

Geoff Huston
December 2017

A Workshop on Internet Economics

In the United States the debate between advocates of market-based resolution of competitive tensions and regulatory intervention has seldom reached the fever pitch that we've seen over the vexed on-again off-again question of Net Neutrality in recent weeks. It seems as if the process of determination of communications policy has become a spectator sport, replete with commentators who laud our champions and demonize our opponents. On the other hand, the way in which we communicate, the manner, richness and reach of our communications shapes our economy and our society, so it's perhaps entirely proper that considerations of decisions relating to its form of governance are matters that entail public debate.

How can we assist and inform that debate? One way is to bring together the various facets of how we build, operate and use the Internet and look at these activities from a perspective of economics. This is the background to a relatively unique gathering, hosted each year by CAIDA, the Centre for Applied Internet Data Analysis, at the University of California, San Diego, at WIE, the Workshop on Internet Economics (<http://www.caida.org/workshops/wie/1712/index.xml>) These are my notes from the 8th such workshop, held in December 2017.

The workshop brought together a diverse collection of practitioners, including network operators, service operators, regulators, researchers and academics. The issues that lie behind many of these thoughts include this topical theme of a post network neutrality network and what it means. We are seeing the rise of tribalism and the re-imposition of various forms of trade barriers that appear to work against the ideals espoused in a world of accessible open communication capability. The traditional view of the public space and communications within that space being seen as part of a society's inventory of public good is being challenged by rapacious private sector actors who are working assiduously to privatise this formerly public space in ways that are insidious, sometimes subtle and invariably perverted. What the role of government in this space should be, and what can and perhaps should be left to market actors to determine is still very unclear. From an historical perspective, the disruptive changes that this rapid deployment of information technologies are bringing to our society are very similar in terms of the breadth and depth of impact to those that occurred during the industrial revolution, and perhaps the thoughts of the nineteenth century economist and revolutionary sociologist Karl Marx are more relevant today than ever before. We are seeing a time of uncertainty when we are asking some very basic questions: what is the purpose of regulation and whom should we apply it to? How can we tell the difference between facts and suppositions?

The initial euphoria of the boundless opportunities created by the Internet are now being balanced by common concerns and fears. Is the infrastructure of software reliable and trustworthy (<http://theatlntc.com/2j6ZP5X>)? Why are we simply unable to create a secure Internet? What is shaping the Internet and is that evolution of forces creating a narrow funnel of self-interest driving the broad agenda? It seems that the Internet is heading down the same path as newspapers, where advertising became the tail that wagged the dog. If advertising on the Internet comes crashing down, what happens then? We are already witnessing the failure of the newspaper model, and will we see a similar failure in the Internet's future? Indeed, what is "the Internet" today? What is in public view and what is private?

As our need for data gets greater how much of the Internet is being withdrawn from public view? And as information is taken away from public purview, to what extent does this leave information asymmetry in its wake? Do consumers really have any realistic valuation on the personal information they exchange for access to online goods and services?

With that in mind, the context is set for this year's workshop.

It's the Economy, Stupid!

In classical public economics, one of the roles of government is to detect, and presumably rectify, situations where the conventional operation of a market has failed. Of course, a related concern is not just the failure of a market, but the situation where the market collapses and simply ceases to exist. Perhaps markets are more than enablers simple transactions between a buyer and a seller. Karl Marx was one of the first to think about the market economy as a global entity, and its role as an arbiter of resource allocation in society. When we take this view, and start looking for potential failure points, one of the signs is that of “choke points” where real investment levels fall and the fall is masked by a patently obvious masquerade of non-truths taking the place of data and facts. Any study of an economy involves understanding the nature of these choke points. Telecommunication services are not an isolated case, but can be seen as just another instance of a choke point in the larger economy. Failure to keep it functioning efficiently and effectively can have implications across many other areas of economic activity.

The telecommunications regulatory regime has, as one of its major objectives, to ensure that the market is functioning effectively. This includes a myriad of functions, including that of ensuring that the market is open to new practices and new entrants, and that the service cost refinements that are accessible through evolutionary innovation are passed on to consumers rather than being retained as a form of cartel or monopoly premium. If we get this right then the economy has a much stronger case for success.

In global terms, the telecommunications market is extremely significant. The sum total of national Gross Domestic Product numbers for the world currently totals some \$80T in US Dollars. When looked at in revenue terms, the revenue level for telecommunications services is some \$2T. By comparison, the total expenditure on advertising is some \$500B. In 2016 Google reported a revenue level of \$90B. If online advertising is the means of funding many of the free online services of the Internet, then the figures today appear to support a case that so far content is still not “king” in the telecommunications environment.

Other parts of the telecommunications space are still extremely lucrative. This observation was illustrated with material gathered from T-Mobile's 2014 filings to the FCC: the cost of data services to the company averaged some \$0.002/Mbyte. On the wholesale market the average wholesale price was \$0.20/Mbyte (or a 100x mark-up), while the retail price was \$20/Mbyte, or a further 100x mark-up over the wholesale price). At least in mobile services there are still highly desirable margins to be made by market actors, and efforts to install a more competitive market discipline appear to have acted more as an incentive for cartel-like collusive practices to maintain extremely high margins rather than a truly competitive pressure to drop prices down to something much closer to the actual cost.

Innovation

Another pressure on markets is that of innovation. New products and services can stimulate greater consumer spending, or realise greater efficiencies in the production of the service. Either can lead to the opportunity of greater revenues for the enterprise.

Innovation occurs in many ways, from changes in the components of a device through to changes in the network platform itself and deployment of arrays of systems. We've seen all of these occur in rapid succession in the brief history of the Internet, from the switch to mobility, the deployment of

broadband access infrastructure, the emerging picture of so-called “smart cities”, and of course the digitisation of production and distribution networks in the new wave of retailing enterprises.

Much of this innovation is an open process that operates without much in the way of deterministic predictability. Governance of the overall process of innovation and evolution cannot concentrate on the innovative mechanism per se, but necessarily needs to foster the process through the support of the underlying institutional processes of research and prototyping. Public funding through subsidies in basic research may be the only way to overcome prohibitively high transaction and adaptation costs in exploring innovative opportunities.

Products, Services, Platforms, and Regulation

From a regulatory perspective, the regulator normally looks at the service, not the facilities. In other words, the topic of regulatory interest is the service, not the network used to deliver the service. For example, a “Broadband Internet Access Service” is defined as something that sends and receives Internet packets within certain service parameters.

However, this distinction is not always preserved, and there are examples where the regulations talk about the facilities rather than the service, such as in the area of defining a duty of interconnection between facilities imposed upon facility operators. The current US Telecommunications Act uses a number of titles to consider various services and service domains, such as Title II’s attention to Broadcast services, Title III for cable services, and so on. While in past years differences services used discrete facility platforms, these days almost every service is just an app running over IP.

Given that the intended subject of the regulation has been the facility operator, while the regulatory constraint has been described as constraints on the service itself, when this is translated into the Internet the outcome could be interpreted as a constraint placed on the underlying ISP relating to the behaviour of applications that run over the top of the ISP network. If the application space is largely a deregulated market, then doesn’t this lead to the observation that these days every service is deregulated?

Maybe the entire issue of whether last mile access ISPs should, or should not, be compelled to work under the provisions of Title II and provide fair access for all services on a fair and non-discriminatory basis is just looking at the issue from the wrong side of the table. Does an entity have a status under these regulations in all of its activities by virtue of providing a regulated service, or is just the regulated service activity the point of constraint?

What is the objective here? Are we trying to protect the consumer, and redress asymmetric information in the marketplace by constraining the service providers and the extent to which customers can be disadvantaged? Or are we attempting to recognise the general benefit from a universally accessible and affordable communications service, and imposing this universal service obligation on the supply side of the industry through regulated obligations? In which case what is a “reasonable” service obligation? Does every household really need the ability to support multiple incoming 4K video streams as a basic communications service?

Communications policy in recent years has been dominated by legal argument taking precedent over economic analysis. However, we are becoming increasingly aware that the toolset for regulation of service behaviours on the Internet are necessarily intertwined with technical analysis.

The development of effective regulation requires professional skills in networking and technology, as distinct from the form of philosophical reasoning used in legal argument. If we are serious about the possibilities offered by “Open Government” then perhaps we need to be equally serious about generating and using “Open Data”. This is influenced by a traditional statistical view of the world, based on the notion of citizen science engaging in government and the development of frameworks based on the use of open data in developing federal policies.

A parallel consideration is data-driven policy making (or evidence based policymaking). Increasingly, it's becoming evident that we need to turn to data analytics in attempting to engage in a broader public role in policy determination. As a case in point, the FCC received some 22,157,658 filings in response to its latest call for comments relating to reconsideration of Net Neutrality provisions. It's no longer feasible to engage teams of readers of these filing equipped with marker pens, and then attempt to integrate the outcomes into some sense of a reasonable summary about the spectrum of received responses. What's being called for is a different form of engagement with the environment of providers and suppliers that can create object open data sets and subject them to analytic techniques that provide information in a repeatable and objective manner. Perhaps these days the FCC and their brethren regulators need more data analytic experts and fewer lawyers!

Products and Lifecycles

The product life cycle can be applied to the maturation of the Internet. We are all well aware of the classic "s-curve" of a product's life cycle, where we see successive waves of adoption from the innovators, then the early adopters seeking early adoption advantage, then the move by the majority, then the extended tail of the late adopters.

What does this mean for business strategies and hence what does this mean for regulators? Innovation is often seen as the lever of competitive access to a market, and there are a number of aspects of this competitive access that has a relationship to innovation in the product lifecycle.

Extensive competition is based around market expansion, and the innovative leverage is based around reducing the impediments to scaling the product. Intensive competition is based around largely static markets, where the intent of innovation nation is to increase the value proposition of the product to the consumer. What is critical in static markets is cost management and quality control. Often this is expressed in supply-side aggregation where providers seek further scale to achieve unit cost reductions.

Is there a theory of technology change? There are evolutionary tweaks that appear as a background constant, and superimposed over this are the occasional major ground-breaking changes. In any case, the impact of the ground-breaking change is irresistible. Not only do such changes sweep all before them, and leave only disruption in their wake, but they invent their own system of self-belief, and create a new orthodoxy of views. For example, the rise and rise of CDNs and their Cloud cousins have profound policy implications, and these implications transcend the technology and create new businesses and unlock new approaches to business. This impacts on the way entities interact, and there will most probably be beneficial opportunities that arise from these changes as well as the potential for increased cost and possibilities for distortions.

Upstream and Downstream Markets

For some time, the topic of two-sided markets has been part of the discourse of Internet Economics, and it looks particularly relevant when looking at the Internet of cloud and content platforms.

In terms of regulatory oversight, we've been looking at the downstream at the expense of looking at the upstream. In other words, the focus of attention has been on the downstream consumer of the service. But ISPs operate as intermediaries, with retail broadband consumers on their downstream, and content distributors, content creators and transit ISPs as their upstream providers. This upstream is a diverse set of relationships and an increasing adoption of an operational model of "winner take all".

The Internet ecosystem is breeding powerful platform operators who can capture large market share by exploiting scale economies, network externalities and high switching costs. These platform operators sit in the so-called "unicorn" sector, and the multi-billion dollar unicorn valuations show several "winner take all" industry actors, including, for example, Airbnb, Amazon, Facebook, Google, and Netflix.

Positive network externalities favour additional subscribers joining the established bandwagon. Massive subscriber populations generate big data that help the unicorns capture the lion's share of advertising revenue, making it possible to fine-tune their data analysis internally and fund acquisition of existing, or potential competitors.

The near-term consumer benefits that accrue from the presence of these unicorns are easy to identify. But the negative upstream impacts require attention. These include workforce dislocation, waning industries and similar social impacts.

The critical question is whether the upstream costs equates to the downstream benefit. Is this a zero-sum gain, or worse, a net cost rather than a net benefit? This raises the question of whether government vigilance is misplaced. Are our obsessions with network neutrality in downstream facing access networks extracting a cost in terms of oversight blindness of upstream costs?

Today's markets are more intertwined here, and regulators should pay heed to calls for more sophisticated data analysis with attention being brought to bear on both upstream relationships and the larger economy as well as downstream consumer impacts. The regulatory risk is not to recognise the critical importance of an open accessible platform, and the way in which incumbents are challenged without requiring a massive level of venture risk. The FCC is a decreasingly small part of the overall picture as more and more activity is not communications per se, but a one of the interactions in the provision of services and service platforms.

A concern here is that the market is so dominated by a small number of incumbents that the market serves the interest of the incumbents exclusively. The opportunity that we would like to see to counteract this concern is that we "normalise" these functions that provide access to the market for new providers and services to the extent that appropriate balances exist without explicit need for regulatory intervention or control. One characterisation of the current situation is that the popular unease with the Internet is not about the size or shape of the providers, but the unease that the market is being gummed up by too much asymmetry of information. If the established market actors simply know too much, then innovative incumbents are stymied by their own relative lack of knowledge, aside from all the more conventional considerations of access to capital, means of production and distribution and access to consumers!

Service Tensions: Content vs Carriage

There is a visible correlation between the rise and rise of the Content Distribution Network (CDN) platform sector, and the corresponding decline in demand for transit services in today's Internet.

Increasingly, the network is not being used to carry user's traffic to service portals, but the service portal has been shifted to sit next to the last mile access networks. However, the CDN providers prefer not to use public Internet infrastructure to support the networks that provision these portals, but are taking the approach of using dedicated communications capacity to build essentially private dedicated feeder systems.

These 'cloud' networked systems are no longer operated as part of the public communications portfolio, but are essentially privately-operated platforms beyond conventional public regulatory oversight, from the perspective of the public communications domain. We are seeing a major part of the original model of the Internet, namely that of the inter-network transit service, being superseded by private networks that operate as CDN feeder systems. Only the last mile access systems are left operating in the public communications realm. It is an open question as to how stable this arrangement will be, and unclear whether we will see further moves to pass parts of the Internet's common access infrastructure into the realm of these private service domains.

Oddly enough, the repeal of Network Neutrality provisions in the United States may well be a catalyst for such a move. The larger of the content providers might perceive that an alternative to paying the

access fees demanded by some of these last mile networks is the option of simply purchasing a controlling interest in the access network itself. In such a world, the CDN becomes the entire Internet.

Network Service Infrastructure

The picture in the Internet carriage provider network has been remarkably steady for over a decade now. Some ISPs are seeing a relatively consistent growth in peak traffic demands within the network growing by a constant 65% CAGR. Some of this is attributed to a growth in the customer base, and a growth in higher capacity access, but the bulk of this growth is content-based growth, where users are increasing their content demands by some 40% CAGR.

Much of the growth appears to come from fibre-connected households where the access capacity is not an impediment to content streaming. The peak delivery rates are now running at access line rate, which is a potential indicator of emerging bottlenecks even in the fibre-based access networks.

It has always been thought that if a network service managed to provision enough capacity in the network it could jump ahead of the demand growth. These figures tend to suggest that wherever that over-provisioned capacity threshold may be, it is probably not even at access speeds of 100Mbps. It tends to suggest that those who confidently assert that individual broadband access speeds of 10Mbps or 25Mbps access speeds are adequate are just out of touch with today's service usage patterns in the Internet. Even 100Mbps is a questionable lower bound for what constitutes "broadband" for today's consumers, and perhaps a service provider with foresight ought to be working on provisioning dedicated broadband access in units of Gbps to each customer in the coming decade.

Data Driven Operations

Much has been said in recent years about using empirical data as the basis of operational decisions. A similar program was taken up by the OECD in its advocacy of data-driven policies. But there's data and there's information, and the task is to determine what data is useful and why. What data would a network operator like to have to assist in the operation an effective and efficient service?

Network operators would like more data on network performance. We are not talking about basic network component behaviours, such as link availability, utilisation levels and packet drop rates. What we are talking about is the performance of applications that operate across the network. Netflix collect this form of data for many access networks, collecting data on streaming speeds, jitter, and loss rates. In other words, the data we are referring to is the view of the network as seen by a transport protocol operating across the network. To collect that data, it helps if you are actually operating one end to the end-to-end transport session, and network operators are not privy to such application level data, so the data is not readily available to network operators merely by virtue of operating the network. What would also help here is create levels of diagnostic data about the components of an end-to-end path. There is a widely held view that many of the issues with the quality of the data streaming experience is attributable to WiFi performance issues in the customer network, and it would be helpful to be able to have data on this to assist customers in understanding why they are experiencing performance issues.

There is a second class of data that is collected by networks, but for various reasons, some of which are as simple as "we've always regarded that as private data" that means that the data is not openly shared. For example, the traffic flows between a network and its neighbour networks. Indeed, one could generalise this and also push out data based on the flows to the destination network as well as the immediately adjacent neighbour network. The other form of data in this class relates to the pricing of various services. Open markets tend to operate more efficiently with open pricing, and most forms of wholesale and other inter-provider agreements include non-disclosure clauses that obscure the true nature of the agreement and the pricing of the wholesale offering. Such price occlusion admits the potential for discrimination and price gouging when market information is asymmetric.

And then there is data that would be useful, but we just don't have. To illustrate this, look at the current state of the DNS. Much of the world relies on the DNS, and much of the DNS relies on the

stable operation of a relatively small number of recursive resolvers and a relatively small cluster of authoritative name servers. One could characterise these servers as forming part of a set of systems that provide essential infrastructure. What is that set? What is their operational status? How fragile are they? The somewhat frustrating observation is that while the DNS is an essential component of the internet, and its efficient and seamless operation is crucial to the internet, we simply have no real data about how it is working. It's not the only application-level activity where operational data is scant, or even non-existent.

More generally, we are seeing a shift to more opaque applications that reveal less and less about how they operate. They are attempting to cloak this operational information from the host network, the host platform, and other applications. For example, the DNS is heading to use TLS, and encrypted network channel, so that DNS queries and responses are encrypted on the wire and cannot be seen, intercepted and potentially altered by the network provider. The predominance of the Chrome browser and its support of QUIC means that we are seeing many TCP sessions disappear behind encrypted UDP packet flows. Pleading for the restoration of an open application behaviour is most likely a plea that will fall on deaf ears today, as in the post-Snowden world we are seeing applications taking extensive measures to protect the online privacy of their users. No doubt if network operators were able to glean this data it may be possible to tune the network's response, but the time for applications and QoS network designers to work together on an open signalling protocol relating to QoS was more than twenty years ago. It's long gone! We just don't share data any more.

Spectrum Management

The radio spectrum that is useable for wireless communication encompasses a relatively broad span of frequencies from 3kHz to 300GHz. This is a common resource, and to limit interference we've devised a framework of assign the exclusive right of use of certain frequencies to certain uses and operators.

This spectrum allocation function has conventionally been undertaken by government administration, attempting to match the propagation properties of certain spectrum bands to the capabilities and service requirements of certain exploitative uses. For example, low bandwidth wide dispersion commercial AM radio operators were assigned bands in the medium wavelength band from approximately 530Khz to 1700Khz. Low frequencies have superior penetrative properties, but have limited bandwidth. Higher frequencies have far superior bandwidth capability, but have progressively limited propagation properties. The range of uses includes broadcasting, emergency communications, defence communications, satellite communications, radio astronomy, and unlicensed use.

As our technology changes, so do our demands for spectrum, and the allocation model was placed under escalating levels of stress as the technology of mobile communications services shifted from specialised use models into mass marketed consumer services. The broad allocations of spectrum to various uses continued, but with the bands assigned to commercial use we shifted to spectrum auctions, intended to improve the efficiency of spectrum use and boost the speed of innovation of use of radio.

The spectrum is fully assigned these days and any sector that wants to expand its use and be assigned more bands can only do so if another sector is willing to release its allocation, or it 'leaks' into unlicensed spectrum for peak periods (https://en.wikipedia.org/wiki/LTE_in_unlicensed_spectrum). Obviously, the mobile communications sector is pushing as hard as it can in the spectrum management agencies for further spectrum allocations, as mobile use continues to expand at unprecedented rates.

In the United States the FCC undertakes the role of spectrum management, and the workshop learned of the current models being used by the FCC to undertake its role. The Commission is now operating a model of "spectrum planning" in response to the observation most spectrum demands are being met by reallocation these days, as distinct from an earlier model of allocation of previously unused spectrum bands as new technologies made exploitation of ever higher frequencies commercially viable. This reallocation process is not a rapid one. Reallocation procedures have taken up to a decade and many

billions of dollars. The aim of the process is to clear a spectrum band of the earlier use and allow a new use to proceed without interference.

This reallocation is not just a response to the rapacious demands of the mobile industry, and their predictable push into cellular broadband data with the successive waves of deployment of higher capacity mobile services, but also a response to the emerging demands of a revitalised satellite industry. We have a new round of proposed MEOs, LEOs and HAPS all trying to co-exist within the same spectrum bands. Smaller spacecraft, and increasing launch capabilities, are fuelling increased level of interest in the use of satellites for certain broadband communications applications. Evidently the currently filed satellite plans encompass the launch of more than 13,000 spacecraft in the near future.

Another factor here is the rise of Software Defined Radio (SDR), which allows highly agile transmitters and receivers. We are seeing the rise of the intensity of use from individual users and applications, the rise in the population of users and services, the increasing diversity of use from terrestrial and vertical systems, and the overall objective is to make the spectrum work harder without imposing unsustainable costs on either the devices or escalating human risk factors.

The last large-scale release of spectrum was in the MHz VHF bands, with the shifting of television broadcasting into the UHF band. This took extraordinary levels of cost and protracted period of time to complete, and it's unlikely we will see a repeat of this form of large scale centrally orchestrated spectrum movement. More recently there have been "double auctions:" where there are bids from operators who are willing to take up vacated spectrum, and bids from incumbents to vacate licensed spectrum, and the double auction attempts to find an efficient reallocation strategy from the various "buy" and "sell" bids.

As reallocation costs escalate, the higher frequency part of the spectrum, which was previously considered to be too expensive to exploit at scale is being opened up. The 5Ghz-7Ghz spectrum is being reallocated for mobile use, and the US has been the first regime to allocate in the spectrum band at 28Ghz. They are also looking at experimental proposals for use in the area of 98Ghz and higher. Much of the current jostling is based around expectations of the rise and rise of the forthcoming 5G services. But how real is this expectation? These are very high frequencies, and they have limited propagation and almost no penetrative qualities. Is a basically line-of-sight service a viable solution for mobile data requirements? Or is the industry being fooled by its own hype here?

Is 5G Fake News?

Much of the public conversation in the United States these days is about "Fake News". There is no doubt that there is a large amount of suppositions and opinions masquerading as truths, but equally there is a large volume of basic truths these days where the temptation to discredit them by labelling them "fake" is overwhelming.

It was reported that the FCC estimates that pervasive 5G services will require an infrastructure investment of some \$275B. The obvious question is whether the mobile data services market is capable of generating revenue to match this investment. The current carrier behaviour suggests a different picture, where US mobile service providers are increasingly investing outside of their core competence to generate market growth. For example, T-Mobile has recently purchased Layer 3 TV (<https://bloom.bg/2zyGbFS>), Verizon has purchased the combination of Yahoo and AOL (called, curiously, "OATH"), and AT&T's efforts to acquire Time Warner continue through the courts. The once incredibly lucrative Average Revenue Per User (ARPU) levels are falling, and the only way mobile providers numbers can raise revenue levels is through extensive bundling of the product with content.

The broad-brush picture of this situation in terms of the capability of the provider side of the sector is pretty bleak as the sector appears to be in a deflationary mode. But what about consumers? Will consumers pay more for a very high speed 5G mobile service? The observation is that over recent years

US household income in all but the top 10% has been relatively flat, so there is limited household capacity to spend more money on Internet access services, 5G or otherwise.

The third part of this picture is the technology itself. It was reported that Google has used machine learning on street view data to find and classify poles as potential 5G base stations. This includes consideration of height, location, trees, and vegetation to generate a projected coverage map. They used a simulation of part of a major city with a 3.5Ghz spectrum service to deliver a 1Gbps 5G service. The study showed that 92 poles would be required to be capable of providing service to 93% of residences. However, because of the limited spectrum at these and lower frequencies there is consideration of using higher frequencies for 5G servers, and there is serious consideration of using millimetre wave services for a retail 5G service offering. How do these 3.8Ghz numbers stack up against an effort to provide the same 5G data service using a 28Ghz carrier? The same simulation, but with a 28Ghz service the same area, would require 1,400 poles, and provide an acceptable level of service to 82% of residences. That's a lot of poles, requiring a lot of power, and that implies that the \$275B estimate may well be on the low side. Maybe 5G really is an over-hyped case of fake news!

Service Infrastructure and FTTH

The Australian experience with high speed wired broadband was evidently somewhat unique. A combination of circumstances, which included a global financial crisis, a strong domestic economy at the time, a growing level of frustration with the incumbent network service providers over their collective determination to concentrate on mobile services to the exclusion of any further investment in the wired customer access network, prompted the government of the day to embark on an ambitious project to replace the copper access network, and some minor relics of a failed CATV deployment two decades earlier with a comprehensive program of Fibre To The Home. Another way of saying this was none of the operators were willing to underwrite such a project on their own, and the wired network was viewed as a derelict legacy of the telephone network that nobody wanted to fix. A digital slum in other words. The subsequent fortunes of this NBN project, including the whims of political favour and disfavour make a wish story all on their own, but here I'd like to contrast this to the experience with FTTH deployment in America.

A summary of the costs in the US are some \$700 to \$1,000 to pass each urban or suburban house on average. If the house elects to connect to the service another \$800 is spent on the tail line to complete the connection. All this has on a par with many other FTTH deployments so far. However, in the US it's an optional decision. Unlike the NBN in Australia where the deployment includes the recovery of the old copper telephone access network and the mandatory conversion of every home to this NBN service, in American context take up of high speed broadband is a consumer choice, and less than 40% of households are taking up the option. This means that each connected service has to cover not only the \$800 tie line plus the network costs of "passing" the house with FTTH, but, assuming that this service is deployed in a revenue neutral or even positive manner, the total revenue from all households that opt to take the service needs to cover the total network deployment costs, effectively including those costs of passing the "missed" houses that do not take up the service. This implies that the "miss rate" has a big impact on service costs. For example, with an average "pass" cost of \$850, and a "miss" rate of 40%, instead of costing some \$1,650 per household that would result from a 100% service take-up rate, each connected service costs some \$2,217 per household. That's a premium of 34% to cover the costs that result from partial take-up of the service. In many local markets, the alternatives to FTTH, namely DOCSIS, xDSL and even LTE services, look a whole lot cheaper, to both the network operator and to the consumer.

High speed broadband deployments in other countries have not experienced such a low take up rate for FTTH. It was reported that Sweden used an up-front payment to cover the initial installation, rather than an extended credit period, and in Norway self-installation of the tie line is a common situation. Other deployments can be highly selective, and the operator only embark on deployments in areas that have an assured target take up rate.

Broadband is highly valued by governments, and if the span and speed of national broadband services are a topic of much national interest at various inter-governmental fora. But consumers evidently do not value high speed broadband in the same enthusiastic light. The workshop heard that surveys of consumers reported that even if the service was completely free the take up rate would still not get much past 70% - 75%. It appears to be the case that where there is established CATV and telephone plant, and the cable infrastructure of coax and copper loops are sunk costs, FTTH is not an irresistibly attractive proposition to any unsubsidised operator.

So why do some network operators deploy FTTH despite these rather dismal numbers? Evidently, in at least one case, it's not an independent business decision, but is part of a larger deal. For example, an arrangement was brokered with the FCC over AT&T's 2015 acquisition of DirectTV (<http://soc.att.com/2k4O0hq>). With this acquisition, AT&T became the largest pay TV provider in the United States and the world, and as part of the FCC's approval of the transaction, AT&T undertook to offer its all-fibre Internet access service to at least 12.5 million customer locations, such as residences, home offices and very small businesses.

While governments are keen to see investment into broadband wireline access networks, the folk who are doing the work need either some significant incentives, or a rather large stick, to get them to play along with the plot!

My Impressions

After two days, I must admit to being a little more confused than before.

I feel that the old definitions of public communications, and the roles and obligations of common carriers in the public sphere were well understood, even if their articulation was often somewhat indirect in the relevant legislative provisions. Even if the various provisions in the US Telecommunications Act made somewhat curious and arbitrary distinctions, the overall objectives seemed to make sense in a world of telephones, television broadcasters, cable TV carriers, radios and newspapers.

But now I'm lost - are we trying to protect the consumer? Are we trying to protect, or even advance, some notional "common national good"? And if that's the case how can streaming multiple 4K video streams into each and every household contribute to a common national good? What are we trying to achieve here?

There are many potential roles for regulation in this space. There is the "emergency last resort" role when you are obliging actors to perform essential public services in the scenario of civil catastrophe or emergency when other service platforms have failed. There is the role of market regulator, attempting to curb the looming excesses of a nascent monopoly. Should the regulatory framework be capable of acting against cartels and other forms of market distortion? Should regulation safeguard some "universal service" obligation, assuming once more that you understand exactly what is meant by "universal" in this context and you have a clear grasp of the "service" being universalised here? Or is the role of regulation a matter of consumer protection by ensuring that providers stand by the offers that they make to consumers? Or is it preserving the accessibility of the environment for some eagerly anticipated, but otherwise undescribed, future disruptive entrant?

The common response, taken by many national regulators, is to take a bit of all of these potential roles, and provide a generous inclusion of political bias and populist appeal. If we want private sector investment in the infrastructure of a public communications platform then we cannot make the investment profile overly antagonistic to private actors. Indeed, if we see each national regulatory regime competing for the same pool of private investment funds there is a certain desire to make the investment profile as attractive as possible. However, a resulting public telecommunications service that is overly exploitative of consumers, or one that imposes unduly high costs on consumers, tends to contribute to an inefficient national economy, with its own attendant risks to national prosperity and

well-being. The common outcome is one where regulation attempts to achieve a little bit of all of these objectives, but in fact is wholeheartedly committed to none of them.

Even if one was able to be clear on exactly what the regulatory objectives may be, consistency in the regulatory role does not appear to be the case. Instead, we are seeing regulatory wavering to match the political climate of the day. and this may well deter both incumbents and challengers from further large scale high risk investment in politically volatile markets. The United States market itself appears to be largely saturated in terms of access to basic services, and the appetite for making further substantial investments in this area looks to be extremely low at present.

The workshop agenda and the material presented at the 2017 WIR workshop is available at <https://www.caida.org/workshops/wie/1712/>.

Author

Geoff Huston B.Sc., M.Sc., is the Chief Scientist at APNIC, the Regional Internet Registry serving the Asia Pacific region.

www.potaroo.net

Disclaimer

The above views do not necessarily represent the views or positions of the Asia Pacific Network Information Centre.