

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

Richard A Steenbergen <ras@turkbergen.com>

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?
 - Probably not.
 - The purpose of this tutorial is to touch on a little of every topic, from the mundane to the unusual.
 - But it helps to have a basic understanding of these topics, 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 and flexible material (relative to copper).
 - Carries tremendous amounts of information (20 Tbps+ today).
 - Can easily carry many different completely independent signals on the same strand, without interference.
 - Can be sent 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.
 - And is typically reserved for systems where the fiber in question is relatively expensive.
 - As with most things in business, cost is the 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 “IMDD”.
 - Which stands for “Intensity Modulation with Direct Detection”.
 - Typically encoded as “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 25GBaud, this technique gets increasing hard to scale.
- “Better than NRZ” systems are becoming more pervasive.
 - QPSK 100G is the basis of most long-haul links today.
 - 100GBASE-PAM4 (Pulse Amplitude Modulation) QSFP28 optics delivering 80km 100G cheaply are starting to ship today, etc.

The Most Basic Distinction in Fiber Types: 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 distortions” 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 50 μm 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/28G 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 distortions
 - Can easily transmit a signal several thousand kilometers (with appropriate amplification), without requiring regeneration.
 - Typically supports distances of ~80km even without amplification.
- SMF has an even broader array of types than MMF.
 - OS1 "tight buffered" for indoor use, OS2 "loose" for buried use.
 - "Classic" SMF can be called "SMF-28" (a Corning product name)
 - But it also comes in many different formulations of Low Water Peak Fiber (LWPF), Dispersion Shifted Fiber (DSF), Non-Zero Dispersion Shifted Fiber, Bend Insensitive Fiber, etc, etc, etc.

Understanding Modal Distortion in MMF



Mode Conditioning Cables

- What happens to a “narrow” laser inside “wide” MMF?
 - It gets bounced around, causing modal distortions.
- This can be improved with a Mode Conditioning Cable
 - A manufactured splice between the SMF and MMF cables, precisely setting the angle of the light sent into the MMF.
 - By controlling the angle, modal distortions can be reduced, allowing greater distances to be achieved over MMF.
 - For example, a 1GE LX over MMF would go from 300M to 550M.



The “What Happens When You...” Table

Transmit Optic Type	Multimode Fiber	Single-mode Fiber
LED/VCSEL Source (GigE SX, FDDI, etc)	Limited by modal distortion, achieves a few hundred meters depending on the exact signal and type of fiber.	Limited by attenuation, diffuse signal doesn't fit into the narrow fiber core. It may actually work, but for a few meters at best.
Laser Source (LX/LR, ER, ZX/ZR, etc)	Limited by modal distortion, but should perform as well or better than an LED source. Not recommended, but it “works” (with a dB hit) if you pass a LR signal through a short bit of MMF (e.g. patch cable)	Achieves maximum distance determined by signal attenuation and other criteria (10km, 40km, 80km, etc).

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, $1/10^{\text{th}}$ 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 $1/10^{\text{th}}$ the signal, -20dB is $1/100^{\text{th}}$ 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.
- There are two main types of dispersion to deal with
 - Chromatic Dispersion
 - Different frequencies of light propagate through a non-vacuum at slightly different speeds. This is why optical prisms work.
 - But if one part of an optical signal travels faster than the other part, the signal will eventually “smear out” over long distances.
 - Polarization Mode Dispersion
 - Caused by imperfection in shape of the fiber (not perfectly round).
 - One polarization of light propagates faster than the other.
 - Older fiber is particularly affected, may get worse with age.

The Effects of Dispersion


- As the signal is dispersed, it is no longer distinguishable as individual pulses at the receiver.



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 (up to 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



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.




Different Types of WDM

- There are several different types of WDM
 - The most common terms are Dense and Coarse.
 - Essentially they both do the same thing in the same way.
 - The only difference is the channel spacing.
 - And sometimes the range of the optical spectrum they cover.



Coarse Wavelength-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 Wavelength-Division Multiplexing

- So what does that make Dense WDM (DWDM)?
 - A much more tightly packed WDM system.
 - Defined by the ITU-T G.694.1 as a “grid” of specific channels.
 - Within C-band, these channel spacings are common:
 - 200GHz – 1.6nm spacing, 20 channels
 - 100GHz – 0.8nm spacing, 40 channels
 - 50GHz – 0.4nm spacing, 80 channels
 - 25GHz – 0.2nm spacing, 160 channels
 - A rough guideline:
 - 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 for high-density 10G systems, before the move to more modern 100G systems at 50GHz spacings.

What Are The Advantages?

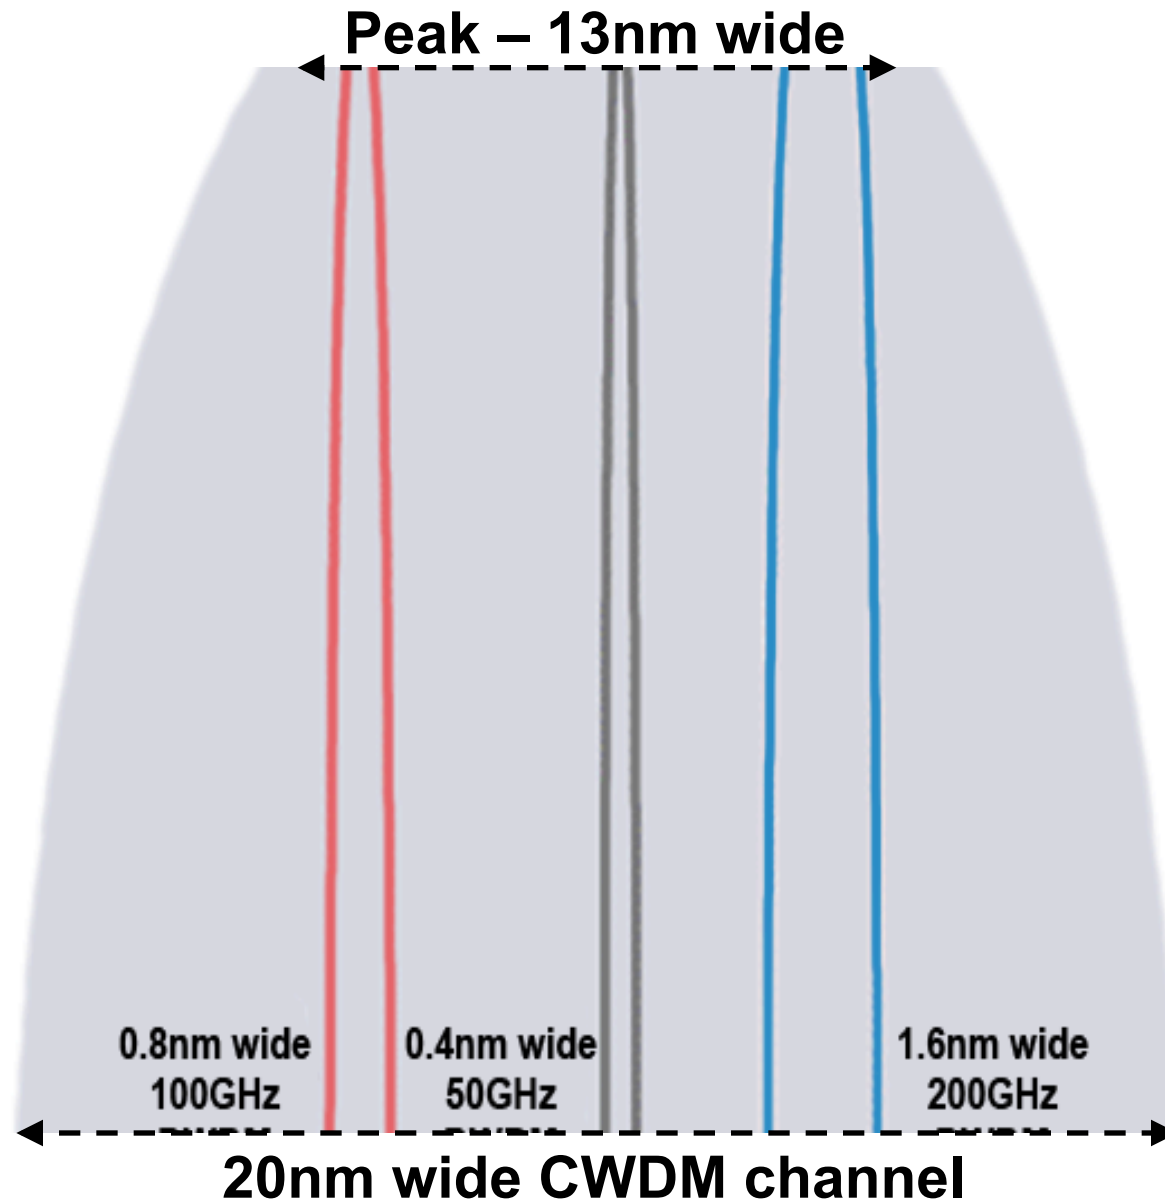
- CWDM

- 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 than other bands.
 - Where Erbium Doped Amplifiers (EDFAs) work.

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
 - Simple combination of two popular windows onto a single strand.
 - 4-lane “Grey” Optics
 - Sometimes it’s much easier/cheaper to implement multi-lane optics
 - 10GE had 10GBASE-LX4 (4x 2.5G channels rather than 1x 10G)
 - 40GE still has SR4/LR4 4x10G, 100G still has SR4/LR4 4x25G
 - These use 600GHz 1295.6nm / 1300.1nm / 1304.6nm / 1309.1nm spacing.
 - Single Strand Optics (BX “bidirectional” standards)
 - E.g. 1310 / 1490nm mux integrated into a pluggable transceiver.



WDM Networking Components



WDM Mux/Demux

- The Mux/Demux
 - Short for “multiplexer”, sometimes called a “filter” or “prism”.
 - The term “filter” is how it actually works, by filtering specific colors.
 - But most people understand a “prism” splits light into the spectrum.
 - A simple device which combines or splits multiple colors of light into a single fiber (called the “common” fiber).
 - Muxes are entirely passive devices, requiring no power.
 - A complete system requires a mux+demux for TX and RX.
 - Most modern devices function the same in both directions, as a mux or demux, so the actual device is the same.
 - Many vendors combine the mux+demux into a single unit for simplicity, but it is really 2 distinct components.

How a Mux Works


- Muxes are actually optical bandpass filters
 - Typically based on Bragg Grating or Dichroic filters.
 - Some frequencies are reflected, the rest are passed through.
 - The channels actually overlap slightly, but have enough isolation to prevent cross-talk interference.

Bragg Grating Effects




The Optical Add/Drop Multiplexer (OADM)

- The Optical Add/Drop Multiplexer (OADM)
 - Selectively Adds and Drops certain WDM channels, while passing other channels through without disruption.
 - Where a mux is used at the end-point of a WDM network to split all of the component wavelengths, an OADM is used at a mid-point, often in a ring.
 - With a well-constructed OADM ring, any node can reach any other node in the ring, potentially reusing the same wavelength multiple times across different portions of the ring.




The ROADM

- The “Reconfigurable Optical Add/Drop Multiplexer”.
 - Essentially a “software tuneable OADM”
 - Essentially a “tunable OADM”, usually in software.
 - Allows you to control which channels are dropped and which are passed through, increasing channel flexibility.
- Some ROADMs are multi-degree
 - Instead of only being able to “pass” or “drop”, there are more than 2 directions of “pass” to choose from.
 - This allows you to build complex star topologies at a purely optical level.




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.



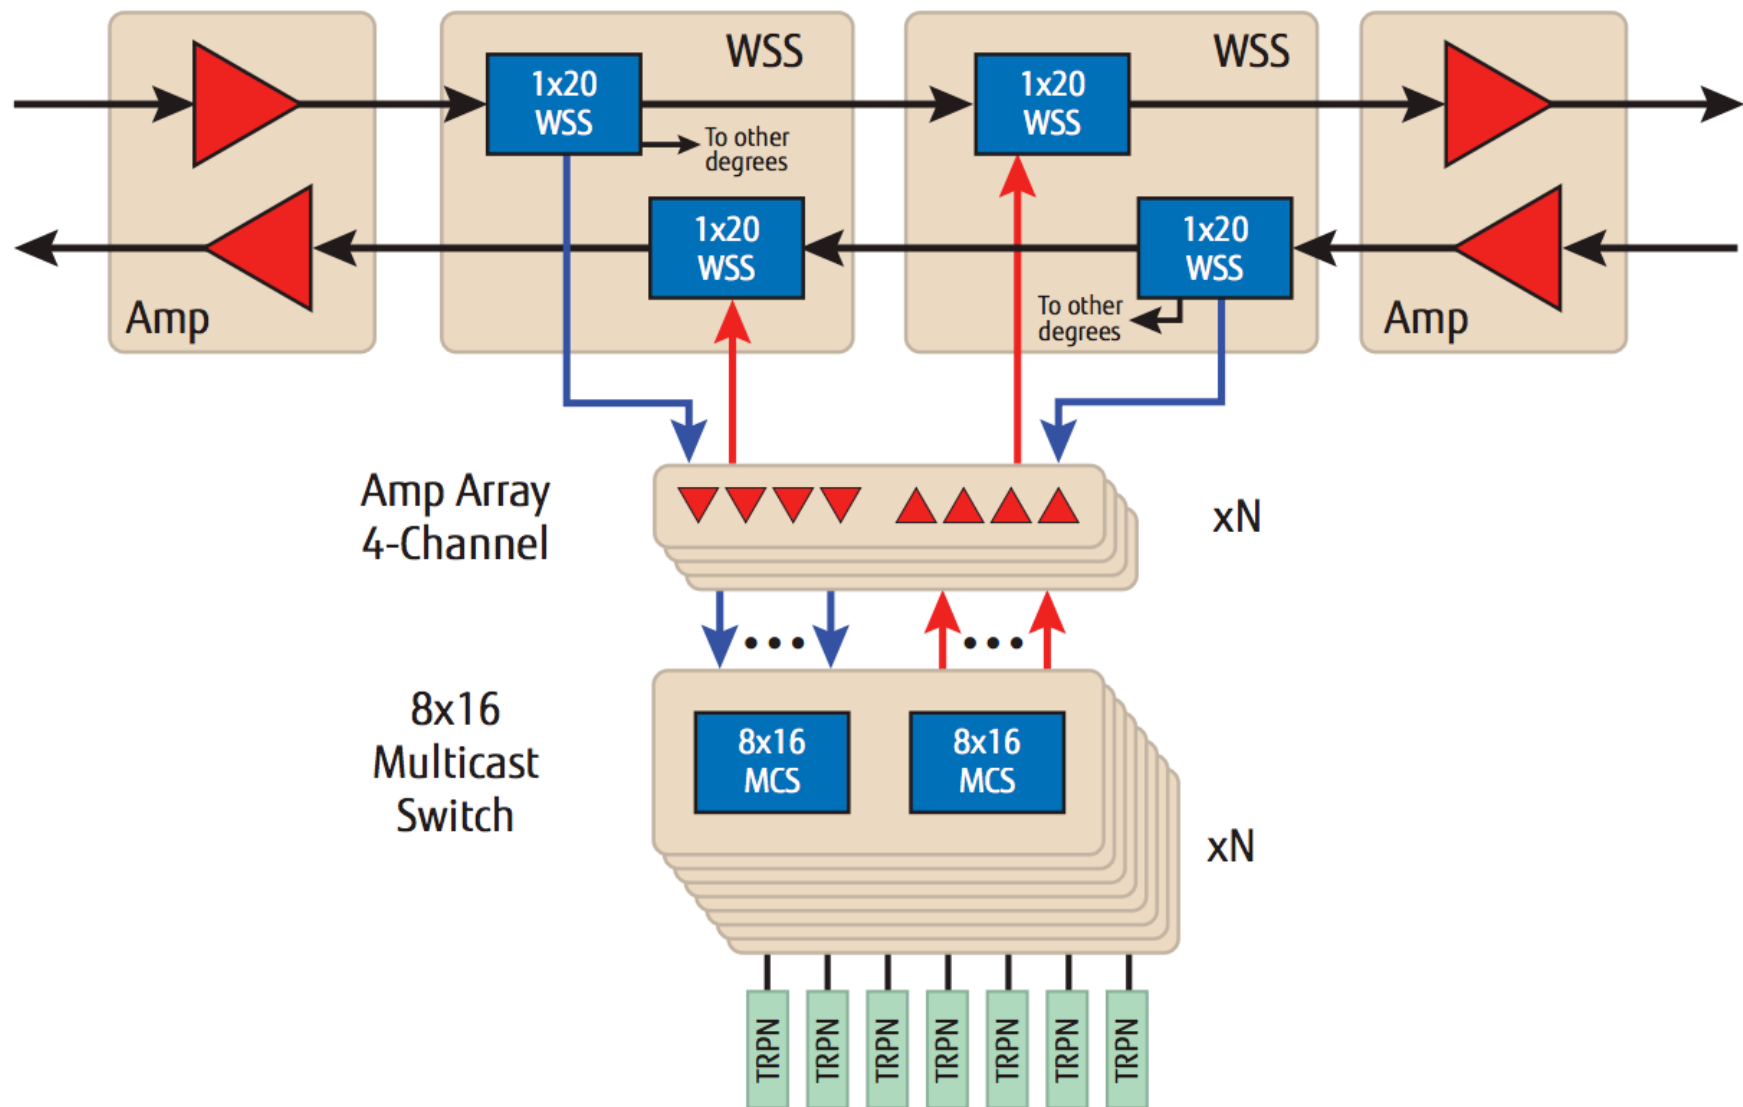
CDC ROADM

- 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 in the ROADM.

Modern Networking and the CDC ROADM

- The newest ROADM technology is called "CDC".
 - Colorless – Any frequency can be dynamically added/dropped to any port, purely by reconfiguring software.
 - Directionless – Channels can be routed out any direction of the ROADM, purely by reconfiguring software.
 - Contentionless – No limitations on reusing the same channel on different directions within the ROADM.
- The goal is to move optical channels entirely with software.
 - Eliminates the need to physically move cables to reconfigure.
 - Allows dynamic bandwidth allocation at an optical level.
 - IETF pushing for vendor interoperability and signaling via mechanisms like PCEP (Path Computation Element Protocol).

Architecture of a CDC ROADM




DWDM Superchannels

- For large bandwidth channels requiring multiple carriers.
- Efficiency can be improved by removing the guardbands.



100 Gbps, PM-QPSK
 $4 \times 12.5 \text{ GHz} = 50 \text{ GHz}$

500 Gbps, PM-QPSK
 $30 \times 12.5 \text{ GHz} = 375 \text{ GHz}$

500 Gbps, PM-16QAM
 $15 \times 12.5 \text{ GHz} = 187.5 \text{ GHz}$



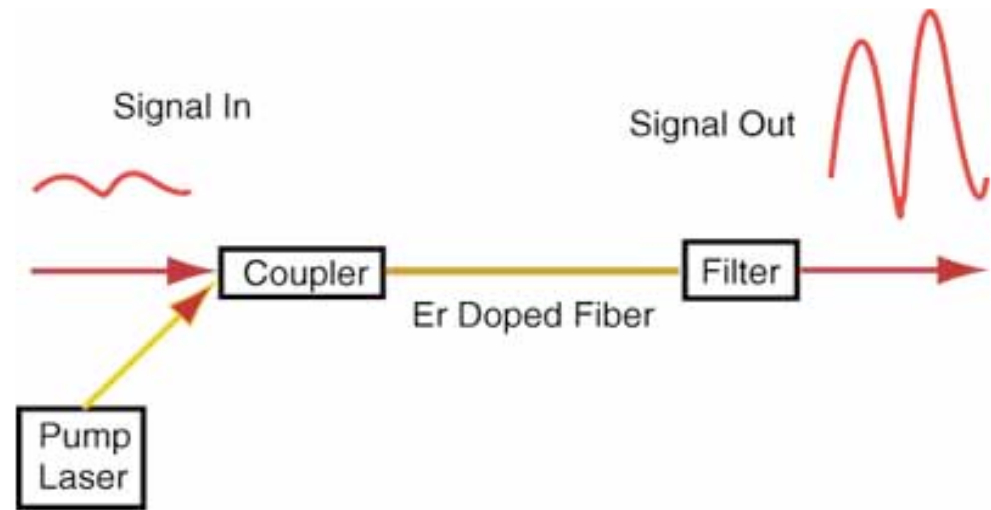
The Evolution of DWDM Channels



Other Optical Networking Concepts

Optical Amplifiers


- Optical amplifiers increase the intensity of a signal
 - There are different types, for different spectrums of light.
 - The most common is the Erbium Doped Fiber Amplifier.
 - Another method is Raman Amplification, typically for ultra long-haul.
 - In an EDFA, a piece of fiber is “doped” with Erbium ions.
 - Additional laser power at 980nm and/or 1480nm is pumped in via a coupler.
 - The interaction between the Erbium and the pump laser causes the emission of light in the C-band spectrum, amplifying the signal.



Optical Switches

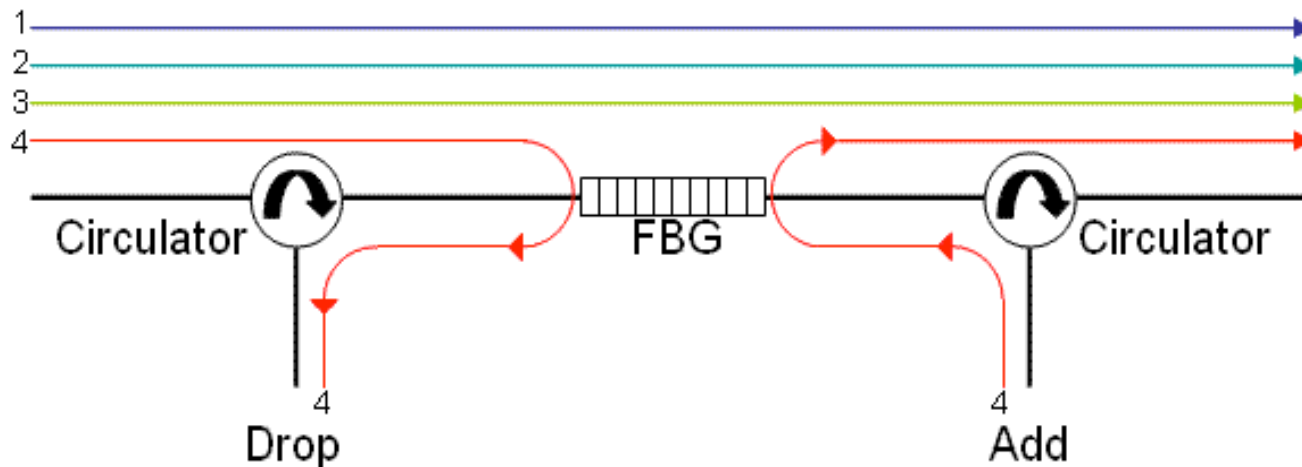
- Optical Switches

- 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.
- Also used inside of complex multi-degree ROADMs, called a WSS (wavelength selectable switch).



Circulator

- A component typically not seen by the end user
 - But used to implement various other common components.
 - Such as muxes, OADMs, and dispersion compensators.
- 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

- Optical Splitters
 - 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 Splitter
 - Often used for “Optical Performance Monitoring”.
 - Tap 1% of the signal and run it to a spectrum analyzer.

Forward Error Correction

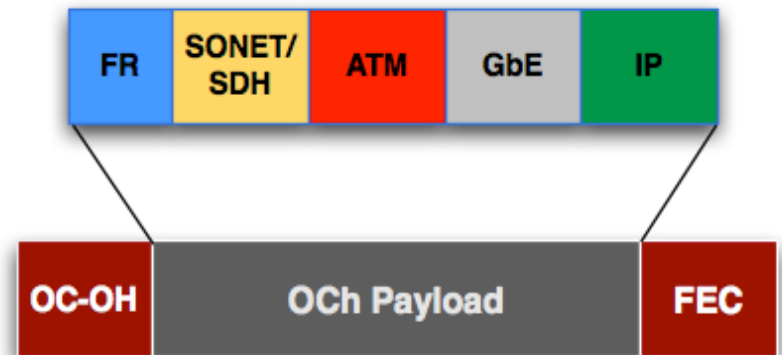
- FEC adds extra/redundant information to the transmission, so the receiver can computationally “recover” from errors.
- In practice, FEC works by improving the required Optical Signal to Noise Ratio (OSNR), allowing improved receiver sensitivity at levels that would otherwise be unusable.
 - Using clever math, padding a 10.325Gbps signal to 11Gbps (7% overhead) can extend a 80km wavelength to 120km or beyond, at the same or better bit error rate.
 - This can really start to matter as older generation DWDM systems are upgraded, since it usually isn't practical to move the amplification huts closer together on a live system.
 - Modern “3rd generation FEC” systems use “Soft Decision” SD-FEC to gain an additional 1-2dB of efficiency, critical in 100G/200G systems.
 - FEC is now integrated into standards like 100GBASE-SR4 as well.

The Benefits of Forward Error Correction



OTN Digital Wrapper Technology (G.709)

- OTN stands for Optical Transport Network
 - A set of standards which allow interoperability and the generic transport of any protocol across an optical network.
 - Implemented as a “wrapper” around another protocol.
 - Why is this needed?
 - So the optical network can be completely transparent.
 - Also, some protocols don't have the same level of troubleshooting capabilities as other protocols.
 - For example, Ethernet is not as good as SONET, because Ethernet wasn't originally designed for the WAN.
 - An OTN wrapper allows the optical troubleshoot with OTN instead.



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 several different types of SMF.
- The most common types are:
 - “Standard” SMF (ITU-T G.652) A.K.A. SMF-28
 - 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 1980s.
- Frequently called “SMF-28”, or simply “classic” SMF.
 - SMF-28 is actually a product name from Corning.
 - 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.


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 will change across different frequencies 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.
 - Running DWDM over DSF causes non-linear interactions.
 - One notable example is called Four Wave Mixing
 - 3 equally spaced wavelengths interact to produce a 4th wavelength.
 - 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”.
 - One spreads the 1550nm band out.
 - The other 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
 - Low-attenuation fiber, at the expense of dispersion.
 - Designed for high-power 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 premise use where the high bend radius of a well designed datacenter may not be practical.
- Modern fibers are usually better than the spec.
 - But much of what's actually in the ground is old fiber.

Dispersion Rates of Commercial Fibers



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:
 - 8-channel CWDM 20nm Mux/Demux: 3.0dB
 - 40-channel DWDM 100GHz Mux/Demux: 3.5dB
 - 80-channel DWDM 50GHz Mux/Demux: 9.5dB

Mismatched Cores



Misaligned Cores

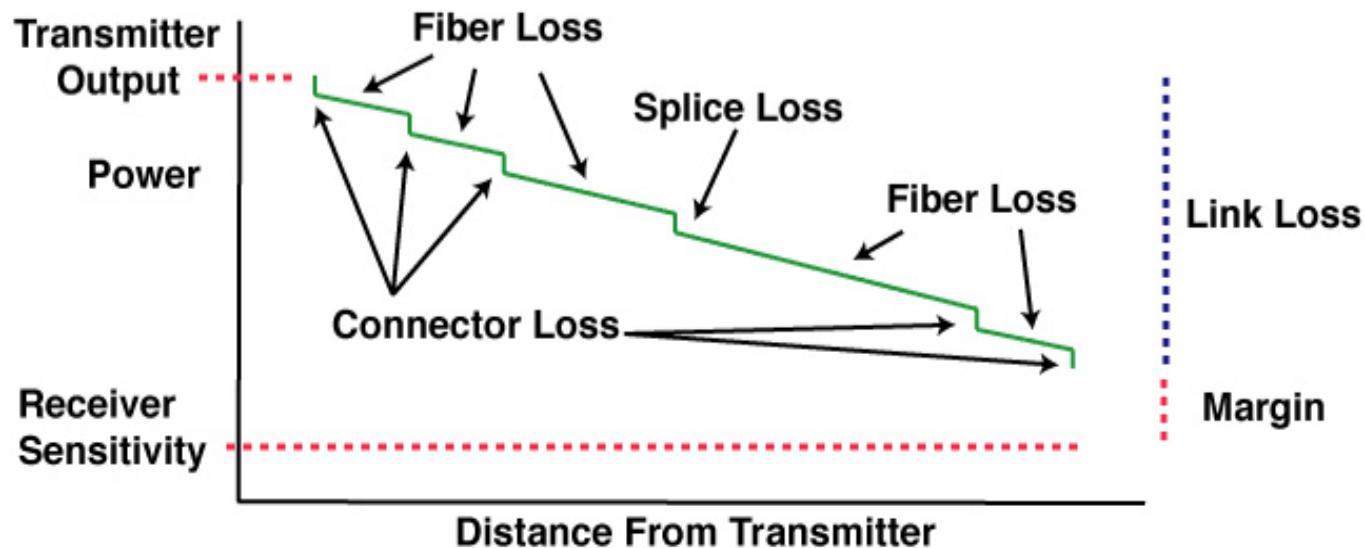


Air Gap Between Fibers



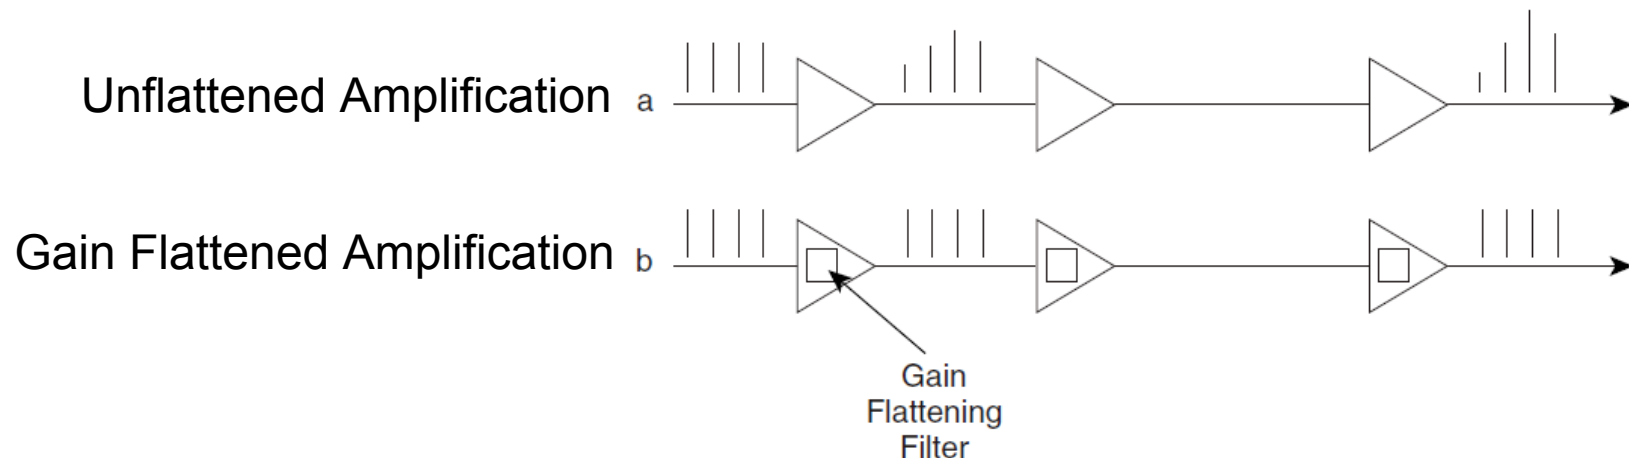
Balling On An (Optical) Budget

- 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 fix will add more splices, etc.



Amplifiers and Power Balance

- Amplifiers introduce their own unique issues.
 - Amplifier gain is not consistent across all wavelengths.
 - The gain must be equalized, or after several amplification stages the power of some channels will be far higher.
 - Mismatched channel powers causes OSNR issues.
 - Care must also be taken when using OADM's, to balance power on passed-thru vs. newly added channels.



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.
 - But the total input power changes as you add channels
 - A single DWDM channel at 10dBm is 0.1mW of input power.
 - 40 DWDM channels at 10dBm is 4mW of power (or 6dBm).
 - 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.
 - A good EDFA needs to monitor system power levels and apply dynamic gain adjustments to maintain a working system.

Dealing with Dispersion

- Dispersion Compensation Unit
 - 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 the signal through this spool reverses the effects of dispersion caused by transmission through the normal fiber.
 - But it also adds extra distance to the normal fiber path, causing additional attenuation, requiring more amplification.
 - Dispersion Compensation spools are typically positioned at optical amplification points for this reason.
 - Circulators can be used to reduce the total amount of fiber needed.




Dealing with Dispersion

- Electronic Dispersion Compensation
 - Dispersion which used to completely ruin a signal can now be compensated for electronically at the receiver.
 - Modern long-haul systems can now handle thousands of kilometers of dispersion compensation.
 - Through sophisticated Digital Signal Processors (DSPs) which compensate for the signal distortion computationally.
 - EDC is being integrated into pluggable optics too.
 - Largely responsible for the 300 meter ranges which can now be achieved over MMF with modern optical standards like 10GBASE-LRM.
 - Technology is getting better all the time too.

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

- As optical impairments (noise, distortion, dispersion, loss of signal, etc) increase...
- The link typically doesn't just outright “die”.
 - It starts taking bit errors, at progressively higher rates.
 - The target maximum Bit Error Rate (BER) is generally 10^{-12} .
 - You can get by with another dBm less signal at 10^{-11} BER.
 - And another dBm less signal after that at 10^{-10} BER.
 - But with exponential progression, the errors gets very bad quickly.

Coherent Optical Technologies

Coherent Optical Technologies


- What exactly are “coherent” optics?
 - A group of advancements in optical technology, which combine to deliver significantly increased optical performance.
 - 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

- What is “Polarization Multiplexing”?
 - Light is (among many other things we don’t fully understand yet) actually 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 made compensating for the changing fiber conditions in real-time practical, allowing dual polarities of light and thus doubling the bandwidth per channel.



Putting It All Together

- Modern long-haul DWDM systems already deliver:
 - 100Gbps transponders with Dual Polarity (DP) Quadrature Phase Shift Keying (QPSK) of 25GBaud signals, to deliver 3000km+ reach.
 - 200Gbps transponders with DP-16QAM at 700km+ reach.
- Further improvements, better DSPs, and better photonics integration onto routers and pluggables, are all expected.

Modulation	Approximate Reach	C-Band Capacity
DP-QPSK	3000 km	9.6 Tbps
DP-8QAM	1500 km	14.4 Tbps
DP-16QAM	700 km	19.2 Tbps
DP-32QAM	350 km	38.4 Tbps
DP-64QAM	175 km	76.8 Tbps

Tools of the Trade


Fiber Optic Power Meter

- Optical Power Meter (or simply a Light Meter)
 - Measures the brightness of an optical signal.
 - Displays the results in dBm or milliwatts (mW).
 - Most light meters 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 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?

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.
 - And the human eye didn't evolve to see infrared.
 - The cornea actually does a very good job of filtering out IR-B light.
 - So an IR-B laser which transmits 10mW of power may still be a Class 1, because that light can't enter the eye.

Optical Networking and Safety


- **Routers**
 - Essentially every single channel laser which can be connected to a router is a Class 1 or Class 1M laser.
 - Even the longest reach 200km+ optics, etc.
- **Optical Amplifiers**
 - Optical amplifiers are capable of putting out enough power to kick a signal into Class 3R (metro) or 3B (long-haul).
- **DWDM Equipment**
 - Total optical power can also increase by muxing together many signals, bringing the total output power into the 3R region even without optical amplification.

Optical Networking and Safety

- So should I be wearing goggles to the colo?
 - Generally speaking, direct router ports are always Class 1 (completely safe under all conditions).
 - Even on DWDM systems, the 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.
 - Extremely high-power DWDM systems have safety mechanisms which detect a fiber cut and cease transmitting a continuous high-power signal until it is repaired.

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.
 - However, most digital cameras can actually see in 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.




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 extremely 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 it’s specified bend radius causes light to leak out.
 - In fact, they even make “macro-bend light meters” which clamp onto the fiber.
 - There are “bend insensitive” fibers for use in residential or office environments which have less bend sensitivity, but they usually trade some performance under normal conditions to achieve this.



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

Can You Mismatch Transceiver Freqs?

- Between certain types of optics, yes.
 - Essentially all optical receivers are wide-band.
 - Though the level of sensitivity may differ for some frequencies.
 - Laser receivers see everything between 1260nm – 1620nm.
 - But they won't be able to see a 850nm LED, for example.
 - 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 it's power reading will be a few dB off.


Can You Mismatch Transceiver Freqs?

- Obscure Optical Networking Trick #738:
 - You may be able to achieve nearly as much distance with a LR/ER (1310nm 10km / 1550nm 40km) pair as with an ER/ER pair.
 - 1550nm has a much lower attenuation rate than 1310nm.
 - 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.
 - The ER optic has a much greater RX sensitivity than the LR.
 - So it will be able to hear the 1310nm signal much better.
- Result:
 - You may only *need* a long reach optic on one side.

Question: Do I Really Need to Clean the Fiber to have it work right?

Do I Really Need to Clean the Fiber?

http://www.fols.org/fols_library/presentations/documents/brown.pdf




Back Reflection = **-67.5 dB**
Total Loss = **0.250 dB**



Back Reflection = **-32.5 dB**
Total Loss = **4.87 dB**

Fiber Contamination and Its Effect on Signal Performance



Clean Connection vs. Dirty Connection

OTDR trace illustration of the significant decrease in signal performance after mating dirty connectors

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 ~ 1.468
 - Speed of light / 1.468 = 204,218,296 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>