1. Introduction

1.1 Distinguishing between Core and Access Networks

In any discussion of broadband networks, it is useful to distinguish between core and access networks. Core networks comprise nodes that are interconnected by high capacity communication links, often in a "mesh" structure so that the failure of any one link does not isolate a node. The nodes may be exchanges or other points where the traffic from multiple of users is aggregated.

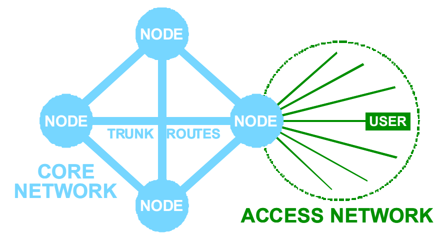

In contrast, access networks refer to the infrastructure that is used to connect individual users to the nearest nodal point in a core network. Links in an access network may be thought of as point-to-point connections between the user and the node – though in practice the traffic from multiple users is often aggregated onto a shared medium approaching the node.

1.2 Core Networks in Australia

In general, there has been a healthy level of investment in core networks in Australia, and there is typically an abundance of high-speed capacity between the major population centres. There has also been considerable investment in network capacity in central business district (CBD) areas, attracted by the comparatively high communication spend of larger corporate and government customers.

Notwithstanding this, at any point in time, there always will be scattered bottlenecks in terms of trunk capacity – where either the infrastructure isn't adequate, or where the absence of competition and monopoly pricing is constraining usage. However, the cost of increasing capacity on trunk routes is comparatively low compared to the cost of upgrading access networks. Where centres are connected by optical fibre, it can be as simple as upgrading the equipment at each end of the fibre.

Changing patterns of usage will put more pressure on core networks (including international links) in the future. In particular, as speed limits are lifted in access networks and the volumes of information transmitted grow (for example, with the increasing exchange of very large video files) it is predictable that more links will become congested and need upgrading. One approach to managing this is to introduce differentiated levels of service, so that users wanting to be assured of consistent, high performance can do so by paying a premium. This is exactly parallel to what has happened with road networks – users wanting faster travel can opt to use toll roads that avoid most of the congestion that occurs on the public road network.

1.3 Access Networks

In contrast to the situation with core networks, investment in access network infrastructure outside of CBD areas has been very patchy since de-regulation. The most significant developments include:

- The Optus hybrid fibre coax (HFC) network, passing around 2.2m homes, mainly in Sydney, Melbourne and Brisbane;
- The Telstra HFC network, passing around 2.5m homes mainly in Sydney, Melbourne and Brisbane; because of the substantial overlap in the "footprints" of the Telstra and Optus HFC rollouts, the total number of Australian homes that has access to one of both of the Telstra and Optus HFC networks is around 2.7m homes only;
- TransACT's fibre-to-the-kerb (FTTK) network, passing some 65,000 homes in Canberra and Neighborhood Cable's HFC network, passing some 90,000 homes in Ballarat, Mildura and Geelong (recently acquired by TransACT);

- A growing number of greenfield estates that are being developed with advanced fibre-to-the-home (FTTH) broadband cabling; whilst these are proliferating, the locations are scattered and (given the pace of land development) it will take some years before any sort of critical mass is achieved.

- A number of wireless broadband networks, notably those of iBurst and Unwired, but including a number of smaller networks typically operated by ISPs.

One other important development in access networks since 2000 that is worth highlighting has been the deployment of ADSL (or its ADSL2 and 2+ variants) in Telstra exchanges following declaration of the unconditioned local loop (ULL) in 1999. This has not (yet) involved any significant upgrades to the national cabling infrastructure – only the installation of different technology to drive the copper in Telstra's network. However, it is particularly noteworthy since for many Australians broadband is synonymous with ADSL.

Today many Australians are already operating at or near the limits of performance that can be achieved over existing infrastructure. Higher performance cannot be unlocked simply by deploying more modern electronics – it requires fundamental re-engineering of network cabling and/or equipment sites to achieve meaningful advances.

2. Access Network Technologies

2.1 Optical Fibre

Optical fibre is the highest-capacity known communications medium. Information is transmitted over optical fibre in the form of light pulses. The ability to simultaneously carry multiple different wavelengths (light colours) gives optical fibre immense capacity – vastly more than is used (or even envisaged) in any end-user application today. Laboratory tests have demonstrated that speeds in excess of 10 Tbps[1] (that is, 1,000,000 Mbps) are possible.

Optical fibre can be driven at virtually whatever performance level and in whatever manner is required by the user. Where the capacity requirements are not high, low-cost commodity equipment can be used. However, by placing more sophisticated (and expensive) equipment at each end of the fibre, the same optical fibre could be driven at performance levels literally thousands of times higher.

Because of its high capacity, optical fibre is the preferred medium for all links in a core network – though other media are occasionally used (typically for short links only) where for whatever reason it is not practical to deploy optical fibre.

Optical fibre is also the ultimate form of connection in an access network. It is widely used in central business districts and for connections to corporate users where one corporate user can aggregate the traffic of many individuals. However, today it is still comparatively rare in residential access networks. This is due in part to the fact that today's residential access networks were deployed when the land was first subdivided and sold – well before the development of optical fibre technologies – and in part to the fact that legacy infrastructure has been able to deliver *adequate* capacity for most residential applications in common use today.

---

[1] ***Network speeds*** are measured in bits-per-second (bps), with the most common measures being a kilo-bit-per-second (abbreviated to kbps, or 1024 bps), a mega-bit-per-second (Mbps, operating at 1024 kbps) a giga-bit-per-second (Gbps, operating at 1024 Mbps) and a tera-bit-per-second (Tbps, operating at 1024 Gbps).

Data volumes are most commonly measured in bytes (or characters) each made up of 8 binary digits (bits, or '0's and '1's), with the most common measures being kilo-bytes (abbreviated to kB, comprising 1024 bytes). Other common measures are the mega-byte (a MB, comprising 1024 kBs) and the giga-byte (a GB, comprising 1024 MBs).

For use in access networks, optical fibre may be deployed in a point-to-point (P2P – not to be confused with "peer to peer") topology, or in a point-to-multipoint (P2MP) topology – often also described as a passive optic network (PON). With P2MP solutions, the fibre extending towards the user branches out so that (typically) 32 or (up to) 64 users receive the same signal. The "branch point" involves no electronics – simply the fusion of fibres. Light sent downstream reaches all end points, and light from any end point reaches the upstream transmission point. However, all information flows are managed so that each user perceives a service that is totally independent of any other users sharing the same fibre.

P2P topologies are the more powerful, since the service to any user is independent of services to all other users and their fibre can be driven at whatever speed is required. However, P2P topologies also carry a price premium, and as a result, P2MP topologies are the more commonly used in residential settings.

Property developers who have recognised the value that good communications infrastructure adds to their land are starting to make arrangements for fibre-to-the-home (FTTH) infrastructure deployed in place of the traditional copper infrastructure. Typically this involves paying a premium in the range from $1500 to $3500 per lot. Although there are now greenfield FTTH projects in every state of Australia, the slow rate at which land is sold and developed in these estates means that it will be quite some time before there is a substantial base of FTTH-connected homes in Australia. Nevertheless, the growing body of such homes will tend to eventually define a new benchmark for broadband performance in Australia.

To date, there have been no significant attempts to replace legacy infrastructure with FTTH in established areas in Australia. In Tasmania, the TasCOLT pilot has seen FTTH deployed past around 1,000 homes in Hobart and Devonport. Whilst there are some overseas models for FTTH deployment in established areas developing, it is difficult to see anything on the horizon in Australia that promises to see FTTH capacity available to large numbers of Australians any time soon.

2.2 Other Technologies

With all other technologies, optical fibre is typically used down to a certain "depth" in the network, and then the final connection to the users is made over some other medium. There are four alternative media commonly used throughout the world, none of which offers the same capacity as optical fibre. As a result, the performance of the end-to-end access link is constrained by the medium used between the end of the fibre and the user. Each of these media and the associated technologies is described in the following sections.

2.2.1. Copper and xDSL Technologies

Where copper is used for the final link to the user, the technology used to drive broadband performance is known generically as digital subscriber line (DSL) technology – often denoted xDSL where "x" can take on a range of values. The most common DSL technology is ADSL and its derivatives, ADSL2 and ADSL2+.

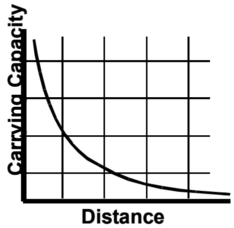

With all xDSL technology, the performance that can be achieved is a function of the length, gauge and condition of the copper. ADSL (asymmetric DSL) – one of the first xDSL variants to gain widespread recognition as a broadband technology – is optimised for the typical copper distances of up to several kilometres that incumbent telephone companies (telcos) have traditionally used to reach customers connected to their exchanges. The performance limit of ADSL is 8 Mbps (downstream) – but this is only achievable at distances below about 1.5 kilometres. DSL variants that are designed for higher performance at shorter loop length include ADSL2+ (an evolution of ADSL that uses more efficient encoding and additional spectrum to achieve downstream speeds of up to 24 Mbps) and VDSL (where V stands for "very-high-speed" and speeds of up to 100 Mbps symmetrical can be achieved).

Some other noteworthy members of the DSL family are:

- HDSL (where H stands for "high bit rate") – uses two good-quality copper pairs to achieve symmetric speeds of 1.5 Mbps (equivalent to an ANSI "T1" circuit) or 2 Mbps (equivalent to an ETSI "E1" circuit) at distances up to around 3 kms. HDSL has largely been eclipsed by SDSL.

- SDSL (where S stands for "symmetric") – uses a single copper pair to achieve symmetric speeds up to 2.3 Mbps at distances up to around 3 kms.

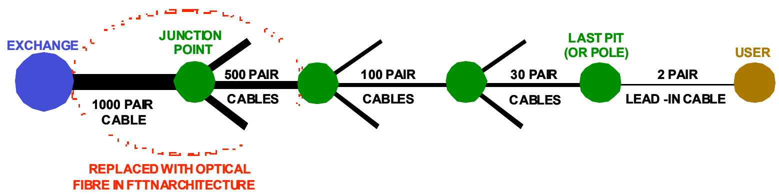

Copper networks are often described as having a "tree and branch" topology. Although each customer is normally connected by a dedicated pair of copper wires, those pairs begin at the exchange in cables that may contain 500 or more pairs. Such cables are cross-connected to multiple smaller cables at each of potentially several junction points (see picture of typical junction point on right) as the circuit gets closer to the user. The final cable that connects a small serving area often contains only 30 pairs or less.

In this sort of network, the cost of replacing cables goes up the closer one gets to the end user. However, by replacing the "upper reaches" of copper with optical fibre, telcos can effectively shorten the length of copper by which a customer is connected and thereby achieve higher performance levels. xDSL equipment is still needed at each end of the copper loop – but rather than the "upstream" end of the loop being located in an exchange, it is located in an outdoor enclosure.

The most common upgrade reflecting this approach is from a fibre-to-the-exchange (FTTX) architecture to a fibre-to-the-node (FTTN) approach – where the node would normally be a kerbside equipment cabinet (such as pictured to the left) located less than a kilometre from the user. At this distance, a node may typically serve around 200-500 residential users. When copper distances are reduced in this manner, ADSL2+ technology can be used to deliver consistently higher performance than would is possible in a FTTX architecture.

An approach which takes optical fibre even deeper into the network is fibre-to-the-kerb (FTTK) – where the nodes are typically located 300m or less from the users. At this distance, the latest VDSL technology is capable of speeds up to 100 Mbps symmetric. This is the architecture that TransACT adopted nearly a decade ago (see picture of TransACT pole-mounted node to the right). The equipment currently driving TransACT's copper pre-dated modern VDSL standards and is limited to 51 Mbps downstream and 1.6 Mbps upstream.

The vast majority of Australians are connected to Telstra's copper network, and many of these currently rely on ADSL (or one of its variants) for their broadband service. The lowest "entry level" ADSL speed that is offered is 256 kbps (downstream) and 64 kbps (upstream) – usually denoted as a 256/64 kbps service. However, many ISPs offer higher performance options such as 512/128 kbps, 512 kbps symmetric, 1024/256 kbps and so on.

The maximum speed that is attainable on a line is a function not only of the line length, but also the gauge (thickness) and condition of the copper. An indication of what is possible can be obtained from databases that record the route and length of the line – but the performance potential can really only be definitively known when a service is activated by placing the modem technology on each end of a line. Even then, performance may vary in some conditions (for example, in wet weather when moisture ingress into cables degrades performance or when many DSL services in the same cable lead to interference).

ADSL2+ is already offered by some Internet Service Providers (ISPs), but users will only experience better-than-ADSL performance levels if they are lucky enough to live within about 1.5 kms copper distance from the nearest exchange.

Some network operators artificially restrict the speeds offered over their lines, usually as an incentive to encourage users wanting higher performance to migrate to a more expensive access plan. However, all modern ADSL technology is adaptive and will automatically settle at the highest speeds that can be reliably maintained over a given copper line.

A FTTN-like architecture is already in use in Telstra's network – and has been the norm in new housing estates for over a decade. However, in addition to an optical fibre connection between the node and the exchange, there is commonly also a copper cable. Customers who want an ADSL service are connected via a copper circuit all the way to the exchange.

As soon as any customers are connected by copper through the node all the way back to the exchange, it is no longer possible to support the higher-performance ADSL standards that could otherwise be offered from the node. This is because an ADSL signal sent from the exchange suffers loss in strength (attenuation) over distance and is weakened by the time it gets to the node. If full-strength signals were injected at the node onto copper pairs in the same cable, they would "swamp" the weakened signals coming from the exchange (see diagram following).

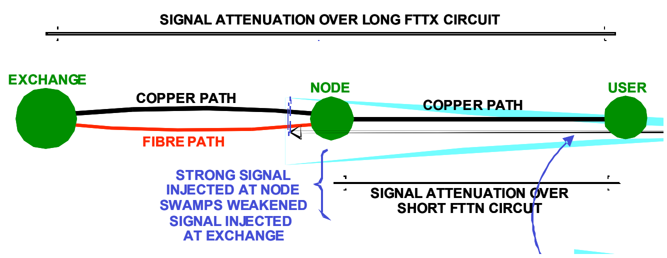

Accordingly converting a traditional FTTX network to a FTTN network has to be an "all or nothing" approach if the benefits of shortened copper loop lengths are to be unlocked.

This consideration is at the heart of competition concerns surrounding discussions about converting Telstra's network to a FTTN architecture. Many third party ISPs have deployed their equipment (DSL Acess Multiplexors or DSLAMs for short) in Telstra exchanges and rent copper loops from Telstra at regulated, wholesale prices. The resultant competition has been very beneficial for growth of the broadband market in Australia. However, if the copper network were re-engineered so that loops stopped at the node, DSLAMs at the exchange (whether owned by Telstra or by third parties) would effectively become "stranded" assets.

The extent to which this presents problems will probably be influenced by the age of equipment. Whilst cabling is a “long-term” asset, the life of the technology used over that cabling is much shorter and invariably is upgraded or replaced periodically. For equipment such as DSLAMS, an operational life less than 5 years would in all likelihood undermine the business assumptions on which its deployment was based. Conversely, any equipment approach the 10-year life would typically be regarded as operating “on borrowed time”. A further consideration is that awareness that FTTN is “coming” is growing, and as plans become clearer, operators will redirect abstain from further investment in areas that have a uncertain or limited future.

Telstra has suggested that it is willing to make the investments needed to upgrade its network to a FTTN architecture, but only if it is *not* compelled to offer capacity to third parties at prices it does not consider to be realistic. The G9 consortium has proposed an alternative of “neutral” ownership of the upper (fibre) reaches of the network, with all parties being allowed access on an equal basis. Telstra’s response has been to threaten the “mother of all legal battles” since the G9 proposal would impinge on its freedom to operate its own copper network in whatever way it sees fit.

The Government has declared its intent to support an upgrade of Telstra’s network to a FTTN architecture, but the specific plans for doing this have yet to be published. It is our view that the most likely form this will ultimately take is for Telstra to be given responsibility for the upgrade, subject to some satisfactory arrangements to ensure that the upgrade does not “wind the clock back” on the competition that has driven innovation in pricing and services over the past decade.

The process of negotiating satisfactory arrangements is likely to take time, and may well frustrate the Government’s target timetable. However, FTTN remains the easiest way to deliver a meaningful performance increase to the majority of Australians.

2.2.2. Coaxial Cable and HFC Technologies

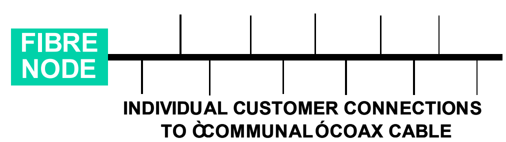

The origins of HFC technology go back to early days of television when it was common practice to install a master (or community) antenna in a prime location selected for good reception of the broadcast signal. Coax cabling was used to reticulate the signal from such an antenna to individual users, giving rise to the “CATV” (Community Antenna Television) acronym. It is ideal for broadcast television because all users share a common signal – in contrast to the architecture of a copper network where each user is connected on a dedicated pair of copper wires.

As CATV networks spread and connected more and more users, amplifiers were introduced to periodically boost the radio frequency (RF) signals that carried television services over longer runs. Unfortunately, the coax cable acts as a good antenna for collecting noise from a wide range of sources – and amplifiers boost noise levels as well as the signal. Accordingly there are limits to how many times a signal can be amplified and still remain usable.

To overcome this, optical fibre (with its low loss and strong immunity to RF noise) was introduced to carry a high quality optical signal through the upper reaches of the network to a node (see example in picture to the left) where it was converted from an optical to an electrical signal and thence reticulated over coax cabling. The resultant HFC networks have become the standard for Cable TV operators.

With the rise of public interest in Internet access from the mid-90's, HFC operators seized the opportunity to re-purpose some of the frequencies previously allocated for television distribution to carry data instead. Since everybody on a HFC network receives the same signal, a special type of modem was needed to filter out individual information from the communal flow. The Data Over Cable System Interface Specification (DOCSIS) standard for cable modems emerged to address this requirement. This standard has now evolved over several generations to underpin broadband data services on HFC networks throughout the world.

The shared nature of the final connection to the user (and hence the bandwidth available on that medium) gives rise to one of the key differentiating characteristics of a HFC network. When there are only a few active users on a coax segment, performance can be "sparkling". However, in busy periods, contention for the available bandwidth rises and performance degrades.

The extent to which this is a problem is determined by the depth to which optical fibre is taken in a HFC network, and the resultant level of sharing. Traditional coax networks were commonly engineered with between 1,000 and 5,000 customers on a common coax segment – and at such high levels of sharing, broadband performance would typically be relatively limited and very variable. More modern networks take fibre deeper, resulting in smaller serving areas and much lower levels of sharing. In the first stages of deregulation, Optus elected to deploy a HFC network as its main means of accessing residential customers. The choice of HFC technology was undoubtedly influenced by the opportunity that Optus and its partners perceived to lead the introduction of Pay TV to Australia. However, in contrast to the 60%+ uptake levels in the USA and Europe, Australians showed only lukewarm interest in paying for subscription TV services and it took a long time before uptake reached double-digit figures. To make matters worse, Telstra embarked upon a rival HFC rollout to broadly the same areas of the market.

Optus passed some 2.2m homes and Telstra some 2.5m homes. However, because of the high level of "overlap", the total number of Australian homes with access to one or both of the Telstra/Optus HFC networks is only 2.7m. With both companies committed to minimum subscriber number guarantees well beyond what was realistic, both lost heavily and by about 1997 all HFC deployments had ceased. Since that time, the only significant rollouts of HFC infrastructure have been by Neighborhood Cable (90,000 homes in Mildura, Ballarat and Geelong) and for TV distribution in a number of greenfield housing estates (notably in Western Australia).

The long-term future of Australia's HFC networks is unclear. Both the Optus and Telstra networks are engineered primarily for TV services, and to keep pace with rising demand for broadband capacity will ultimately require upgrading the networks by pushing fibre deeper in order to reduce cell sizes. With cell sizes of just a few tens of customers, the performance potential is excellent.

If Telstra's copper network is upgraded to a FTTN architecture with the capacity to deliver video services, the value of a separate HFC network will decline over time. Indeed, the network may ultimately become more of an operational liability than an asset – especially if confronted with the need for further investment in order to keep the network viable.

In contrast, one might expect that Optus would be interested in upgrading its HFC network in order to maintain access to some 2.2m residences without dependence on using any Telstra network assets. However, in the wake of the unfortunate history of its original HFC investment, there are no obvious signs that the Company has an appetite for further investment in the HFC network.

2.2.3. Electricity Wires and BPL Technologies

Broadband over Power Lines (BPL) is a comparatively new technology that has generated a lot of interest. Modem chipsets are used to send/receive a broadband signal over ordinary electrical wiring.

The technology is used in two comparatively different ways. As a form of *in-building* communication, end-users able to purchase modems that essentially bridge an Ethernet service from one power outlet to another. More sophisticated configurations can be used to reticulate a broadband signal throughout larger buildings – such as apartment blocks.

BPL can also be used as an *access* technology – using electricity supply wiring as the means for getting into the home. It is in this context that it has generated most interest – the ability to access customers without the massive costs involved in deploying new wiring generates has natural appeal, particularly amongst utilities who are interested in new revenue opportunities outside of the regulated returns in their traditional businesses.

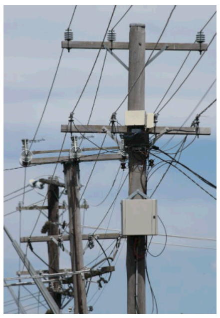

From an architectural perspective, BPL is fundamentally similar to HFC in that a shared signal is reticulated to a community of users – in this case, over power lines rather than a coax cable segment. The signal "injection" point (see example on power pole in diagram on left) is most commonly located at the low voltage (LV) transformer, since transformers naturally inhibit the propagation of signals upstream and into adjacent "cells". For optimum performance, optical fibre would be used to carry the signal to the injection point – and thus BPL could also be described as a fibre-to-the-transformer (FTTT) architecture. However, other means of signal delivery (including BPL over the medium voltage feeder lines) are available if the performance objectives are not high.

The number of customers connected to a LV transformer varies from one or two (in rural areas) to as many as 500 (in urban areas). More typically, the average in urban areas is between 50 and 100 residential customers. The level of signal sharing depends on how many users in an LV cell subscribe to the service.

The capacity of the signal is the other key parameter that determines the performance that can be given to individual users. At this stage, BPL is not yet a standards-based technology. Whilst efforts are underway through the IEEE, it is likely to be some time before standards and interoperability become a reality. In the interim, there are three vendor groups vying to have their technology embraced as the foundation for standards. The company currently offering the highest performance is the Spanish supplier, DS2.

The DS2 chipset – heavily favoured in Australia and Europe – is often cited as a 200 Mbps chipset. However, this is a measure of raw performance at the physical layer rather that the throughput that can be achieved. In optimum conditions, a DS2 chipset can deliver about 85 Mbps throughput – but in real-world conditions, the signal commonly needs to be regenerated[2] to traverse the distances that are common in a 240V electricity network. Simplistically, the amount of capacity that is available to the community in a LV cell can be divided by n+1 where n is the number of times the signal needs to be regenerated. For most practical purposes, therefore, it is more realistic to expect capacity of 30-50 Mbps shared amongst all users in an Australian LV cell.

BPL has somewhat greater potential in the USA's 110V power environment. Lower voltages mean shorter reach, and as a result, LV cells in urban USA typically comprise only 5-10 homes. The need for signal regeneration is also eliminated with short LV distances, and the low levels of sharing mean that a more credible broadband service can be delivered to each user.

As a broadband access technology, BPL faces some significant challenges before it is likely to have a material impact on the Australian broadband landscape:

---

[2] Regeneration can be accomplished by frequency division multiplexing (FDM – that is, using separate frequencies, one to receive and another to retransmit), or by time division multiplexing (TDM – that is, receiving for a period before requesting the transmitter to pause whilst the information received is relayed).

- the most "natural" parties to own and operate a BPL network are arguably the electricity utilities – but they tend to show great caution in diversifying into new areas; most recently, Tasmania's Aurora Energy has scrapped its exploration of BPL technology, despite several years' investment in successful trials; this appears to be first and foremost a case of the utility returning to "core business";

- until standards are settled, the use of in-home BPL products will typically generate interference to and conflict with any BPL access technologies that are in use in the same neighbourhood;

- in aerial deployments, interference levels are of concern to the ham radio community and protocols for avoiding the contended frequencies have yet to be established in Australia;

- the amount of capacity that is available can underpin a good service, but only at low uptake levels; beyond around 15-20% uptake, performance would be degraded in periods of high activity and users would typically find an ADSL2+ service more attractive in terms of consistent performance; unless technology breakthroughs allow LV cells to be subdivided into smaller isolated cells each capable of receiving a full-strength BPL signal, high levels of sharing will limit the potential for BPL to become a mainstream access technology.

BPL may find a natural home in supporting emerging new communication requirements of the utilities themselves. Today's electricity supply networks are largely "dumb" networks, with utilities having little or no visibility into what happens beyond their major zone substations. This stands in striking contrast to modern telecommunications networks where every active element is monitored and information and communications technology (ICT) is deeply embedded in the network.

Electricity grids can be modernised and new levels of efficiency achieved by monitoring all the active elements in the grid – in particular, every low voltage transformer and every supply meter. Such monitoring relies on a communications fabric, and BPL offers abundant capacity in a form that is naturally aligned with the architecture of the grid. Growing global pressure on greenhouse gas emissions could see utilities embracing Smart Grid technology in the future. If this leads to BPL being deployed, there may be potential for surplus capacity to be used for general broadband services.

2.2.4. Radio Waves and Wireless Technologies

With radio solutions, the final connection to the customer is via radio waves to/from a "base station" or transmission tower.

The physics of radio transmission are complex, but performance can be characterised by a few key parameters:

- The radio frequency (RF) band in which the signal is transmitted. As a simple rule of thumb, the higher the frequency the shorter the reach and the greater the dependence on clear line-of-sight between the transmission and reception point.

- The width of the frequency band used in transmission. As a simple rule of thumb, the higher the frequency the greater the bandwidth – simply because capacity is related to the number waves per second. A one mega-herz (MHz) band has a thousand times as many cycles as a one kilo-herz (kHz) band.

- The shape of the beam. This may be omni-directional (360^o), sectored (typically 60^o or 120^o) or quite narrow (<10^o). A good parallel is a water spray which may also be sprayed omni-directionally, in a general direction or in a tightly focused jet stream. The broader the coverage area, the more likely that the signal will be shared amongst a large number in the reception area. Very narrow beams are used for point-to-point (P2P) wireless links, offering high capacity but minimal coverage. P2P links are not especially relevant to a consideration of access network technologies except as an option to cabling for satisfying "backhaul"[3] requirements.

- The transmission power. The stronger the signal, the greater the reach – just as a loud noise carries further than a soft noise. The reach of a signal is essentially determined by frequency and transmission power.

- Whether the spectrum being used is licensed (and hence protected from "open slather" interference) or unlicensed. Unlicensed spectrum is freely available for all to use, but as a result there can be high noise levels that limit performance levels.

The wireless technologies of greatest relevance to a discussion of broadband access are in Australia WiFi, 3G, WiMAX and iBurst and satellite. A user's preference amongst these different technologies is most likely to be shaped by:

- coverage in the areas of interest to a user;
- the device that the users wants to use (3G phone, notebook computer with wireless access card, etc);
- the level of performance required;
- any need for mobile access.

WiFi

In its most common configuration, WiFi is a short-reach technology (nominally 100 metres), with relatively high capacity (11 Mbps, 54 Mbps, 108 Mbps or even higher). WiFi operates in an unlicensed spectral band, and many homes operate their own in-home wireless network. In this regard, the limited reach is a decided benefit – if the signals reached further, the levels of interference from neighbours would largely cripple the service since it is only possible to accommodate three non-overlapping signals in the frequency band used by WiFi. The comparatively small coverage area of a WiFi base station is known as a "hot spot".

In addition to private hotspots, various operators host public hot spots where anyone can access the Internet either freely, or for a usage fee. There are many different ways of collecting usage fees including subscription, "scratchy cards" (that contain a period of credit and where the password is revealed by scratching away some masking paint), credit card, charging to a mobile phone account and so on. The proliferation of public hot spots has given rise to a business for "hot spot aggregators" who offer a single payment mechanism that is honoured at typically thousands of hot spots around the world.

WiFi hot spots offer excellent performance in their limited coverage areas, but are not well suited to mobile or nomadic users who want sustained coverage whilst on the move. There have been some efforts around the world to provide blanket coverage of larger areas, but most of these have been abandoned because of technical difficulties in provide quality coverage or because of problems in funding the necessary infrastructure.

---

[3] Backhaul refers to the carriage of aggregated traffic from a node (such as a wireless base station) into the core network.

One of the key strengths of WiFi technology is the fact that support is integrated into so many user devices. Most modern notebook computers and many personal digital assistants have inbuilt WiFi technology, and new devices (such as "dual mode" phones that use WiFi where available and otherwise default to the mobile phone networks) are starting to proliferate. This is a significant advantage not shared by any of the other wireless access technologies.

3G and the Mobile Phone Networks

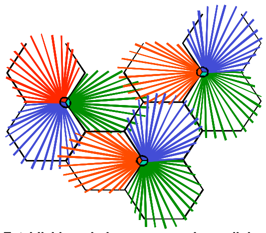

Establishing wireless coverage in a cellular network using non -interfering frequency bands

The 3G networks are an evolution of the global systems for mobile telephony (GSM) networks to support broadband data access. GSM networks have a cellular structure, with transmission towers distributed through a serving area to provide blanket coverage. Each tower (see picture on right) typically supports one or more arrays of three antennae, each with a 120^{o} beam that is carefully focused to cover a designated area (as illustrated in the diagram on the left).

Australian GSM networks operate in the 900 Mhz and 1.8 GHz frequency bands. Additional antennae can be deployed (creating a finer cellular structure) to provide extra capacity if usage levels reach saturation point. In essence, this equates to pushing optical fibre (for backhaul) deeper into the network.

The first stage in introducing support for data users was the generalised packet radio service (GPRS) – allowing long-held connections with traffic charged by volume rather than time. Today GPRS access is available wherever there is a GSM signal. As such, it is available in the more populated areas and along major highways. However, the GPRS service operates at approximately dial-up speeds and hence is not a true broadband service.

3G (3rd Generation) technology is a generic term referring to more powerful broadband data capabilities. In Australia, 3G has been progressively introduced into Australia's GSM networks, utilising a 2.1 Ghz frequency band from the same towers as are used for GSM. To rationalise infrastructure costs, Telstra and Hutchisson share infrastructure and Optus and Vodafone share infrastructure. Thus there are in essence two rival networks supporting 3G services. The basic 3G services operate at 384 kbps – a little above the 256 kbps that represents the lowest performance level for broadband delivered via ADSL.

Given their origins, GPRS, 3G and related services building on mobile telephony infrastructure are intrinsically well-suited to supporting mobile users who move through the coverage areas of different base stations. Connections are seamlessly handed over to adjacent cells, and connectivity can be maintained whilst traversing large distances.

Unlike GPRS, 3G service is *not* available throughout the GSM coverage area. Coverage tends to be provided near major airports and in the central areas of Australia's larger cities. However, the technology used for connection typically supports "fallback" to the lower GPRS standard when moving out of a 3G coverage area. As such service can be maintained – even though performance levels may vary.

New standards such as high speed downlink packet access (HSDPA) offer even higher levels of broadband performance – with theoretical speeds from to 1.8 to 14.4 Mbps. Whilst these speeds sound impressive, it must be remembered that the capacity of the signal is shared amongst all in the coverage area – and practical speeds achieved are typically much lower. Accordingly the only way to support large numbers of high-speed users is to divide the serving area into ever-smaller cells so that the number of users sharing a signal is reduced.

Telstra is currently in the process of decommissioning its CDMA network and will re-use the spectrum that it used (850 Mhz band) for supporting 3G services. The lower frequency give extended reach – something that is particularly helpful in rural and low-density areas, allowing good coverage from a more sparse array of base stations than could be achieved with the 2.1 GHz frequency band.

WiMAX & iBurst

In contrast to 3G with its origins in mobile telephony, WiMAX and iBurst have been designed from the outset as wireless broadband delivery technologies. Eliminating the requirement to support voice calls allows more capacity for broadband services. Both technologies support a level of mobility between adjacent coverage areas.

iBurst uses a proprietary protocol and operates in a licensed 1.9 GHz frequency band. The company (Personal Broadband Australia Pty. Ltd.) driving its deployment aims to eventually provide coverage to around 75% of Australians, concentrating on the major population centres and business districts. Connections speeds up to 1 Mbps are currently supported.

WiMAX has been designed to operate in multiple frequency bands, with initial interest concentrating on the unlicensed 5.8 GHz and licensed 2.3 and 3.4 GHz bands. Connection speeds to 12 Mbps can be supported with current technology.

In Australia, the licensed bands are owned (predominantly) by Unwired (in city areas) and Austar (in rural areas). The Opel (Optus/Elders) consortium plans to use freely available 5.8 GHz spectrum, relying on the fact that the airwaves are not unduly crowded in the rural and regional areas that are the prime targets for its Wimax deployment. Little has been publicised about its deployment plans at this stage, though it is understood that Opel will be deploying a fixed variant of the WiMAX standards rather than the variant that supports mobility.

The Unwired Group launched its service in the Sydney area using a proprietary technology in the 3.4 GHz range. However, in 2005 it announced a substantial investment from Intel and plans to migrate future deployments to WiMAX technology.

WiMAX is of particular interest because of Intel's intentions to include WiMAX support into the Centrino chipsets that are widely used to support WiFi connectivity in laptops from early 2008. The proliferation of WiMAX-capable devices will undoubtedly lead to a ready-made market for carriers offering WiMAX access.

Satellite

Satellite technology represents the "ultimate" in wireless coverage, with the signal being bounced off satellite positioned in geostationary orbit at an altitude of 36,000 kms. The resultant beam can be shaped to cover an area as large as the entire Australian continent.

There are two key downsides for satellite-based broadband:

- The distance that the signal must travel from an earth station to the satellite and then down to the receiver will be 72,000 kms or longer, depending on how distant the transmitting earth station and the receiver are from the optimum point immediately "below" the satellite. Even though RF waves travel at the speed of light, this means a delay (or latency) of at least a quarter of a second over and above the other delays in the end-to-end communications path. As a result, satellite-based services are ill-suited to real-time applications such as VoIP and video-conferencing.

- Although a satellite may deliver many tens or even hundreds of Mbps of broadband capacity, the large serving areas mean potentially high levels of sharing – and thus the capacity available to an individual user can be both limited and highly variable.

For these reasons, satellite is likely to remain the broadband access technology of last resort, used principally in remote areas where it is simply not feasible to bring cabled technologies within reach of the user.

Intrinsic Limitations of all Wireless Technologies

Many observers believe that wireless technologies are on the threshold of rendering the need for cabled broadband obsolete – a "silver bullet" that will solve pressures on network infrastructure. This belief leads to the comfortable view that no action is needed upgrade Australia's cabled network infrastructure to meet the needs of the future.

Unfortunately this is a naïve view. Whilst computing power has continued to double every 18-24 months (Moore's Law) for more than 40 years, telecommunications is governed a law of limitation (Shannon's Law) that describes the maximum capacity that can be obtained through a noisy channel such as the airwaves. Speeds will continue to increase as more computing power can be thrown at the task of isolating signals against a noisy background – but there is no scientific evidence to support the view that there will be a fundamental breakthrough in wireless capacities.

In a historical comment on cabled and wireless communication, Alcatel has noted that wireless technologies have consistently lagged behind cabled technologies by an order of magnitude for many years. The illustration below summarises the essential reason why this is so. A cabled connection can be compared to delivering water to its intended destination by hose with surgical accuracy – versus a wireless connection where the water is taken by a hose to sprinkler and then sprayed in the general direction of the destination.

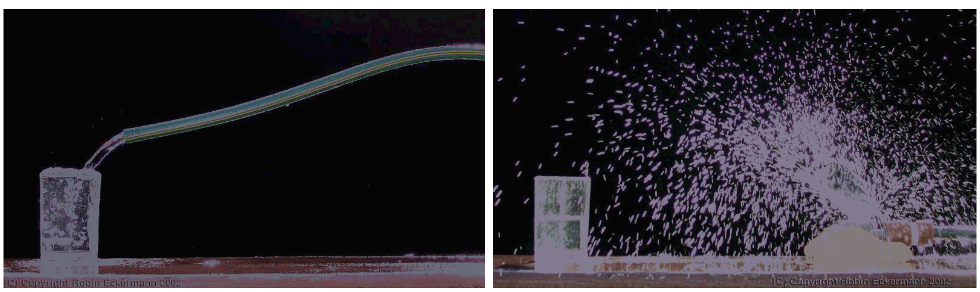

Wired versus Wireless Access Technologies

It would be wrong to discount the importance of wireless technologies on the grounds that they do not ultimately offer the same capacity as cabled alternatives. In certain parts of Australia, wireless access will remain the most economically viable mean of delivering broadband for some time to come. Furthermore, in regional and more remote areas, the inherently lower capacity can be much less of an issue for two reasons. Firstly, the airwaves are typically less cluttered with interfering signals, and hence performance is less prone to degradation. Secondly, at low population densities a signal is typically shared amongst a smaller community of users and hence there is more for each user.

Unique Opportunities Created by Wireless

Even more importantly, wireless technologies create opportunities that simply cannot be replicated on cabled networks – the ability to support "untethered" access by nomadic and mobile users. This creates the opportunity to fundamentally re-engineer and streamline activities to access information when and where needed, and to capture information when and where it is generated.

There are good examples throughout Australia and overseas where government agencies have pioneered new approaches that harness the unique benefits of wireless access[4].

2.3 Common Theme

It should be apparent from the preceding sections that there is a common theme in access network capacity. No matter what the access technology, the key to performance lies in pushing optical fibre closer to the customer:

- In the case of the DSL family of technologies, performance increases hyperbolically as the copper distances become shorter;
- In the case of all other access technologies (HFC where the signal is shared amongst everyone on the same coax segment, BPL where the signal is shared amongst everyone on the same power segment or Wireless where the signal is shared amongst everyone in the beam), the performance increases as the level of sharing is reduced – but this means breaking the serving area into smaller cells, each of which needs to be fed with a signal.

Thus the task of improving Australia's broadband capacity is first and foremost and cabling challenge – not a technical challenge. Furthermore, if the process of pushing optical fibre deeper into the network is taken to its logical limit, the result is FTTU (or fibre-to-the-user).

This applies even with wireless technologies. For example, to provide blanket coverage with WiFi technology and performance levels (and setting aside issues of interfering signals), it would be necessary to deploy a grid of base stations such that no point was more than 100m from a base station. Needless to say, connecting such a dense array of base stations back into core networks represents a substantial cabling challenge! Furthermore, even this configuration, each signal would typically be shared amongst multiple users[5].

Steps such as FTTN will certainly lift capacity – but as pressure for ever higher performance grows, architectures that take fibre all the way to the user will inevitably attract more and more interest.

---

4 See for example http://www.dbcde.gov.au/__data/assets/pdf_file/24256/Government_use_of_wireless_broadband.pdf

5 The level of sharing in a cell obviously varies considerably depending on the size of the cell and the density of users within the service area. For a FTTN network (where distances are limited to around 1 km) in an urban environment such as the more established areas of Canberra, the number of reachable homes would typically range from about 150 to 500. For TransACTs FTTK network (where distances were limited to 300m), the number of reachable homes ranged from around 20 to 100, with an average of around 55.