

Multiplexing, Modulation and Detection

What is the best mix?

January 2010

Abstract

This document discusses the different challenges that are found in high capacity terrestrial and submarine systems, and the technical solutions that are appropriate.

1 INTRODUCTION

Telecommunications capacity demand has always increased from one year to the next, giving an obvious technical challenge to get more capacity per fibre pair, while also meeting economic and operational constraints. In the past technology has risen to the challenge with higher symbol rates and closer wavelength spacing. For both technical and economic reasons this is unlikely to continue and the latest move to get more capacity consists of technology to transmit more information per wavelength by using new types of multiplexing, better error correction and modulation schemes and more sensitive detection, rather than continuing to ever higher line-rates.

Unsurprisingly, so far there is no single perfect solution and technical conferences continue to debate the merits of different mixes of modulation, multiplexing and detection. This paper tries to explain some of the issues and show how there are different "best" solutions depending on the target application. To avoid covering too many types of systems, it considers only high-capacity optically-amplified fibre systems with overall lengths greater than 100 km.

In the following sections it considers how best to multiplex traffic in a restricted optical bandwidth. It looks at Amplitude and Phase modulation schemes, where there is little dispute that phase modulation is preferred, but the choice of the number of levels and the signal format is less obvious. The type of detection, direct or coherent, is the final selection which has an important impact on the design and ultimate performance.

Selecting the best combination depends on the type of system. Currently, submarine and terrestrial systems are quite different, although their designs seem likely to converge in the future. Equally the key issues for long and short systems are not the same.



While Terrestrial and Submarine Systems both consist of a line of optical amplifiers connected with single-mode fibre, they take quite different approaches in terms of the types of fibre that they use, as shown in the following diagram, where the numbers represent Chromatic Dispersion CD in ps/nm/km.

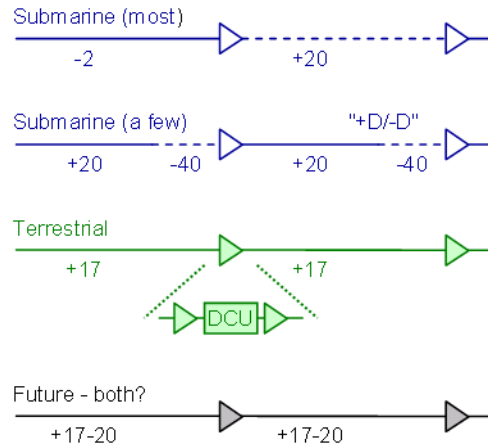


Figure 1 - Different Fibre Maps

With the exception of the "Future – both?" they are designed to get the accumulated chromatic dispersion to be small, while avoiding long lengths of low dispersion fibre. Low dispersion and WDM is a bad mix because it results in adjacent wavelengths travelling at roughly the same speed, thus giving the maximum non-linear interaction between them. One of the best ways to minimise non-linear effects is to ensure that chromatic dispersion is relatively high; large core fibre (which reduces the power density) is an additional approach.

Terrestrial systems use high positive dispersion fibre, with each In-Line Amplifier (ILA) consisting of two optical amplifiers with a Dispersion Compensation Unit (DCU) between them. In a submarine system space and electrical power are limited and amplifiers expensive (due to the need for extensive reliability screening) so a single amplifier is used, with a mixture of mainly negative dispersion fibre balanced by a repeater section of positive dispersion fibre every few sections. A recent alternative ["Submarine (newer)" in the diagram, but also known as "+D/-D" or "Dispersion Managed Fibre (DMF)"] has been to mix positive and negative dispersion within a single repeater section; this is a great solution from a transmission perspective, although it is more expensive and needs care when doing a repair.

An interesting speculation is that in the future both submarine and terrestrial systems could be built entirely from high dispersion large core fibre (which would minimise non-linear effects) with Digital Signal Processing (DSP) being used to compensate for the huge dispersion that would accumulate. There are already



some very encouraging laboratory results, and it will be interesting to see how the industry reacts. This approach should also be applicable to new build terrestrial systems, giving the potential to simplify the intermediate amplifiers, which will no longer require Dispersion Compensation Units (DCUs).

Existing submarine systems, with their relatively low dispersion fibres and potentially long spans, however, represent cases where non-linearity is significant, so upgrading to add extra capacity is likely to need a different solution.

Capacity, line-rate, symbol rate

For simple binary transmission these are all roughly the same, but with multi-level signaling the two are quite different. In this paper 10/40/100G refers to the approximate information capacity; the rate at which symbols are transmitted will be shown as X Gbaud.

2 MULTIPLEXING

Recently 10G wavelength separation has dropped from 50 GHz to 25 GHz (in the latest systems), but this process has probably reached its limit because of the spectral width of the signals. The figure shows two RZ spectra which are spaced at 33 GHz.




Figure 2 - Spectrum of 10 Gbaud RZ pulses

For long span systems Return to Zero (RZ) pulses give a significant performance benefit and are almost always used. The diagram shows how at 33 GHz spacing there is already a small overlap between adjacent RZ spectra, which becomes greater at 25 GHz spacing. The diagram is somewhat theoretical and real spectra can be made a bit narrower, making 25 GHz spacing quite possible with good optical filters. However, it serves to show that this is close to the limit for RZ pulses; Non Return to Zero (NRZ) pulses have a somewhat narrower spectrum, but they suffer significant penalties for long spans and are generally used only in Terrestrial systems and short Submarine links. Getting greater capacity by



increasing the line-rate is not possible because this would make the spectrum broader. (There are formats – e.g., Phase Shaped Binary Transmission [PSBT] which operate well in band-limited systems – but the general principle remains true.)

It is, however, possible to transmit two signals on the *same* wavelength using two orthogonal polarisations, which have very little interaction with each other. Even after long distances there will still be two orthogonal polarisations, but the states will vary with time, so the technical challenge is to find a way to track the polarisation states and recover the two data signals. Possible ways of doing this are described later.

Launching two polarisations offers the possibility of transmitting twice the number of optical carriers in a given bandwidth, but this doesn't necessarily double the traffic capacity, as one needs to maintain a good Optical Signal to Noise Ratio (OSNR). Since amplifier noise is generally fixed, the total output power must be roughly doubled if the total number of channels is doubled. For terrestrial systems amplifier power can be increased (by adjustment or replacement), but for most subsea systems the amplifier power is fixed, so capacity increases are possible only for those cases where transmission was limited by amplifier bandwidth.

3 ERROR CORRECTING CODES

It could be argued that Forward Error Correction (FEC) is the most effective single technology, as it adds several dB for relatively little cost.

So far all the codes used apply correction after the receiver has produced a one or zero based on whether the signal is above or below a threshold value. It is easy to calculate the gain for a perfect code of this type; the number of errors that can be corrected depends on the number of different parity checks that are performed. Increasing the number of checks adds to the overhead, but the graph below shows that increasing the overhead (and thus the complexity) of the code adds only small improvements. The squares show examples of actual codes that are used.



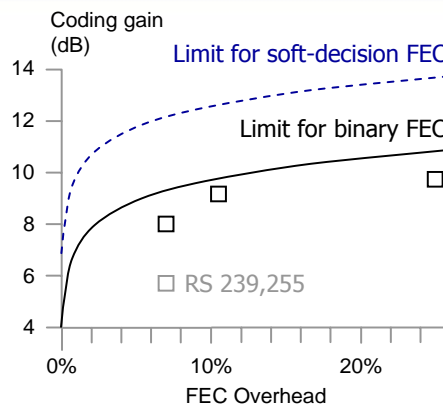


Figure 3 - FEC gain vs. Overhead

Increasing the overhead also has the undesirable effect of increasing the line-rate and the spectral width of the signal.

A potential solution to this is the use of "soft decision" correction, where the receiver no longer produced a stream of binary digits. Instead a very high speed Analogue to Digital Converter (ADC) produces a multi-bit signal which gives the amplitude of the detected signal. This helps because a signal which is close to the decision threshold is more likely to be on the "wrong" side of it than one which is distant. With this extra information, and using appropriate coding and decoding, it is possible in theory to produce results 3 dB better than the binary process. In practice, the decoding process is *extremely* complex and currently no-one has yet shown soft-decision solutions at Gigabit rates which offer significant advantages over binary codes, but the potential rewards make this an attractive area for development.

4 MODULATION

4.1 Line-rates and Multi-levels

Firstly it is worth remembering that symbol rate and information transfer rate are not the same, although they are close for binary signals with no Forward Error Correction (FEC). The use of FEC adds an "Overhead" of check bits, which makes the line-rate higher, typically by 7-25%. However, the line rate can be reduced by using multi-level symbols, a technique commonly used in electrical and radio systems when bandwidth is limited. The diagram below shows how phase modulation with four different phases allows each symbol to carry two bits of information.



Quaternary (phase)




Figure 4 - Example of multi-level format

Many non-optical systems use multiple levels of amplitude and phase, with digital TV for example using up to 64 different levels, but until now most long-span fibre optic systems have used simple binary transmission because signaling speed has not been an issue. Previous capacity increases were achieved by moving from line-rates of 2.5 Gbit/s to 10 Gbit/s and reducing the channel spacing from 0.8 nm (100 GHz) down to 0.4 nm (50GHz) or in some cases 0.3 nm. Why not continue to line-rates 40 Gbit/s, or even 100 Gbit/s and even closer spacing?

Dispersion and cost are two reasons. Increasing the symbol rate makes the symbol duration shorter; it also makes its spectrum broader. The narrower symbols are obviously much more sensitive to pulse spreading and the broader spectrum means that the same chromatic dispersion causes more pulse spreading.

It is relatively simple to compensate for chromatic dispersion, but compensating for Polarisation Mode Dispersion (PMD) is somewhat more difficult and expensive, mainly because PMD is a time-varying effect and requires adaptive compensation. Submarine fibres have PMD values of around 0.05 ps/√km, with an absolute worst case of 0.1 ps/√km. 75% of terrestrial fibres fall between 0.1 and 0.5 ps/√km, but values as high as 5 ps/√km are not uncommon and peak values can exceed 20 ps/√km. Penalties start to become significant once the dispersion exceeds around 10% of a bit period, as shown on the following graph, where the circled value is the PMD in ps/√km. Note, however, this is an approximation and that it is possible to add PMD mitigation devices to extend the range; this needs to be done for individual wavelengths and is thus potentially expensive.



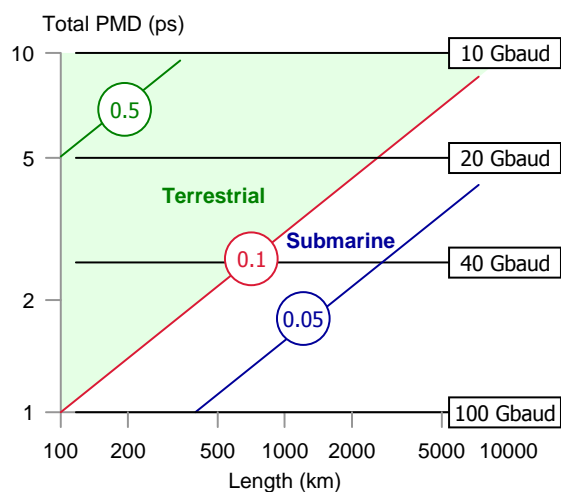


Figure 5 - PMD limits without mitigation

From this graph it is clear that binary signaling at 100 Gbaud will be very challenging for almost any distance other than Metro spans, while 40 Gbaud is limited to shorter spans, with the limits depending on the actual PMD values. However, 10 Gbaud and 20 Gbaud binary signals can propagate over very long spans of submarine-grade fibre.

In addition to the transmission problems that dispersion poses at higher line-rates, components that work at high speeds are usually more expensive than lower speed devices, which usually consume less power and are of well proven reliability. (This is partly linked to the volume and duration of production, so it is likely to change with time.)

Given these factors it might seem obvious that transmission at $N \times 10$ Gbaud would be the best solution, but one has to balance this against other factors. 10 Gbaud implies multiple levels to transmit 40G or 100G – a lot of levels for 100G! This in turn makes both transmitter and receiver architecture more complex, increasing component counts – and usually cost as well – while reducing reliability. In certain cases (a 20 Gbaud example is shown in Figure 8) the higher line-rate may be propagate better than the lower one, so the choice is not straightforward.

4.2 Phase modulation

Early systems used On/Off or Amplitude Shift Keying (OOK or ASK), but it is well known that modulating phase rather than amplitude leads to better receiver performance – typically around 3 dB can be gained – and people are increasingly



moving to Phase Shift Keying (PSK) or Differential Phase Shift Keying (DPSK). Differential Phase Shift Keying uses a phase inversion from one bit to the next to indicate a zero: a one is signaled by keeping the phase constant. The cleverness of this approach becomes clear when one looks at how they are detected.




Figure 6 - Binary DPSK detection

The figure above shows how adding the DPSK signal to a one bit delayed version of itself results in an amplitude modulated signal which is easily converted into an electrical signal by photodiodes. This detection scheme is relatively simple to implement and produces significantly better performance than an ASK signal with the same Optical Signal to Noise Ratio (OSNR).

However, since the signal and the shifted signal both contain noise it should be possible to improve on the detection process using coherent detection, where the signal is mixed with a low noise local oscillator of the same frequency and polarisation. In this case zeros and ones are signaled by different phases.




Figure 7 - Binary PSK detection



Coherent PSK detection is harder to implement because the local oscillator needs to match frequency, phase and polarisation of the incoming signal quite precisely, yet the phase and polarisation of the signal at the end of a system vary with time.

(An interesting solution to this problem is shown in Figure 12 - Digital coherent detection). Because of these complexities, most suppliers have implemented DPSK schemes and are now considering coherent detection, although at least one supplier has implemented it. Moving from ASK to DPSK gains around 3 dB; moving from DPSK to coherent PSK gains less than 1 dB, so developing coherent detection represents a technical/economic decision. In Section 5 it becomes clear that in the case of long non-linear systems DPSK detection may actually be better than coherent detection.

Early 10 Gbaud experiments rapidly confirmed that DPSK, which involves a relatively simple detection scheme, did indeed give better performance, but they also showed that in regions where chromatic dispersion is low the performance was not so good for spans of several thousand km. This is because in these regions the effect of cross-phase modulation are worst and a phase modulated signal is naturally sensitive to phase noise. However, 20G DPSK (with its broader spectrum) is much less affected, as shown by the following graph.

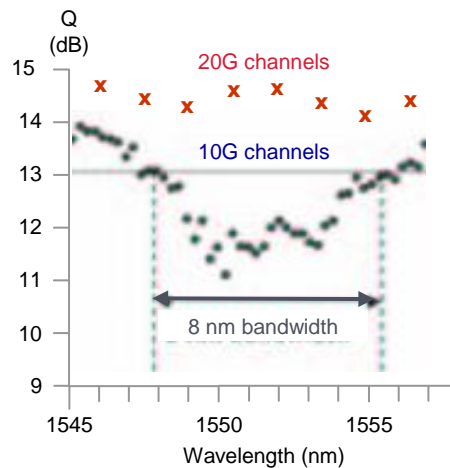


Figure 8 - Comparison of 10 Gbaud and 20 Gbaud DPSK

The graph shows the results from two different publications [References 1. Heinrich Hertz for 20 Gbaud; 2. Alcatel-Lucent for 10 Gbaud]. The offset is because the systems are not *exactly* the same, but it is clear that all the 20G channels offer roughly the same performance, while the central 10 Gbaud DPSK channels are noticeably worse than the others. Alcatel-Lucent authors suggest that the solution is to use ASK channels in the low dispersion region – not ideal, since it requires two types of spare.



As mentioned earlier, dispersive effects are less significant for shorter systems and this effect is significant only for long span system, e.g., long submarine systems. (Xtera uses 20G for long repeated systems, but 10 Gbaud and 40G rates for shorter submarine systems and terrestrial systems.)

4.3 Multi-level Systems

While transmitting multiple levels offers the possibility of increasing the number of bits per symbol this can only be achieved by also increasing the average power per symbol. This can be seen in the following graphic.




Figure 9 - Power for 2, 4 and 8 level schemes

Adjacent phase levels must be roughly the same distance apart to allow detection in the presence of noise. This in turn means that as the number of levels is increased, so the signal power also increases, and a good approximation is that one requires $2x$ ($= 3\text{dB}$) more power each time the number of levels is doubled. (This is effectively the same as saying that each bit requires a fixed amount of energy or a certain number of photons.) In a linear system this is not a problem, but most existing submarine systems are not linear and simulation and experiment suggest that for spans of >4000 km, two levels are better than four.

5 DETECTION

Compared with amplitude modulation, using phase modulation reduces the minimum required OSNR by around 3 dB for differential detection, and around 4 dB for coherent detection.

The small extra gain offered by coherent detection is clearly of value, but it is important to appreciate that coherent detection can be seriously affected by non-linear effects. For example the effect of self-phase modulation and cross-phase modulation is to create very high frequency phase noise which gets worse if the system is long or if the chromatic dispersion is close to zero. For medium length terrestrial systems the effect is minimal because the dispersion is high. For long submarine spans, however, dispersion is low and coherent detection may have



little benefit and can be worse than differential detection when non-linear effects cause occasional rapid phase shifts.

In coherent detection the local oscillator tracks the average phase and is affected by the sudden phase shift, while differential detection uses the phase of the previous pulse, which will also be shifted by a similar amount. [Reference 3]

It is also interesting to look at the components needed to implement differential and coherent detection systems.




Figure 10 - Differential detection

In a simple differential detection system the first step is to adjust the incoming signal so that a simple polarisation splitter will separate the two signals. After this each signal passes through a splitter, one bit delay and combiner, and is then detected by photodiodes in the receivers. The polarisation adjuster needs to be adaptive and represents a significant new feature, albeit not a particularly complex one.




Figure 11 - Analogue coherent detection

To achieve coherent detection a local oscillator signal is added to the two polarisations. The key difficulty is controlling the local oscillator to match both the frequency and phase of the incoming signal. This is quite hard and an interesting alternative is shown in the next diagram.




Figure 12 - Digital coherent detection



Here there is no polarisation adjustment, so the two receivers each receive a combination of the two polarisations. Very high speed Analogue to Digital Converters (ADCs) produce digital samples of the incoming signals which are then fed to a very high speed Digital Signal Processor (DSP) chip, which computes two different (but related) linear combinations of the samples to re-create the two polarisations. While this is more complex, as the polarisation state will change, it removes the need for polarisation adjustment, and offers some other interesting possibilities. For example it can compensate for minor drifts in the local oscillator and can also digitally re-shape the pulses to compensate for the effects of pulse spreading caused by chromatic dispersion. At the expense of even more complex processing it could provide some degree of mitigation for non-linear distortions.

To achieve this requires some very high speed electronics. The ADCs must digitise in real time and sample at least two times per bit, i.e. >20 Giga samples/sec for 10G signals. The DSP must be able to process two such streams and perform polarisation rotation, oscillator drift correction and dispersion compensation. The first two of these processes must be adaptive, as the parameters of these corrections will be time-varying. Dispersion compensation is also best as an adaptive process to simplify initial set-up and to handle the effect of repairs which may change the overall dispersion.

The high speed and amount of processing make the DSP chips required significantly more complex than those used in more routine applications – one paper [Reference 4] describes a CMOS ASIC with 20 million gates and a consumption of >20 watts – and it is likely that they will remain large and relatively high consumption devices. However, the removal of dispersion compensating fibre/devices and the optical amplifiers needed to compensate for their loss means that ultimately the DSP approach should result in overall lower power consumption. Whether there is a benefit will depend on the cost/consumption/reliability of the units that are being replaced by the DSP chips.

6 SUMMARY / CONCLUSIONS

It seems clear that different solutions are required for different types of systems. System length and non-linearity represent two key system parameters, with non-linearity becoming less problematic if the fibre's chromatic dispersion or core size is high.

Existing terrestrial systems use fibre with relatively high dispersion and operate over moderate distances, as do unrepeaters submarine links. Repeated submarine systems, however, use fibres with low chromatic dispersion and are generally long enough for non-linear effects to be important.

Linear systems can use multi-level formats (such as Quadrature phase-shift keying (QPSK)) and coherent detection to improve performance, but when non-linear



effects become significant the extra power needed for multiple levels increases symbol distortion and the best format choice is binary. Because non-linear effects can produce high-speed phase shifts, coherent detection over long distances is likely to offer little (if any) gain over direct detection, which has the benefit of simplicity.

The following table attempts to summarize the solutions for different systems.

Fibre Cd	Short (<3000 km)	Long (>3000 km)
High	Terrestrial NRZ QPSK Coherent Det.	Both in future? NRZ QPSK Coherent Det. EDC
Low	Submarine (existing systems)	
	NRZ QPSK / DPSK Coherent Det.	RZ DPSK Direct Det.

Figure 13 - Comparison of system types

(EDC = Electronic Dispersion Compensation)

Terrestrial and short submarine links suffer little from non-linear effects and can use similar solutions, e.g., NRZ format, multi-level signals and coherent detection, although the submarine cases will need more attention to dispersion compensation.

Existing long submarine systems need to use the RZ format, binary signals and are unlikely to benefit from coherent detection.

In the future it is possible that all systems could be based on high dispersion fibre, with Electronic Dispersion Compensation handling the high chromatic dispersion.

Nearly all types of fibre optic link are suitable for dual polarisation transmission, but this will only provide higher capacity when the link was previously limited by bandwidth rather than noise.

Digital Signal Processing offers the potential to replace polarisation adjustment and



optical dispersion compensation; it also relaxes local oscillator stability requirements in coherent detection and may be able to mitigate some non-linear effects. Given these benefits, its relatively high complexity and power consumption should be justified.

Finally, there is still room for improvement. Soft decision decoding offers the potential for a further 2 dB or so of coding gain which should translate into at least 50% more capacity for a given system.

Acknowledgments

The author would like to thank Joerg Schwartz, Paul Farrugia, Hervé Février, Benoit Kowalski and others for the help they have provided me while writing this white paper.

In the event that you have questions or comments, don't hesitate to contact:

Tony Frisch SVP Repeated Solutions

Telephone	+44 1335 344 394
Mobile	+44 7963 650 079
E-mail	tony.frisch@xtera.com

References

- 1 "DPSK Transmission Experiments Over DSMF and NZDSF", Philippe Plantady, et al.
- 2 "Wavelength Dependent Performance of 20 Gbit/s RZ-DPSK for Upgrades of Non-Slope-Matched Submarine Transmission Systems", Lutz Molle, et al.
- 3 "A performance comparison of differential and coherent detection ...", Renaudier, et al. OFC 2007.
- 4 "DSP, A Disruptive Technology for Optical Transceivers", K. Roberts, ECOC 2009.

Xtera Communications specializes in network infrastructure that delivers the highest capacity, reach and value. Providing solutions for enterprise and telecom companies, Xtera offers an extensive portfolio of optical and IP networking solutions for submarine, long-haul, regional, metropolitan and enterprise applications. With deployments across five continents, Xtera's optical transport solutions help service providers expand and accelerate their market reach with new deployments and extend the life of existing network assets with cost-effective upgrades. Our IP networking solutions for WAN traffic management deliver efficient, flexible network optimization. Xtera's innovative technologies offer exceptional quality and performance, driving customer success. For more information, visit www.xtera.com.

Xtera and the Xtera logo are registered trademarks of Xtera Communications, Inc. All other trademarks are the property of their owners.

Copyright © 2010 Xtera Communications, Inc. All rights reserved. Information in this document is subject to change without notice. Xtera assumes no responsibility for any errors that may appear in this document.

