Avoiding the bandwidth price dilemma using optical networking

David Payne, Russell Davey.
BTexact, Adastral Park, Martlesham Heath, Ipswich, Suffolk, UK.

Keywords:
PON, access, network, optical, fibre

Abstract
As operators of telecommunications try to move into a truly broadband future a major problem arises concerning the mismatch between the cost of providing bandwidth and the consequential growth in revenues. The fundamental problem is that the 80% learning curve of the electronic technology that networks are built upon is not fast enough to sustain the enormous growths that could arise. This paper analyses this problem and proposes a possible solution that tries to use the power of optical networking technology to eliminate tiers, nodes and port cards from the network as a way of massively reducing the unit costs of bandwidth.

Introduction
After more than two decades since it was first seriously proposed, fibre to the home/office (usually generically referred to as FTTH) is back on the agenda as an access technology for serious consideration. The candidate architecture most likely to succeed is the passive optical network (PON) [1] probably based on FSAN specifications initially [2}, although Ethernet solutions (EPON) may also play a role.

The problem for FTTH and indeed alternative fibre to the street node solutions that bring fibre closer to the customer has been the high capital investment required for any large-scale deployment [3]. This has led to deferment of adopting FTTH and FTTStreet node in favour of more incremental broadband solutions based on DSL technology that reuses as much as possible, the huge historical investment made in copper pairs. However the high capital costs of fibre solutions only exacerbates a more fundamental problem that faces operators and indeed the whole telecommunications industry when trying to move into a truly broadband future.

The bandwidth cost problem
The problem facing the industry is that the cost of increasing network capacity to meet the inevitable demand for bandwidth driven by broadband services can far exceed the subsequent growth in revenues! Bandwidth growth has exceeded revenue growth for years but has previously been balanced by the normal price decline that occurs in any industry as product volumes increase. The problem for the telecommunications industry as it enters a broadband future is that many of the possible service and growth scenarios, produce such large growths in bandwidth that traditional price decline alone will not be sufficient to keep growth costs in line with growth in revenue.

Traditional networks are built upon a large base of electronic equipment which only price declines in line with the usual price declines seen within the electronics industry. Historically electronic goods and products have declined with an 80% learning curve ie. for every doubling of product volume the unit price drops to 80% of the previous price. Over the past decade bandwidth has followed the same relationship, as network capacity doubles the costs of providing that capacity has fallen in line with the price decline of the supporting electronic technology. Importantly there is no evidence that the overall situation is changing!

Overall revenues from the customer base grow relatively slowly. Although for any particular service, revenues can grow quite fast, it is usually via revenue substitution from within the telecommunications business rather than truly new revenue entering the sector. On a macro-economic scale this is perfectly reasonable, GDPs of countries, and indeed the world, only grow at modest rates (typically ~1-3%). If a large sector of the economy such as telecommunications revenues grow at rates much greater than GDP over an extended period there will be significant consequential economic change ie. Increased spend in telecomms will mean a down turn in some other sector of the economy. Although such changes do occur they generally occur slowly. In the UK growth in telecomms spend has indeed exceeded growth in GDP for several years implying underlying changes in spending patterns, however expecting significant and sustainable changes in the spend ratios over extended periods of time is unrealistic.

Year 2006/07	Pragmatic Internet	Optimistic + Moderate Video	Very Optimistic & Video Centric
Broadband Internet customers	6 million	7 million	10 million
Video/Vod customers	0.15 million	5 million	7 million
Average Internet session time/day	70 mins	80 mins	100 mins
Average Internet Session Bandwidth	110 kb/s	500 kb/s	1100 kb/s
Average Video session time/day	20 mins	22 mins	90 mins
Average Video Session Bandwidth	2 Mb/s	3 Mb/s	5.3 Mb/s

Table 1 Assumptions used in modelling traffic scenarios.

The relative overall slow growth of revenues for telecomms compared to the potential for very rapid increase in bandwidth demand, implies that the cost per unit of bandwidth must decline very rapidly if returns on investment are to be maintained, much more rapidly for example than traditional price declines of products in the electronics industry.

This problem of the cost of bandwidth provision outstripping the increase in revenues is illustrated by the three service scenarios categorised as: Pragmatic Internet, Optimistic + Moderate Video and Very Optimistic & Video Centric. The main service and usage assumptions behind the scenarios are shown in Table1. Note these are scenarios not forecasts, the intent is merely to illustrate the impact that different service growth and usage demand patterns could have on network capacity requirements and subsequently the bandwidth price decline required to maintain return on investment assuming historical increases in revenues.

The corresponding growth in network traffic for these scenarios is shown in Figure 1.

It can be seen from the results in Figure 1 that there could be a huge range in the bandwidth demands that a future network may need to meet. The much bigger problem is that the revenue growth will not match the cost of servicing the bandwidth growth, particularly of the optimistic scenarios.

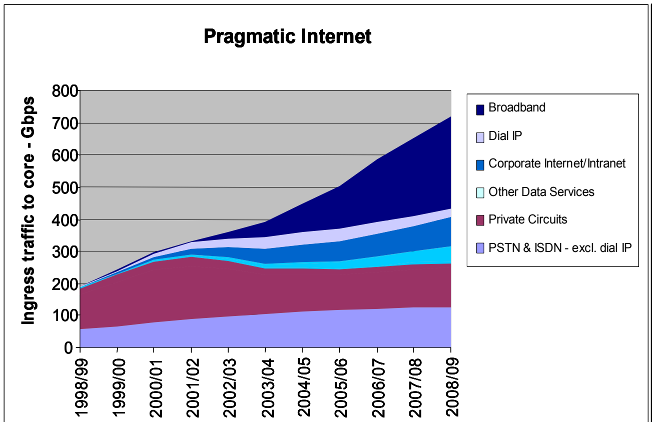

This scenario can be considered a pragmatic internet scenario. It only relies on ADSL and Cable Modem technology. Streamed video services are a minority and most video entertainment is delivered through conventional systems eg, terrestrial, satellite or cable. Internet growth is assumed to be fairly robust but not particularly ambitious. Customers are assumed to be split between cable modem operators and ADSL operators.

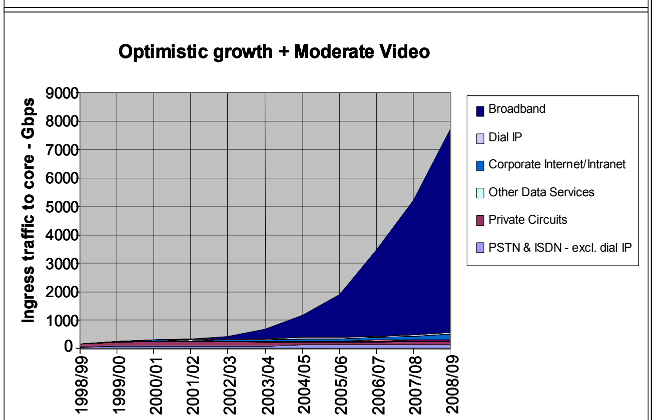

This scenario assumes a much greater use of content rich internet applications which pushes up average session bit rates. It implies that ADSL, where capable, will enable customers to operate up to ~2Mb/s and also some VDSL and/or fibre solutions are also rolled out. The other major change from the "pragmatic" scenario is that a large proportion of customers also use their broadband connection for some streamed video applications, although it is not the main entertainment video delivery system used by customers.

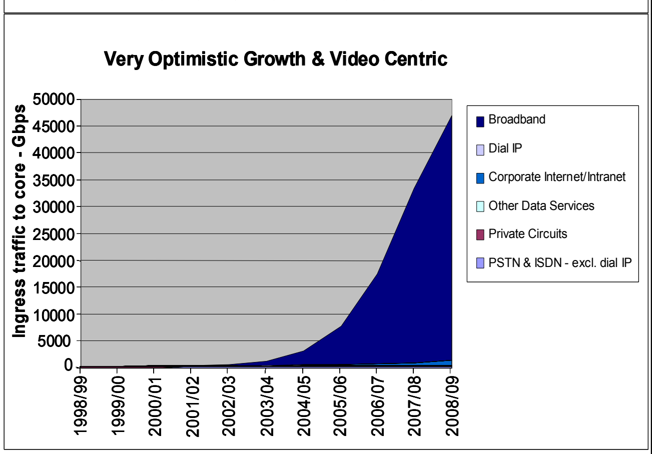

This is a much more optimistic scenario but without stretching the demands of plausibility of the types and usage of the services postulated. The main difference is the use the network to deliver high quality video entertainment services and fast file transfers for business and e-commerce. Even in this scenario the Video usage is only equivalent to about 1 feature film per day implying most entertainment video would still be delivered by conventional means. However network delivered video is assumed to be a personalised (eg. VoD) type service rather than broadcast/multi-cast and is transmitted to individual customers.

Figure 1. Projected traffic growth for the service scenarios listed in Table 1

Unless ways are found of significantly increasing the price decline of bandwidth, this problem will severely limit the growth that could occur in the future.

The following analysis illustrates this problem.
Ideally the industry would like to maintain the return on any capital expenditure (ROCE) required to expand network capacity during a period of sustained growth otherwise there is a risk of profitability suffering or even going out of business! The interesting question to ask of the scenarios modelled is: what price decline or learning curve for bandwidth must be achieved to enable the growths shown whilst maintaining ROCE?

If all growths are expressed as a compound annual growth rate (CAGR), there is a simple relationship linking revenue growth, bandwidth growth and the learning curve for the price decline per unit of bandwidth to maintain ROCE, given by:

$$1+G_R \geq (1+G_B)^{1+L}$$

Where: GR is the CAGR for Revenues
GB is the CAGR for Bandwidth
and L = Log(L%/100)/Log(2)
Where L% is the learning curve (expressed traditionally as a percentage).

A learning curve is defined as the percentage decline in price of a product as the product volume doubles, an 80% learning curve will mean that the price of a product at a volume V will decline to 80% of that price at volume 2xV. In this case the product volume is taken as bandwidth.

This is the macro-economic condition that needs to be met for a sustainable business in the light of unprecedented demand for bandwidth [3]! This relation ship is plotted in figure 2 with contours of constant revenue growth.

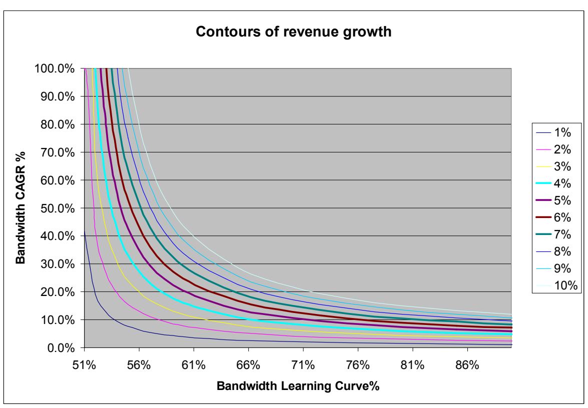

Figure 2 Relationship between revenue & bandwidth growths and the price learning curve for unit bandwidth

Typically it is reasonable to expect time averaged revenue growths in the range 4% to 7% CAGR. Given that electronic systems have learning curves of ~80% (historically transmission systems have been following such a learning curve for the past decade at least!) the curves in figure 2 would suggest that bandwidth growths of ~10% can be sustained. Note this is a macro-economic argument applied to the total system size it does not apply to individual products and services or sub-networks within the main network. These could grow at much faster rates by substitution of other products or markets or they could be small niches and simply cause too small a perturbation to be noticeable in the bigger picture. In the scenarios illustrated above we are looking at large scale changes to the capacity of the network and though simple, this model gives a useful guide to the bounds on growths that can be expected to be viable, given a set of revenue constraints.

The envelopes of bandwidth growth for the scenarios illustrated in figure 1 do not exhibit constant CAGRs due to the many individual parameters and growth and service usage functions that are in model used. However if an approximation is fitted to the regions which start where bandwidth from broadband begins to "take off" (from ~2002/3), then approximate CAGRs can be obtained. It should be noted that in the very optimistic scenario, saturation of some parameters were beginning to occur in the latter years in the model and the classical "S" curve is beginning to manifest itself, to get a good fit to CAGR the last year was omitted. Using these fitted bandwidth growth rates, assuming a nominal 5% revenue growth and the relationships shown in figure 2, then the following estimates of the required learning curves for bandwidth price decline can be estimated:

Scenario	CAGR fit	Bandwidth learning Curve
Pragmatic Internet	12%	~70%
Optimistic & Moderate Video	62%	~54%
Very Optimistic & Video Centric	128%	~52%

Table 2 Learning curves to meet given network bandwidth growths

The "Pragmatic Internet" scenario is on the bounds of possibility with electronic centric systems and network architectures not very different from today's. So the current strategies followed by operators deploying ADSL or Cable Modem technology should just be viable and will probably maintain ROCE particularly if some additional revenue growth can be obtained from the richer service sets that can be offered.

However, the case is very different as we go beyond this scenario to the higher bandwidth scenarios that assume greater than a few hundred kilobits per second session rates and greater use of personalised video services. It can be seen that the price declines required to meet the bandwidth growths in these scenarios, are much faster than the traditional 80% learning curves typically arising in the electronics industry. In the case of the optimistic and video centric scenarios the required price declines are so fast that it can be safely claimed they will not be met, no technology has ever price declined at this rate! From this perspective it can concluded that future network architectures that continue to use traditional electronic solutions will not be able to price decline sufficiently fast to be able to maintain operating margins and profitability. This will apply regardless of any protocol changes such as converging onto an IP/MPLS platform!

The follow on question for optical technology and optical networking is, can it produce faster unit bandwidth price declines than traditional network developments? Considered from a purely technology perspective there is no reason to suppose that the price of optical technology itself will decline any faster than electronic technology. Many of the manufacturing techniques used within the optical communications technology industry have evolved from or are developments of processes used within the electronics industry.

interfaces required. This network proposition is illustrated in figure 3.

The proposed network is based on the passive optical network (PON) principles but with optical amplification to boost the power budget increasing bandwidth, geographical range and optical split compared to a conventional PON. It can be pedantically argued that this is no longer a *passive* optical network because of the use of an "active" optical amplifier within the PON. However the original access proportion of the network remains passive and the term PON will continue to be used in this paper (for the purist the P could become Photonic).

The long reach enabled by the amplifiers, is used to bypass the usual backhaul network which would conventionally be SDH combined with Metro WDM systems. The use of amplifiers also allows the possible by-pass and ultimate removal of the local exchange or remote concentrator site. The increased capacity of the system coupled with the

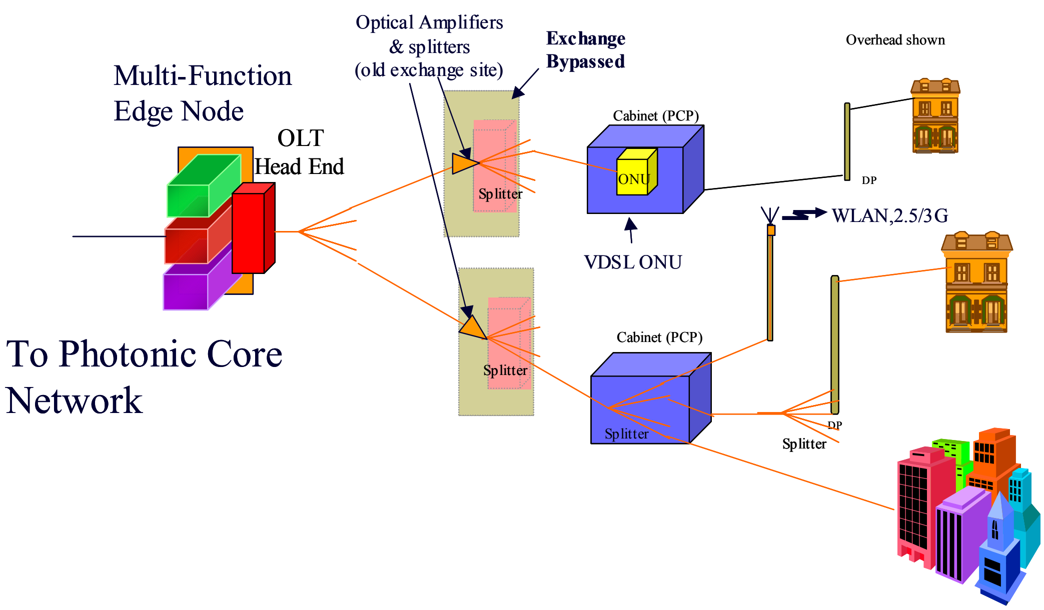

Figure 3 Integrated optical access and Backhaul (Metro) network

increased power budget also enables a much greater optical split to be deployed, allowing more customers to be able to share the capacity and costs of the PON and the OLT located in the core edge node.

Opportunities using optical networking

The opportunity offered by optical technology is the prospect of new radical approaches to network architectures aimed at exploiting the capacity, reach and flexibility of optical networking to minimise the numbers of tiers, nodes and interface ports required in a large scale network. That is network architectures that enable the removal of significant proportions of the electronics from within the network!

A key to this is the design of a network that combines the access network and outer-core/backhaul/metro network into one common high capacity network that can bypass and eventually eliminate the current access nodes [4]. At the same time it minimises the number of customer ports and network

This combined access and backhaul network terminates on core edge nodes that marshal and groom traffic onto a photonic inner core network that interconnects these core edge nodes via wavelength channels. For the UK it is envisage that there would be about 100 of these core edge nodes. This could lead to a highly simplified network that has the potential for significantly reducing unit costs and enabling scenarios with very high bandwidth growth. Such a future network is illustrated in Figure 4. The long reach access network operates at 2.5 or 10Gb/s and could serve ~500 (possibly 1000) customer sites on each amplified PON [5]. The traffic is terminated onto a multi-function node with functionality for IP, ATM and TDM traffic and services. Note this does not imply that this functionality is all in one box, although this is one option, but that the functionality is co-located at the same geographical node. It could also be terminated on an all IP/MPLS node if that was the desired solution. However the optical solution is relatively agnostic to the higher layer functionality.

Traffic from the long reach access networks terminating on the core edge node is marshalled into transport containers (eg. this could be concatenated VC4-n containers over SDH or any other adequate transport protocol) and placed onto a network would require protection paths and mechanisms because it now provides the outer core network connectivity. These protection paths could be up to twice as long as the primary paths. The deep reach access network effectively by-passes the metro or outer-core collector network and can significantly reduce the cost of back haul from the access node to the first core node.

Long Reach Access

A long reach access network such as an amplified PON can provide a very flexible and highly functional traffic

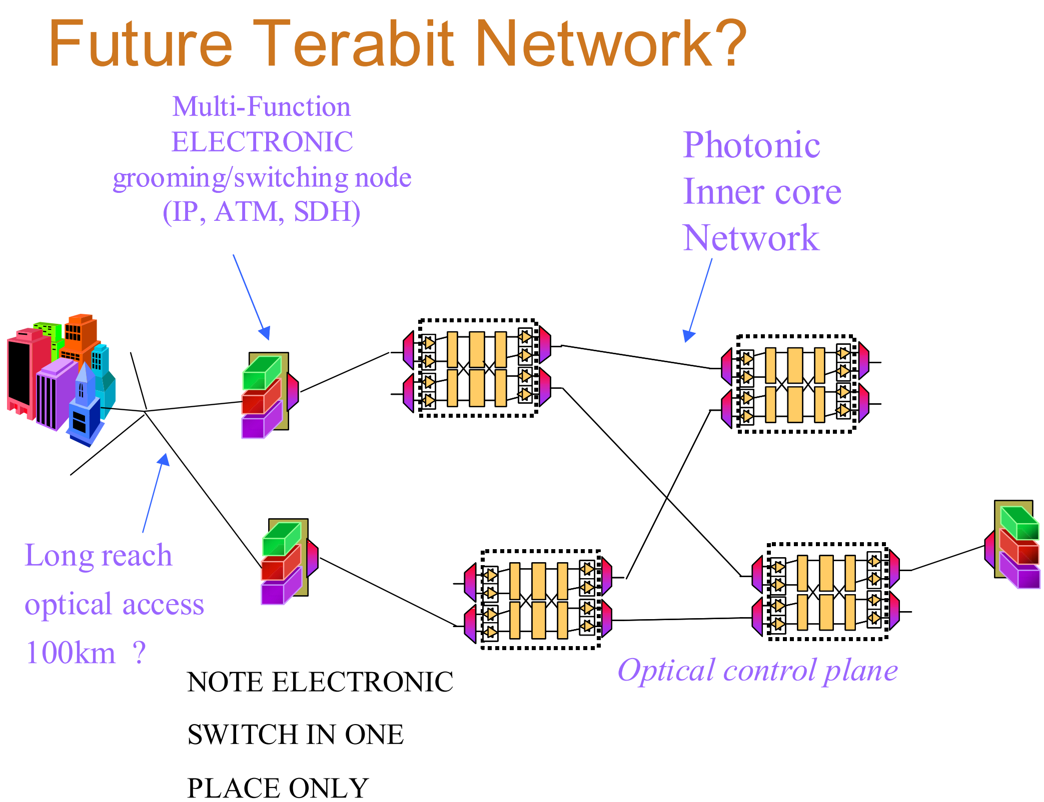

Figure 4 An End to End network vision using Long Reach Access with photonic core networks

wavelength destined for another core edge node. The only circuit manipulation and routing that takes place in the core is at wavelength granularity. This is to avoid optical to electronic conversions followed by electronic processing and then electronic to optical conversion for onward transmission. This function is now only performed, as far as possible, in the edge nodes. Note that in practice many of the optical nodes in the core will be geographically co-located with an electronic edge node but logically the architecture is as shown. The optical core may also have an optical control plane to facilitate set-up and tear-down of connections across the core.

This architecture would require a PON standard to be developed with higher split, longer reach and much higher capacity than the current generation of passive optical networks. The 100km reach is required because such a

engineering capability. With the right functionality built into the transport protocol, bandwidth can be allocated to customers on demand, broadcast and multi-cast services can be offered with minimum bandwidth and switching functionality required in the core. It can act as a distributed concentrator and service switch and simplify the switch and routing functions required in the edge node.

Sophisticated monitoring and diagnostic tools could also be implemented that could give early warnings of problems before they become traffic affecting and also provide terminal performance, identification and location diagnostics.

Security is always an issue with shared access media and needs to be addressed carefully to ensure adequate security both from eavesdropping and malicious damage to the network integrity perspective. There are many approaches that have been considered during the development of passive optical networks and the issue is more of choosing the most cost effective solutions rather than a lack of potential solutions.

Other advantages that the deep reach access architecture can offer are:

Symmetrical capacity. Although asymmetrical capability could easily be offered, large capacity PONs that serve all customer types will probably need to be symmetrical. The proportion of capacity allocated for broadcast services can be a relatively modest fraction of the total capacity and the bulk of the capacity may well need to have greater symmetry for future service requirements.

It can serve all customer types and services. This could include residential customers, multi-dwelling units, SMEs and even large business sites. Larger business sites that are already served with point-to-point fibre systems may continue to use the existing fibre and topologies, although the long reach access technology could supplement or replace these systems if greater service capability was required or local exchange by-pass was implemented.

The combined access and outer-core network that long reach access architectures span, could provide guaranteed QoS parameters by suitable choice of transport protocol and control system. Real time streamed services and PSTN quality voice (with all the associated delay constraints) together with common packet protocols such as Ethernet and IP will need to be supported.

The high capacity of the system coupled with traffic management and dynamic bandwidth assignment produces the lowest cost per unit of bandwidth per customer. As discussed above this will be essential if the potential bandwidth demands of the future are to be met economically.

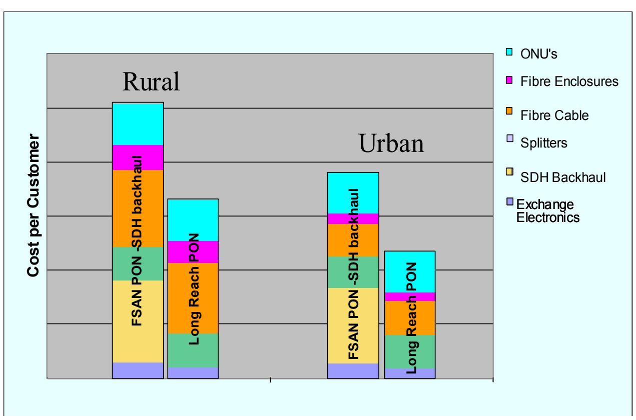

Figure 5 Relative costs of the long reach access system and conventional PON plus metro

The major saving associated with this architecture is the large reduction in backhaul costs that become possible. The costs for the access portion remain similar to conventional PONs because the structure and deployment of fibre and the associated infrastructure are very similar in both cases. The saving in backhaul costs arise because the Metro WDM or SDH/SONET equipment is by passed with a simple node consisting of optical power splitters and a few optical amplifiers. This is illustrated in Figure 5 which shows the results from a cost model comparing the long reach access network against a more conventional FSAN based PON with an SDH backhaul or metro network. The costs of the components were based on currently available prices which where then price declined for volume roll out using nominal 80% learning curves. The low cost of the backhaul/metro part of the long reach access network raises the interesting possibility that this solution might also be viable in the more rural and low density serving areas.

Summary

The scenario modelling described above shows that the future growth and cost of network capacity, driven by a truly broadband future with true triple play services, could easily out strip any consequential net revenue growth into the telecommunications sector. The evidence for this conclusion is based on the historical price decline of network bandwidth and the historical growth in revenues when applied to possible future service scenarios. It is also concluded that the price decline of electronics and optical technology alone will not solve this problem for the industry as it moves forwards into the broadband future and implies that either the growth cannot happen or new radical approaches to the network architecture must be found.

The opportunity offered by optical technology is to enable new architectures that significantly reduce the number of tiers, nodes and port or interface cards required within the network. This use of optical technology can radically reduce the amount of equipment required per unit of bandwidth and this can produce bandwidth price declines much greater than that from just historical equipment price declines alone.

The architecture being developed, to exploit this advantage from optical technology, combines the access and backhaul or metro network into one integrated optical network. This by-passes the conventional SDH/Metro-WDM backhaul networks and also enables access exchange/RCU site by-pass in the future. This integrated access and back haul network is terminated on a relatively small number of core edge nodes (~100) which in turn are interconnected via a photonic core network operating at wavelength granularity.

This “long reach access” with “photonic core” architecture offers the following advantages:

- Large reduction in Backhaul costs.
- Could also be viable for rural geotypes.
- Minimum costs per unit of bandwidth per customer.

- Provides traffic multiplexing, concentration and grooming function.
- Provides flexible bandwidth assignment.
- Could provide symmetrical capacity.
- Could provide guaranteed QoS.
- Can serve all customer types and services.
- Can deliver true triple play and beyond!

The issue is; can the significant level of capital investment required to realise this future architecture generate a sufficiently fast return from the expected operational cost savings and revenue up-lift opportunities it can generate? Or do we face a future of only moderate broadband growth and the threat of declining revenues?

References

1. J R Stern, J W Balance, D W Faulkner, S Hornung, D B Payne. “Passive optical networks for telephony applications and beyond”, Electronics Letters 19th November 1987, Vol. 23 No. 24 pp 1255-1257.
2. http://www.fsanweb.org/default.asp
3. D B Payne, R P Davey. “The future of fibre access systems”, BT Technology Journal Vol.20 No.4 October 2002
4. D P Shea, A Ellis, D B Payne, R P Davey. “10Gb/s PON with 100km Reach and x1024 Split”, Proceedings ECOC 2003, Rimini.
5. D P Shea, RP Davey, A Lord, J E Mitchell. “Design aspects of long reach optical access networks”, Proceedings ISSLS 2004, Edinburgh.