LOW LATENCY CABLES

Alasdair Wilkie (Hibernia Networks)
email: alasdair.wilkie@hibernianetworks.com

Hibernia Networks
International Exchange Centre
Clonshaugh Industrial Estate
Dublin 17
Ireland

Abstract: Low Latency cables are a new and interesting addition to the submarine cable community where they are used to support the lowest latency routes between city pairs i.e. New York and London. Whilst the cables are limited in their routing, there can be only one "Lowest Latency" path on any one route. In particular they present a series of challenges that go against the grain with regards to the standard way of thinking with respect to route engineering and improvements need to be made in laying techniques to overcome the dangers of untried routes. This paper will provide an insight to the rationale for low latency cable and explain how Hibernia Atlantic have overcome the challenges and embarked on the lowest latency path from London to New York including the new Express cable.

1. Why Low Latency?

The main requirement for Low Latency routes has come from High-Frequency Traders, sometimes known as Algorithmic Traders.

Although there are other customers who are looking for lower latency i.e.:-

- cloud computing
- video on demand systems
- replication systems
- film editing

2. Euclidean, Haversine and Great Circle Distances.

The principle behind a low latency cable is the shortest route, and as every school boy or girl knows the shortest distance between two points is a straight line. So looking at London to New York, a straight line using Euclidean geometry between two points in space is given by:-

$$\frac{d^2}{R^2} = 2 - 2cos\Theta_1 cos\Theta_2 \cos(\phi_1 - \phi_2) - 2\sin\Theta_1 sin\Theta_2$$

Where :-

d = distance
R = distance from origin
Θ = Latitude in radians
ϕ = Longitude in radians

This gives a distance of 5401 km but has the disadvantage of tunnelling through the earth to a depth of greater than 600 km.

This therefore is the first impractical solution that we come across.

The next solution is to use the Haversine formula to calculate the Great Circle route between two points on a sphere, this is given by :-

$$haversin\frac{d}{R} = haversin\,(\Theta_2 - \Theta_1) + \cos\Theta_1 \cos\Theta_2\, haversin(\phi_2 - \phi_1)$$

Where :-

$$\text{haversin}(\Theta) = \frac{1 - \cos\Theta}{2}$$

This gives us a distance of 5577km from London to NewYork but the Haversine formula is for a sphere and the Earth is an oblate spheroid. Although up until fairly recently this would have been an acceptable method, however in 1975 Vincenty[1] derived a formula since modified by Karney[2] based on work by Bessel which has significantly improved the accuracy of Great Circle distances.

The Vincenty formula uses an iterative calculation until the difference between the λ is < 10^{-12} or 0.6mm accuracy. The full formula can be found in Ref 1.

Using the modified Vincenty formula the Great Circle distance from London to New York is 5594km this therefore is the theoretical shortest distance that could be achieved for a cable from London to New York. This still doesn't take into consideration any hills or valleys that will change the length of the route.

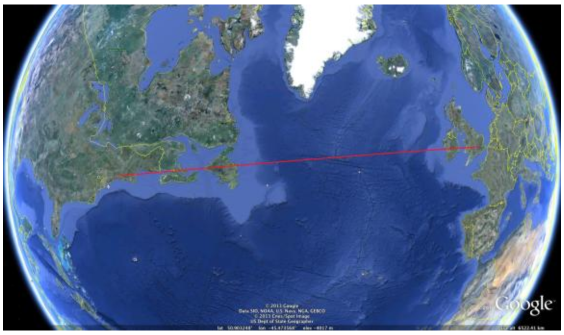

Great Circle route from London to New York.

3. Latency Calculations

A simplistic calculation for latency and one that works for short routes particularly on land is that there is 1ms of Round Trip Delay (RTD) for each 100km of route. However when applied to transoceanic cable routes this simplistic calculation is inaccurate and a more empirical formula is required.

Breaking down the route into constituent components we see that the following are the major contributors to latency

- Fibre
- Amplifiers
- Equalisation
- Transponders

Dealing with each of these in turn we have

Fibre

Light travels in a vacuum at 299,792.458 km/s, however in a fibre the glass slows the transmission of light down and this is reflected in the refractive index of the glass which for standard fibre types commonly used in long haul telecommunications is in the region of 1.463 to 1.5.
Using a range of these index figures on a hypothetical 1000km link gives the following latencies for each:-

IOR	Latency for 1000km (ms)	%Δ
1.463	4.88	Ref
1.4682	4.90	+0.36
1.5	5.00	+2.53

So even for a small decrease in the refractive index a useful improvement in latency can be achieved.

For our calculation the unit latency per km is calculated as follows:-

$$RTD_{Fibre}\ (ms\ per\ km) = \frac{IOR}{c}$$

Where c is the speed of light in a vacuum in km per ms.

Amplifiers

The majority of amplifiers on the market today have internal fibre lengths varying from 30 to 60m meters in each amplifier. About 15m of this is for the erbium doped fibre and the rest is for splicing and storage. The more amplifiers the more fibre and the greater latency, so reducing amplifier count downwards and reducing the fibre in the amplifiers will again have a positive effect on the latency.
Amplifier latency is designated as RTD_{AMP}

Equalisation

There are two areas of equalisation that need to be considered, one is the inline equalisation on the submerged plant and the other is the equalisation of the signals in the terminal station.

The inline equalisers are a similar prospect to the amplifiers and since their numbers are usually very small these can be ignored for all practical purposes.

Equalisation in the terminal station, is done in one of three ways, fibre equalisation, Bragg gratings or for coherent systems electronic equalisation.

In-line fibre equalisation should be avoided in low latency systems as not only will it add significant latency but since it is carried out on a per channel basis no two channels will have the same latency.

Bragg gratings or electronic equalisation can both be used as these add insignificant amounts of latency to the system and importantly the latency is the same for all channels.

Transponders

Transponders at the end of each link or segment provide the final interface to the customer's network. The majority of interfaces are still at 10 Gbit/s even with the advent of 40Gbit/s and 100Gbit/s line rates. Transponder latency varies from ≈50µs to ≈200µs per card. Transponder latency per transponder pair is designated by RTD_{TX}

So to add all of the above together we have an empirical formula for a simple cable landing station to cable landing station system as follows:-

$$RTD_{Total} = lRTD_{Fibre} + xRTD_{Amp} + RTD_{TX}$$

Where:-

l is the total fibre length
x is the number of amplifiers

In the empirical formula above we have ignored any equalisation units either land or marine based, however this will provide a good approximation of the achievable latency.

4. Route Considerations

With all cable routes, and in particular ones that contain constraints such as low latency, the consideration for a combined land and sea route is paramount. This also means that the integration for the land and sea routes must be taken into account, every interface can add latency.

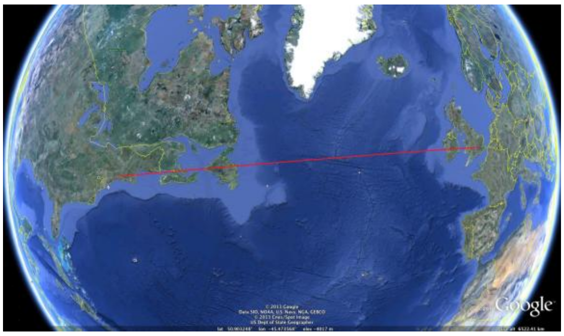

The route above from London to New York would pose some serious challenges with regards to building this route, the number of land to sea interfaces as well as the limitations regarding the land routes would make this an economically unjustifiable route to construct.

If the route is split between land and marine, in the knowledge that the majority of marine routes can follow a great circle route if required, or at least close to one, then we can understand how far from the great circle route we can deviate before we compromise the overall latency of the system. All fibre routes below are based on using the walking distance between two points on maps.google.com[3], which whilst not perfect is an easy way to start.

Looking specifically at the section of the great circle route above, then there are two particular issues that make this route economically impossible. One is the direct route from London to just north of Llan-non on the Welsh coast, this direct route passes through the Brecon Beacon mountains as well as the centre of many large towns and cities. The great circle distance from London to Llan-non is 303km whereas the shortest usable road route is 344km.

The same issue is also true of looking at the route across Ireland. The landing point on the above route is close to Glendoyne with an exit point near Pollnaclogha. The great circle route between these points is 221km but the shortest path for a cable would be 255km.

Doing the same for the other land sections we have the following (all distances in km):-

Country	Great Circle Route	Shortest Duct Route
UK	303	344
Ireland	211	255
Newfoundland, Canada	414	525
PEI & Nova Scotia, Canada	156	196
USA	356	400

If these increases in distance are added up we get an increase in length from the great circle route of 280km this is for duct, to realistically look at fibre length then 10% must be added to the duct length to get the fibre length, making the variation 308km.

Taking this as a staring point the redesign using a primarily marine route can take place.

5. Route Redesign

Since submarine cables can invariably be installed in almost a great circle route then increasing the amount of submarine cable in the route between London and New York is the most effective way of keeping the latency low.

Doing this between Weston-super-Mare in the UK and Shirley in Long Island provides a great circle route of 5323km so this is the next starting point for low latency reduction.

The submarine route shall be split into three main areas, the deep water section and the European and North American Continental shelves.

For the European continental shelf, the obvious starting point is somewhere at Weston-super-Mare, UK. The next limiting point is Fastnet rock on the south west of Ireland. From Fastnet a great circle route can be drawn straight to Newfoundland.

Putting these waypoints into a map format we can see that he beginnings of a route would look something like this.

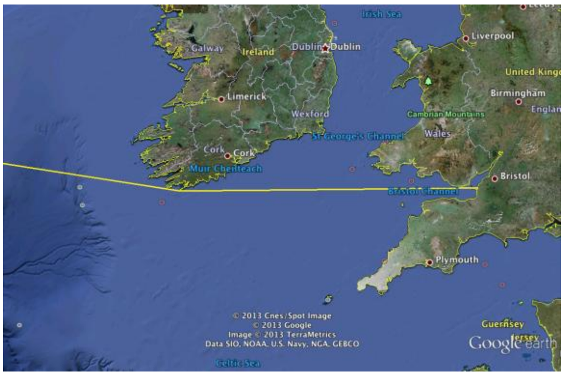

For the North American continental shelf after passing the southern Avalon Peninsula in Newfoundland the next way point is the southern tip of Nantucket Island, before heading to the beach near Shirley, NY.

Putting all of these waypoints into a trans Atlantic route we have a route of 5364km which is only 41km longer than the great circle route between the two points.

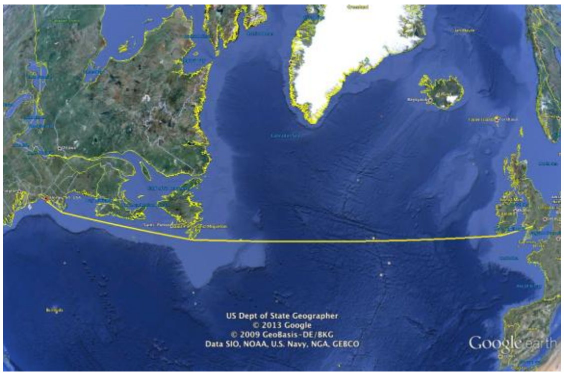

This is the starting point for the route before the detailed engineering can take place.

The detailed engineering is the same as would be done on an ordinary route so the following must be take into account:-

- Windfarms and other offshore renewable energy
- Oil and Gas Exploration blocks
- Severe slopes
- Marine Conservation Zones and other protected areas
- Other cables and pipelines
- Limitations imposed by governments within 12 mile limit and the EEZ
- Traffic separation zones
- Anchoring areas
- Fishing areas

This list is not exhaustive and there will be other limitations on the route which come up during the desktop study or survey.

6. Final Route for Hibernia Express

Due to the fact that Hibernia Networks already have a cable from Halifax to New York via Lynn, which is fairly close to the great circle route from London to New York, it was decided to bring the cable into the current landing station near Halifax. This also had the advantage of avoiding the very shallow Georges Bank and making the regulatory burden less by not bringing a cable into the USA.

The other areas which have been avoided are Ballards Bank to the south of the Avalon peninsula due to steep slopes. Irish territorial waters for licensing reasons, Welsh territorial waters for special areas of conservation and the Atlantic Array windfarm area in the Bristol Channel.

Putting all of these facts into consideration and ensuring that the ICPC recommendations for pipeline and cable crossings are taken into account a survey route has been found.

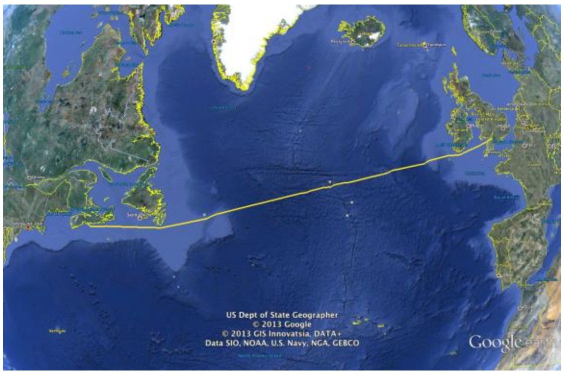

This route together with the existing route from Halifax to New York and a new route from Weston-Super-Mare to London will provide Hibernia Networks with the latency required.

7. REFERENCES

[1] Vincenty, T. (April 1975a). "Direct and Inverse Solutions of Geodesics on the Ellipsoid with application of nested equations". *Survey Review* XXIII (176): 88–93.

[2] Karney, C. F. F. (2013). "Algorithms for geodesics". *J. Geodesy* **87** (1): 43–55

[3] Maps.google.com © Google