September, 2013

The Future of Copper

By Paul S. Cawte
http://pscomms.wordpress.com/2013/09/19/the-future-of-copper/

The death of copper access networks has long been predicted, but in many countries they remain alive and well and in fact still reign as the dominant access technology. Why hasn't fibre with all its alleged benefits taken over? The answer is simple, where copper is already deployed it can be upgraded cost effectively to support high speed broadband services, whereas the cost to replace copper with fibre is more often than not prohibitive.

The story is different for new network deployments where fibre is an attractive option. Also in certain geographies where housing is dense (e.g. large apartment blocks) it can make sense to upgrade the connection to the building to fibre, but even in these cases internal distribution may well remain as copper.

With current and planned future developments in copper access technology this general situation is unlikely to change. It is true that fibre will be deployed deeper into the access network to reduce the length of copper connections (called drops) to the home, but in the majority of cases the final copper connection will be retained.

There are two types of copper access network that are widely deployed:

- Twisted pair networks, historically deployed by phone companies to deliver phone services, and
- Coaxial, cable TV (CATV) networks, historically deployed by cable TV service providers.

Both of these technologies have well developed upgrade paths that will allow them to remain competitive with fibre for a considerable time yet.

This paper looks at the trends in development of both twisted pair and CATV copper access networks, showing how they will remain competitive with fibre access. Following these trends, the paper then predicts what could become the standard architecture for copper access, with an incremental upgrade path to pure fibre.

Twisted Pair Access

Twisted pair networks were originally deployed by phone companies to provide telephone services. In these networks, each customer of the phone company has a dedicated twisted pair cable from their home back to the exchange building where the telephone exchange was housed. Each twisted pair connection consists of a pair of copper wires or conductors twisted together along their length. This twisting is important as it reduced the impact of electrical interference on the phone service and is similarly important for the broadband Internet services that are now delivered over the same cables.

The connection from the customer to the exchange building was normally several kilometres or more in length. To manage the large volume of connections required back at the exchange building a number of cross-connect points where established to combine the individual twisted pair cables from the customer onto multi-pair cables (e.g. 50 or 100 pairs) carrying the connections from a group of homes and then onto multi-100 pair cables for connection back to the exchange building. One typical configuration, such as that deployed by British Telecom (BT), consists of two levels of cross-connection:

- Primary cross-connect point (PCP): which is the first level of cross-connection after the exchange building. This for example might provide service to roughly 1000 homes with the distance to the home being between 500-1000m. Normally the PCP is housed in a street side cabinet. As it is an entirely passive electrical cross-connect point it does not need a power connection to it.
- Distribution point (DP): the point at which the individual customer drop connections are cross-connected onto larger multi-pair cables. The nature of the distribution point will depend very much on the local geography it is serving. For example:
    - In rural areas it might be a joint mounted on a telephone pole from which the drop cables are hung,
    - In urban areas it might be in a small street side cabinet or even in an underground joint,
    - For MDUs (multi-dwelling units) such as apartment blocks the distribution point will often be housed within the building itself.

  Given the diversity of locales, DPs can range significantly in size, but they may typically serve 10-50 homes and are located 100-200m from the home.

Even though when twisted pair access networks were deployed, broadband Internet or even just the Internet was not envisioned, these networks have proved remarkably adaptable and with some ingenuity have been upgraded to support increasingly high-speed Internet access. This can be seen through a series of technological and access network architecture evolution steps as described in the following sub-sections.

Dial-up access

In the early days of the Internet and other data services, consumer access was achieved via dial-up access over a standard phone line as voice band data transmitted and received by a modem. This approach saw successive iterations in modem technology from 300 baud audio coupled modems to directly connected 56kbit/s V.90 and V.92 modems. Whilst this represented a number of generations in modems the access network remained largely unchanged other than an increased number of 2nd lines from customers ordering a dedicated phone line to use for data services.

ISDN

Basic rate ISDN was intended as a replacement for traditional analogue phone lines. It was designed to allow two phone calls or 128kbit/s of data to be transmitted over a standard copper access connection. Basic rate ISDN never found widespread adoption (except in Germany), as it was expensive and too late to the market to gain widespread acceptance. In most markets its major use was a backup circuit for mid-range business routers at the time connected via 2Mbit/s leased lines or frame relay circuits.

ADSL

The next major enhancement in Internet access over twisted pair was using asymmetrical digital subscriber line (ADSL) technology. This solution made use of frequency splitters on the twisted pair line, so that the low frequencies (below 3.1kHz) were used for voice and the higher frequencies used for broadband Internet service. To enable ADSL services a DSLAM (digital subscriber line access multiplexer) was deployed in the exchange building and an ADSL modem in the customers home. At both ends a filter or splitter was required to separate the voice signal from the broadband data signal.

In most cases changes in the copper access network between the customer and the exchange building were not required other than some tidying up and enforcement of good cabling practices. However, in the exchange building significant re-cabling was required to incorporate the filters from which a connection was now required to both the exchanged (for voice) and DSLAM (for Internet services).

Although, as has been indicated, little change was required to the copper plant, this was not the case for customers on very long connections (or loops). The performance of ADSL diminishes as the loop length increased and for connections much above 6km a reliable ADSL service could not be provided. For such customers, either the DSLAM had to be installed closer to their home (for example in a street side cabinet) or ADSL service was not available.

ADSL2+

To improve the performance of ADSL, ADSL2+ was introduced. ADSL2+ achieved its greater performance by using higher transmission frequencies. This increased the downstream performance from 8Mbit/s (for ADSL) to up to 24Mbit/s. However, these high data rates were only available for customers living very close to the exchange. For longer loops the performance of ADSL2+ was not much better than ADSL and was seen as increasingly sub-standard compared to alternative offerings from cable TV companies.

VDSL2 and FTTN/FTTB

To provide the next step wise upgrade in performance VDSL2 was introduced. By using even higher frequencies on the copper cable than ADSL2+, VDSL2 offered higher data rates, however, these were only available on short loops. To overcome this, VDSL2 envisaged a number of deployment scenarios:

- In the exchange building, much as for ADSL. This was only applicable in those places were the loop lengths were short (<1km).
- In a street side cabinet, with a fibre connection from the exchange building to the cabinet. This is often called fibre to the node (FTTN). In BT's case for example, they deployed MSANs (multi-service access nodes) at the PCP cabinet. The practical difficulties of this are:
    - Installing a new cabinet which may meet with resistance from local residents or local authorities

  - Obtaining a power connection for the MSAN
  - Installing the fibre backhaul to the exchange building.

  In BT's case they offer an up to 80Mbit/s service using this FTTN architecture.

- In fibre to the curb (FTTC) architectures the MSAN is installed closer to the home (e.g. <300m). This can provide higher data rates but of course requires a much larger number of cabinets. Most operators seem to be focussing on FTTN deployments, with FTTC remaining a future option.
- In countries with high density accommodation (e.g. high rise apartment blocks) fibre to the building (FTTB) is very applicable. In this deployment scenario an MSAN is deployed in an equipment room within the apartment building, with VDSL2 distribution over the existing in-building twisted pair copper. This is an attractive solution as compared to FTTN space and power for the cabinet are less of a problem, and it is just the fibre connection that is required to the building.

In terms of actual deployments, the current major use cases of VDSL2 are FTTN and FTTB.

Vectoring

The performance of VDSL2 is limited by cross-talk with other twisted pair cables carrying VDSL2 within the same cable binder. This means that as the penetration of VDSL2 service increases the performance achieved by all users will diminish somewhat. The next stage of VDSL2 evolution is to cancel this cross-talk using an approach called vectoring. Vectoring makes use of the considerable processing power available in DSPs (digital signal processors) and the knowledge of what has been transmitted on every VDSL2 pair to largely cancel out the effects of cross-talk. With vectoring the performance achieved is close to that achieved when just a single VDSL2 pair is in use. This means that vectoring is able to achieve close to perfect cross-talk cancellation. Major operators are now in the process of upgrading their MSAN deployments to support vectoring. In the UK, for example, BT plan to launch an up to 100Mbit/s product.

Bonding and Phantom Mode

Once vectoring has been implemented there are two further steps that can be taken to increase performance. One approach is the bonding of two or more pairs. Whilst operators typically do not have enough copper pair capacity in their network to provide this for everybody, a figure of around 10-20% of customers is realistic. Using 2-pair bonding and vectoring a doubling in performance can be achieved. Moreover, an additional virtual or phantom pair can be created between the 2-pairs. This phantom pair allows an additional signal to be carried, so that 2-pair bonding, plus vectoring, plus phantom mode would yield a 2.5-3 times improvement in performance over a single pair, potentially yielding a download rate of 250 to 300Mbit/s. Currently bonding is not widely deployed and consequently consumer grade, bonding capable modems are not available. Phantom mode has been demonstrated in lab deployments but is not yet supported by commercial equipment.

G.fast and FTTdp

Whilst bonding with or without phantom mode might remain somewhat of niche solution for example restricted to business customer connections, the next significant step in twisted pair access evolution will be the introduction of G.fast and a move to fibre to the distribution point (FTTdp) architectures. FTTdp is effectively a form of FTTC, with active equipment deployed at the distribution point. This removes the copper distribution cables and just leaves the final drop cable to the customer as copper. G.fast is being developed specifically for these architectures, incorporating a number of specific capabilities making it easier to deploy. The G.fast standard is still under development and is consequently not finalised. It is, however, expected to have the following capabilities:

- Targeted for short loops less than 200m
- Designed to support rates up to 1Gbit/s at 100m and 500Mbit/s at 200m
- Capable of being reverse power fed (i.e. from the home), so that an operator does not need to provide a power connection to the DP (something that would be quite difficult given the number of DPs).
- Designed for a low number of users per DP (e.g. 8 to 24)
- Equipment is likely to be manufactured in a range of units designed for easy deployment, such as:
    - Ruggedised, pole mounted for aerial delivery
    - Sealed and passively cooled for underground deployment
    - Ruggedised, wall mounting for use on the side of the street
    - Etc.

In principle deploying fibre to the DP should be relatively straightforward. The main challenge is installing the fibre to the DP and simply the very large number of DPs that will need to be upgraded. Whilst this author does not know the exact number for the UK, but it will be in the high hundreds of thousands, if not over one million.

It is clear that rolling out fibre to the DP will be a major undertaking, nevertheless this is substantially less difficult that upgrading all homes to have direct fibre connections. One big challenge in the UK is that a significant number of the drop cables are directly buried and replacing these with fibre would be a mammoth undertaking.

Coax, CATV Access

Cable TV coaxial access networks are substantially different from twisted pair telephony access networks. In a cable TV network the coaxial cable is shared between many users, whereas in a phone network each customer has their own twisted pair line. However, coaxial cables are much better conductors of high frequency signals that twisted pair cables so that a much great bandwidth can be sent over the coax – i.e. the whole TV spectrum.

Originally CATV networks just transmitted TV channels over a shared cable – i.e. logically they were more or less the same as a broadcast TV network, but with aerial transmission replaced by transmission over a coaxial cable. The first networks were one way only, transmitting signals from the head-end to the home, and consisted entirely of coaxial cable that branched and split along the way. Wide-band RF amplifiers were used to maintain the signal strength as it was reduced by splitting and due to losses along the length of the cable.

As optical technology improved, operators were able to improve the performance and reliability of their CATV networks by introducing hybrid fibre coax (HFC) architectures. This involved replacing large sections of trunk coax with optical fibre – thus introducing fibre to the node architecture into CATV HFC networks. Operators also started to introduce a return path – i.e. a transmission path from the home to the head-end. This allowed them to introduce pay-per view services and collect telemetry information from the optical node and amplifiers.

At this stage CATV networks still just carried large numbers of TV channels with some interactivity achieved via pay-per view or near video on demand (N-VOD). With the deregulation of the telecoms sector vendors and operators spotted the opportunity to gain additional revenues by providing telephone services over cable TV networks and a number or proprietary cable phone systems were developed and deployed. Early deployments suffered from reliability problems due to the challenges of operating over cable TV networks. This was particularly true of the return path. Nevertheless, when operators and vendors gained more experience with the technology and the operation of it a number of successful deployments were achieved, such as those using the Arris Cornerstone Voice product. Deployments, however, remained relatively limited due to the high cost of the technology, which was not helped by its proprietary nature.

Although the early cable voice systems were not generally a commercial success the experience gained enabled Internet access using cable modems to be a success. Whilst early solutions were again proprietary, the DOCSIS (data over cable service interface specification) architecture was standardised by CableLabs. This found widespread adoption throughout the cable industry and was one of the key drivers of the broadband Internet revolution – at the time a number of PTTs (public telephone companies) had not committed to widespread ADSL deployment and it was the success of the cable industry in proving the customer demand for broadband Internet service that forced the telco operators to aggressively role out ADSL service.

DOCSIS works by allocating one or more of the TV channels for data transmission. It also of course requires a return path. A large part of the work of deploying DOCSIS was the installation or upgrading of the return path to allow for reliable data transmission.

The capacity on a DOCSIS network is shared between all the users within the segment. Operators can adjust the segment size and the number of channels they allocated to DOCSIS to vary the capacity. In the US, where the TV channels are 6MHz wide each downstream channel typically carries 42.88Mbit/s whereas on the European 8MHz wide channels 55.62Mbit/s can be carried. In the upstream a variety of channel widths and modulation rates are supported. However, a typical configuration is an upstream channel capacity of 10.24Mbit/s.

The degree to which these channels are shared depends upon the segmentation. There are different types of segmentation in a CATV network:

- Optical segmentation – this corresponds to the optical node size. Other segmentations result from the combination of one or more nodes for different upstream and downstream channels.
- Downstream data – is the number of nodes over which a single or group of downstream data channels is shared. A historically common configuration was to share the downstream data channels over four optical nodes.
- Upstream data – is the number of nodes over which one or multiple upstream data channels are shared. Due to the challenge of noise ingress into a cable return paths this most commonly corresponds to a single optical node.
- Downstream TV – is the number of nodes over which the TV channel line up is shared. This can be very large, for example, a whole city.
- Downstream VoD – is the number of nodes over which video on demand channels are shared. This needs to be relatively small otherwise a large number of channels will need to be allocated so that there is always capacity available for a user to access their desired video on demand.

The DOCSIS standard has gone through a number of iterations, increasing performance, functionality and data rate:

- DOCSIS 1.0 – was the initial version
- DOCSIS 1.1 – introduced quality of service (QoS) mechanisms required to support voice and improved security. DOCSIS 1.1 was deployed widely and in many cases enabled cable operators to provide competitive primary line voice services.
- DOCSIS 2.0 – introduced more upstream modulation options and the advanced PHY (physical layer). This included ingress cancellation, which considerably improved the robustness of upstream communication.
- DOCSIS 3.0 – introduced channel bonding. Prior to DOCSIS3.0 a modem could only use a single upstream and downstream channel. This created an absolute limit on service of up to 40Mbit/s in the US and 50Mbit/s in Europe. DOCSIS 3.0 allowed this limitation to be overcome by combining channels. Thus, for example, if four channels are combined, then an up to 100Mbit/s service can easily be provided.
- DOCSIS 3.1 – removes the 6 or 8MHz channel structure and allows a single block of up to 200MHz of spectrum to be utilised. The increased efficiency gained from removing this channel structure and improved modulation allows for a 50% increase in throughput. DOCSIS 3.1 also allows the split between upstream and downstream spectrum to be changed. This is important, as the lack of available upstream bandwidth is becoming a significant limiting factor to overall DOCSIS performance.

Network operators are exploiting the new capabilities offered by the latest DOCSIS standards to introduce increasingly high speed broadband Internet services. At the same time they are seeing more of their video traffic shift to IP delivery. Historically in a cable TV network video traffic has been separate from

IP traffic and has been carried over dedicated DVB traffic streams, which did not incur the over-head of IP transport. This is changing now and increasingly video on demand (VoD) is delivered over IP and is a mixture of "curated" premium content and web-based content such as YouTube. The advantage of moving to IP is that content can be viewed on devices other than set-top boxes, so for example an IP based VoD server could equally serve content to a high-end set-top box, a low cost IP TV adapter and an iPad.

As more traffic moves to IP delivery and is delivered on demand rather than as part of a broadcast TV service an increasing proportion of an operator's spectrum will be used for on-demand content (Internet and video). To support this, operators either need to convert more of the traditional broadcast spectrum to on-demand (reducing the number of broadcast TV channels) or reduce the segmentation size or more likely a combination of both.

There are other strong reasons why a CATV network operator will want to further reduce the node size. CATV networks are wide-bandwidth, shared analogue transmission networks. This makes them difficult to maintain and subject to ingress noise. By taking fibre deeper into the network:

- node sizes are reduced resulting in less ingress, and
- the number of active elements is reduced. This reduction in active elements improves network performance allowing more higher order modulation to be used.

The ideal architecture for performance reasons is fibre to the last amplifier, effectively turning the last amplifier into an optical node. In networks where a large bank of taps are installed next to this last amplifier to connect the customer drops, this is equivalent to fibre to the distribution point in a twisted pair access network. Where the taps are more distributed this is not quite equivalent to fibre to the distribution point.

With DOCSIS 3.1, using 500MHz of the available spectrum a data rate of 5Gbit/s can be achieved. This would allow a 1Gbit/s service to be provided, with the level of segmentation reduced over time as the average usage rate creeps up. From this it can be seen that both twisted pair and HFC networks will in the future be capable of supporting 1Gbit/s service to consumer by a series of incremental upgrade steps.

Fibre to the DP Architectures

As has been described in the previous sections both telco, twisted pair networks and CATV, coaxial networks are likely to evolve to fibre to the distribution point (FTTdp) architectures, or in the case of CATV networks something close to it. This makes sense as:

- The largest cost of a full upgrade to fibre to the home (FTTH) is replacing all of the customer drop cables.
- Network performance can be increased dramatically by shortening the length of the copper cable, and in the case of CATV networks dramatically reducing the segment size.

- Network operational costs (related to reliability) are reduced by: removing copper trunk cables, reducing the number of active elements and moving from shared media approaches closer to star topologies.

It is clear that for twisted pair networks the preferred technology will be the final standardised version of G.fast. In this author's opinion, what is less clear is the future of CATV network standards. Whilst CableLabs does have a plan for the evolution of HFC networks into passive optical networks, this author remains sceptical as to the likely success of this approach. Instead, in this paper the author sets out an alternative, but contentious, conjecture that both twisted pair and CATV networks will converge on the same technology. This is that:

- Cable TV networks will likely evolve more slowly to fibre to the DP architectures than twisted pair networks. This is because coaxial cable is higher performing than twisted pair meaning that from a pure bandwidth perspective there is less need to take fibre quite so deep. However, as has been indicated there are operational reasons why an operator would want to take fibre deeper.
- If one assumes that G.fast equipment is already deployed in large quantities in twisted pair networks then when a CATV operators comes to deploy a fibre to the distribution point architecture, why would they not consider using G.fast, but with the final connection to the home over coaxial cable rather than twisted pair. Given the large volumes of connections in existing twisted pair network this is likely to be a lower cost alternative than deploying a CATV specific solutions (from the point of view of both network equipment and customer premise equipment).

Therefore, in this author's opinion, a cable specific variant of G.fast is likely to emerge making use of the superior performance characteristics of coaxial cable, whereas as twisted pair specific variant might make use of the two pairs typically present on a drop cable. In this scenario one might imagine that a 2Gbit/s service can be provided over either a 2 pair twisted pair drop or coaxial drop cable.

The Final Shift to Fibre

As has been indicated, for the majority of users, copper connections to the home will remain the most cost effective approach to 2020 and beyond. Eventually, however, if access rates carry on increasing at the current rate then the highest speed access lines will exceed what can be delivered over copper. Nielsen's Law states that a high-end user's connection speed will increase by 50% per year. Note that this refers to the connection speed and not the average bit rate consumed by the user, which is normally much less. Nielsen's law is not a real law, but an empirical fit to historic trends. Whilst it is clear that access speeds will continue to increase over the next years, it is not immediately obvious how a connection of >1Gbit/s would be utilised by the user. Much as computers are doing it could well be that connections become good enough for the vast majority of normal users.

Nevertheless, if we assume operators continue delivering connections in line with Nielsen's law then they will need to introduce a fibre overlay of their FTTdp networks to serve their high-end customers with a native fibre to the home (FTTH) service. Clearly, if this is the eventual plan then this should be borne in mind when operators deploy their FTTdp solutions and sufficient fibre strands should be included to support the FTTH overlay. It is not yet clear at this stage in the development of the technology what the best optical architecture will be for these overlays. Given they will be for premium customers, GPON will not be adequate. Even 10G-PON might not be considered adequate for this premium tier, leaving the options of point-to-point fibre or a future PON technology such as wavelength based PONs.

Over time, the optical overlay will become the increasingly preferred method to deliver service and the copper FTTdp infrastructure will be withdrawn from service on a DP-by-DP basis as the customers naturally migrate to the higher speed connections. Given that the fibre overlay will eventually be the preferred connection approach for all customers, a fair amount of consideration will need to be given to the optical architecture to ensure that whilst it can cost effectively support a premium tier it will also scale to support 100% pentration in a very cost effective manner. This will be the challenge the future fibre access technologies should try and address.