
Architecting a National Optical Fiber Open-Access Network: The Australian Challenge

Vijay Sivaraman, The University of New South Wales
Craig Russell and Iain B. Collings, CSIRO ICT Centre
Adam Radford, Cisco Systems, Australia

Abstract

In this article we discuss the challenges faced when building a national optical access network that is ubiquitous (i.e., reaches all households nationwide) and open access (i.e., allows any provider to offer their services over it). Our work is inspired by a public fiber access network being built in Australia called the National Broadband Network (NBN) that will deploy fiber to 93 percent of premises, providing broadband access at potential data rates of 100 Mb/s and above. We highlight how the nationwide open access nature of the network creates new architectural challenges, in terms of both fiber layout as well as interconnection with private providers competing to offer retail services. We identify some of the technical choices related to the support of advanced features such as quality of service, multicast, and reliability, and discuss potential challenges providers may face when rolling out their services in such a network.

There is steady growth worldwide in the deployment of fiber access networks for residential broadband connectivity. Fiber penetration rates (including fiber-to-the-home [FTTH] and fiber-to-the-building [FTTB]) are particularly high in Asian countries such as South Korea (52 percent), Japan (34 percent), Hong Kong (33 percent), and Taiwan (25 percent), according to the FTTH Council AP, 2010. Almost all access fiber deployments are private undertakings (e.g., KT in Korea and NTT in Japan). In this article, we focus on the architectural challenges in building a *public* (i.e. taxpayer funded) optical access network. The requirements of a public network are fundamentally different to that of a private network in at least two ways:

- It should provide ubiquitous coverage nationwide, so as not to disadvantage a section of the population.
- It should be an “open access” wholesale network to allow fair competition among private providers to offer their retail services to consumers.

Our intent is to show that these two requirements create new challenges in the architecture of the fiber access network, in areas ranging from physical layout and interconnectivity to support for advanced network capabilities and new service offerings.

Our work is motivated by a public fiber access network currently being built in Australia. The network is being built by the Australian government and is called the National Broadband Network (NBN) [1]. It will take 10 years to build, and is expected to cost approximately A\$40 billion overall. The network will provide fiber access to around 93 percent of premises, with the remaining (predominantly in regional and rural areas) covered by fixed wireless and satellite. In this article we use the Australian NBN as a case study to highlight the archi-

tectural challenges surrounding construction and operation of a public fiber access network. However, we note that many of these challenges also apply to other public fiber access networks such as Singapore’s NextGen NBN [2], although the scale of these challenges will vary depending on factors such as network size and associated operational model.

Our contributions in this article are twofold. First, we highlight how the nationwide and open access nature of a public network creates challenges related to physical layout of the fiber access network and its interconnection mechanism with backhaul networks belonging to private providers competing to offer retail services. Second, we identify technical choices in the support of advanced network capabilities such as quality of service, multicast, and reliability, and discuss potential challenges faced by providers in delivering their services using these capabilities. Throughout our discussions we use the Australian NBN as an example to illustrate the architectural choices and their impact.

The remainder of the article is organized as follows. We briefly describe the context of the Australian NBN that motivates this work. We discuss general challenges related to fiber open access network architectures, and the interconnection between the access and backhaul networks, respectively. Issues related to network capabilities are discussed, and those related to network services are highlighted. The article is then concluded.

Context: The Australian Scenario

The Australian federal government has stated that it aims to be among the world’s leading digital economies by 2020, with specific goals [3] focusing on online participation by Australian households and businesses, smarter management of

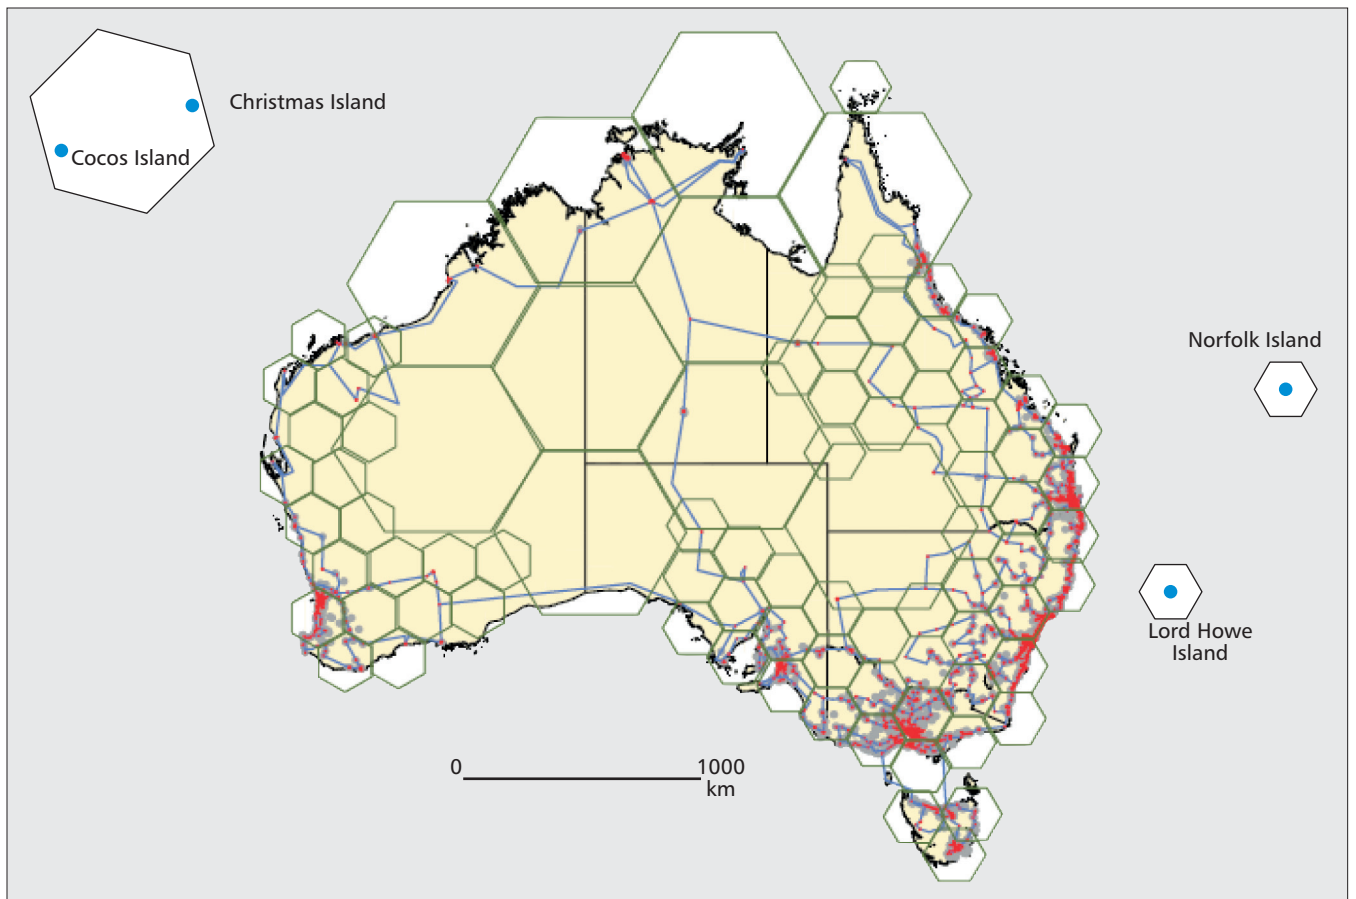


Figure 1. NBN coverage: 93 percent fiber (red), 4 percent fixed wireless (grey/black), 3 percent satellite (hexagons). Source: [1]

the environment and infrastructure, improved health and aged care services, expanded online education, increased tele-commuting, improved online government service delivery and engagement, and greater digital engagement in regional Australia. Consequently, it is building an open-access broadband network (the NBN) which is primarily targeting residential users; though all premises will also be connected, including small and big businesses, schools and hospitals.

The NBN will provide fiber-to-the-premises (FTTP) to 93 percent of Australian households (which can typically avail of download speeds of 100 Mb/s), with an additional 4 percent being covered by fixed wireless and the remaining 3 percent via satellite, as shown in the coverage map in Fig. 1.

The NBN has a staged product roadmap that includes five product drops [4], starting with basic telephony and high-speed broadband at speeds of 100 Mb/s. This is expected to progress to entertainment capability, including triple-play services (voice, video, and data). Small to medium businesses will then be supported with multi-site connectivity, secure and reliable access, progressing to larger enterprises that require speeds of 1 Gb/s and transparent VLAN services. Finally, mission-critical sites such as hospitals will be provided with high availability and access diversity.

The wholesale pricing plan provides an entry-level plan of 12 Mb/s downstream and 1 Mb/s upstream service that is uniformly priced across fiber, wireless, and satellite access, at all locations across the country (so as not to disadvantage remote areas). For the 93 percent of premises that will have fiber access, many are expected to take up a 100 Mb/s downstream and 40 Mb/s upstream service. In addition to the above wholesale services at the user side (user-network interface, UNI), the retail service provider also pays for the network side (network-network inter-

face, NNI) connectivity. If 70 percent of the approximately 11 million Australian premises take up connectivity to the NBN, it is expected to yield revenue of A\$7.6 billion by 2025.

From a technology perspective, the NBN is deploying a gigabit passive optical network (GPON) in a modular architecture shown in Fig. 2. The local network supplies 2.5 Gb/s downlink and 1.25 Gb/s uplink capacity on the access fiber. These rates are expected to be shared among 24 households using passive optical splitters/couplers. Each local exchange (called fiber distribution area, FDA) terminates optical fibers from 200 premises. The FDAs are in turn interconnected hierarchically in a regular pattern, with 16 FDAs aggregating (via redundant fiber topology) to a fiber serving area module (FSAM), 24 of which aggregate to a fiber serving area (FSA). An FSA connects via a transit network to a point of interconnect (POI), which is where retail service providers (RSPs), which operate their own long-haul networks, connect with the NBN. Currently there are expected to be 121 POIs nationwide. Physical dimensions of components of the network will vary across the continent; however, the distance from the fiber aggregation node (FAN) to the edge of the FSA is less than 15 km (rectilinear).

All the network assets, including equipment and fiber, from the premises to the POI will be owned and operated by NBN Co., the wholly owned government entity that has been established to deploy and operate the NBN. RSPs will purchase from the NBN provisioned Ethernet transport service between the POI and the customer premises at the wholesale prices mentioned above. This decoupling of the network infrastructure from the service allows an RSP to access any customer in the country by connecting to the (relatively small number of) POIs, while also allowing a customer to choose any (or even several) RSPs for network services.

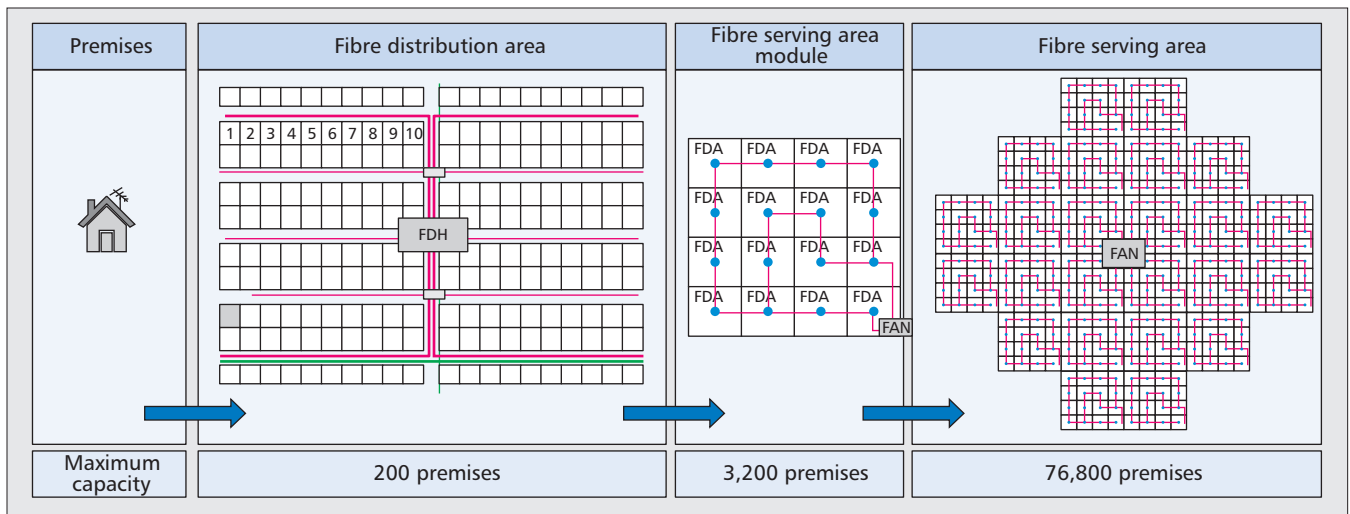


Figure 2. Hierarchical topology of the NBN access fiber. Source: [1].

In the following sections we will discuss the architectural options in designing any national fiber open-access network. In each section we will provide specific details of the relevant components of the Australian NBN as a case study example.

Physical Architecture Challenges

In this section, we highlight some of the architectural challenges relating to the physical layout of a nationwide fiber access network. We focus on two specific aspects: we discuss access technology and access fiber layout options in the local network, and cost-effective replication of the fiber infrastructure in the distribution network.

Local Network

Within the context of a continental-scale network, the local network connects households to the local exchange. In what follows we examine the options to achieve this connectivity.

Medium: Fiber, Copper, or Wireless? — Fiber is a very attractive medium as it offers virtually unlimited bandwidth capacity at very low error rates, while having less weight, lower energy-per-bit, and longer lifespan than copper. The primary attraction of wireless is the lower build-out cost (since no parts need to be installed in the field, other than at user premises and the local exchange), and the potential to support mobile devices (smart phones and pads). A national network will typically need to choose an appropriate mix of access medium technologies based on factors such as geography, population distribution, bandwidth needs of services, and cost constraints. In the Australian context, it is especially the case, since the country is geographically vast but has high concentration of population in a few regions. The Australian NBN is aiming to serve as many as 93 percent of premises by fiber (shown by red dots in Fig. 1). Fixed wireless (that provides wireless connectivity to a household, not to a mobile device) will serve 4 percent of Australian premises (shown by black/grey dots in Fig. 1), and the remaining 3 percent of premises will be served by satellite (the satellite cells are shown as hexagons in the coverage map of Fig. 1).

Fiber Reach: Curb or House? — This is an especially important aspect of a public network. It is a balance of technology and pragmatism (in terms of the ability to actually gain access to properties in the publicly funded context). A fiber access network can have fiber built to the node/curb (called FTTN/FTTC), and then use the local copper loop to access households. This is generally a cheaper option as it avoids

reconnecting every household (although it is subject to the local loop being made available by the provider who owns it). However, FTTN/FTTC requires active equipment in the street, which creates operational challenges with regard to battery backup, flood damage, and so on. We note that the Australian NBN has chosen to deploy fiber to the premises (FTTP) throughout, and avoided the use of copper. Even though this is the more expensive option, it removes the (often aging, particularly in Australia) copper from the local loop, and provides much higher bandwidths that scale well into the future.

FTTP Technology: P2P or PON? — There is a choice between having a dedicated point-to-point (P2P) fiber to each household, or a PON whereby one transceiver at the optical line terminal (OLT) equipment in the local exchange communicates with several optical network terminal (ONT) devices, one in each household, over an infrastructure consisting only of passive splitters and couplers. Experience from other countries shows that 80 percent of the cost of building last mile networks relates to the civil engineering construction costs, and electronics and cable component only account for 20 percent of the total cost. Nevertheless, the absence of active components in the field gives the PON architecture higher longevity and reliability in a more cost efficient manner.

There are two dominant PON standards: GPON, which has larger deployment in Europe and North America, offers 2.5 Gb/s downstream and 1.25 Gb/s upstream, and supports non-native transport protocols including asynchronous transfer mode (ATM) and time-division multiplexing (TDM), while also supporting Ethernet. The other standard, Ethernet PON (EPON), is more widely deployed in Asia (Japan, Korea, Taiwan), and supports 1 Gb/s in each direction. We note that the Australian NBN has chosen GPON technology, which caters better to asymmetric traffic patterns wherein customers consume more data than they produce. A typical deployment in which 24 households share a GPON would give an average download bandwidth of approximately 100 Mb/s to each household. However, the NBN topology has spare fibers to support premises that require higher speeds via P2P fiber.

Fiber layout: Underground or Overhead? — One of the major challenges in fiber rollout relates to the high cost of running fiber from the curb to each household. While underground deployment is preferred due to higher longevity and lower chance of damage, trenching to lay the fiber can be very expensive. In the Australian NBN case, they are planning to utilize existing ducts and pits where possible. The remaining

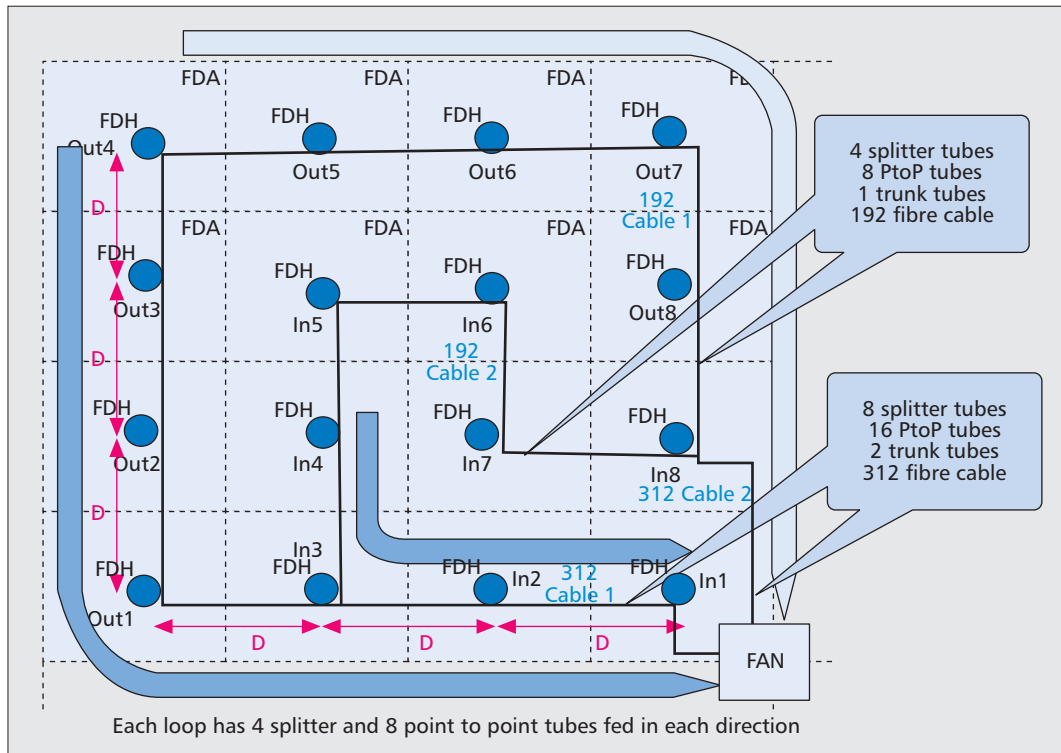


Figure 3. FSAM layout showing redundant interconnection of FDHs. Source: [1]

premises will have aerial cabling, co-located with existing electricity distribution networks. Each household will have an ONT device at which the fiber terminates. This device will come equipped with Ethernet and telephony ports.

Challenges within Premises — We believe that some of the practical challenges that will arise within or near the premises in connecting to the network will include:

- **Power supply:** From where will the ONT draw power, and with what battery capacity should it be equipped to provide critical telephony services during power outages? There is also debate about whether to provide battery backup for the ONT given that many households use cordless phones that rely on the mains power supply.
- **Connectors:** Should the fibers be pre-terminated in the factory before they are installed (to reduce cost) or should the connectors be attached by the field personnel during installation (to reduce excess fiber coils etc.)?
- **Wiring:** How will the wiring be done from the ONT to other outlets in the house?
- **Multi-dwelling units:** How will the wiring (fiber or copper) run from the ONT to the various apartments in the building (and who is responsible for it)?
- **Wireless:** If many neighboring households operate a wireless access point, will the resulting WiFi interference seriously impact usability of the network?
- **Security:** Will encryption (bidirectional or downlink only) be enabled on the PON infrastructure to protect the privacy of user data, or is it best left to the application layer (for residential) or network layer (for multi-site business)?

Distribution Network

When building an open-access network it is important to realize economies of scale to reduce costs particularly if it is a public network intended to serve millions of premises. The most effective way to achieve this is to develop a hierarchical modular design that can then be replicated cost-effectively throughout the network. Several factors influence the number

of levels in the hierarchy and the dimensions of each level, such as density of households, geographical distances (which are constrained by optical reach and splitting ratios), redundancy requirements (i.e. how many premises become disconnected due to a fiber cut), and ease of management and troubleshooting.

In the case of the Australian NBN, their modular arrangement, shown in Fig. 2, has the following dimensions: Four houses (2 each side of the road), each referred to as a geocoded national address file (GNAF), connect to a network access port (NAP). A NAP contains 12 fibers, 2 provisioned per house and an additional 4 spares for future use (e.g. business services). Fifty NAPs connect to a fiber distribution hub (FDH), which serves a fiber distribution area (FDA) of 200 GNAFs. The FDH contains the splitter for PON services. Only one fiber per GNAF is spliced at the FDH, although there are two provisioned. Point-to-point services can be delivered by bypassing the splitter.

16 FDH are aggregated in a fiber serving area module (FSAM) and connect back to a fiber aggregation node (FAN). The FSAM has two 312-fiber cables containing 26 tubes (12 fibers each). There are 8 splitter tubes (PON), 16 P2P tubes, and 2 trunk tubes. These are arranged in a redundant inner and outer loop, as shown in Fig. 3, to protect against fiber cuts. The inner loop is half the outer loop (13 tubes). One splitter tube is delivered to the FDA in the loop. Up to 24 FSAMs are connected to a single FAN at the center of a fiber serving area (FSA). The FAN is where the optical line termination (OLT) resides. As mentioned above, the FAN has $2 * 312$ fibers delivered to each of the 24 FSAMs. Since each FSAM has 16 FDAs, each serving 200 households, there are up to $24 * 16 * 200 = 76,800$ houses in an FSA. There will be approximately 980 FSAs in the final network. Each FSA connects (possibly via a transit network) to a point of interconnect (POI), where the access network hands off traffic to retail service providers (also referred to as RSPs or access seekers). The interconnectivity with access seekers is discussed in detail in the next section.

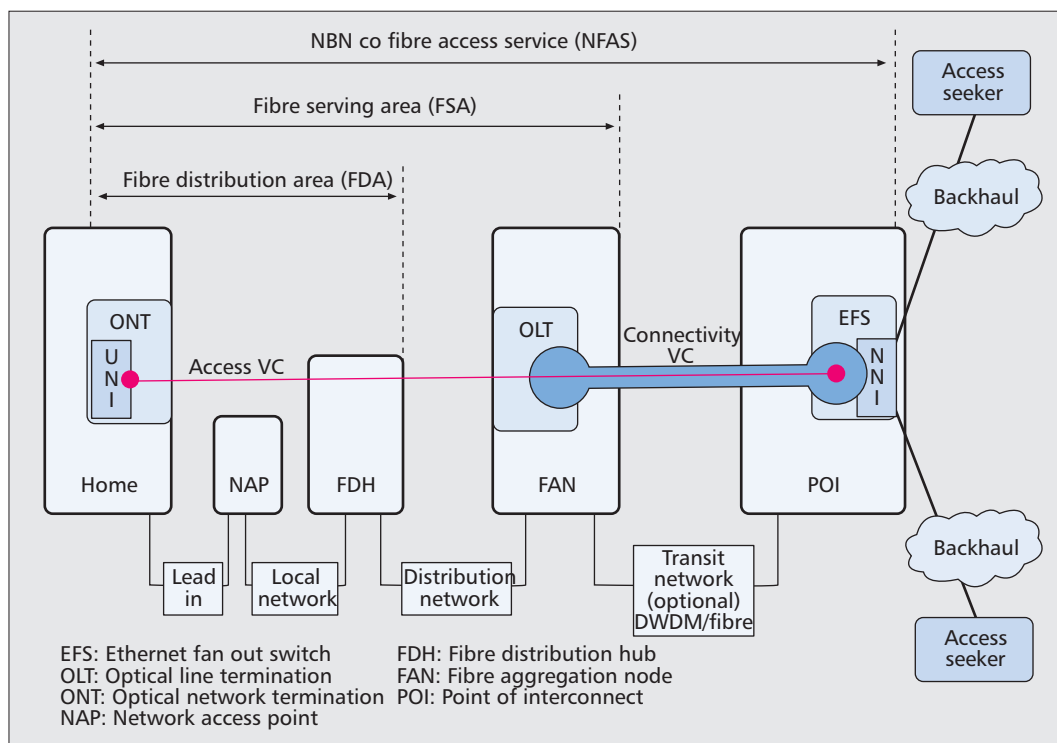


Figure 4. Layer 2 architecture and POIs between the NBN and access seekers (RSPs).

Interconnection Architecture Challenges

Public networks provide wholesale infrastructure, but do not offer retail services, for two reasons: first, the complexity of developing and managing services for end consumers is too high, and second, a monopoly in retail services would stifle innovation. For these reasons, a public fiber access network needs to be open access, that is, allow fair competition among retail service providers in accessing potential customers. Likewise, an open access network gives consumers maximum choice in selecting their RSPs, and indeed, a single household may be served by multiple access seekers, say, one offering video services and another offering general Internet access. We discuss two important architectural choices that are central to meeting these challenges.

Interconnection Interface

The first challenge concerns the OSI layer at which the network hands off traffic to an Access Seeker (or RSP). If traffic is handed off at layer 1 (e.g. dark fiber or entire wavelength), the Access Seeker that gains access to that fiber (or wavelength) would not only have a monopoly on that household, but also on all premises connected on that PON, thus preventing competition. If, on the other hand, the network hands off traffic to the access seeker at layer 3 (i.e., IP packets), the network will have to act as the first-hop router to the customer premises equipment, and take on all the associated complexities of a layer 3 network, such as address assignment, implementing routing protocols, managing peering with access seekers, migration to IPv6, and implementing IP multicast protocols. Therefore, layer 2 (Ethernet frames) connectivity between the public access network and the access seeker is most suitable as it combines flexibility (ability to virtualize and share the fiber infrastructure among the access seekers) with low complexity (since IP addressing and routing do not need to be managed). These reduce the barrier of entry for new RSPs and should result in greater competition.

We note that the NBN uses the IEEE 802.1ad S-TAG/C-TAG structure (generally referred to as provider bridging, stacked VLANs, or simply q-in-q tagging) to transport Ethernet frames between the access seeker at the NNI and the user

premises at the UNI. In the most common case of traditional unicast data services, this two-level VLAN addressing scheme uses an inner C-TAG to address an individual access virtual circuit between the NNI and a UNI (called access VC), and an outer S-TAG to address a connectivity virtual circuit between the NNI and an FSA (called connectivity VC). This architecture, as shown in Fig. 4, therefore enables multiple access seekers to simultaneously deliver services to an individual premises, and it scales well (up to 2^{24} user-provider relationships per FSA). However, it also offers additional flexibility. By allowing the S-TAG to address individual end-user UNI ports, the access seeker can use the C-TAG for its own purposes on a per-user basis (e.g., priority tagged frames for quality of service differentiation or for offering VLAN transparent services to enterprises). Furthermore, multicast services can be deployed from a single NNI to multiple FSAs by having the S-TAG address an FSA and using medium access control (MAC) forwarding within the FSA.

Points of Interconnect

Another important architectural challenge concerns the number of points of interconnection (POIs) between the public access network and private backhaul networks (belonging to RSPs). In regional areas, which have low population densities, the number of users served by an aggregation point would typically be small, and an RSP who connects their equipment there would have to bear a much higher cost per user. To prevent this from disadvantaging smaller RSPs, traffic from several aggregation points in remote areas needs to be aggregated, via a transit network, into an aggregation node site that incorporates a POI. For such cases the cost of the transit network borne by the public purse has to be balanced against an RSP's ease of access to customers in remote areas.

In the context of the Australian NBN, these trade-offs have led to a decision to aggregate 967 FSAs to 121 POIs nation-wide. This is an average of approximately eight FSAs per POI, but the number will vary greatly between metropolitan and regional areas. In a metropolitan area, a FAN is likely to cover a large number (as high as 76,800) premises, and hence there will be one POI for a few (as low as 2) FANs. On the other hand, the POI for a regional

area will cover many FANs — for example, the Toowoomba region in southeast Queensland will incorporate 43 FANs to a POI. Roughly a third of the POIs are expected to be in metropolitan areas and the remaining in regional Australia. Furthermore, the Australian NBN has recently announced plans to build 100 Gb/s dense wavelength-division multiplexing (DWDM) services that provide nearly 1Tb/s of access between a POI and FSA.

For a large continent such as Australia with vast tracts of sparsely populated areas, the large number of POIs will entail high cost in building transit networks. We therefore think it is likely that a market will emerge for private transit providers offering wholesale connectivity between POIs. It is even conceivable that an RSP could use multiple such transit providers. This may lead to additional complexity of inter-connection between the RSPs and the transit providers, since each transit provider may support a different transport technology (e.g., MPLS vs. IP) or may not allow VLAN transparency, thereby necessitating mechanisms for VLAN ID translation.

Challenges in Network Capabilities

Ethernet bitstream service provision is essential in any public open access network. In the hierarchical design discussed earlier, Ethernet will need to run from the POI to an end-user premises. The base service from the RSP will then at least need to include high-speed broadband and telephony, which can evolve to also deliver video services and enterprise-grade products. In what follows we briefly review network capabilities required for public open access networks, and discuss some challenges service providers may face when building retail offerings using these capabilities.

Ethernet Bitstream Product

The basic principles of Ethernet bitstream products were discussed above. We believe this will present access seekers with some new challenges. Among these will be increased traffic volume, since public fiber open access networks will increase the bandwidth to a household by a factor of 20 to 50 (and users will also have access to significantly greater uplink capacity than currently available). Access seekers will have to adequately provision their networks to deal with this (and decide on how much oversubscription they operate with). In order to optimize (and guarantee) their service delivery, they will also need to understand the nature of the traffic they carry, for example, how much of it is local vs. international, P2P vs. client-server, social networking vs. cloud, and so on. With the increased bandwidth available and the new service offerings (e.g., multicast and quality of service) to users, old traffic models may no longer be useful in predicting traffic volumes and types. Access seekers that offer video content distribution and IPTV services will also have to determine placement and interconnectivity of their content servers to provide a good user experience. In order to deal with flash crowds (e.g., major sporting events), they may also require dynamic provisioning capabilities.

Enterprise Product

To assist RSPs deliver high quality and reliability solutions to businesses the following enhanced features are required: optical interface options at the UNI; higher symmetrical speed offerings up to 1 Gb/s; carrier Ethernet VLAN transparency options; access VC protection options; an enhanced range of service provisioning and restoration options; and operations, administration, and maintenance and enhanced reporting options. Again, we believe this will present access seekers some new challenges since these services will be delivered over the same infrastructure as residential services and they will need to ensure that oversubscription rates will not impact service delivery of enterprise products.

Also, the question of security and privacy of enterprise data will need to be considered since the public access network may or may not be responsible for any encryption or security guarantees.

Class of Service

The wholesale access network needs to provide a sufficient variety of quality of service (QoS) classes that RSPs can ensure appropriate performance for the higher-layer applications and services they offer. In the case of the Australian NBN Fiber Access Service (NFAS), the product supports four traffic classes at the Ethernet layer. This feature represents a significant improvement in QoS delivery and application performance differentiation in the Australian telecommunications market, particularly for residential customers who have historically only had best effort service offerings. The specification of the different traffic classes is via two parameters, committed information rate (CIR) and peak information rate (PIR), and it is each access seeker's responsibility to request the appropriate values of these parameters to guarantee application performance to its end users. The four traffic classes available are (with example applications): TC_1 (voice), specified by CIR; TC_2 (interactive streaming and real-time video), specified by CIR; TC_3 (premium data), specified by CIR and PIR; and TC_4 (best effort data), specified by PIR. The scheduling mechanism used for these classes is strict priority for TC_1 and TC_2 (with TC_1 at the highest priority) and weighted fair queuing for TC_3 and TC_4.

The NFAS product specification refers access seekers to Internet Engineering Task Force (IETF) RFC 4594 [5] for recommendations on how to map IP DSCP values to the four IEEE 802.1p PCP values accepted at the UNI and NNI (these bits are located in the VLAN tags). While supporting only four traffic classes and two parameter values for bandwidth specification may be considered coarse compared to other models of quality of service provisioning, it should be more than sufficient for the vast majority of residential customers since it allows for voice, interactive streaming and real-time video, and best effort. This is a dramatic improvement on current broadband offerings. We are of the opinion that the NBN architecture and class of service (CoS) provisioning provides enough flexibility to allow access seekers to offer a variety of service guarantees without being overly prescriptive as to how this should be done. This should enable unprecedented levels of innovation in the service provider market and facilitate competition. It is not clear, however, how access seekers will manage and control traffic contention and congestion on their circuits, particularly for the expected growth of new applications and traffic types, as alluded to earlier. This represents an opportunity for further research on how to optimize bandwidth specification, traffic contention and congestion management, and quality of service guarantee.

Multicast Capability

Multicast capability is mostly required by RSPs engaged in video content distribution and IPTV. Fortunately the public network need only support multicast at layer 2; it is up to RSPs to manage, dimension, and implement their own IP-based multicast architectures. Layer 2 multicast is a fairly simple technology, requiring switches to forward multicast Ethernet frames away from the source and toward receivers based on information learned through IGMP snooping. It is possible to use IGMPv3 snooping, thereby enabling RSPs to offer source specific multicast, as in the Australian NBN. However, there are many challenges for the RSPs in dimensioning their multicast services, based on wholesale pricing and demand, that may become apparent only when these services are deployed. The uptake of these services should stimulate further research into more efficient/optimal ways for access seekers to cache content in their IP networks (as mentioned above).

Resilience/Reliability

Resilience can be built into several components of the fiber access network: in the fiber serving nodes (within an FSA) via redundant links between GPON OLTs and the fan-out switches. Furthermore, access seekers can achieve high availability across the NNI via several mechanisms. Single chassis resilience can be achieved by logically bundling a number of parallel physical links and using IEEE 802.1ad link aggregation. Alternatively, chassis diversity is available for a single NNI group at a POI whereby interfaces are grouped into an active/standby relationship. Finally, site diversity is available across different POIs using a similar active/standby relationship for interfaces. At the user premises (UNI), redundant connectivity entails higher cost and is likely to be a business-grade product.

Services and Impact

In this section we briefly comment on the possible impact of a public fiber open access network. For a detailed study in the particular case of the Australian NBN, see [6].

Retail Service Providers

It is widely understood that retail providers of Internet access will benefit from a larger customer base at higher access speeds, and will have to prepare their network for increased traffic volumes and types. More important, we believe that a public fiber open access network will undoubtedly allow retail service providers to specialize and develop new services, such as video content distribution, gaming, etc. Due to its open-access and wholesale nature, such a network has the potential to enable service providers to reach consumers who are interested in these niche services.

Other Service Providers: OTT, Cloud, CDN

Providers of over-the-top (OTT) audio (e.g. Skype) and video (e.g. YouTube) services, who have long argued for network neutrality, should benefit from an increased choice of RSPs that will be technically available to them to reach their customers. Cloud computing services should also receive a boost — for example, Apple's iCloud and Amazon's cloud-drive each offer 5GB of storage capacity for free that become much more usable over a 100 Mb/s access connection. Operators of content distribution networks (CDNs) should also be able to deliver better customer experience by placing their content caches at or near the POIs.

E-Health and E-Education Services

It is well recognized that a public open access network will provide opportunities to significantly improve healthcare services, particularly to rural and regional areas that suffer through lack of access to medical specialists. For example, building on CSIRO technologies [7], one could envisage delivering improved home-based care systems, access intelligent portals for medical image analysis, and use telemedicine to bring specialist healthcare to remote areas. In a similar way to health services, public fiber open-access networks are expected to enable more people to access more educational materials, via schools, universities and vocational training programs. In particular, rural areas will be able to participate in education that previously has only been available in urban centers.

E-Government Services

There are many government services that a public open-access network has the potential to facilitate. Broadly, the relevant government can benefit in at least three ways: it has the potential to reap the benefits of increased productivity from the private sector (via increased online business activity); it can communicate more easily with the community (education campaigns, reports and studies), informing the community (compli-

ance requirements, entitlements and services offered), and servicing the community (e-tax and telehealth); and state and local governments also potentially could improve complementary public services and social inclusion in areas ranging from health, education, housing, transportation, business regulation, etc. These represent ways in which the digital divide, whereby people living in rural communities have limited access to services commonplace in urban areas, can be bridged.

Conclusions

We have provided an overview of the challenges surrounding the architecture of a public optical fiber access network that provides ubiquitous and open access coverage. We have highlighted the role of physical fiber layout in providing ubiquity of coverage, and of interconnection between access and backhaul networks in maximizing participation. We have identified technology choices related to the support of advanced capabilities such as quality of service, multicast, and reliability in the network, alongside discussing the challenges faced by providers in developing innovative services using these capabilities. We have used as a case study the Australian National Broadband Network (NBN), a publicly owned fiber access infrastructure project currently underway to connect Australian premises at speeds of 100 Mb/s and above by 2020.

References

- [1] NBN Co., "National Broadband Network – Information Pack," <http://www.nbnco.com.au>, Apr 2011.
- [2] Singapore Govt., "Next Gen NBN," <http://www.ida.gov.sg/Infrastructure/20090717105113.aspx>.
- [3] Australian Govt. Dept. of Broadband, Communications and the Digital Economy, "National Digital Economy Strategy," <http://www.nbn.gov.au/the-vision/digitaleconomystrategy>, 2011.
- [4] NBN Co., "Wholesale Access Service: Product and Pricing Overview for Access Seekers," <http://www.nbnco.com.au>, Dec 2010.
- [5] J. Babiarz *et al.*, "Configuration Guidelines for DiffServ Service Classes," IETF RFC 4594, Aug 2006.
- [6] Report by Allen Consulting Group Pty Ltd for the Australian Govt. Dept. of Broadband, Communications and the Digital Economy, "Quantifying the Possible Economic Gains of Getting More Australian Households online," http://www.dbcde.gov.au/_data/assets/pdf_file/0004/135508/Quantifying_the_possible_economic_gains_of_getting_more_Australian_households_online.pdf, Nov. 2010.
- [7] CSIRO ICT Centre, "Enabling a revolution in healthcare through e-health innovation," <http://www.csiro.au/science/healthcare-revolution-through-ehealth-innovation.html>.

Biographies

VIJAY SIVARAMAN (vijay@unsw.edu.au) is an associate professor at the University of New South Wales in Sydney, Australia, and works part-time at the CSIRO. He obtained his Ph.D. from the University of California at Los Angeles, and has worked at Bell Labs and a Silicon Valley startup manufacturing optical switch-routers. His research interests include optical networking, packet switching and routing, network design, and QoS.

CRAIG RUSSELL (Craig.Russell@csiro.au) is a principal research engineer in the Wireless and Networking Technologies Laboratory of the CSIRO ICT Centre. He obtained his Ph.D. in applied mathematics from Macquarie University, and holds a CCIE certification from Cisco Systems. He has worked for a tier 2 telecommunications carrier, and is currently doing network architecture for the Australian Square Kilometer Array (ASKAP) and CSIRO's wireless broadband system for rural Australia.

IAIN B. COLLINGS (Iain.Collings@csiro.au) is research director of the Wireless and Networking Technologies Laboratory at the CSIRO ICT Centre. He holds a Ph.D. from the Australian National University. His main expertise is in design and analysis of digital communication systems, and he has authored over 220 refereed journal and conference papers. He is a recipient of the IREE Neville Thiele Award and IEEE ComSoc Stephen O. Rice Best Paper Award for *IEEE Transactions on Communications*.

ADAM RADFORD (aradford@cisco.com) works for Cisco Systems as a consulting systems architect, working with their largest customers across the Australia and New Zealand (ANZ) region on strategy and direction. He has spent over 20 years in the computer networking industry in a variety of roles, from research, implementation, design, and architecture, and has been the technology lead for both Cisco Networkers and Technology Solutions Roadshow for the past four years.