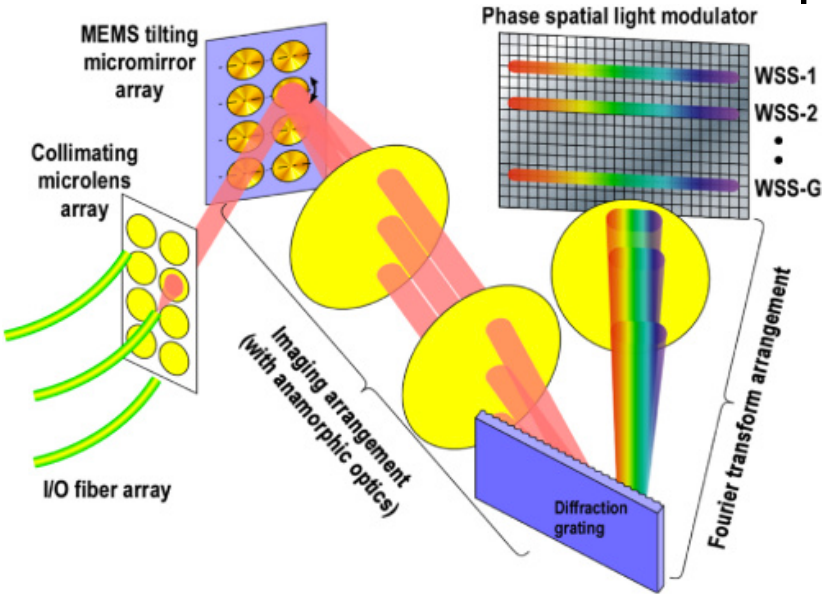

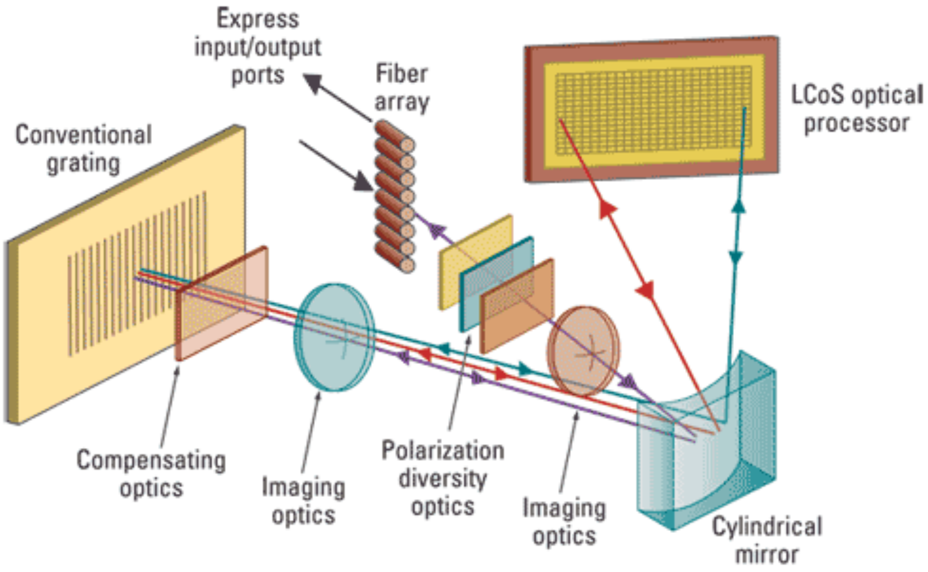

Circulator

- A component typically not seen by the end user.
  - But frequently used inside other popular components.
    - Bragg grating based components, like OADMs and small muxes.
    - Dispersion compensation spools, amplifiers, etc.
    - Very useful when building single-strand bidirectional systems too.
  - A circulator has 3 fiber ports.
    - Light coming in port 1 goes out port 2.
    - Light coming in port 2 goes out port 3.

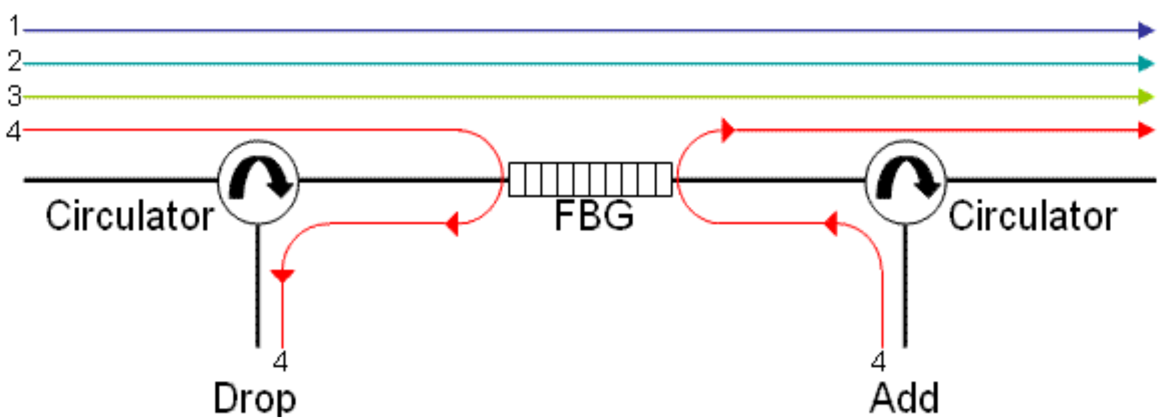

Splitters and Optical Taps

- Do exactly what they sound like they do, split a signal.
- Common examples are:
  - A 50/50 Splitter
    - Often used for simple “optical protection”.
    - Split your signal in half and send down two different fiber paths.
    - Use an optical switch with power monitoring capabilities on the receiver, have it automatically pick from the strongest signal.
    - If the signal on one fiber drops, it switches to the other fiber.
  - A 99/1 or 98/2 Splitter
    - Often used for “Optical Performance Monitoring”.
    - Tap a small % of the signal, and run it to a spectrum analyzer.

Forward Error Correction

- FEC adds extra/redundant information to the transmission, so the receiver can computationally “recover” from any errors.
- In practice, FEC works by lowering the required OSNR, which can help an otherwise unusable signal function normally.
  - Using clever math, padding a 10.325Gbps signal to 11Gbps (7% overhead) can extend a signal from 80km to 120km or beyond.
  - This can really matter when upgrading older DWDM systems.
    - Since it usually isn’t practical to move amp sites closer on a live system.
  - FEC has evolved significantly as well.
    - 1st Gen – RS-FEC – 6% overhead for ~6dB of net coding gain.
    - 2nd Gen - EFEC – 7% overhead for ~8-9dB of net coding gain.
    - 3rd Gen – SD-FEC – 20-25% overhead for 10-11dB coding gain.
      - It might not seem worth it, but a 1-2dB gain in OSNR can hugely increase optical reach.
  - FEC is key to many standards like 100GBASE-SR4 as well.

The Benefits of Forward Error Correction

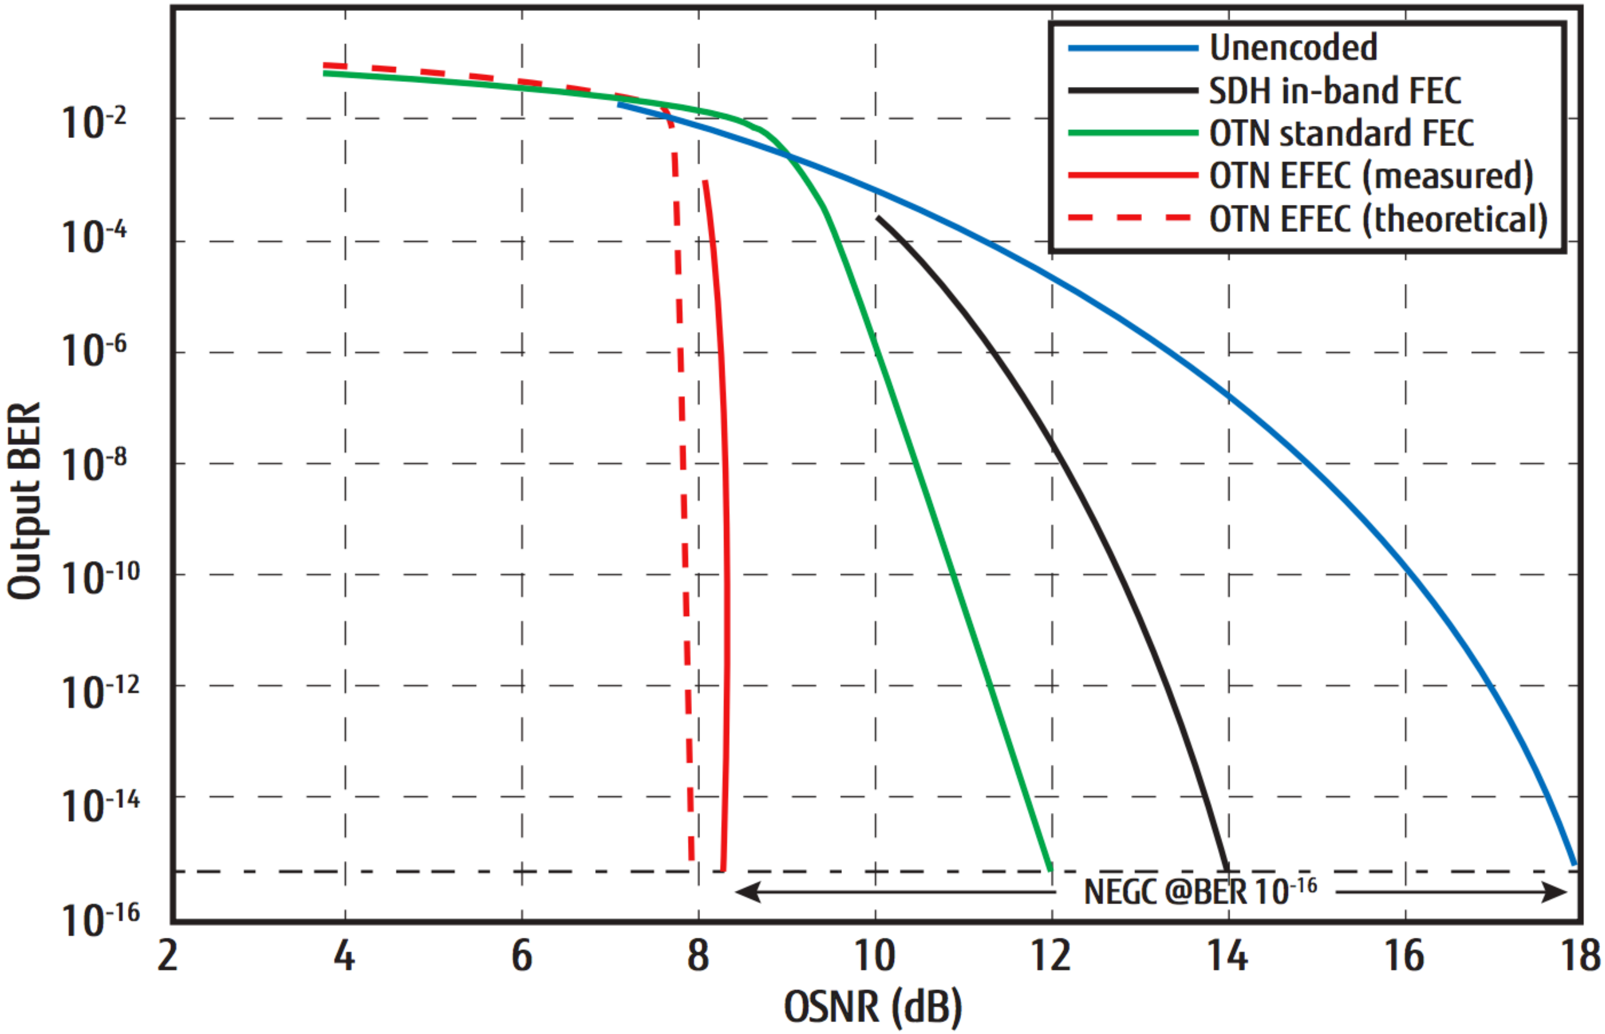

OTN Digital Wrapper Technology (G.709)

OTN stands for Optical Transport Network

- Replacement for SONET/SDH, with support for optical networking.
  - A standard for the generic transport of any protocol across a common optical network, with TDM mux/demux capabilities.
  - Implemented as a "wrapper" around other protocols.
- Why is this needed?
  - Pure optical channels only make sense for high-speed protocols.
    - Example: A single 100GE service, delivered over a 100G wave.
  - Low speed services still need to be aggregated.
    - Example: 10x10GE services on a 100G wave.
  - OTN technology lets the optical network be completely transparent to underlying protocols.
  - Can also help with troubleshooting.

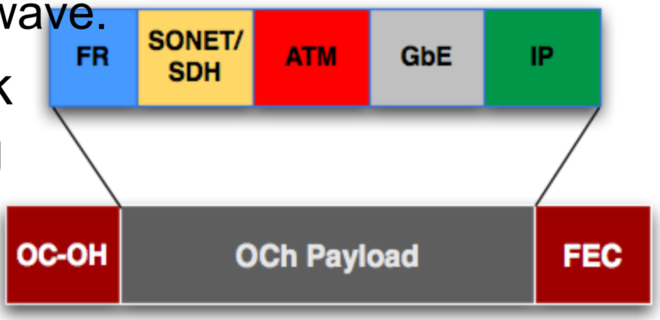

Types of Single Mode Optical Fiber

Types of Single-Mode Fiber

- We’ve already discussed how single-mode fiber is used for essentially all long-reach fiber applications.
- But there are also many different types of SMF.
- The most common types are:
  - “Standard” SMF (ITU-T G.652) A.K.A. SMF-28
  - Full Spectrum (Low Water Peak) Fiber (ITU-T G.652.C/D)
  - Dispersion Shifted Fiber (ITU-T G.653)
  - Low-Loss Fiber (ITU-T G.654)
  - Non-Zero Dispersion Shifted Fiber (ITU-T G.655)
  - Bend Insensitive Fiber (ITU-T G.657)

“Standard” Single-Mode Fiber (G.652)

- One of the original fiber cables.
  - Deployed widely throughout the 1990s.
- Frequently called “SMF-28”, or simply “classic” SMF.
  - SMF-28 is actually a Corning product name.
  - Also called NDSF (Non-Dispersion Shifted Fiber).
- Optimized for use by the 1310nm band.
  - Has the lowest rate of dispersion here.
  - Originally deployed before the adoption of WDM.
- Ironically, has come full circle to again being the best choice for modern high-speed DWDM systems.

Low Water Peak Fiber (G.652.C/D)

- Modified G.652, designed to reduce water peak.
  - Water peak is a high rate of attenuation at certain frequencies due to OH- hydroxyl molecule within the glass.
  - This high attenuation makes certain bands “unusable”.

Absorption of Light by Hydrogen at Various Wavelengths

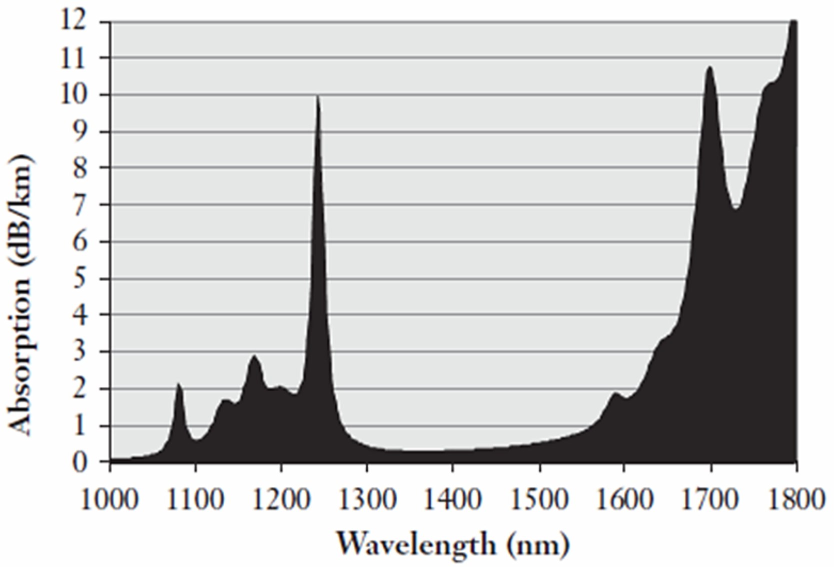

Attenuation of Standard vs. Low Water Peak Fiber

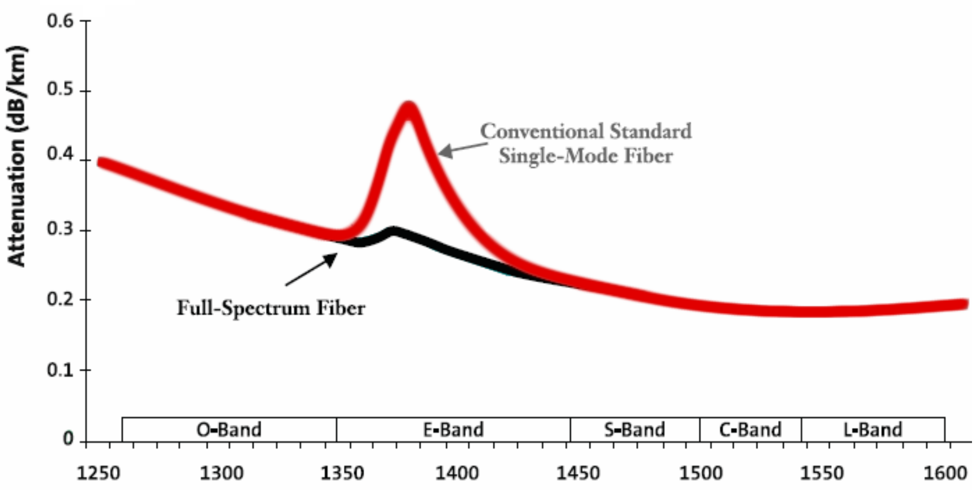

Dispersion Shifted Fiber (ITU-T G.653)

- An attempt to improve dispersion at 1550nm.
  - The rate at which chromatic dispersion occurs changes depending on the frequency of light.
    - The point of lowest dispersion in G.652 occurs at 1300nm.
    - But this is not the point of lowest attenuation, which is around 1550nm.
  - DSF shifts the point of lowest dispersion to 1550nm too.
- But this turned out to cause big problems.
  - Worked well for simple, single channel systems.
  - But running DWDM signals over DSF caused huge amounts of non-linear interactions at high power.
    - The worst of which is Four Wave Mixing (FWM).
  - As a result, this fiber is rarely used today.

Non-Zero Dispersion Shifted Fiber (G.655)

- Similar concept to Dispersion Shifted Fiber
  - But the zero point is moved outside of the 1550nm band.
  - This leaves a small amount of dispersion, but avoids the non-linear cross-channel interactions cause by DSF.
- To manage dispersion, NZDSF comes in 2 types
  - NZD+ and NZD-, with opposite dispersion "slopes".
    - The "transmission fiber" still spreads out 1550nm just a bit.
    - Then "compensation fiber" compresses it in the opposite direction.
  - By switching between the two slopes, the original signal can be maintained even over extremely long distances.

Other Single-Mode Fiber Types

- G.654
  - Ultra low attenuation, high power capable fiber.
  - Designed for ultra-long reach systems like undersea cables.
- G.657
  - Bend Insensitive fiber (reduced sensitivity at any rate).
  - Uses a higher refractive index cladding than normal fiber.
  - Designed for patch cable use, where a perfect bend radius may not be practical.
- Modern fibers are often better than these specs.
  - But much of what’s actually in the ground is old fiber.

Dispersion Rates of Commercial Fibers

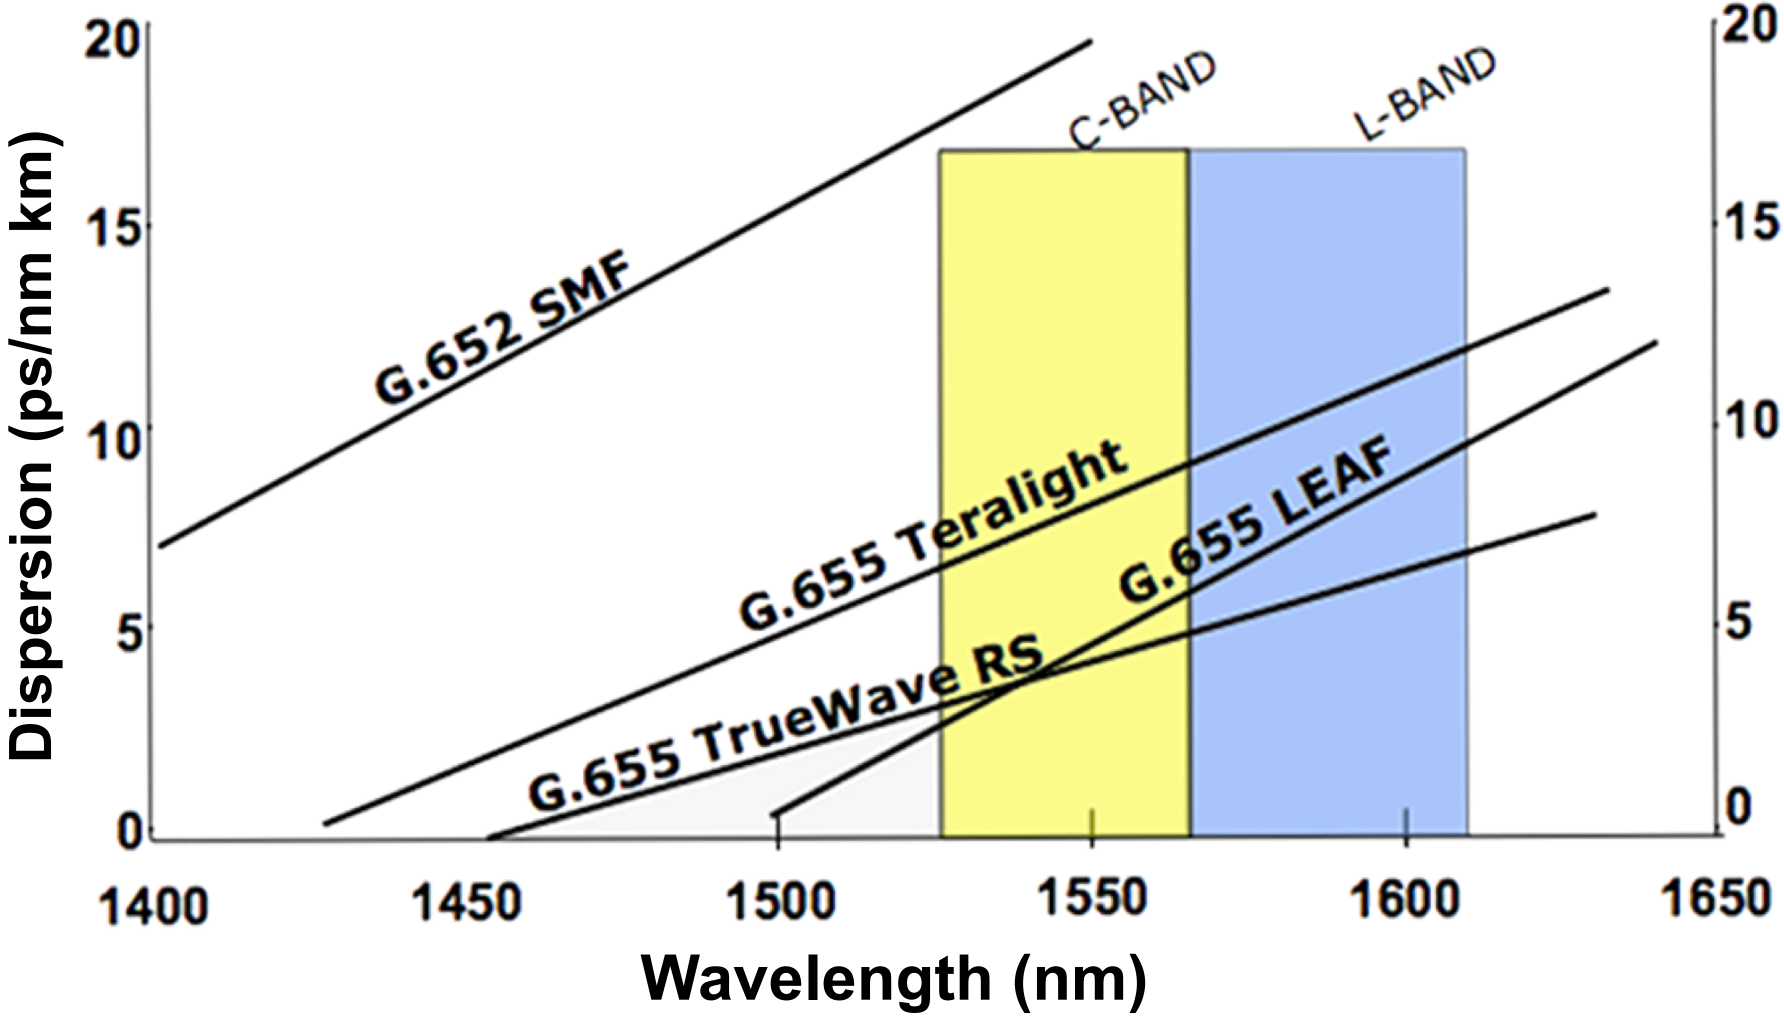

Non-Linear Impairments

Non-Linear Impairments

- Might be better described as “high power problems”.
  - If you don’t transmit at high powers, you’ll never see them.
    - But if you care about reach, you’ll probably be trying to push this.
    - What is “high power”? “Depends”, but usually above +4dBm / channel.
- Non-Linear Impairments can be categorized as:
  - Stimulated Scattering
    - Stimulated Brillouin Scattering (SBS)
    - Stimulated Raman Scattering (SRS)
  - Kerr Effect
    - Intense light causes changes to the refractive index of the fiber.
    - Four Wave Mixing (FWM), Self-Phase Modulation (SPM), Cross-Phase Modulation (XPM)

Stimulated Brillouin Scattering (SBS)

- SBS is the first major impairment to high launch powers
  - Excessive power transmitted into the fiber causes acoustic vibration at an atomic level within the lattice structure of the glass.
  - These vibrations set up Bragg grating effects, causing reflections.
  - Past a certain point, power is reflected back rather than forwards.
  - This limits power, causes errors, and can damage the transmitter.
- SBS is highly dependent on the "power density" in the fiber.
  - Wider linewidths spread the optical power out over more freq.
  - SBS suppression techniques include "dithering" to a wider signal.
  - Coherent helps quite a bit here, higher baud rates do too.
- SBS impact also largely requires long distances of fiber.
  - Putting high power through a very short span may not hurt you.
  - Typical "effective length" maxes out at around 20km.

Stimulated Raman Scattering (SRS)

- SRS is related to the SBS phenomenon.
  - Used intentionally, this is what makes Raman amplification work.
  - Unintentionally, it causes power transfer from one wave to another
- Tighter channel spacing actually REDUCES SRS effects.
  - But adding more total channels increases them.

Example max launch powers, in G.655 NZDSF fiber:

Channel Count	200GHz Spacing	100GHz Spacing	50GHz Spacing
8	15 dBm / ch	18 dBm / ch	21 dBm / ch
16	8.6 dBm / ch	11.6 dBm / ch	14.7 dBm / ch
32	2.5 dBm / ch	5.5 dBm / ch	8.5 dBm / ch
40	0.5 dBm / ch	3.6 dBm / ch	6.6 dBm / ch
80	-5.5 dBm / ch	-2.5 dBm / ch	0.5 dBm / ch

Four Wave Mixing (FWM)

- Regularly spaced signals can interact with each other, to create harmonics in other frequencies.
- The closer they're spaced, the worse the effects.
- Transmission rate independent behavior.
- Uneven channel spacing can reduce the effects.
- FWM is most prevalent in low dispersion fibers.

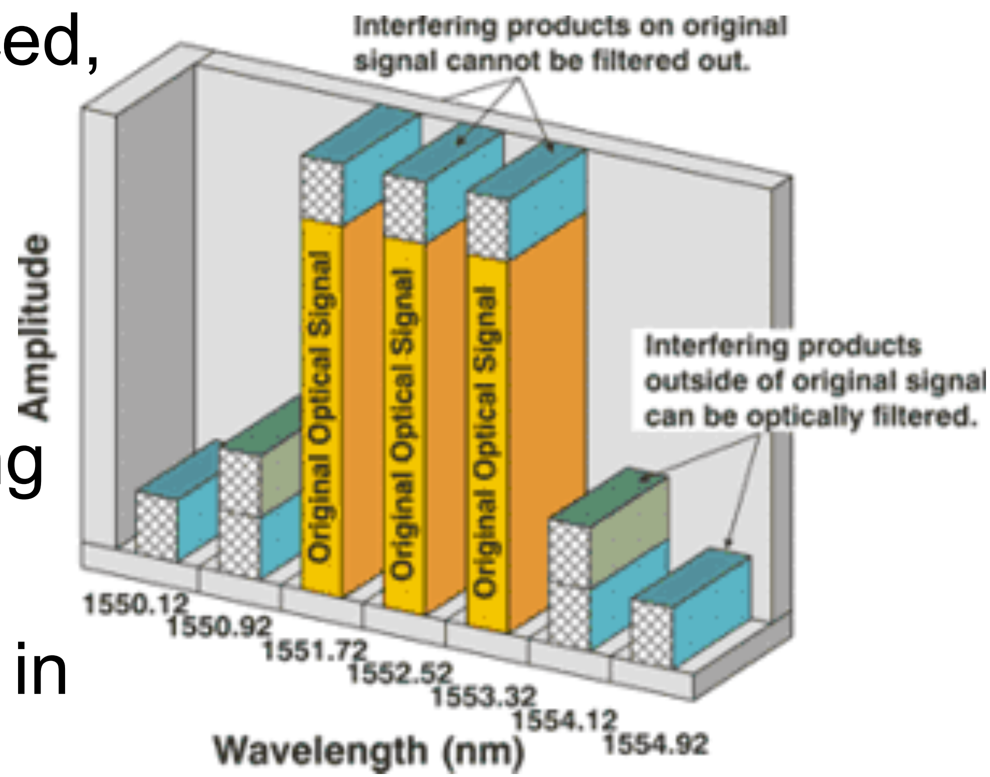

Four Wave Mixing Efficiency

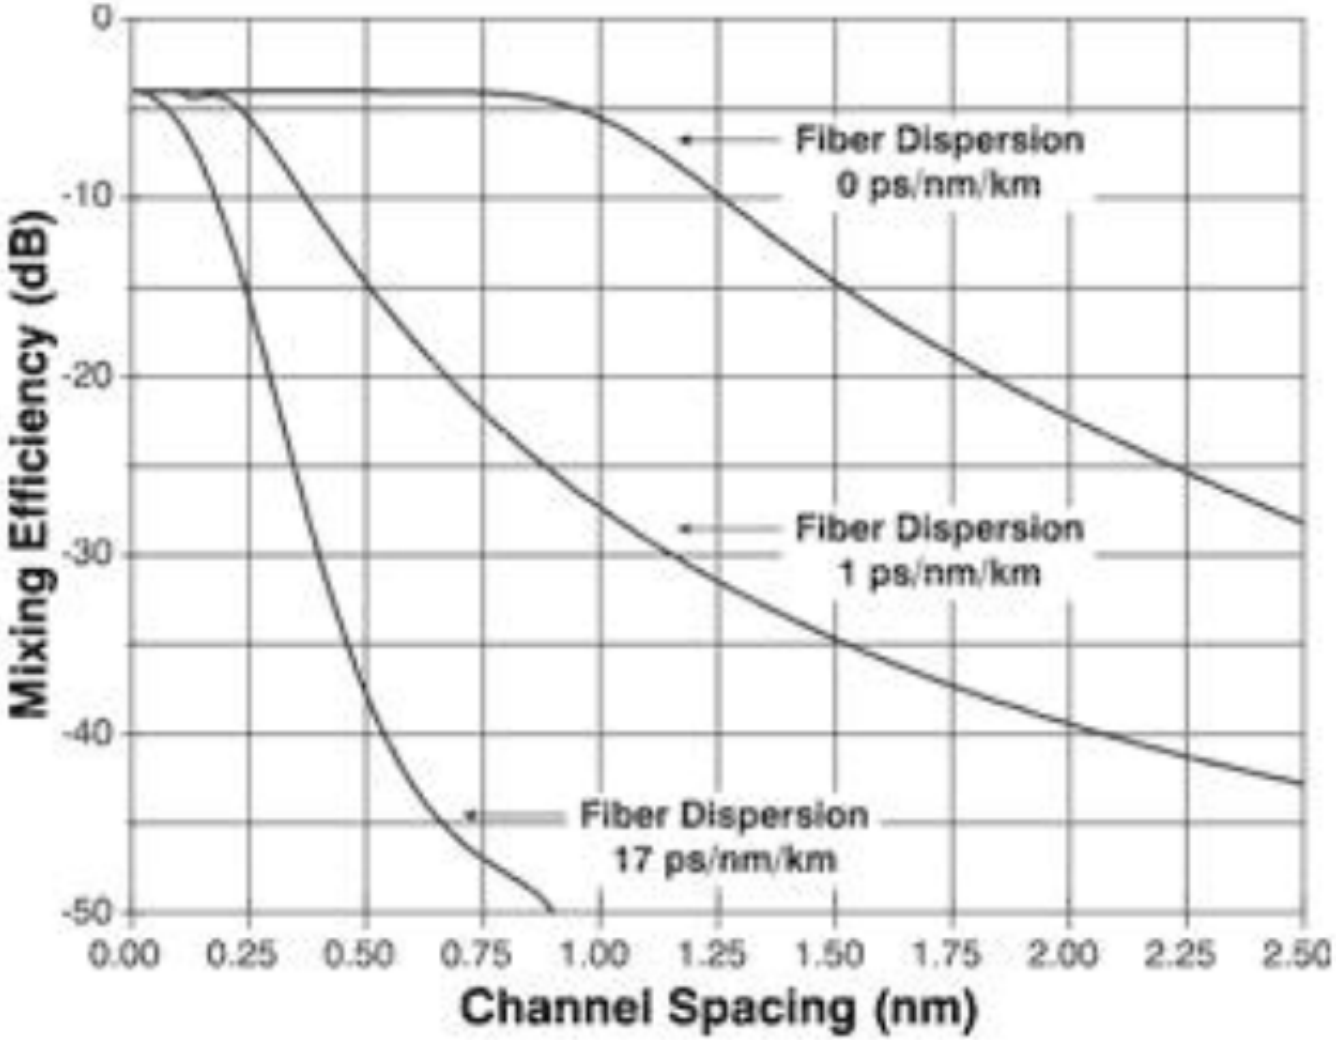

Four Wave Mixing Examples

DWDM Channels		Fiber Chromatic Dispersion Coefficient		
		2 ps/(nm·km)	5 ps/(nm·km)	10 ps/(nm·km)
Number of Channels	Channel Spacing (GHz)	Max Signal Power (dBm)	Max Signal Power (dBm)	Max Signal Power (dBm)
8	10	–11	–6	–4
	25	–3	1	4
	50	3	7	10
	100	9	13	15
16	10	–13	–10	–6
	25	–5	–1	1
	50	0	4	8
	100	6	10	14
32	10	–14	–10	–6
	25	–6	–1	1
	50	0	4	8
	100	6	10	13

Interchannel Effects (XPM, SPM)

- Cross-Phase Modulation (XPM)
  - One wavelength of light can affect the phase of another.
  - Can cause inter-channel cross-talk on DWDM systems.
  - Also caused by mixing NRZ and Coherent systems.
    - Coherent systems actually modulate on phase, so neighboring NRZ channels cause XPM penalties in coherent channels.
    - A 100GHz (minimum) to 200GHz (best) guard band helps this.
  - High CMD helps prevent XPM.
- Self-Phase Modulation (SPM)
  - Occurs when the change in signal power between a 0 and 1 is so strong that it triggers Kerr effect.
  - Low CMD helps prevent SPM.

Non-Linear Threshold Examples

Non-Linear Effect	Max Launch (SMF28)	Max Launch (NZDSF)	Channel Count	Channel Spacing	Line Width
FWM	15 dBm	13 dBm	8	100 GHz	
FWM	13 dBm	10 dBm	32	100 GHz	
SPM	12 dBm	10 dBm	1	N/A	
XPM	15 dBm	11 dBm	8	100 GHz	
SBS	7 dBm	5 dBm	N/A	N/A	10MHz
SBS	15 dBm	13 dBm	N/A	N/A	200MHz
SRS	19 dBm	18 dBm	8	100 GHz	
SRS	5 dBm	3.5 dBm	40	100 GHz	

Nonlinear Effects and Effective Area

- ALL nonlinearities are related to the power “density”.
  - A larger fiber (technically a larger “Mode Field Diameter”) spreads the power over a larger area, reducing peak intensity.
  - This measurement is called a fiber’s “Effective Area” (A_{eff}).
    - If not specified in the fiber specs, use MFD and $\pi * r^2$
  - The quickest way to improve all nonlinearities at once is to use fiber with a larger effective area.
  - Some common examples:
    - Standard G.655 NZ-DSF – 50 μm^2
    - LEAF or TrueWave XL NZ-DSF– 75 μm^2
    - Standard G.652 “SMF28”-based NDSF – 80 μm^2
    - Submarine Fiber (e.g. Corning Vascade) - 150 μm^2
- One tradeoff: Larger Effective Area = Less Raman Gain

Coherent Optical Technologies

Coherent Optical Technologies

- What exactly are "coherent" optics?
  - A group of advancements in optical technology, which combined to deliver vastly improved performance over Direct Detect.
    - Named after the ability to track optical phase ("phase coherence").
  - Specifically, coherent technologies generally consist of:
    - Polarization multiplexing.
    - High-order phase modulation techniques.
    - Using a laser as a local reference oscillator on the receive side.
    - Advanced Digital Signal Processors (DSPs) which are necessary to tie all of these together, recombine the signals, and compensate for impairments.
  - These technologies combined to deliver:
    - Significantly improved spectrum efficiency (went from 1.6 Tbps to 9.6 Tbps+)
    - True 100G/200G and beyond optical signals, not just Nx10G signals.
    - High-bandwidth optical signals which are usable over long distances.
    - Eliminating the need for physical Dispersion Compensation Units.

Improved Modulation Techniques

- Historically optical systems used “IM-DD” modulation.
  - Simplistic “bright for a 1, dim for a 0” type modulation.
  - This yields only 1 data bit per “symbol”, or modulation change.
  - 10GE meant modulating the light 10 billion times/sec, or 10 Gigabaud.
- But adding bandwidth by increasing clock cycles has limitations.
  - For years, the industry was not able to break through the “10G barrier” caused by increasing chromatic and polarization dispersion impairments.
  - Technology advanced only by packing the channels tighter (160 channels in C-band), and throwing more Nx10G’s at the problem.
- Improving the modulation technique yields more bits per symbol.
  - Quadrature Phase Shift Keying (QPSK) delivers 2 bits per symbol.
  - 8 Quadrature Amplitude Modulation (8QAM) delivers 3 bits per symbol.
  - 16 Quadrature Amplitude Modulation (16QAM) delivers 4 bits per symbol.
  - Etc, etc.

Polarization Multiplexing

- Light is (among many other things we don't fully understand yet) a wave of electromagnetic energy, propagating through space.
- In 3-Dimensional space (e.g. a cylindrical fiber), you can send two independent orthogonal signals which propagate along a X and Y axis, without interfering with each other.
- Modern DSPs have it possible to compensate for changing fiber conditions in real time, effectively doubling bandwidth.

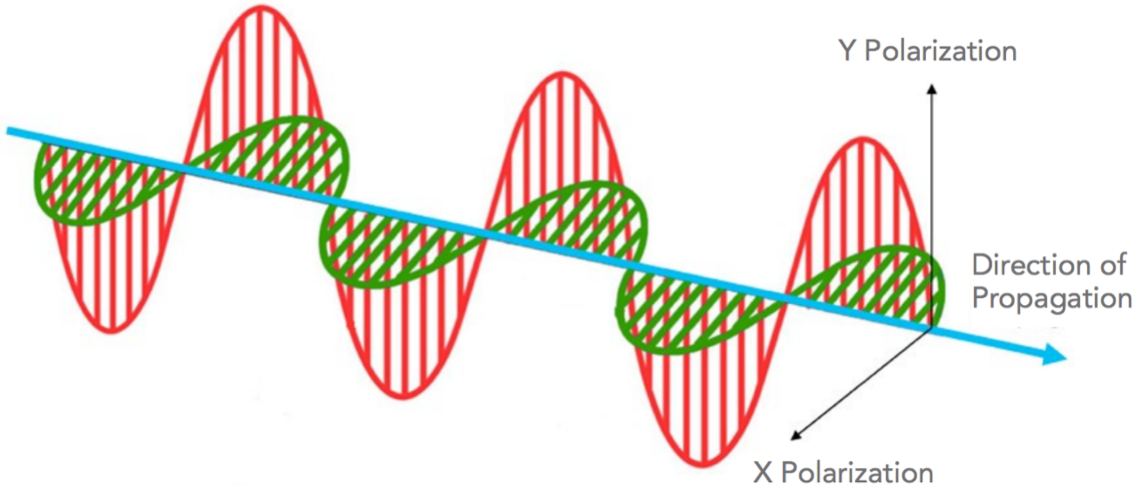

Polarization, Baud, and Modulation

- Total transponder bandwidth is a combination of:
  - Polarization – Today dual polarization, to double capacity.
  - Baud – Higher baud needs wider channel sizes, better DACs.
  - Modulation – Higher modulation needs better OSNR levels.

Data Rate	Baud Rate	Polarities	Modulation Format	Channel Size	Raw BW (with FEC)	Efficiency (bits/s/Hz)	OSNR Required
100G	32G	2	DP-QPSK	37.5GHz	128G	2	10.5 dB
150G	32G	2	DP-8QAM	37.5GHz	192G	3	16.0 dB
200G	32G	2	DP-16QAM	37.5GHz	256G	4	19.5 dB
200G	56G	2	DP-8QAM	62.5GHz	224G	3	17.5 dB
400G	56G	2	DP-32QAM	62.5GHz	560G	5	23.0 dB
200G	64G	2	DP-QPSK	75GHz	256G	4	14.5 dB
400G	64G	2	DP-16QAM	75GHz	512G	4	21.0 dB
600G	64G	2	DP-64QAM	75GHz	768G	6	25.0 dB

More About Coherent

- Other cool features of Coherent technology
  - Need for dispersion compensation all but eliminated.
    - Coherent DSPs eat CMD for lunch - 200,000 ps/nm or more.
    - In fact, Coherent systems work BETTER with CMD.
  - Tunable laser works as a local oscillator on RX side.
    - Coherent can “lock on” to any frequency, and ignore the rest.
    - So in many cases you don’t even need a full filtering mux.
    - Enables colorless/directionless ROADM designs.
- But there are some downsides too.
  - Many components, and expensive / power hungry DSP.
  - Very difficult to integrate into high-density “pluggables”.
    - As of 2017, even 16nm DSPs only fit in CFP2 (called CFO2-DCO).

Engineering an Optical Network

Insertion Loss

- Even the best connectors and splices aren't perfect.
  - Every time you connect two fibers together, you get loss.
  - The typical budgetary figure is 0.5dB per connector.
    - Actual loss depends on your fiber connector and mating conditions.
- Insertion loss is also used to describe loss from muxes.
  - Since it is the "penalty you pay just for inserting the fiber".
  - Some real-life examples:
    - 40-channel DWDM 100GHz Mux/Demux: 3.5dB
    - 80-channel DWDM 50GHz Mux/Demux: 9.5dB
      - Effectively just 2x 100GHz muxes (even+odd) plus an interleaver.

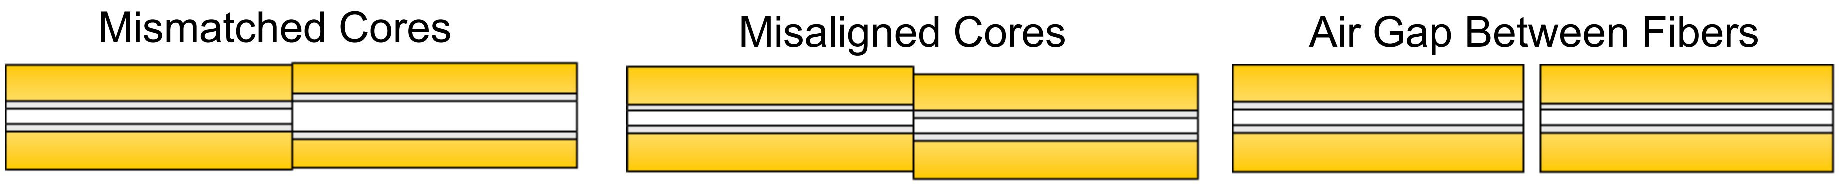

Balling On An (Optical) Budget

- To plan your optical network, you need a budget.
  - When an optic says “40km”, this is only a guideline.
  - Actual distances can be significantly better or worse.
  - It’s also smart to leave some margin in your designs.
    - Patch cables get bent and moved around, optic transmitters will cool with age, a fiber cut and repaired will add more loss, etc.

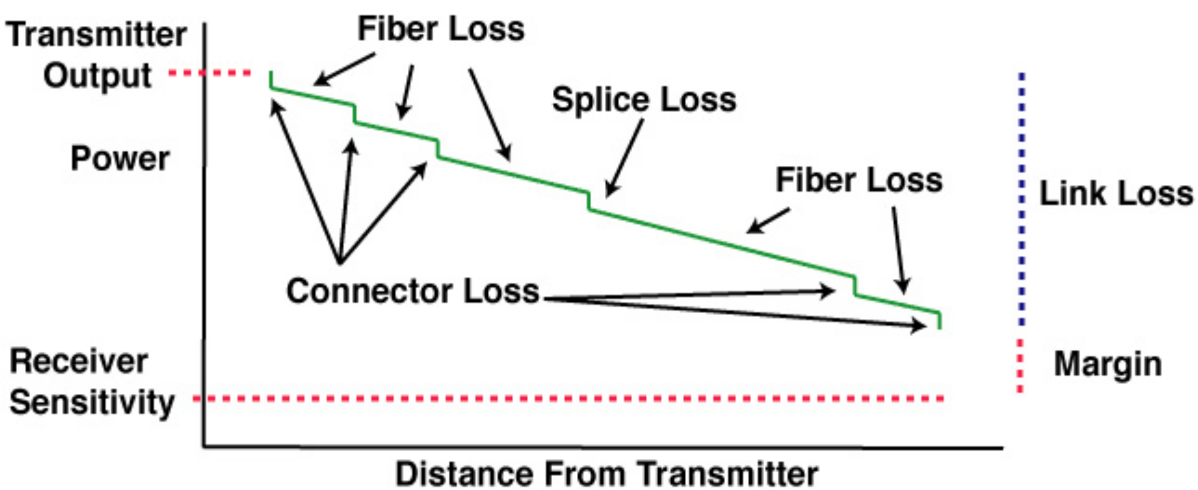

PC vs UPC vs APC

- Beware of the different types of ferrule connectors.
  - (Ultra) Physical Contact
    - Blue Connectors
    - PC - < -30dB Back Reflection
    - UPC - < -55dB Back Reflection

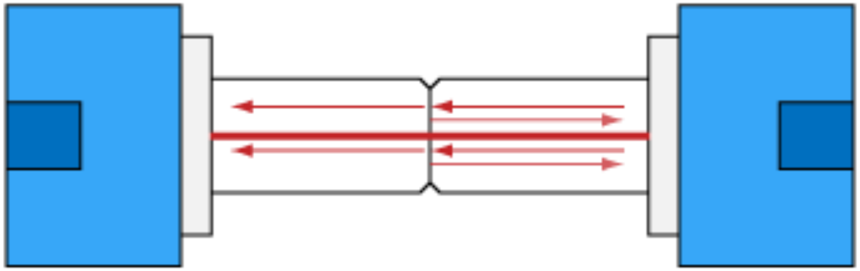

  - Angled Physical Contact
    - Green Connectors
    - 8˚ angle on the ferrule
    - < -65dB Back Reflection
    - Incompatible with PC / UPC!
    - Useful for high power applications

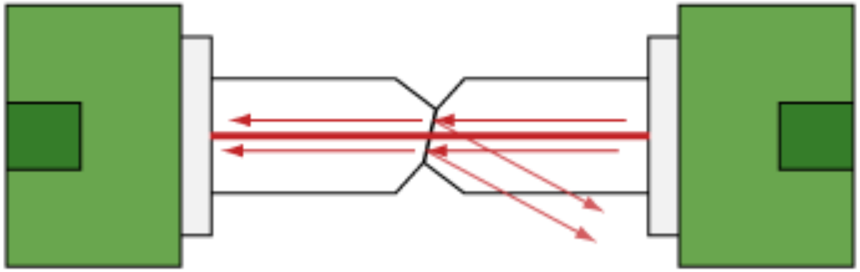

  - Why? When disconnected, even UPC reflects massively.
    - On a high powered amplifier, reflections could cause damage.

Dispersion Compensation Units

- Essentially just big a spool of fiber in a box.
  - Designed to cause dispersion in the opposite direction (with the opposite “slope”) as the transmission fiber used.
  - Passing through this spool reversed the effects of dispersion caused by the transmission fiber.
  - But also adds fiber distance (typically 20% overhead).
  - Usually deployed at amp sites.
  - Best in the middle of a 2-stage amp with mid-stage access.
  - Circulators can reduce the total amount of fiber needed.

Dealing with Dispersion

- Electronic Dispersion Compensation
  - Dispersion which used to completely ruin a signal is now be compensated for electronically at the receiver.
  - Example: 10GBASE-LRM 300 meters over MMF
- Dispersion is worst for Direct Detect systems.
  - PAM4 requires EXTREME CMD compensation.
  - Tolerances of +/- 100 ps/nm, tunable DCM required.
- While coherent systems eat dispersion for lunch
  - They're capable of reading phase information.
  - And use sophisticated Digital Signal Processors (DSPs) to compensation computationally.

Re-amplifying, Reshaping, and Retiming

- Signal Regeneration (Repeaters)
  - Different types are described by the “R’s” that they perform.
  - 1R – Re-amplifying
    - Makes the analog signal stronger (i.e. makes the light brighter)
    - Typically performed by an amplifier.
  - 2R – Reshaping
    - Restores the original pulse shape that is used to distinguish 1’s and 0’s.
  - 3R – Retiming
    - Restores the original timing between the pulses.
    - Usually involves an O-E-O conversion.

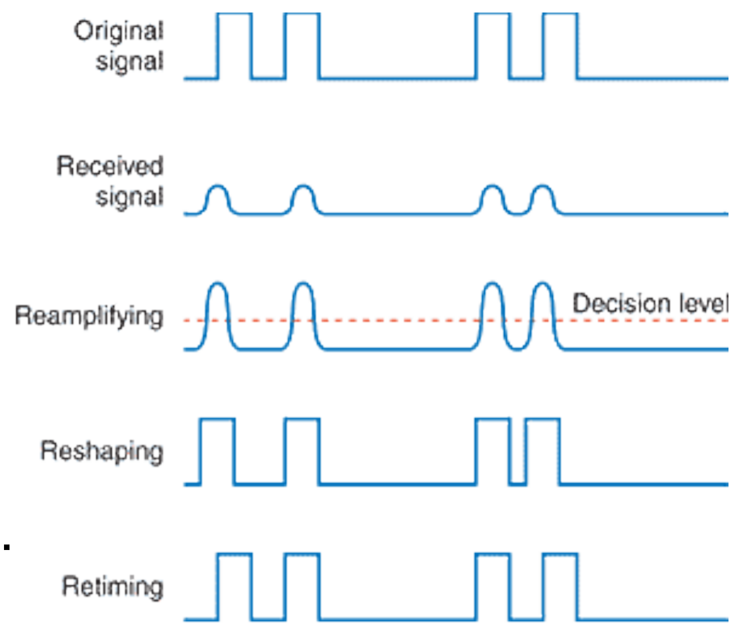

Bit Error Rates (BER)

- As optical impairments add up, links don't just "die".
  - They start taking bit errors, at progressively higher rates.
  - The probability that this will happen is the Bit Error Rate.
- For 99% confidence (100 bit error samples), test:

Date Rate	BER 10^{-9}	BER 10^{-11}	BER 10^{-12}	BER 10^{-13}
100 Gbps	1 sec	2 min	21 min	3 hr 29 min
40 Gbps	3 sec	6 min	53 min	8 hr 47 min
10 Gbps	13 sec	21 min	3 hr 30 min	1d 10 hr 58m
1 Gbps	2 mins	3 hr 30 min	1d 10 hr 58 min	14d 13 hr 33m

OSNR(dB) and Bit Error Rates

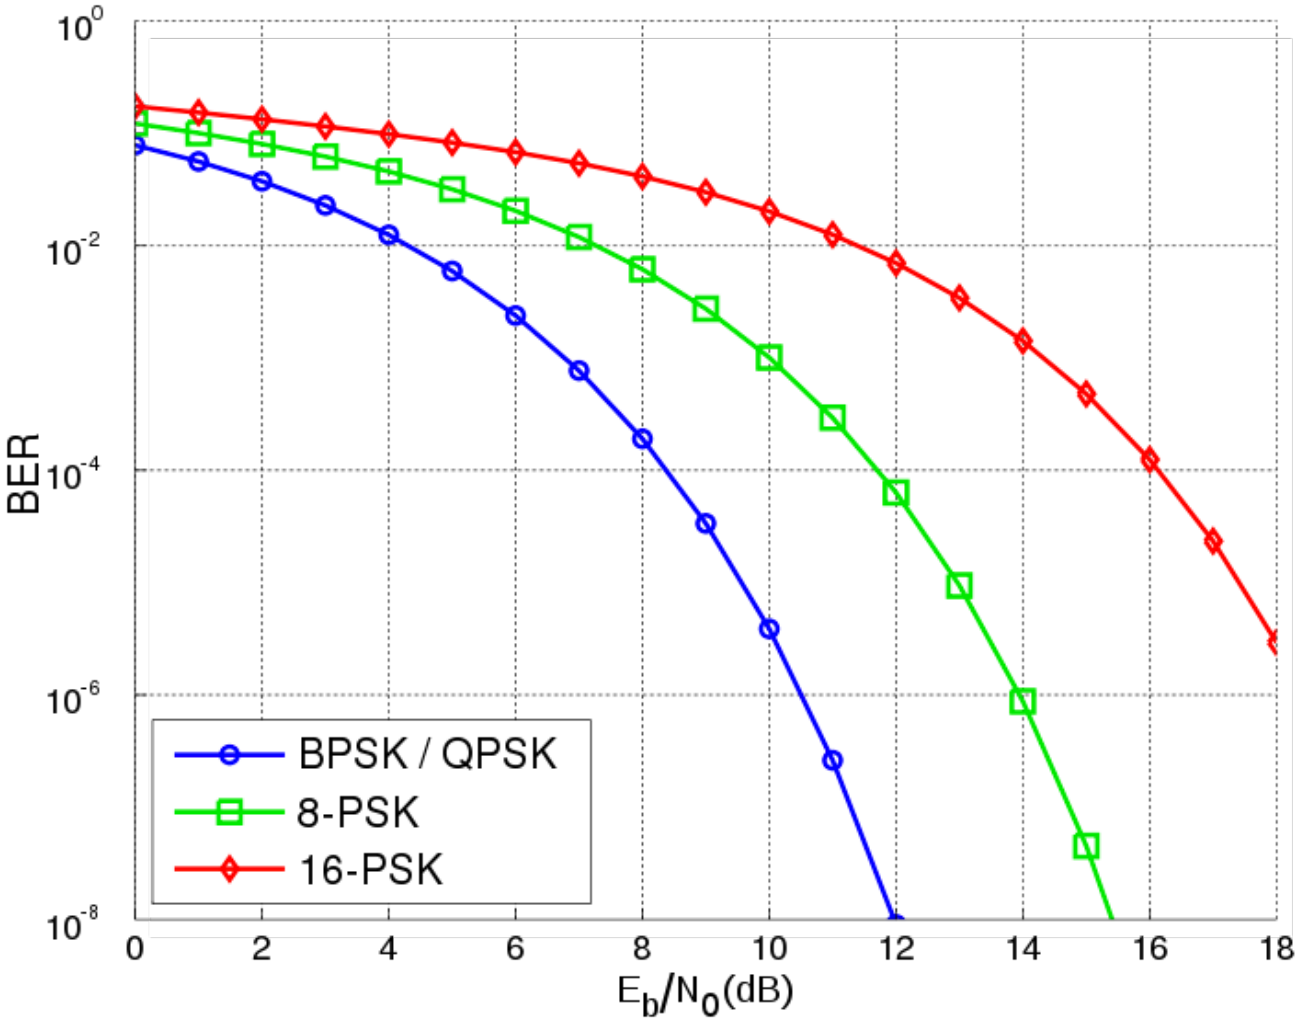

Tools of the Trade

Optical Power Meter (or Light Meter)

- Measures the brightness of an optical signal.
- Displays the results in dBm or milliwatts (mW).
- Most light meters also include a "relative loss" function, as well as absolute power meter.
  - Designed to work with a known-power light source on the other end, to test the amount of loss over a particular fiber strand.
  - These results are displayed in dB, not dBm.
  - Frequently the source of much confusion in a datacenter, when you use the wrong mode!
  - If I had a nickel for every time someone told me they just measured a +70 signal on my fiber…

Optical Time-Domain Reflectometer (OTDR)

- An OTDR is a common tool for testing fiber.
- Injects a series of light pulses into a fiber strand.
- Analyzes the light that is reflected back.
- Used to characterize a fiber, with information like:
  - Splice points, and their locations.
  - Overall fiber attenuation.
  - Fiber breaks, and their locations (distance from the end-point).

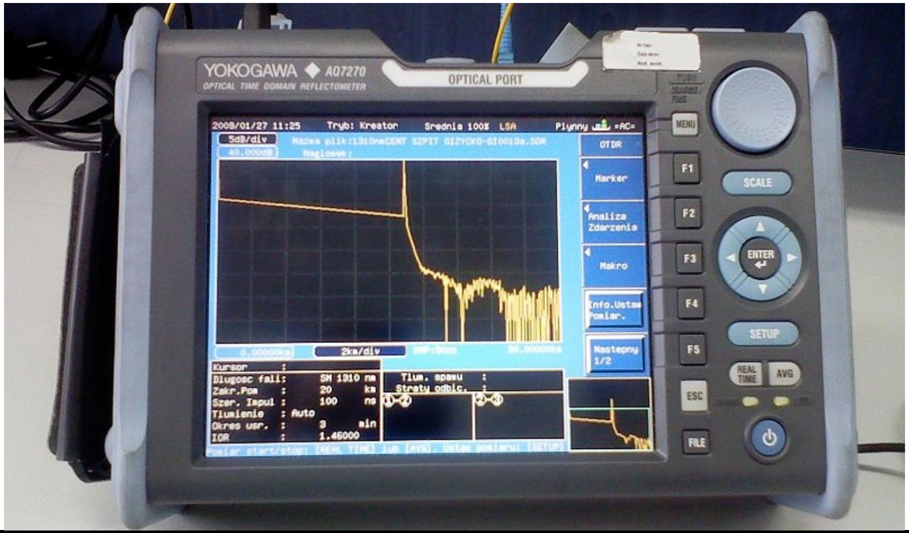

Example OTDR Output

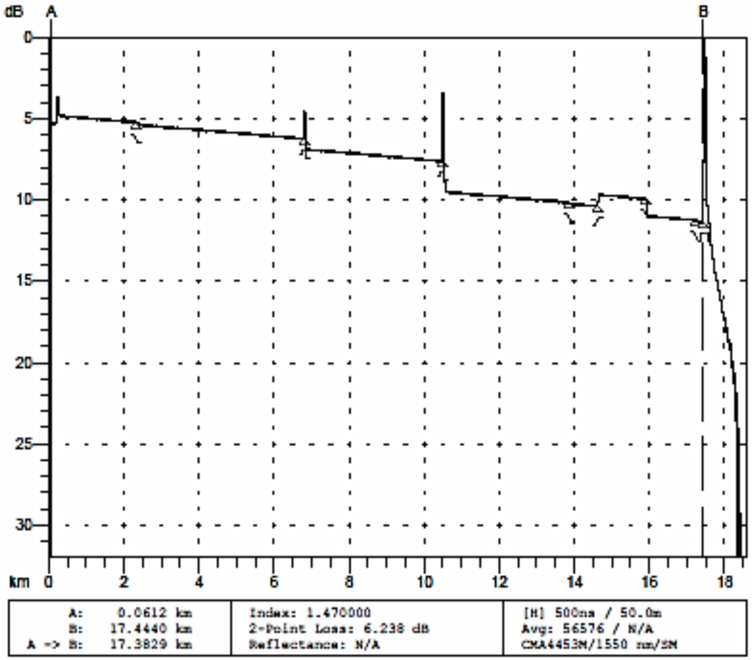

Feature #/Type	Location (km)	Event-Event (dB)	(dB/Km)	Loss (dB)	Refl (dB)
1/N	2.3310	??	??	0.12	
2/R	6.8035	0.91	0.203	0.64	-58.19
3/R	10.4907	0.72	0.196	1.86	-48.24
4/N	13.8639	0.70	0.206	0.06	
5/N	14.6205	0.14	0.188	-0.71	
6/N	15.9114	0.26	0.205	1.06	
7/N	17.2350	0.25	0.193	0.08	
8/E	17.4491	0.05	0.211	>3.00	>-33.55S

Overall (End-to-End) Loss: ??

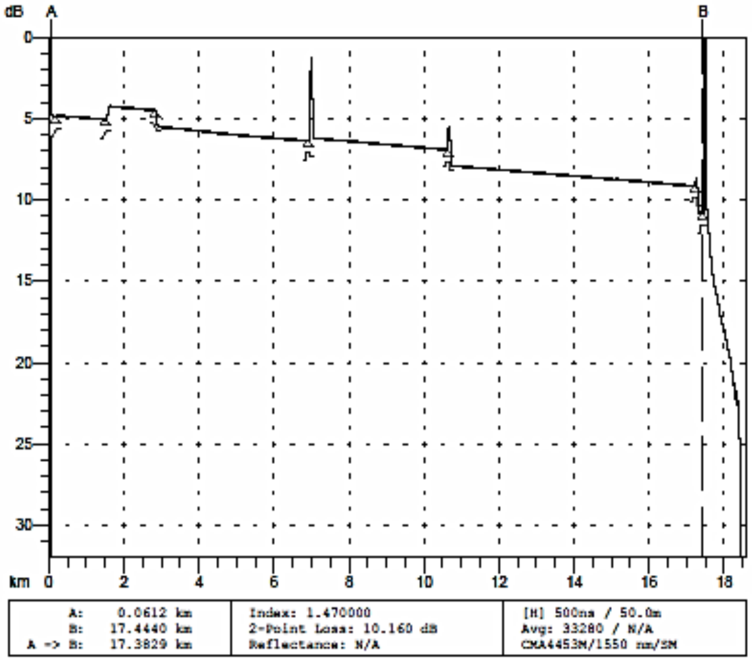

Feature #/Type	Location (km)	Event-Event (dB)	(dB/Km)	Loss (dB)	Refl (dB)
1/N	0.1937	0.02	0.121	-0.06(2P)	
2/N	1.5194	0.24	0.184	-0.82	
3/N	2.8327	0.26	0.197	0.99	
4/R	6.9421	0.90	0.219	-0.21	>-46.37
5/R	10.6396	0.75	0.203	0.96	-56.69
6/R	17.2269	1.28	0.194	1.61	-61.90
7/E	17.4512	0.04	0.184	>3.00	>-34.48S

Overall (End-to-End) Loss: 5.97 dB

Question: Can I really blind myself by looking into the fiber?

Or - Beware of Big Scary Lasers

Laser Safety Guidelines

- Lasers are grouped into 4 main classes for safety:
  - Class 1 – Completely harmless during normal use.
    - Either low powered, or laser is inaccessible while in operation.
    - Class 1M – Harmless if you don't look at it in a microscope.
  - Class 2 – Only harmful if you intentionally stare into them
    - Ordinary laser pointers, supermarket scanners, etc. Anyone who doesn't WANT to be blinded should be protected by blink reflex.
  - Class 3 – Should not be viewed directly
    - Class 3R (new system) or IIIA (old system)
      - Between 1-5mW, "high power" Internet purchased laser pointers, etc.
    - Class 3B (new system) or IIIB (old system)
      - Limited to 500mW, requires a key and safety interlock system.
  - Class 4 – Burns, melts, destroys Alderaan, etc.

Laser Safety And The Eye

- Networking lasers operate in the infrared spectrum
  - Infrared can be further classified as follows:
    - IR-A (700nm – 1400nm) – AKA Near Infrared
    - IR-B (1400nm – 3000nm) – AKA Short-wave Infrared
  - Laser safety levels are based on what can enter the eye.
    - Remember, the human eye didn’t evolve to see infrared.
    - The cornea actually does a good job of filtering out IR-B light.
    - So IR-B has much higher safety limits than visible light.

Max power (continuous, without auto-shutdown features) for IR-B:

Class 1	Class 3R	Class 3B	Class 4
10 dBm	17 dBm	27 dBm	Above 27dBm

Optical Networking and Safety

- Routers
  - Essentially every single-channel laser that can be connected to a router is a Class 1 or Class 1M laser.
  - Even “long reach” 200km+ optics are no exception:
    - A multi-lane optic can have the highest output, e.g. 40G LR4 = 8mW
- Optical Amplifiers
  - Can easily have output powers of 3R (metro) or 3B (long-haul).
  - Raman amplifiers are almost always Class 4.
  - But they all have Automatic Power Reduction/Shutdown too.
- DWDM Equipment
  - Total output power is the sum of all muxed input signals.
  - This can put the total output power into the 3B territory even without amplification, and often has no auto-shutdown feature.

Optical Networking and Safety

- So should I be wearing goggles in the colo?
  - Generally speaking, your standard client optics are always Class 1 (completely safe under all conditions).
  - Even on amplified/DWDM systems, light rapidly disperses as soon as it leaves the fiber and travels through air.
  - Wavelengths above 1400nm are IR-B, and are mostly blocked by the human eye. Most high power optics and long-reach systems are in this range.
  - High-power systems are legally required to have auto-shutdown safety mechanisms if they detect a cut.
- But, don't hold a DWDM mux directly to your eye.
- And be extra careful with a fiber microscope.

Why Look Into The Fiber Anyways?

- Can you even see the light at all?
  - No, the human eye can only see between 390 – 750nm.
  - No telecom fiber signal is directly visible to the human eye.
- But, I looked at 850nm and I saw red?
  - What you're seeing are the sidebands of an imperfect signal generation, not the main 850nm signal itself.
  - Many digital cameras can see infrared.
  - One trick to check for light in a fiber is to hold it up to your camera phone.
    - You can try this on your TV's remote control.
    - Except newer/nicer ones filter IR, for picture quality. iPhone started blocking IR as of 4S/5.

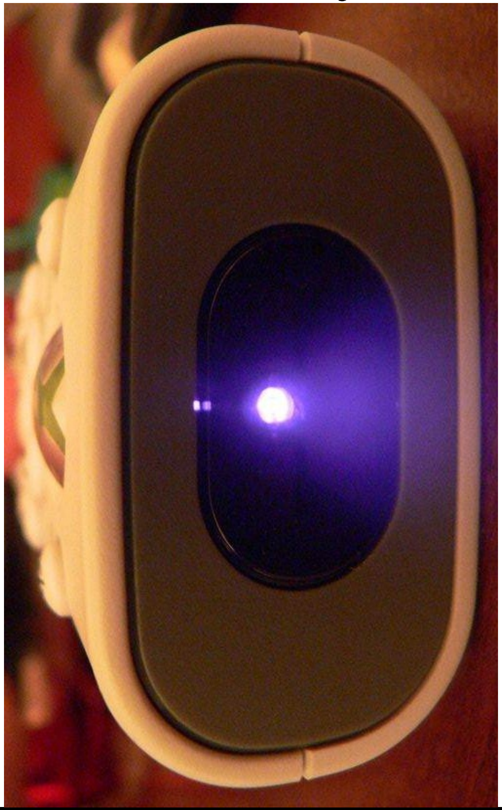

Question: Can optical transceivers be damaged by over-powered transmitters?

Damage by Overpowered Transmitters?

- Well, yes and no.
  - Actually, most optics transmit at roughly the same power.
    - The typical output of 10km vs 80km optics are within 3dB.
  - Long reach optics achieve their distances by having more sensitive receivers, not stronger transmitters.
    - 80km optics may have a 10dB+ more sensitive receiver than 10km
    - These sensitive receivers are what are in danger of burning out.
  - There are two thresholds you need to be concerned with.
    - Saturation point (where the receiver is "blinded", and takes errors).
    - Damage point (where the receiver is actually damaged).
    - The actual values depend on the specific optic.
    - But generally speaking, only 80km optics are at risk.

Tx and Rx Optical Power Ranges

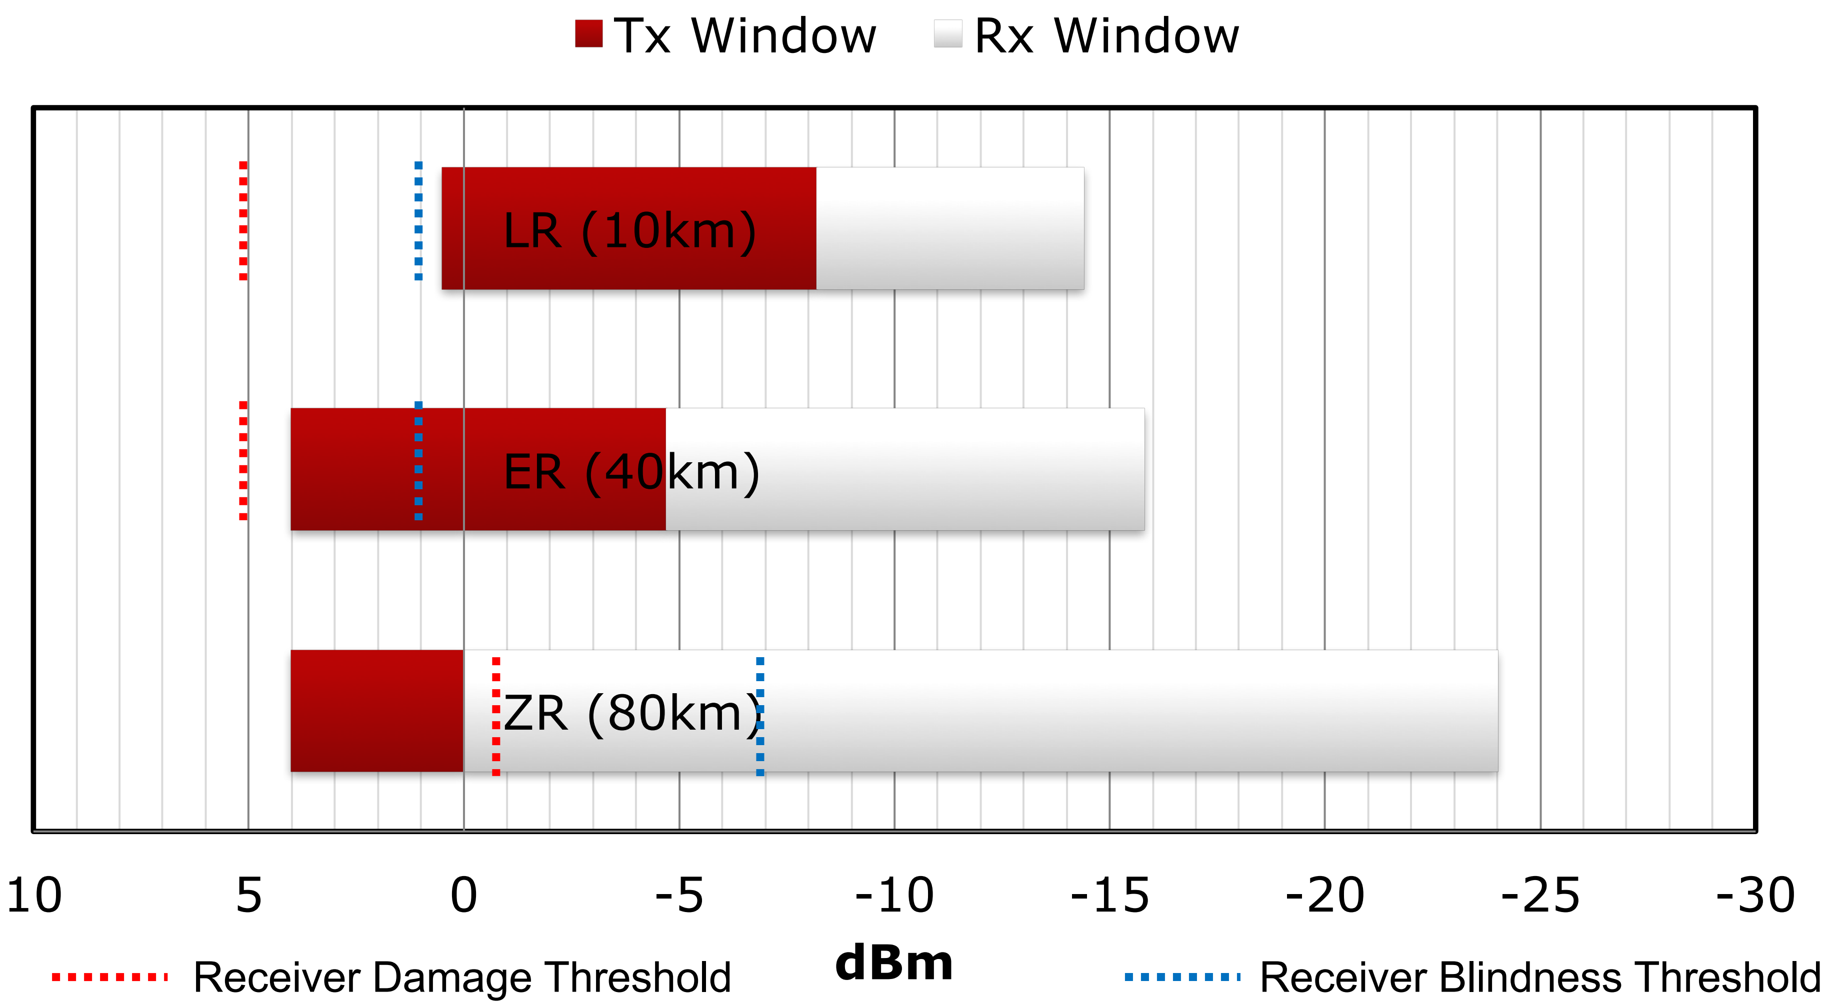

Question: Do I really need to be concerned about bend radius?

Is Bend Radius Really A Concern?

- Yes, bend radius is a real issue.
  - Remember that total internal reflection requires the light to hit the cladding below a "critical angle".
  - Bending the fiber beyond its specified bend radius causes the light to "leak" out.
  - There are "bend insensitive" fibers, though they usually trade some level of performance for this.
  - These are pretty useful in datacenter applications, when humans don't do the right thing.

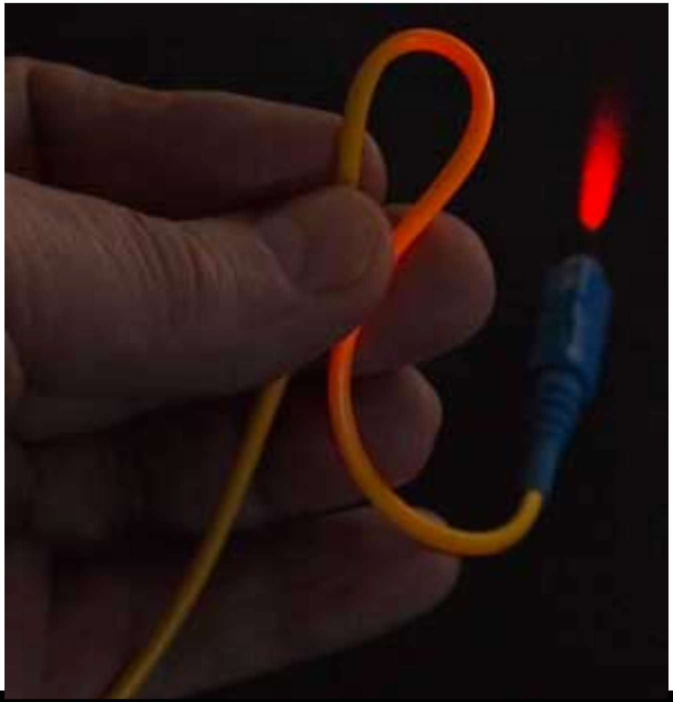

Practical Bend Radius Examples (SMF)

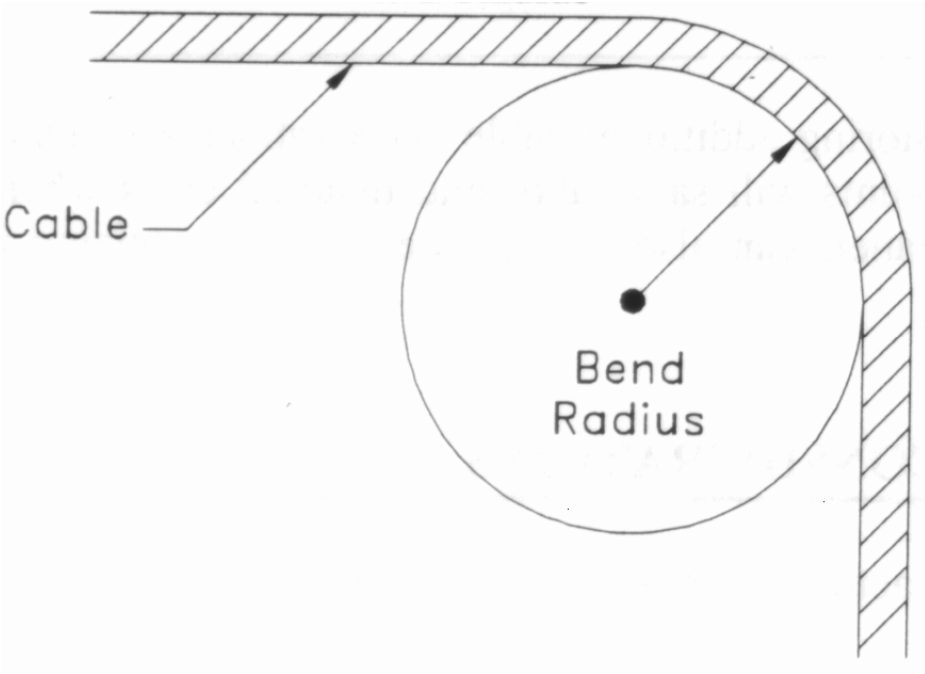

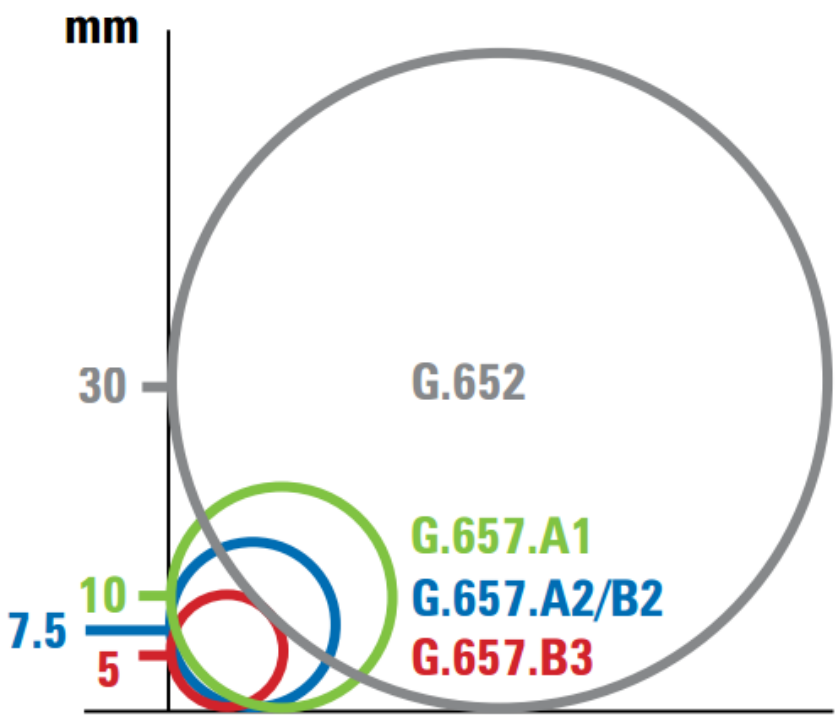

Question: Can two transceivers on different wavelengths talk to each other?

Can You Mismatch Transceiver Freqs?

- Between certain types of optics, yes.
  - All optical receivers have wideband photodetectors.
    - Laser receivers “see” everything between 1260nm – 1620nm.
    - But they won’t be able to see a 850nm LED, for example.
    - Coherent receivers can even “lock on” to one specific frequency.
  - Many DWDM networks are build around this premise.
    - By using one wavelength going A->B and other going B->A, you can achieve a bidirectional system over a single fiber strand.
    - The DWDM filters (muxes and OADMs) provide hard cut-offs of certain frequencies, but the transceivers can receive any color.
  - The only “gotcha” is optical power meters will be wrong.
    - A meter that is calibrated to read a 1310nm signal will see a 1550nm signal just fine, but its power reading will be a few dB off.

Can You Mismatch Transceiver Freqs?

- You can also mismatch frequencies for added reach.
  - You can achieve nearly as much distance with an LR/ER pair (1310nm 10km / 1550nm 40km) as with an ER/ER pair.
    - The ER transmits at 1550nm, which has a lower rate of attenuation.
      - Around 0.2dB/km vs 0.35dB/km, depending on fiber type.
      - So the LR side receives a much stronger signal than the ER side.
    - And the ER optic has a much greater RX sensitivity than the LR.
      - So it will be able to hear the 1310nm signal much better than an LR optic would in the same position.
- Result:
  - You may only *need* a long reach optic on one side.

Question:
Do I Really Need to Clean the Fiber?

Do I Really Need to Clean the Fiber?

- Dirt can actually DAMAGE fiber permanently.
  - A mating force of 2.2lb, over a 200µm surface area...
  - Is 45,000 lbs per square inch of pressure.
  - This can permanently pit and chip your fiber cables!
- Buy a cheap cleaning kit!

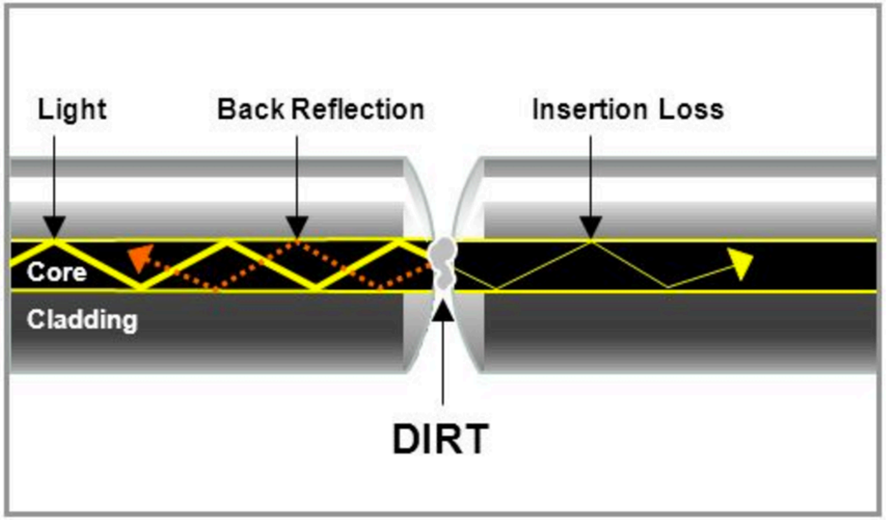

Other Misc Fiber Information

How Fast Does Light Travel In Fiber?

- Ever wondered how fast light travels in fiber?
  - The speed of light is 299,792,458 m/sec
  - SMF28 core has a refractive index of 1.4679
  - Speed of light / 1.4679 = 204,232,207 m/sec
  - Or roughly 204.2 km/ms, or 126.89 miles/ms
  - Cut that in half to account for round-trip times.
    - So, approximately 1ms per 100km (or 62.5 miles) of RTT.
- Why do you see a much higher value in real life?
  - Remember, fiber is rarely laid in a straight line.
  - It is often laid in rings which take significant detours.
  - Dispersion compensation can add extra distance too.

Send questions, comments, complaints to:

Richard A Steenbergen <ras@turkbergen.com>