

February 2020
Geoff Huston

Deep Sea Diving

Last month I attended the New Zealand Network Operators' Group meeting (NZNOG'20 <http://www.nznog.org/nznog-2020>). One of the more interesting talks for me was given by Cisco's Beatty Lane-Davis on the current state of subsea cable technology. There is something quite compelling about engineering a piece of state-of-the-art technology that is intended to be dropped off a boat and then operate flawlessly for the next twenty-five years or more in the silent depths of the world's oceans! It brings together advanced physics, marine technology and engineering to create some truly amazing pieces of networking infrastructure.

A Potted (and Regionally Biased) History of Subsea Cables

On the 5th August 1856, after a couple of false starts, the Atlantic Telegraph Company completed the first trans-Atlantic submarine telegraph cable. It was a simple affair with seven copper conductor wires, wrapped with three coats of the new wonder material, *gutta-percha* (or as we know it today, rubber). This was further wrapped in tarred hemp and an 18-strand helical sheath of iron wires. It didn't last long, as the cable company's electrical engineer, Dr Wildman Whitehouse, had a preferred remedy for a fading signal by increasing the voltage on the circuit (as distinct from William Thompson's (later Lord Kelvin) chosen remedy of increasing the sensitivity of his mirror galvanometer receivers). A 2KVDC power setting proved fatal for the insulation of the cable, which simply ceased to function from that point onward.




Figure 1 – Isambard Kingdom Brunel's SS *Great Eastern*, the ship that laid the first lasting transatlantic cable in 1866

Over the ensuing years techniques improved, with the addition of in-line amplifiers (or repeaters) to allow the signal to be propagated across longer distances, and progressive improvements in signal processing to improve the capacity of these systems. Telegraph turned to telephony, valves turned to transistors and polymers replaced rubber, but the basic design remained the same: a copper conductor sheathed in a watertight insulating cover, with steel jacketing to protect the cable in the shallower landing segments.

In the Australian context the first telegraph system, completed in 1872, using an overland route to Darwin, and then short undersea segments to connect to Singapore and from there to India and the UK.



Figure 2 – Erecting the first telegraph pole for the Overland Telegraph Line in Darwin in 1870

Such cables were used for telegraphy, and the first trans-oceanic telephone systems were radio-based; it took some decades for advances in electronics to offer a cable-based voice service.

One of the earlier systems to service Australia was commissioned in 1962. COMPAC supported 80 x 3Khz voice channels linking Australia via New Zealand, Fiji, and Hawaii to Canada, and from there via a microwave service across Canada and then via CANTAT to the UK. This cable used valve-based undersea repeaters. COMPAC was decommissioned in 1984 at the time when the ANZCAN cable was commissioned.

ANZCAN followed a similar path across the Pacific, with wet segments from Sydney to Norfolk Island, then to Fiji, Hawaii and Canada. It was a 14Mhz analogue system with solid state repeaters spaced every 13.5km.

This was replaced in 1995 by the PACRIM cable system, with a capacity of 2 x 560Mhz analogue systems. COMPAC had a service life of 22 years, and ANZCAN had a service life of 11 years. PACRIM had a commercial life of less than two years as it was already superseded by 2.5Ghz submarine circuits on all-optical systems, and it was woefully inadequate for the explosive demand of the emerging Internet.

These days there are just under 400 submarine cables in service around the world, and 1.2 million km of cable. Telegeography's map (<https://www.submarinecablemap.com>) is a comprehensive resource that shows this cable inventory.

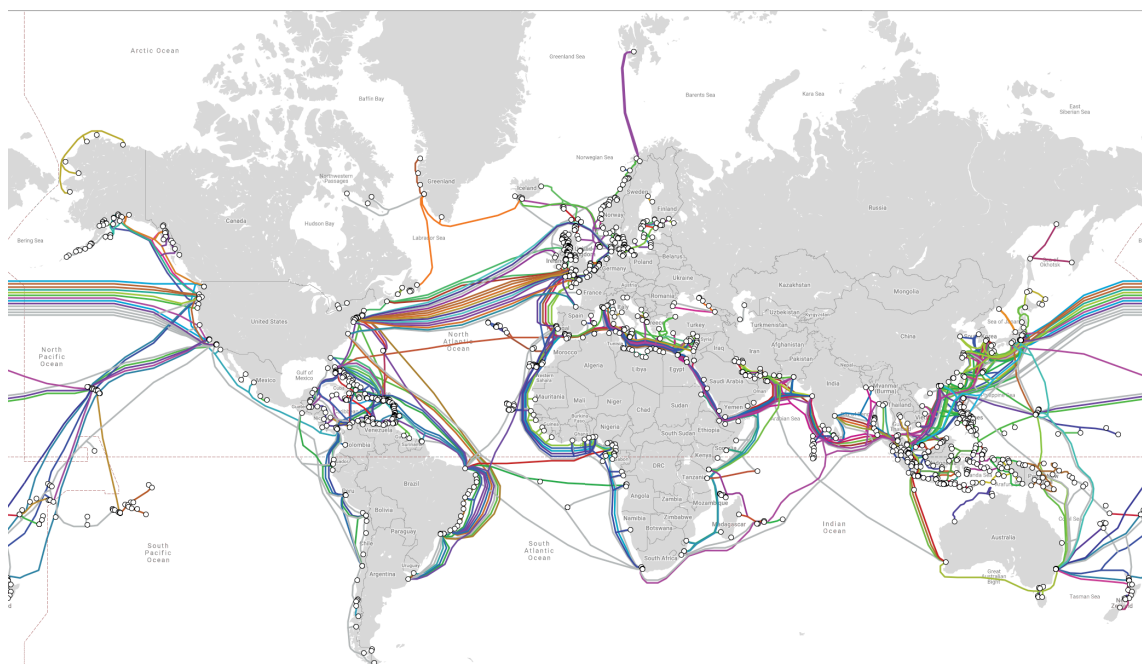


Figure 3 – Telegeography's map of subsea cables

Ownership

The first cable systems were incredibly expensive undertakings as compared to the size of the economies that they serviced. The high construction and operating costs made the service unduly prohibitive to most potential users. For example, when the Australian Overland Telegraph was completed, a 30-word telegram to the UK cost the equivalent of three week's average wages. Early users were limited to the press and government agencies as a result.

The overland telegraph system used 'repeater' stations where human operators recorded incoming messages and re-keyed the message into the next cable segment. Considering that across Asia few of these repeater station operators were native English speakers, it is unsurprising that the word error rate in these telegrams was up to one third of the words in a message. So the service was both extremely expensive and highly error prone.

What is perhaps most surprising is that folk persisted in spite of these impediments, and, in time, the cost did go down and the reliability improved.

Many of these projects were funded by governments and used specialised companies that managed individual cables. In the late 19th century there was the *British Indian Submarine Telegraph Company*, the *Eastern Extension Australasia and China Telegraph Company* and the *British Australia Telegraph Company*, among many others. The links with government were evident, such as the commandeering of all cable services out of Britain during WWI. The emergence of the public national telephone operator model in the first part of the twentieth century was mirrored by the public nature of the ownership of cable systems.

The model of consortium cable ownership was developed within the larger framework, where a cable was floated as a joint stock private company that raised the capital to construct the cable as a debt to conventional merchant banking institutions. The company was effectively 'owned' by the national carriers that purchased capacity on the cable, where the share of ownership equated to the share of purchased capacity.

The purchased capacity on a cable generally takes the form of the purchase of an indefeasible right of use (IRU), giving the owner of the IRU exclusive access to cable capacity for a fixed term (commonly for a period of between 15 and 25 years, and largely aligned to the expected service life of the cable). The IRU conventionally includes the obligation to pay for a proportion of the cable's operational costs.

When the major customer for undersea cable systems was the national carrier sector, the costs of each of the IRUs were normally shared equally between the two carriers who terminated either end of the IRU circuits. (this was part of the balanced financial settlement regime of national carriers, where the cost of common infrastructure to interconnect national communications services was divided equally between the connecting parties). While this model was developed in the world of monopoly national carriers, the progressive deregulation of the carrier world did not greatly impact this consortium model of cable ownership for many decades. Part of the rationale of this shared cost model of half circuits constructed on jointly owned IRUs was one where carriers cooperated on the capital and recurrent costs of facility construction and operation and competed on services.

One of the essential elements of this bureaucratic style of ownership was that capacity pricing of the cable was determined by the consortium, presenting any carrier or group of carriers to undue the other members of the consortium. The intention was to preserve the market value of the cable by preventing undercutting and dumping. The actual outcome was a classic case of supply side rationing and price fixing where the capacity of the cable was released into the market in small increments, ensuring that demand always exceeded available capacity over the lifetime of the cable, and cable prices remained buoyant.

The internet construction boom in the 1990s coincided with large scale deregulation of many telco markets. This allowed other entities to obtain cable landing rights in many countries, which led to the entrance of carrier wholesalers into the submarine cable market, and also led to the concept of "wholly

owned capacity" on cable system. Early entrants in these markets were wholesale providers to the emerging ISP sector, such as Global Crossing, for example, but it was not long for the large content enterprises, including Google, Facebook and others to enter this submarine cable capacity market with their own investments. The essential difference in this form of operation is the elimination of price setting, allowing cable pricing to reflect prevailing market conditions of supply and demand.

Cables

The basic physical cable design that's used today for these subsea systems is much the same as the early designs. The signal bearer itself has changed from copper to fibre, but the remainder of the cable is much the same. A steel strength member is included with the signal bearers and these signal cables are wrapped in a gel to prevent abrasion.

Surrounding the signal bundle is a copper jacket to provide power, then an insulating and waterproofing sheath (a polyethylene resin), then, depending of the intended location of the particular cable segment, layers of protection. The shallower the depth of the cable segment and the greater the amount of commercial shipping, the higher the number of protection elements, so that the cable will survive some forms of accidental snagging of the cable (Figure 5).




Figure 4 – Cable cross-section

Normally the cable is laid on the seabed, but in areas of high marine activity the steel sheathed cable might be laid into a ploughed trench, and in special circumstances the cable may lay in a trough cut out from a seabed rock shelf.

The cable laying technique has not changed to any significant degree. An entire wet segment is loaded on a cable laying ship, end-to-end-tested, and then the ship sets out to traverse the cable path in a single run. The speed and position of the ship is carefully determined so as to lay the cable on the seabed without putting the cable under tensile stress. The ship sails the lay path in a single journey without stopping, laying the cable on the seabed, whose average depth is 3,600m, and up to 11,000m at its deepest. The cable is strung out during laying up to 8,000m behind the lay ship.



Figure 5 – Loading a Cable Ship



Figure 6 – A Cable Laying Ship



Figure 7 – Cable Repeater

Cable repair is also a consideration. It takes some 20 hours to drop a grapnel to a depth of 6,000m, snag the cable and retrieve one end, and that depth is pretty much the maximum feasible depth of cable repair operations. Cables in deeper trenches are not repaired directly but spliced at either side of the trench. The implication is that when very deep-water cable segments fail repairing the cable can be a protracted and complex process.

While early cable systems provided simple point-to-point connectivity, the commercial opportunities in using a single cable system to connect many end points fuelled the need for the provision of *Branching Units*. The simplest form of an optical branching unit is to split up the physical fibres in the cable core. These days its more common to see the use of reconfigurable optical add-drop multiplexers (ROADMs). These units allow individual or multiple wavelengths carrying data channels to be added and/or dropped from a common carriage fibre without the need to convert the signals on all of the wavelength division multiplexed channels to electronic signals and back again to optical signals. The main advantages of the use of ROADMs is the deferral of planning of entire bandwidth assignment in advance, as ROADMS allow reconfiguring capacity in the system in response to demands. The reconfiguration can be done as and when required without affecting traffic already passing the ROADM.



Figure 8 – Cable Branching Unit

The undersea system is typically referred to as the *wet segment*. and these systems interface to surface systems at cable stations. These stations house the equipment that supplies power to the cable. The power configuration is DC, and long haul cable systems are powered by systems that typically use 10KVDC feeds at both ends of the cable. The cable station also typically includes the wavelength termination equipment and the Line Monitoring Equipment

Optical Repeaters

Optical “Repeaters” are perhaps a misnomer these days. Earlier electrical repeaters operated in a conventional repeating mode, using a receiver to convert the input analogue signal into a digital signal, and then re-encoding the data into an analogue signal and injecting it into the next cable segment.

These days optical cable repeaters are photon amplifiers that operate at full gain at the bottom of the ocean for an anticipated service life of 25 years. Light (at 980nm or 1480nm) is pumped into a relatively short erbium-doped fibre segment. The erbium ions cause an incoming light stream in the region around 1550nm to be amplified. The pump energy causes the erbium ions to enter a higher energy state, and when stimulated by a signal photon, the ion will decay back to a lower energy level, emitting a photon at the energy level of the stimulated state, but with a light frequency equal to the triggering incoming signal. This emitted amplified signal conveniently shares the same direction and phase as the incoming light signal. These are called “EDFA” units (Erbium Doped Fibre Amplifiers). This has been totally revolutionary for subsea cables. The entire wet segment, including the repeaters are entirely agnostic with respect to the carrier signal. The number of lit wavelengths, the signal encoding and decoding, and the entire cable capacity is now dependant on the equipment at the cable stations at each end of the cable. This has extended the service life of optical systems, where additional capacity can be scavenged from deployed cables by placing new technology in the cable stations at either end, leaving the wet segment unaltered. The wet plant is agnostic with respect to the cable carrying capacity in these all optical systems.




Figure 9 – EDFA amplification

The subsea optical repeater units are designed to operate for the entire operational life of the cable with any further intervention. The design includes an element of redundancy in that if a repeater fails then the cable capacity may be degraded to some extent but will still operate with viable capacity.

The EDFA units have a bias in amplification across the operational frequency range and it's necessary to add a passive filter to the amplified signal in order to generate a flatter power spectrum. This allows the cumulative sum of these in-line amplifiers to produce an outcome that maximises signal performance for entire spectrum of the band used on the cable. Over extended distances this is still insufficient, and cables may also use active units, called "Gain Equalisation Units". The number, spacing and equalisation settings used in these units is part of the customised design of each cable system.

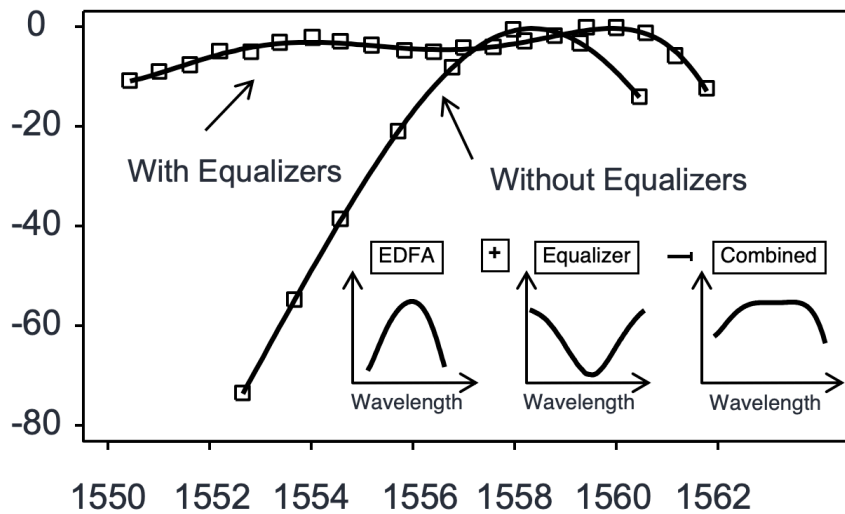


Figure 10 – EDFA Gain Equalization

In terrestrial systems amplifier control can be managed dynamically, and as channels are added or removed the amplifiers can be reconfigured to produce optimal gain. Subsea amplifiers have no such dynamic control, and they are set to gain saturation, or always on “maximum”. In order to avoid overdriving the lit channels, all unused spectrum channels are occupied by an ‘idler’ signal.

Repeaters are a significant cost component of the total cable cost, and there is a compromise between a ‘close’ spacing of repeaters, every 60km or so, or stretching the inter-repeater distance to 100km and making significant savings in the number of repeaters in the system. On balance it is the case that the more you are prepared to spend on the cable system the higher the cable carrying capacity.

The observation here is that a submarine cable is not built by assembling standard components and connecting them together using a consistent set of engineering design rules, but by customising each component within a bespoke design to produce a system which is built to optimise its service outcomes for the particular environment where the cable is to be deployed. In many respects every undersea cable project is built from scratch.

Cable Capacity and Signal Encoding

The earliest submarine cable optical systems were designed in the 1980's and deployed in the late 80's. These first coaxial cable systems used electrical regeneration and amplification equipment, and amplifiers were typically deployed every 40km on the cable. The first measure that was used to increase cable capacity was to use a system that had been the mainstay of the radio world for many years, namely frequency division multiplexing (FDM). The first optical transmission electrical amplification cables used FDM to create multiple voice circuits over a single coaxial cable carrier. These cables supported a total capacity of 560Mb divided into some 80,000 voice circuits.

Plans were underway to double the per coax-bearer cable capacity of these hybrid optical / electrical systems when all-optical EDFA systems were introduced, and the first deployments of EDFA submarine cable systems were in 1994. These cables used the same form of frequency-based optical multiplexing, where each optical cable is divided into a number of discrete wavelength channels (lambdas) in an analogous sharing framework termed WDM (Wave Division Multiplexing).

As the carrier frequency of the signal increased, coupled with long cable runs, the factor of chromatic dispersion became more critical. Chromatic dispersion describes the phenomenon that light travels at slightly different speeds at different frequencies. What this means is that a square wave input will be received as a smoothed wave, and at some point, the introduced distortion will push the signal beyond the capability of the receiver digital signal processor (DSP) to reliably decode. The response to chromatic dispersion is the use of negative dispersion fibre segments, where the doping of the negative dispersion cable with germanium dioxide is set up to compensate for chromatic dispersion. It's by no means a perfect solution, and while it's possible to design a dispersion compensation system that compensates for dispersion at the middle frequency of the C-band, the edge frequencies will still show significant chromatic dispersion when using long cable runs.

This first generation of all-optical systems used simple on/off keying (OOK) of the digital signal into light on the wire. This OOK signal encoding technique has been used for signal speeds of up to 10Gbps per lambda in a WDM system, achieved in 2000 in deployed systems, but cables with yet higher capacity per lambda are infeasible for long cable runs due to the combination of chromatic dispersion and polarisation mode dispersion.

At this point coherent radio frequency modulation techniques were introduced into the digital signal processors used for optical signals, combined with wave division multiplexing. This was enabled with the development of improved digital signal processing (DSP) techniques borrowed from the radio domain, where receiving equipment was able to detect rapid changes in the phase of an incoming carrier signal as well as changes in amplitude and polarization.

Using these DSPs it's possible to modulate the signal in each lambda by performing phase modulation of the signal. Quadrature Phase Shift Keying (QPSK) defines four signal points, each separated at 90-degree phase shifts, allowing 2 bits to be encoded in a single symbol. A combination of 2-point polarisation mode encoding and QPSK allows for 2 bits per symbol. The practical outcome is that a C-band based 5Thz optical carriage system using QPSK and DWDM can be configured to carry a total capacity across all of its channels of some 25Tbps, assuming a reasonably good signal to noise ratio. The other beneficial outcome is that these extremely high speeds can be achieved with far more modest components. A 100G channel is constructed as 8 x 12.5G individual bearers.

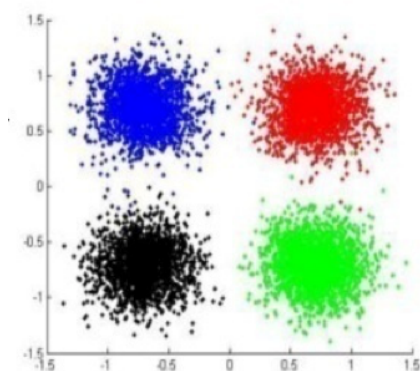


Figure 11 – Phase-Amplitude space mapping of QPSK keying

This encoding can be further augmented with amplitude modulation. Beyond QPSK there is 8QAM that adds another four points to the QPSK encoding, adding additional phase offsets of 45 degrees and at half the amplitude. 8QAM permits a group coding of 3 bits per symbol but requires an improvement in the signal to noise ratio of 4db. 16QAM defines, as its name suggests 16 discrete points in the phase amplitude space which allows the encoding of 4 bits per symbol, at a cost of a further 3db in the minimum acceptable S/N ratio. The practical limit of increasing the number of encoding points in phase amplitude space is the signal to noise ratio of the cable, as the more complex the encoding the greater the demands placed on the decoder.

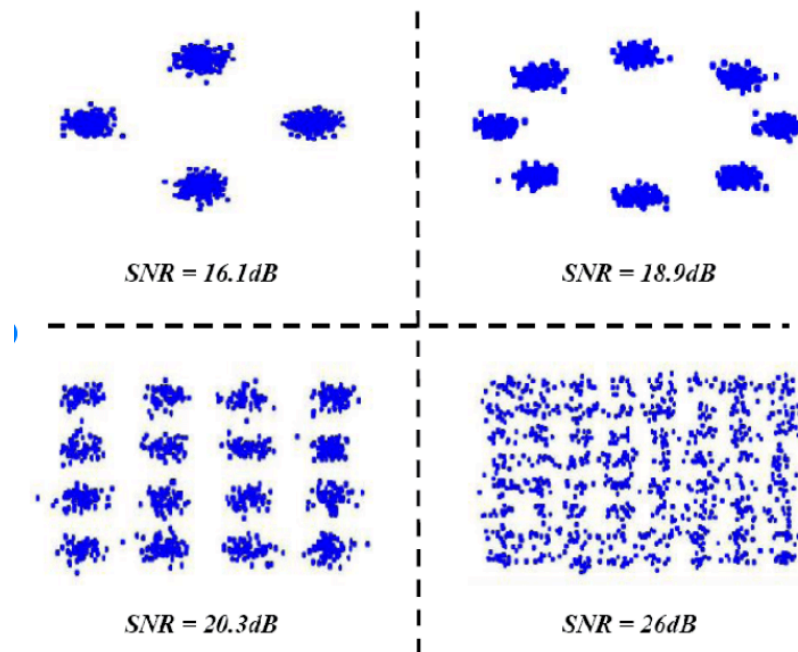


Figure 12 - Adaptive Modulation Constellations for QPSK, 8PSK, 16QAM AND 64QAM

The other technique that can help is extracting a dense digital signal from a noise-prone analogue bearer is the use of Forward Error Correcting codes in the digital signal. At a cost of a proportion of the signal capacity, the FEC codes allow the detection and correction of small number of errors per FEC frame. The current state of the art in FEC is Polar code, where the channel performance has now almost closed the gap to the Shannon limit which sets the bar for the maximum rate for a given bandwidth and a given noise level.

It's now feasible and cost efficient to deploy medium to long cable systems with capacities approaching 50Tbps in a single fibre. But that's not the limit we can achieve in terms of capacity of these systems.

There are two frequency bands available to cable designers. The conventional band is C-band, which spans wavelengths from 1,530nm to 1,565nm. There is an adjacent band, L-band, which spans wavelengths from 1,570nm to 1,610nm. In analogue terms there is some 4.0 to 4.8 THz in each band. Using both bands with DWDM and QPSK encoding can result in cable systems that can sustain some 70Tbps per fibre through some 7,500km of cable.

Achieving this bandwidth comes at a considerable price, as the EDFA units operate either in C or L bands so cable systems that have both C and L bands need twice the number of EDFA amplifiers (which of course requires double the injected power at the cable stations. At some point in the cable design exercise it may prove to be more cost efficient to increase the number of fibre pairs in a cable than it is to use ever more complex encoding mechanisms, although there are some limitations to the total amount of power that can be injected into long distance submarine cables, so typically such long distance systems have no more than 8 fibre pairs.

A different form of optical amplification is used in the Fibre Raman Amplifier (FRA). The principle of FRA is based on the Stimulated Raman Scattering (SRS) effect. The gain medium is undoped optical fiber and power is transferred to the optical signal by a nonlinear optical process known as the Raman effect. An incident photon excites an electron to the virtual state and the stimulated emission occurs when the electron de-excites down to the vibrational state of glass molecule. The advantages of FRAS are variable wavelength amplification, compatibility with installed single mode fibre, and its suitability to extend EDFAs. FRAs require very high pump power lasers and sophisticated gain control, however a combination of EDFA and FRA may result in a lower average power over a span. FRAs operate over a very broad signal band (1280nm – 1650nm).

However, all this assumes that the glass itself is a passive medium that exhibits distortion (or noise) only in linear ways. Twice the cable, twice the signal attenuation and twice the degree of chromatic dispersion and so on. When large amounts of energy, in the form of photons, are pumped into glass it displays non-linear behaviour, and as the energy levels increase the non-linear behaviours become significantly more evident.

The Optical Kerr effect is seen with laser injection into optical cables where the intensity of the light can cause a change in the glass' index of refraction, which in turn can produce modulation instability particularly on phase modulation used in optical systems. Brillouin scattering results from the scattering of photons caused by large scale, low-frequency phonons. Intense beams of light through glass may induce acoustic vibrations in the glass medium that generate these phonons. This causes remodulation of the light beam as well as modification of both light amplification and absorption characteristics.

The result is that high density signalling in cables requires high power which results in non-linear distortion of the signal, particularly in terms of phase distortion. The trade-offs are to balance the power budget, the chromatic and phase distortion and the cable capacity. Some of the approaches used in current cable systems include different encoding for different lambdas to optimise the total cable carrying capacity.




Figure 13 – Non-linear distortion of a QPSK keyed signal

Current cable design since 2010 largely leaves dispersion to the DSP and does away with dispersion compensation segments in the cable itself. The DSP now performs feedback compensation. This allows for greater signal coherency, albeit at a cost of increased complexity in the DSP. Cable design is now also looking at larger diameter of the glass core in SPF fibre. The larger effective size of the glass core reduces the non-linear effects of so much power being pumped into the glass, leading to up to ten times the usable capacity in these larger core fibre systems. Work is also looking at higher powered DSPs that can perform more functions. This relies to a certain extent on Moore's Law in operation. Increasing the number of gates in a chip allows for greater functionality to be loaded onto the DSP chip and still have viable power and cooling requirements. Now only can DSPs operate with feedback compensation, they can also perform adaptive encoding and decoding. For example, a DSP can alternate between two PSK encodings for every symbol pair. For example, if every symbol was encoded with alternating 8QAM and 16QAM, then the result is an average of 7 bits per symbol, reducing the large increments between the various PSK encoding levels. The DSP can also test the various phase amplitude points and will reduce its use of symbols that have a higher error probability. This is termed *Probabilistic Constellation Shaping* or PCS. This can be combined with FEC operating at around 30% to give a wide range of usable capacity levels for a system.

Taking all this into account, what is the capacity of deployed cable systems today? The largest capacity cable in service at the moment is the MAREA cable, connecting Bilbao in Spain to Virginia Beach in the US, with a service capacity of 208Tbps.

Futures

We are by no means near the end to the path in the evolution of subsea cable systems, and ideas on how to improve the cost and performance abound. Optical transmission capability has increased by a factor of around 100 every decade for the past three decades and while it would be foolhardy to predict that this pace of capability refinement will come to an abrupt halt, it also has to be admitted that sustaining this growth will take a considerable degree of technological innovation in the coming years.

One observation is that the work so far has concentrated on getting the most we can out of a single fibre pair. The issue is that to achieve this we are running the system in a very inefficient power mode where a large proportion of the power gets converted to optical noise that we are then required to filter out. An alternative approach is to use a collection of cores within a multi-core fibre and drive each core at a far lower power level. System capacity and power efficiency can both be improved with such an approach.

The refinements of DSPs will continue, but we may see changes to the systems that inject the signal into the cable. In the same way that vectored DSL systems use pre-compensation of the injected signal in order to compensate for signal distortion in the copper loop, it may be possible to use pre-distortion in the laser drivers, or possibly even in the EDFA segments, in order to achieve even higher performance from these undersea systems.

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.

Author

Geoff Huston AM, B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net